

Stochastic Optimization for Performative Prediction

Celestine Mendler-Dünner* Juan C. Perdomo* Tijana Zrnic* Moritz Hardt†
 {mendler, jcperdomo, tijana.zrnic, hardt}@berkeley.edu

University of California, Berkeley

Abstract

In performative prediction, the choice of a model influences the distribution of future data, typically through actions taken based on the model’s predictions.

We initiate the study of stochastic optimization for performative prediction. What sets this setting apart from traditional stochastic optimization is the difference between merely updating model parameters and deploying the new model. The latter triggers a shift in the distribution that affects future data, while the former keeps the distribution as is.

Assuming smoothness and strong convexity, we prove rates of convergence for both greedily deploying models after each stochastic update (greedy deploy) as well as for taking several updates before redeploying (lazy deploy). In both cases, our bounds smoothly recover the optimal $O(1/k)$ rate as the strength of performativity decreases. Furthermore, they illustrate how depending on the strength of performative effects, there exists a regime where either approach outperforms the other. We experimentally explore the trade-off on both synthetic data and a strategic classification simulator.

1 Introduction

Prediction in the social world is often *performative* in that a prediction triggers actions that influence the outcome. A forecast about the spread of a disease, for example, can lead to drastic public health action aimed at deterring the spread of the disease. In hindsight, the forecast might then appear to have been off, but this may largely be due to the actions taken based on it. Performativity arises naturally in consequential statistical decision-making problems in domains ranging from financial markets to online advertising.

Recent work [14] introduced and formalized *performative prediction*, an extension of the classical supervised learning setup whereby the choice of a model can change the data-generating distribution. This perspective leads to an important notion of stability requiring that a model is optimal on the distribution it entails. Stability prevents a certain cat-and-mouse game in which the learner repeatedly updates a model, because it no longer is accurate on the observed data. Prior work established conditions under which stability can be achieved through repeated risk minimization on the full data-generating distribution.

When samples arrive one-by-one over time, however, the learner faces a new challenge compared with traditional stochastic optimization. With every new sample that arrives, the learner has to decide whether to deploy the model, thereby triggering a drift in distribution, or

* Equal contribution. † MH is a paid consultant for Twitter.

to continue to collect more samples from the same distribution. Never deploying a model avoids distribution shift, but forgoes the possibility of converging to a stable point. Deploying the model too greedily could lead to overwhelming distribution shift that hampers convergence. In fact, it is not even clear that fast convergence to stability is possible at all in an online stochastic setting.

1.1 Our contributions

In this work, we initiate the study of stochastic optimization for performative prediction. Our main results are the first convergence guarantees for the stochastic gradient method in performative settings. Previous finite-sample guarantees [14] had an exponential dependence on the dimension of the data distribution.

We distinguish between two natural variants of the stochastic gradient method. One variant, called *greedy deploy*, updates model parameters and deploys the model at every step, after seeing a single example. The other, called *lazy deploy*, updates model parameters on multiple samples before deploying a model. We show that both methods converge to a stable solution. However, which one is preferable depends both on the cost of model deployment and the strength of performativity.

To state our results more precisely we recall the formal setup of performative prediction. In performative prediction, we assume that after deploying a model parameterized by θ , data are drawn from the distribution $\mathcal{D}(\theta)$. The distribution map $\mathcal{D}(\cdot)$ maps model parameters to data-generating distributions.

Given a loss function $\ell(z; \theta)$, a *performatively stable* model θ satisfies the fixed-point condition,

$$\theta \in \operatorname{argmin}_{\theta'} \mathbb{E}_{z \sim \mathcal{D}(\theta)} \ell(z; \theta').$$

Performative stability expresses the desideratum that the model θ minimizes loss on the distribution $\mathcal{D}(\theta)$ that it entails. Once we found a performatively stable model, we therefore have no reason to deviate from it based on the data that we observe.

The stochastic gradient method in this setting operates in a sequence of rounds. In each round k , the algorithm starts from a model θ_k and can choose to perform $n(k)$ stochastic gradient updates where each data point is drawn i.i.d. from the distribution $\mathcal{D}(\theta_k)$. After $n(k)$ stochastic gradient updates, the algorithm deploys the new model parameters θ_{k+1} . Henceforth, the data-generating distribution is $\mathcal{D}(\theta_{k+1})$ and the algorithm proceeds to the next round. For greedy deploy, $n(k) = 1$ for all k , whereas for lazy deploy $n(k)$ is a hyperparameter we can choose freely.

To analyze the stochastic gradient method, we import the same assumptions that were used in prior work on performative prediction. Apart from smoothness and strong convexity of the loss function, the main assumption is that the distribution map is sufficiently Lipschitz. This means that a small change to the model parameters (in Euclidean distance) leads to small change in the data-generating distribution (as measured in the Wasserstein metric).

Our first main result shows that under these assumptions, greedy deploy achieves the same convergence rate as the stochastic gradient method in the absence of performativity.

Theorem 1.1 (Greedy deploy, informal). *If the loss is smooth and strongly convex and the distribution map is sufficiently Lipschitz, greedy deploy converges to performative stability at rate $O(1/k)$, where k is the number of model deployment steps.*

Generally speaking, the Lipschitz parameter has to be smaller than the inverse condition number of the loss function for our bound to guarantee convergence. The exact rate stated in Theorem 3.2 further improves as the Lipschitz constant tends to 0.

In many realistic scenarios, data are plentiful, but deploying a model in a large production environment is costly. In such a scenario, it makes sense to aim to minimize the number of model deployment steps by updating the model parameters on multiple data points before initiating another model deployment. This is precisely what lazy deploy accomplishes as our next result shows.

Theorem 1.2 (Lazy deploy, informal). *Under the same assumptions as above, for any $\alpha > 0$, lazy deploy converges to performative stability at rate $O(1/k^\alpha)$ provided that $O(k^{1.1\alpha})$ samples are collected between deployments k and $k + 1$.*

In particular, this shows that any distance from optimality $\delta > 0$ can be achieved with $(1/\delta)^c$ model deployments for an arbitrarily small $c > 0$ at the cost of collecting polynomial in $1/\delta$ many samples.

Our main theorems provide upper bounds on the convergence rate of each method. As such they can only draw an incomplete picture about the relative performance of these methods. Our empirical investigation therefore aims to shed further light on their relative merits. In particular, our experiments show that greedy deploy generally performs better than lazy deploy when the distribution map has a small Lipschitz constant, i.e., the performative effects are small. Conversely, lazy deploy fares better when the distribution map is less Lipschitz. These observations are consistent with what our theoretical upper bounds suggest.

1.2 Related work

Perdomo et al. [14] introduced the performative prediction framework and analyzed algorithms for finding stable points that operate at the population level. While they also analyze some finite-sample extensions of these procedures, their analysis relies on concentration of the empirical distribution to the true distribution in the Wasserstein metric, and hence requires exponential sample complexity. In contrast, our analysis ensures convergence even if the learner collects a single sample at every step.

There has been a long line of work [3, 4, 5, 6, 10] within the learning theory community studying concept drift and learning from drifting distributions. Our results differ from these previous works since in performative prediction, changes in distribution are not a passive feature of the environment, but rather an active consequence of model deployment. This introduces several new considerations, such as the conceptual idea of performative stability, which is the main focus of our investigation.

Our work draws upon ideas from the stochastic convex optimization literature [7, 13, 15, 17, 18, 19]. Relative to these previous studies, our work analyzes the behavior of the stochastic gradient method in performative settings, where the underlying objective changes as a response to model deployment.

Lastly, we can view instances of performative prediction as special cases of reinforcement learning problems with nice structure, such as a Lipschitz mapping from policy parameters to the induced distribution over trajectories (see [14] for further discussion). The variants of the stochastic gradient method we consider can be viewed as policy gradient-like algorithms [1, 9, 20, 21] for this setting.

2 Preliminaries

We start by reviewing the core concepts of the framework of performative prediction. Afterwards, we set the stage for our analysis of stochastic algorithms by first considering gradient descent at the population level. In doing so, we highlight some of the fundamental limitations of gradient descent in performative settings.

2.1 The framework of performative prediction

Throughout our presentation, we focus on predictive models f_θ that are parametrized by a vector $\theta \in \Theta \subseteq \mathbb{R}^d$, where the parameter space Θ is a closed, convex set. The model or classifier, f_θ , maps instances $z \in \mathbb{R}^m$ to predictions $f_\theta(z)$. Typically, we think of z as being a feature, label pair (x, y) . We assess the quality of a classifier f_θ via a loss function $\ell(z; \theta)$.

The key theme in performative prediction is that the choice of deployed model f_θ influences the future data distribution and hence the expected loss of the classifier f_θ . This behavior is formalized via the notion of a *distribution map* $\mathcal{D}(\cdot)$, which is the key conceptual device of the framework. For every $\theta \in \Theta$, $\mathcal{D}(\theta)$ denotes the distribution over instances z induced by the deployment of f_θ . In this paper, we consider the setting where at each step, the learner observes a single sample $z \sim \mathcal{D}(\theta)$, where f_θ is the most recently deployed classifier. After having observed this sample, the learner chooses whether to deploy a new model or to leave the distribution as is before collecting the next sample.

We adopt the following Lipschitzness assumption on the distribution map. It captures the idea that if two models make similar predictions, then they also induce similar distributions.

Definition 2.1 (ε -sensitivity [14]). A distribution map $\mathcal{D}(\cdot)$ is ε -sensitive if for all $\theta, \theta' \in \Theta$:

$$W_1(\mathcal{D}(\theta), \mathcal{D}(\theta')) \leq \varepsilon \|\theta - \theta'\|_2,$$

where W_1 denotes the Wasserstein-1, or earth mover's distance.

The value of ε indicates the strength of performative effects; small ε means that the distribution induced by the model f_θ is not overly sensitive to the choice of θ , while large ε indicates high sensitivity. As an extreme case, $\varepsilon = 0$ implies $\mathcal{D}(\theta) = \mathcal{D}(\theta')$ for all $\theta, \theta' \in \Theta$ and hence there are no performative effects, as in classical supervised learning.

Given how the choice of a model induces a change in distribution, a naturally appealing property of a predictive model in performative settings is that it achieves minimal risk on the distribution that it induces. This solution concept is referred to as *performative stability*.

Definition 2.2 (Performative stability). A predictive model $f_{\theta_{\text{ps}}}$ is *performatively stable* if

$$\theta_{\text{ps}} \in \arg \min_{\theta} \mathbb{E}_{z \sim \mathcal{D}(\theta_{\text{ps}})} \ell(z; \theta).$$

We refer to θ_{ps} as being performatively stable, or simply stable, if $f_{\theta_{\text{ps}}}$ is performatively stable.

Performative stability captures an equilibrium notion in which a prediction induces a shift in distribution by the environment, yet remains simultaneously optimal for this new distribution. These solutions are referred to as stable since they eliminate the need for retraining. Besides eliminating the need for retraining, there are cases where performatively stable solutions also have good predictive power on the distribution they induce. More specifically, stable points can imply a small *performative risk*, $\text{PR}(\theta) \stackrel{\text{def}}{=} \mathbb{E}_{z \sim \mathcal{D}(\theta)} \ell(z; \theta)$, in the case of a strongly convex loss and a reasonably small sensitivity parameter ε (Theorem 4.3, [14]).

To illustrate these abstract concepts, we instantiate a simple traffic prediction example with performative effects which will serve as a running example throughout the paper.

Example 2.3 (ETA estimation). Suppose that each day we want to estimate the duration of a trip on a fixed route from the current weather conditions. Let $x \in \{0, 1\}$ denote a binary indicator of whether the current day is sunny or rainy, and suppose that $\Pr\{x = 1\} = p \in (0, 1)$. Let f_θ denote the deployed model which predicts trip duration y from x . Assume y behaves according to the following linear model:

$$y = \mu + w \cdot x - \varepsilon \cdot (f_\theta(x) - \mu),$$

where $\mu > 0$ denotes the usual time needed to complete the route on a sunny day, $w > 0$ denotes additional incurred time due to bad weather, and $-\varepsilon \cdot (f_\theta(x) - \mu)$ denotes the performative effects, for some $\varepsilon \in (0, 1)$. Namely, if the model predicts a faster than usual time to the destination, more people want to take the route, thus worsening the traffic conditions and making y large. If, on the other hand, the model predicts a longer trip, then few people follow the route and the resulting y is smaller. Suppose that the model class is all predictors of the form $f_\theta(x) = x\theta_1 + \theta_2$, where $\theta = (\theta_1, \theta_2)$ and $\theta_1 \in (0, w), \theta_2 \in (0, 2\mu)$. It is not hard to see that the distribution map corresponding to this data-generating process is ε -sensitive.

Assume that we measure predictive performance according to the squared loss, $\ell((x, y); \theta) = \frac{1}{2}(y - \theta_1 x - \theta_2)^2$. Then, a simple calculation reveals that the unique performatively stable solution, satisfying Definition 2.2, corresponds to

$$\theta_{\text{PS}} = \left(\frac{w}{1 + \varepsilon}, \mu \right).$$

In fact, one can show that θ_{PS} is simultaneously optimal in the sense that it minimizes the performative risk, $\theta_{\text{PS}} = \arg \min_{\theta} \text{PR}(\theta) = \arg \min_{\theta} \mathbb{E}_{(x, y) \sim \mathcal{D}(\theta)} \ell((x, y); \theta)$.

2.2 Population-level results

Before analyzing optimization algorithms in stochastic settings, we first consider their behavior at the population level. Throughout our analysis, we make the following assumptions on the loss $\ell(z; \theta)$, which hold for broad classes of objectives. To ease readability, we let $\mathcal{Z} = \cup_{\theta \in \Theta} \text{supp}(\mathcal{D}(\theta))$.

(A1) (*joint smoothness*) A loss function $\ell(z; \theta)$ is β -jointly smooth if the gradient¹ $\nabla \ell(z; \theta)$ is β -Lipschitz in θ and z , that is for all $\theta, \theta' \in \Theta$ and $z, z' \in \mathcal{Z}$ it holds that,

$$\|\nabla \ell(z; \theta) - \nabla \ell(z; \theta')\|_2 \leq \beta \|\theta - \theta'\|_2 \quad \text{and} \quad \|\nabla \ell(z; \theta) - \nabla \ell(z'; \theta)\|_2 \leq \beta \|z - z'\|_2.$$

(A2) (*strong convexity*) A loss function $\ell(z; \theta)$ is γ -strongly convex if for all $\theta, \theta' \in \Theta$ and $z \in \mathcal{Z}$ it holds that

$$\ell(z; \theta) \geq \ell(z; \theta') + \nabla \ell(z; \theta')^\top (\theta - \theta') + \frac{\gamma}{2} \|\theta - \theta'\|_2^2.$$

For $\gamma = 0$, this is equivalent to convexity.

We will refer to $\frac{\beta}{\gamma}$, where β is as in (A1) and γ as in (A2), as the condition number.

In this paper we are interested in the convergence of optimization methods to performative stability. However, unlike classical risk minimizers in supervised learning, it is not a priori clear that performatively stable solutions always exist. We thus recall the following fact regarding existence.

¹Gradients of the loss ℓ are always taken with respect to the parameters θ .

Fact 2.4 ([14]). Assume that the loss is β -jointly smooth (A1) and γ -strongly convex (A2). If $\mathcal{D}(\cdot)$ is ε -sensitive with $\varepsilon < \frac{\gamma}{\beta}$, then there exists a unique performatively stable point $\theta_{\text{PS}} \in \Theta$.

We note that it is not possible to reduce sensitivity by merely rescaling the problem, while keeping the ratio γ/β the same; the critical condition $\varepsilon\beta/\gamma < 1$ remains unaltered by scaling.²

The upper bound $\varepsilon < \gamma/\beta$ on the sensitivity parameter is not only crucial for the existence of unique stable points but also for algorithmic convergence. It defines a regime outside which gradient descent is not guaranteed to converge to stability even at the population level.

To be more precise, consider *repeated gradient descent* (RGD), defined recursively as

$$\theta_{k+1} = \theta_k - \eta_k \mathbb{E}_{z \sim \mathcal{D}(\theta_k)} [\nabla \ell(z; \theta_k)], \quad k \geq 1, \quad \text{where } \theta_1 \in \Theta \text{ is initialized arbitrarily.}$$

As shown in the following result, RGD need not converge to a stable point if $\varepsilon \geq \frac{\gamma}{\beta}$. Furthermore, a strongly convex loss is necessary to ensure convergence, even if performative effects are arbitrarily weak.

Proposition 2.5. Suppose that the distribution map $\mathcal{D}(\cdot)$ is ε -sensitive. Repeated gradient descent can fail to converge to a performatively stable point in any of the following cases, for any choice of positive step size sequence $\{\eta_k\}_{k \geq 1}$:

- (a) The loss is β -jointly smooth (A1) and convex, but not strongly convex (A2), for any $\beta, \varepsilon > 0$.
- (b) The loss is β -jointly smooth (A1) and γ -strongly convex (A2), but $\varepsilon \geq \frac{\gamma}{\beta}$, for any $\gamma, \beta, \varepsilon > 0$.

On the other hand, if $\varepsilon < \gamma/\beta$ we prove that RGD converges to a unique performatively stable point at a linear rate. Proposition 2.6 strengthens the corresponding result of Perdomo et al. [14], who showed linear convergence of RGD for $\varepsilon < \gamma/(\gamma + \beta)$. Proofs can be found in Appendix C.

Proposition 2.6. Assume that the loss is β -jointly smooth (A1) and γ -strongly convex (A2), and suppose that the distribution map $\mathcal{D}(\cdot)$ is ε -sensitive. Let $\varepsilon < \frac{\gamma}{\beta}$, and suppose that $\theta_{\text{PS}} \in \text{Int}(\Theta)$. Then, repeated gradient descent (RGD) with a constant step size $\eta_k = \eta \stackrel{\text{def}}{=} \frac{\gamma - \varepsilon\beta}{2(1 + \varepsilon^2)\beta^2}$ satisfies the following:

- (a) $\|\theta_{k+1} - \theta_{\text{PS}}\|_2 \leq \left(1 - \frac{\eta(\gamma - \varepsilon\beta)}{2}\right) \|\theta_k - \theta_{\text{PS}}\|_2$, where $0 < \frac{\eta(\gamma - \varepsilon\beta)}{2} < 1$.
- (b) The iterates θ_k of RGD converge to the stable point θ_{PS} at a linear rate, $\|\theta_{k+1} - \theta_{\text{PS}}\|_2 \leq \delta$ for $k \geq \frac{2}{\eta(\gamma - \varepsilon\beta)} \log\left(\frac{\|\theta_1 - \theta_{\text{PS}}\|_2}{\delta}\right)$.

Together, these results show that γ/β is a sharp threshold for the convergence of gradient descent in performative settings, thereby resolving an open problem in [14]. Having characterized the convergence regime of gradient descent, we now move on to presenting our main technical results, focusing on the case of a smooth, strongly convex loss with $\varepsilon < \gamma/\beta$.

²The reason is that the notion of *joint* smoothness we consider does not scale like strong convexity when rescaling θ . For example, rescaling $\theta \mapsto 2\theta$ (thus making $\varepsilon \mapsto \varepsilon/2$) would downscale the strong convexity parameter and the parameter corresponding to the usual notion of smoothness in optimization by a factor of 4, however the smoothness in z would downscale by a factor of 2.

3 Stochastic optimization results

We introduce two variants of the stochastic gradient method for optimization in performative settings (i.e. stochastic gradient descent, SGD), which we refer to as *greedy deploy* and *lazy deploy*. Each method performs a stochastic gradient update to the model parameters at every iteration, however they choose to deploy these updated models at different time intervals.

To analyze these methods, in addition to (A1) and (A2), we make the following assumption which is customary in the stochastic optimization literature [7, 16].

(A3) (*second moment bound*) There exist constants σ^2 and L^2 such that for all $\theta, \theta' \in \Theta$:

$$\mathbb{E}_{z \sim \mathcal{D}(\theta)} \left[\|\nabla \ell(z; \theta')\|_2^2 \right] \leq \sigma^2 + L^2 \|\theta' - G(\theta)\|_2^2, \text{ where } G(\theta) \stackrel{\text{def}}{=} \arg \min_{\theta'} \mathbb{E}_{z \sim \mathcal{D}(\theta)} \ell(z; \theta').$$

Given the operator $G(\cdot)$, performative stability can equivalently be expressed as $\theta_{\text{PS}} \in G(\theta_{\text{PS}})$.

3.1 Greedy deploy

A natural algorithm for stochastic optimization in performative prediction is a direct extension of the stochastic gradient method, whereby at every time step, we observe a sample $z^{(k)} \sim \mathcal{D}(\theta_k)$, compute a gradient update to the current model parameters θ_k , and deploy the new model θ_{k+1} (see left panel in Figure 1). We call this algorithm *greedy deploy*. In the context of our traffic prediction example, this greedy procedure corresponds to iteratively updating and redeploying the model based off information from the most recent trip.

While this procedure is algorithmically identical to the stochastic gradient method in traditional convex optimization, in performative prediction, the distribution of the observed samples depends on the trajectory of the algorithm. We begin by stating a technical lemma which introduces a recursion for the distance between θ_k and θ_{PS} .

Lemma 3.1. *Assume (A1), (A2) and (A3). If the distribution map $\mathcal{D}(\cdot)$ is ε -sensitive with $\varepsilon < \gamma/\beta$, then greedy deploy with step size η_k satisfies the following recursion for all $k \geq 1$:*

$$\mathbb{E} \left[\|\theta_{k+1} - \theta_{\text{PS}}\|_2^2 \right] \leq \left(1 - 2\eta_k(\gamma - \varepsilon\beta) + \eta_k^2 L^2 \left(1 + \varepsilon \frac{\beta}{\gamma} \right)^2 \right) \mathbb{E} \left[\|\theta_k - \theta_{\text{PS}}\|_2^2 \right] + \eta_k^2 \sigma^2.$$

Similar recursions underlie many proofs of SGD, and Lemma 3.1 can be seen as their generalization to the performative setting. Furthermore, we see how the bound implies a strong contraction to the performatively stable point if the performative effects are weak, that is when $\varepsilon \ll \gamma/\beta$.

Using this recursion, a simple induction argument suffices to prove that greedy deploy converges to the performatively stable solution (see Appendix D). Moreover, it does so at the usual $O(1/k)$ rate.

Theorem 3.2. *Assume (A1), (A2) and (A3). If the distribution map $\mathcal{D}(\cdot)$ is ε -sensitive with $\varepsilon < \frac{\gamma}{\beta}$, then for all $k \geq 0$ greedy deploy with step size $\eta_k = \left((\gamma - \varepsilon\beta)k + 8L^2/(\gamma - \varepsilon\beta) \right)^{-1}$ satisfies*

$$\mathbb{E} \left[\|\theta_{k+1} - \theta_{\text{PS}}\|_2^2 \right] \leq \frac{M_{\text{greedy}}}{(\gamma - \varepsilon\beta)^2 k + 8L^2},$$

where $M_{\text{greedy}} = \max \{ 2\sigma^2, 8L^2 \|\theta_1 - \theta_{\text{PS}}\|_2^2 \}$.

<u>Greedy Deploy</u>	<u>Lazy Deploy</u>
<p>Input: step size sequence $\{\eta_k\}_{k=1}^\infty$ Deploy initial classifier $\theta_1 \in \Theta$ For each $k = 1, 2, \dots$</p> <ul style="list-style-type: none"> - Observe $z^{(k)} \sim \mathcal{D}(\theta_k)$ - Update model parameters: $\theta_{k+1} = \theta_k - \eta_k \nabla \ell(z^{(k)}; \theta_k)$ - Deploy θ_{k+1} 	<p>Input: step size sequence $\{\eta_{k,j}\}_{k,j=1}^\infty$ Deploy initial classifier $\theta_1 \in \Theta$ For each $k = 1, 2, \dots$</p> <ul style="list-style-type: none"> - Set $\varphi_{k,1} = \theta_k$ - For each $j = 1, \dots, n(k)$: <ol style="list-style-type: none"> 1. Observe $z_j^{(k)} \sim \mathcal{D}(\theta_k)$ 2. Update model parameters: $\varphi_{k,j+1} = \varphi_{k,j} - \eta_{k,j} \nabla \ell(z_j^{(k)}; \varphi_{k,j})$ - Deploy $\theta_{k+1} = \varphi_{k,n(k)+1}$

Figure 1: Stochastic gradient method for performative prediction. Greedy deploy publishes the new model at every step while lazy deploy performs several gradient updates before releasing the new model.

Comparing this result to the traditional analysis of SGD for smooth, strongly convex objectives (e.g. [15]), we see that the traditional factor of γ is replaced by $\gamma - \varepsilon\beta$, which we view as the effective strong convexity parameter of the performative prediction problem. When $\varepsilon = 0$, there are no performative effects and the problem of finding the stable solution reduces to that of finding the risk minimizer on a fixed, static distribution. Consequently, it is natural for the two bounds to identify.

3.2 Lazy deploy

Contrary to greedy deploy, lazy deploy collects multiple data points and hence takes multiple stochastic gradient steps between consecutive model deployments. In the setting from Example 2.3, this corresponds to observing the traffic conditions across multiple days, and potentially diverse conditions, before deploying a new model.

This modification significantly changes the trajectory of lazy deploy relative to greedy deploy, given that the observed samples follow the distribution of the last *deployed* model, which might differ from the current iterate. More precisely, after deploying θ_k , we perform $n(k)$ stochastic gradient steps to the model parameters, using samples from $\mathcal{D}(\theta_k)$ before we deploy the last iterate as θ_{k+1} (see right panel in Figure 1).

At a high level, lazy deploy converges to performative stability because it progressively approximates *repeated risk minimization* (RRM), defined recursively as,

$$\theta_{k+1} = G(\theta_k) \stackrel{\text{def}}{=} \arg \min_{\theta' \in \Theta} \mathbb{E}_{z \sim \mathcal{D}(\theta_k)} \ell(z; \theta') \quad \text{for } k \geq 1 \quad \text{and } \theta_1 \in \Theta \text{ initialized arbitrarily.}$$

Perdomo et al. [14] show that RRM converges to a performatively stable classifier at a linear rate when $\varepsilon < \gamma/\beta$. Since the underlying distribution remains static between deployments, a classical analysis of SGD shows that for large $n(k)$ these “offline” iterates $\varphi_{k,j}$ converge to the risk minimizer on the distribution corresponding to the previously deployed classifier. In particular, for large $n(k)$, $\theta_{k+1} \approx G(\theta_k)$. By virtue of approximately tracing out the trajectory of RRM, lazy deploy converges to θ_{PS} as well. This sketch is formalized in the following theorem. For details we refer to Appendix E.

Theorem 3.3. Assume (A1), (A2), and (A3), and that the distribution map $\mathcal{D}(\cdot)$ is ε -sensitive with $\varepsilon < \frac{\gamma}{\beta}$. For any $\alpha > 0$, running lazy deploy with $n(k) \geq n_0 k^\alpha$, $k = 1, 2, \dots$ many steps between deployments and step size sequence $\eta_{k,j} = (\gamma j + 8L^2/\gamma)^{-1}$, satisfies

$$\mathbb{E}[\|\theta_{k+1} - \theta_{\text{PS}}\|_2^2] \leq c^k \cdot \|\theta_1 - \theta_{\text{PS}}\|_2^2 + \left(c^{\Omega(k)} + \frac{2}{k^{\alpha \cdot (1-o(1))}} \right) \cdot M_{\text{lazy}},$$

where $c = \left(\varepsilon \frac{\beta}{\gamma}\right)^2 + o(1)$ and $M_{\text{lazy}} = \frac{3(\sigma+\gamma)^2}{\gamma^2(1-c)}$. Here, $o(1)$ is independent of k and vanishes as n_0 grows; n_0 is chosen large enough such that $c < 1$.

3.3 Discussion

In this section, we have presented how varying the intervals at which we deploy models trained with stochastic gradient descent in performative settings leads to qualitatively different algorithms. While greedy deploy resembles classical SGD with a step size sequence adapted to the strength of distribution shift, lazy deploy can be viewed as a rough approximation of repeated risk minimization.

As we alluded to previously, the convergence behavior of both algorithms is critically affected by the strength of performative effects ε . For $\varepsilon \ll \gamma/\beta$, the effective strong convexity parameter $\gamma - \varepsilon\beta$ of the performative prediction problem is large. In this setting, the relevant distribution shift of deploying a new model is negligible and greedy deploy behaves almost like SGD in classical supervised learning, converging quickly to performative stability.

Conversely, for ε close to the convergence threshold, the contraction of greedy deploy to the performatively stable classifier is weak. In this regime, we expect lazy deploy to perform better since the convergence of the offline iterates $\varphi_{k,j}$ to the risk minimizer on the current distribution $G(\theta_k)$ is unaffected by the value of ε . Lazy deploy then converges by closely mimicking the behavior of RRM.

Furthermore, both algorithms differ in their sensitivity to different initializations. In greedy deploy, the initial distance $\|\theta_1 - \theta_{\text{PS}}\|_2^2$ decays polynomially, while in lazy deploy it decays at a linear rate. This suggests that the lazy deploy algorithm is more robust to poor initialization. While we derive these insights purely by inspecting our upper bounds, we find that these observations also hold empirically, as shown in the next section.

In terms of the asymptotics of both algorithms, we identify the following tradeoff between the number of samples and the number of deployments sufficient to converge to performative stability.

Corollary 3.4. Assume (A1), (A2), and (A3), and that $\mathcal{D}(\cdot)$ is ε -sensitive with $\varepsilon < \frac{\gamma}{\beta}$.

- To ensure that greedy deploy returns a solution θ^* such that, $\mathbb{E}[\|\theta^* - \theta_{\text{PS}}\|_2^2] \leq \delta$, it suffices to collect $\mathcal{O}(1/\delta)$ samples and to deploy $\mathcal{O}(1/\delta)$ classifiers.
- To achieve the same guarantee using lazy deploy, it suffices to collect $\mathcal{O}(1/\delta^{\frac{\alpha+1}{(1-\omega)\alpha}})$ samples and to deploy $\mathcal{O}(1/\delta^{\frac{1}{\alpha}})$ classifiers, for any $\alpha > 0$ and some $\omega = 1 - o(1)$ which tends to 1 as n_0 grows.

We see from the above result that by choosing large enough values of n_0 and α , we can make the sample complexity of the lazy deploy algorithm come arbitrarily close to that of greedy deploy. However, to match the same convergence guarantee, lazy deploy only performs $\mathcal{O}(1/\delta^{\frac{1}{\alpha}})$ deployments, which is significantly better than the $\mathcal{O}(1/\delta)$ deployments for greedy deploy.




Figure 2: Convergence of lazy and greedy deploy to performative stability for varying values of ε . We use $n(k) = k$ for lazy deploy. The results are for the synthetic Gaussian example with $\mu = 10$, $\sigma = 0.1$.

This reduction in the number of deployments is particularly relevant when considering the settings that performative prediction is meant to address. Whenever we use prediction in social settings, there are important social costs associated with making users adapt to a new model [12]. Furthermore, in industry, there are often significant technical challenges associated with deploying a new classifier [2]. By choosing $n(k) = n_0 k^\alpha$ appropriately, we can reduce the number of deployments necessary for lazy deploy to converge while at the same time improving the sample complexity of the algorithm.

4 Experiments

We complement our theoretical analysis of greedy and lazy deploy with a series of empirical evaluations³. First, we carry out experiments using synthetic data where we can analytically compute stable points and carefully evaluate the tradeoffs suggested by our theory. Second, we evaluate the performance of these procedures on a strategic classification simulator previously used as a benchmark for optimization in performative settings by [14].

4.1 Synthetic data

For our first experiment, we consider the task of estimating the mean of a Gaussian random variable under performative effects. In particular, we consider minimizing the expected squared loss $\ell(z; \theta) = \frac{1}{2}(z - \theta)^2$ where $z \sim \mathcal{D}(\theta) = \mathcal{N}(\mu + \varepsilon\theta, \sigma^2)$. For $\varepsilon > 0$, the true mean of a distribution $\mathcal{D}(\theta)$ depends on our revealed estimate θ . Furthermore, for $\varepsilon < \gamma/\beta = 1$, the problem has a unique stable point. A short algebraic manipulation shows that $\theta_{\text{PS}} = \frac{\mu}{1-\varepsilon}$. As per our theory, both greedy and lazy deploy converge to performative stability for all $\varepsilon < 1$.

Effect of performativity. We compare the convergence behavior of lazy deploy and greedy deploy for various values of ε in Figure 2. We choose step sizes for both algorithms according to our theorems in Section 3. In the case of lazy deploy, we set $\alpha = 1$, and hence $n(k) \propto k$.

We see that when performative effects are weak, i.e. $\varepsilon \ll \gamma/\beta$, greedy deploy outperforms lazy deploy. Lazy deploy in turn is better at coping with large distribution shifts from strong performative effects. These results confirm the conclusions from our theory and show that the

³Code is available at <https://github.com/zykls/performative-prediction>.




Figure 3: Convergence of lazy and greedy deploy to performative stability. Results are for the strategic classification experiments with $\varepsilon = 100$. (left panel) convergence as a function of the number of samples. (right panel) convergence as a function of the number of deployments.

choice of whether to delay deployments or not can indeed have a large impact on algorithm performance depending on the value of ε .

Deployment schedules. We also experiment with different deployment schedules $n(k)$ for lazy deploy. As described in Theorem 3.3, we can choose $n(k) \propto k^\alpha$ for all $\alpha > 0$. The results for $\alpha \in \{0.5, 1, 2\}$ and $\varepsilon \in \{.2, .6, .9\}$, are depicted and compared to the population-based RRM algorithm, in Figure 4 in the Appendix. We find that shorter deployment schedules, i.e., smaller α , lead to faster progress during initial stages of the optimization, whereas longer deployments schedules fare better in the long run while at the same time significantly reducing the number of deployments.

4.2 Strategic classification

In addition to the experiments on synthetic data, we also evaluate the performance of the two optimization procedures in a simulated strategic classification setting. Strategic classification is a two-player game between an institution which deploys a classifier f_θ and individual agents who manipulate their features in order to achieve a more favorable classification.

Perdomo et al. [14] introduce a credit scoring simulator in which a bank deploys a logistic regression classifier to determine the probability that an individual will default on a loan. Individuals correspond to feature, label pairs (x, y) drawn from a Kaggle credit scoring dataset [8]. Given the bank’s choice of a classifier f_θ , individuals solve an optimization problem to compute the best-response set of features, x_{BR} . This optimization procedure is parameterized by a value ε which determines the extent to which agents can change their respective features. The bank then observes the manipulated data points (x_{BR}, y) . This data-generating process can be described by a distribution map, which we can verify is ε -sensitive. For additional details we refer to Appendix A.

At each time step, the learner observes a single sample from the distribution in which the individual’s features have been manipulated in response to the most recently deployed classifier. This is in contrast to the experimental setup in [14], where the learner gets to observe the entire distribution of manipulated features at every step. While we cannot compute the stable point analytically in this setting, we can calculate it empirically by running RRM until convergence.

Results. The inverse condition number of this problem is much smaller than in the Gaussian example; we have $\gamma/\beta \approx 10^{-2}$. We first pick ε within the regime of provable convergence, i.e., $\varepsilon = 10^{-3}$, and compare the two methods. As expected, for such a small value of ε greedy deploy is the preferred method. Results are depicted in Figure 5 in the Appendix.

Furthermore, we explore the behavior of these algorithms outside the regime of provable convergence with $\varepsilon \gg \gamma/\beta$. We choose step sizes for both algorithms as defined in Section 3 with the exception that we ignore the ε -dependence in the step size schedule of greedy deploy and choose the same initial step size as for lazy deploy (Theorem 3.2). As illustrated in Figure 3 (left), lazy significantly outperforms greedy deploy in this setting. Moreover, the performance of lazy deploy significantly improves with α . In addition to speeding up convergence, choosing larger sample collection schedules $n(k)$ substantially reduces the number of deployments, as seen in Figure 3 (right).

Acknowledgements

We wish to acknowledge support from the U.S. National Science Foundation Graduate Research Fellowship Program and the Swiss National Science Foundation Early Postdoc Mobility Fellowship Program.

References

- [1] Alekh Agarwal, Sham M Kakade, Jason D Lee, and Gaurav Mahajan. Optimality and Approximation with Policy Gradient Methods in Markov Decision Processes. volume 125 of *Proceedings of Machine Learning Research*, pages 64–66. PMLR, 09–12 Jul 2020.
- [2] Algorithmia. 2020 State of Enterprise Machine Learning. 2020. https://info.algorithmia.com/hubfs/2019/Whitepapers/The-State-of-Enterprise-ML-2020/Algorithmia_2020_State_of_Enterprise_ML.pdf.
- [3] Peter L. Bartlett. Learning with a Slowly Changing Distribution. In *Proceedings of the Conference on Computational Learning Theory (COLT)*, pages 243–252, 1992.
- [4] Peter L. Bartlett, Shai Ben-David, and Sanjeev R. Kulkarni. Learning Changing Concepts by Exploiting the Structure of Change. *Machine Learning*, 41(2):153–174, 2000.
- [5] Rakesh D Barve and Philip M Long. On the Complexity of Learning from Drifting Distributions. *Information and Computation*, 138(2):170–193, 1997.
- [6] Omar Besbes, Yonatan Gur, and Assaf Zeevi. Non-Stationary Stochastic Optimization. *Operations Research*, 63(5):1227–1244, 2015.
- [7] Léon Bottou, Frank E Curtis, and Jorge Nocedal. Optimization Methods for Large-Scale Machine Learning. *SIAM Review*, 60(2):223–311, 2018.
- [8] Kaggle. Give Me Some Credit. <https://www.kaggle.com/c/GiveMeSomeCredit/data>, 2012.
- [9] Sham M Kakade. A Natural Policy Gradient. In *Advances in Neural Information Processing Systems*, pages 1531–1538, 2002.

- [10] Anthony Kuh, Thomas Petsche, and Ronald L Rivest. Learning Time-Varying Concepts. In *Advances in Neural Information Processing Systems (NIPS)*, pages 183–189, 1991.
- [11] John Miller, Chloe Hsu, Jordan Troutman, Juan Perdomo, Tijana Zrnic, Lydia Liu, Yu Sun, Ludwig Schmidt, and Moritz Hardt. WhyNot, 2020.
- [12] Smitha Milli, John Miller, Anca D. Dragan, and Moritz Hardt. The Social Cost of Strategic Classification. In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT*)*, page 230–239. Association for Computing Machinery, 2019.
- [13] Eric Moulines and Francis R Bach. Non-Asymptotic Analysis of Stochastic Approximation Algorithms for Machine Learning. In *Advances in Neural Information Processing Systems (NIPS)*, pages 451–459, 2011.
- [14] Juan C. Perdomo, Tijana Zrnic, Celestine Mendler-Dünner, and Moritz Hardt. Performative Prediction. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2020.
- [15] Alexander Rakhlin, Ohad Shamir, and Karthik Sridharan. Making Gradient Descent Optimal for Strongly Convex Stochastic Optimization. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 1571–1578, 2012.
- [16] Benjamin Recht and Stephen J. Wright. *Optimization for Modern Data Analysis*. 2019. Preprint available at <http://eecs.berkeley.edu/~brecht/opt4mlbook>.
- [17] Herbert Robbins and Sutton Monro. A Stochastic Approximation Method. *The Annals of Mathematical Statistics*, pages 400–407, 1951.
- [18] Tom Schaul, Sixin Zhang, and Yann LeCun. No More Pesky Learning Rates. In *Proceedings of the International Conference on Machine Learning (ICML)*, volume 28, pages 343–351, 2013.
- [19] Shai Shalev-Shwartz, Ohad Shamir, Nathan Srebro, and Karthik Sridharan. Stochastic Convex Optimization. In *Proceedings of the Conference on Computational Learning Theory (COLT)*, 2009.
- [20] Richard S Sutton, David A McAllester, Satinder P Singh, and Yishay Mansour. Policy Gradient Methods for Reinforcement Learning with Function Approximation. In *Advances in Neural Information Processing Systems*, pages 1057–1063, 2000.
- [21] Ronald J Williams. Simple Statistical Gradient-Following Algorithms for Connectionist Reinforcement Learning. *Machine learning*, 8(3-4):229–256, 1992.

A Additional evaluations and details on experimental setup




Figure 4: Convergence to performative stability of lazy deploy for the synthetic Gaussian example with $\mu = 10, \sigma = 0.1$. (top row) We show convergence of lazy deploy as a function of the number of samples collected for various values of ε . (bottom row) We plot convergence in the same setting, but now as a function of the number of deployments. For comparison we add greedy deploy (red) and RRM (dashed, gray line). The stars indicate the value attained at the end of our simulation (50k SGD updates).




Figure 5: Convergence of lazy and greedy deploy to performative stability. Results are for the strategic classification experiments with $\varepsilon = 0.001$. (left panel) convergence as a function of the number of samples collected. (center panel) convergence as a function of the number of deployments. (right panel) excess performative risk with respect the the stable classifier θ_{PS} as a function of stochastic gradient updates.

Input: base distribution \mathcal{D} , classifier f_θ , cost function c , and utility function u
Sampling procedure for $\mathcal{D}(\theta)$:

1. Sample $(x, y) \sim \mathcal{D}$
2. Compute best response $x_{\text{BR}} \leftarrow \arg \max_{x'} u(x', \theta) - c(x', x)$
3. Output sample (x_{BR}, y)

Figure 6: Distribution map for strategic classification (Perdomo et al. [14]).

Algorithm parameters. If not stated otherwise we use the following step size schedules proposed by our theory:

- greedy deploy (Theorem 3.2): $\eta_k = \frac{c_\eta}{k+k_0}$, where $c_\eta = \frac{1}{\gamma-\varepsilon\beta}$ and $k_0 = \frac{8L^2}{(\gamma-\varepsilon\beta)^2}$.
- lazy deploy (Theorem 3.3): $\eta_{k,j} = \frac{c_\eta}{j+k_0}$, where $c_\eta = \frac{1}{\gamma}$ and $k_0 = \frac{8L^2}{\gamma^2}$.

In Figure 3 we experiment with $\varepsilon = 100$ which is outside the regime of our convergence guarantees. Therefore we adapt the ε -dependence of the step size in greedy deploy. In particular, we pick $c_\eta = \frac{100}{\gamma}$ and $k_0 = \frac{8L^2}{\gamma^2}$ for both algorithms. The factor 100 was found empirically to reduce runtime. The deployment schedule $n(k)$ for lazy deploy is parameterized by α as $n(k) = n_0 k^\alpha$, where we choose $n_0 = 1$ for our experiments.

Confidence intervals. We repeat all our experiments 30 times and plot the mean μ_s and the shaded area $\mu_s \pm z \frac{s}{\sqrt{n}}$ where s denotes the standard deviation computed over the runs and $z = 1.645$. The value of z is chosen to ensure 90% coverage assuming Gaussian errors in the data.

A.1 Synthetic Gaussian experiments

The distribution map for the synthetic example is given by $\mathcal{D}(\theta) = \mathcal{N}(\mu + \theta\varepsilon, \sigma^2)$ where we use $\mu = 10$ and $\sigma = 0.1$ for our experiments. The SGD updates take the following form:

- greedy deploy: $\theta_{k+1} = \theta_k + \eta_k(z^{(k)} - \theta_k)$ where $z^{(k)} \sim \mathcal{D}(\theta_k)$.
- lazy deploy: $\varphi_{k,j+1} = \varphi_{k,j} + \eta_{k,j}(z_j^{(k)} - \varphi_{k,j})$, where $z_j^{(k)} \sim \mathcal{D}(\theta_k)$.

We initialize all optimization procedures at the risk minimizer $\theta_1 = \mu$ to mitigate effects of bad initialization and to instead focus on the effects of performativity.

A.2 Strategic classification simulator

For these experiments, we use the same experimental setup used by Perdomo et al. [14] as implemented in the WhyNot library [11]. We include all the relevant details for the sake of completeness.

The distribution map for this strategic classification example is described in Figure 6. The base distribution \mathcal{D} is a subsampled version of the Kaggle dataset [8] with $d = 10$ features and $n = 18357$ examples. Labels are binary variables $y \in \{0, 1\}$ and indicate whether an individual defaulted on a loan or not. We preprocess the data and normalize features to have zero mean and unit standard deviation. Out of the ten features, three are treated as strategic features. These are dimensions 1, 6, 8 corresponding to features such as the number of open credit lines.

The empirical distribution on these 18k points is considered to be the true distribution. To run our stochastic optimization experiments, we simply sample a single example from the data set according to the data-generating process described in Figure 6.

Individual utilities $u(\theta, x) = -\theta^\top x$ are linear and the costs are quadratic $c(x', x) = \frac{1}{2\varepsilon} \|x' - x\|^2$. Together, these lead to an ε -sensitive distribution map as shown in [14].

The loss of the institution is a logistic loss with ℓ_2 regularization:

$$\frac{1}{n} \sum_{i=1}^n \left[\log(1 + \exp(x_i^\top \theta)) - y_i x_i^\top \theta \right] + \frac{\lambda}{2} \|\theta\|^2$$

This loss is γ -strongly convex and $\beta = \max\left\{2, \frac{1}{4n} \sum_{i=1}^n \|x_i\|_2^2 + \gamma\right\}$ jointly smooth [14]. We fix $\lambda = 10^3/n$ for all experiments. When evaluated on the base distribution, the objective has parameters $\beta = 4.72$, $\gamma = 0.054$ which yields $\frac{\gamma}{\beta} = 0.011$.

B Technical lemmas

Lemma B.1 (Kantorovich-Rubinstein). *A distribution map $\mathcal{D}(\cdot)$ is ε -sensitive if and only if for all $\theta, \theta' \in \Theta$:*

$$\sup \left\{ \left| \mathbb{E}_{Z \sim \mathcal{D}(\theta)} g(Z) - \mathbb{E}_{Z \sim \mathcal{D}(\theta')} g(Z) \right| \leq \varepsilon \|\theta - \theta'\|_2 : g : \mathbb{R}^p \rightarrow \mathbb{R}, g \text{ 1-Lipschitz} \right\}.$$

Lemma B.2 (Lemma C.4 in [14]). *Let $f : \mathbb{R}^n \rightarrow \mathbb{R}^d$ be an L -Lipschitz function, and let $X, X' \in \mathbb{R}^n$ be random variables such that $W_1(X, X') \leq C$. Then*

$$\|\mathbb{E}[f(X)] - \mathbb{E}[f(X')]\|_2 \leq LC.$$

Lemma B.3 (First-order optimality condition). *Let f be convex and let Ω be a closed convex set on which f is differentiable, then*

$$x_* \in \arg \min_{x \in \Omega} f(x)$$

if and only if

$$\nabla f(x_*)^\top (y - x_*) \geq 0, \quad \forall y \in \Omega.$$

Lemma B.4 (Theorem 3.5 in [14]). *Suppose the loss function is γ -strongly convex (A2) and β -jointly smooth (A3). Then, for all $\theta, \theta' \in \Theta$, it holds that,*

$$\|G(\theta) - G(\theta')\|_2 \leq \varepsilon \frac{\gamma}{\beta} \|\theta - \theta'\|_2.$$

Lemma B.5. *Let $s \in (0, 1)$, and fix $\alpha > 0$, then,*

$$\sum_{k=1}^t k^{-\alpha} s^{t-k} \leq \frac{s^{t(1-2^{-1/\alpha})}}{1-s} + \frac{2t^{-\alpha}}{1-s}.$$

Proof. Denote by $a_k \stackrel{\text{def}}{=} k^{-\alpha}$. Let $M_t = \max\{m \in \mathbb{N} : a_m > 2a_t\}$. We decompose the sum depending on M_t as follows:

$$\sum_{k=1}^t a_k s^{t-k} = \sum_{k=1}^{M_t} a_k s^{t-k} + \sum_{k=M_t+1}^t a_k s^{t-k}.$$

We bound the first term trivially, by applying the fact that $a_k \leq 1$. For the second term, we use the fact that $a_k \leq 2a_t$ for $k > M_t$. We thus get:

$$\sum_{k=1}^t a_k s^{t-k} \leq \sum_{k=1}^{M_t} s^{t-k} + 2a_t \sum_{k=M_t+1}^t s^{t-k} \leq \frac{s^{t-M_t}}{1-s} + \frac{2a_t}{1-s}.$$

Since $a_k = k^{-\alpha}$, then $M_t \leq \frac{t}{2^{1/\alpha}}$, and so

$$\frac{s^{t-M_t}}{1-s} + \frac{2a_t}{1-s} \leq \frac{s^{t(1-2^{-1/\alpha})}}{1-s} + \frac{2a_t}{1-s}.$$

■

C Population-level results: proofs

C.1 Proof of Proposition 2.5

Let $\Theta = \mathbb{R}$, and let $z \sim \mathcal{D}(\theta)$ be a point mass at $1 + \varepsilon\theta$. This distribution map is clearly ε -sensitive. Furthermore, define the loss as,

$$\ell(z; \theta) = -\beta z \theta + \frac{\gamma}{2} \theta^2,$$

where $\beta \geq \gamma$ is an arbitrary positive scalar. Note that this objective is convex in θ and β -jointly smooth. Furthermore, it has a unique performatively stable point $\theta_{\text{PS}} = \frac{\beta/\gamma}{1-\varepsilon\beta/\gamma}$ whenever $\varepsilon \neq \frac{\gamma}{\beta}$; when $\varepsilon = \frac{\gamma}{\beta}$, there is no stable point. Repeated gradient descent has the dynamics:

$$\begin{aligned} \theta_{k+1} &= \theta_k - \eta_k \mathbb{E}_{z \sim \mathcal{D}(\theta_k)} \nabla \ell(z; \theta_k) \\ &= \theta_k - \eta_k (\gamma - \varepsilon\beta) \theta_k + \eta_k \beta \\ &= (1 - \eta_k (\gamma - \varepsilon\beta)) \theta_k + \eta_k \beta. \end{aligned}$$

If $\gamma = 0$, then the loss $\ell(z; \theta)$ is convex. Furthermore, for any values of $\varepsilon, \beta > 0$ and any positive step size sequence $\{\eta_k\}_{k=1}^{\infty}$, it holds that $1 + \eta_k \varepsilon \beta > 1$ meaning that RGD diverges.

To prove the second part of the statement, if $\gamma > 0$, then the loss is γ -strongly convex. Furthermore, if $\varepsilon > \gamma/\beta$, then for any step size sequence $\{\eta_k\}_{k=1}^{\infty}$, $1 - \eta_k (\gamma - \varepsilon\beta) > 1$ and RGD again diverges. When $\varepsilon = \frac{\gamma}{\beta}$, there is no stable solution and hence RGD does not converge to stability.

C.2 Proof of Proposition 2.6

This proof is essentially a consequence of Lemma 3.1. By following the steps of Lemma 3.1, we get

$$\begin{aligned} \|\theta_{k+1} - \theta_{\text{PS}}\|_2^2 &\leq \|\theta_k - \theta_{\text{PS}}\|_2^2 - 2\eta_k (\mathbb{E} \nabla \ell(z^{(k)}; \theta_k))^\top (\theta_k - \theta_{\text{PS}}) + \eta^2 \|\mathbb{E} \nabla \ell(z^{(k)}; \theta_k)\|_2^2 \\ &\stackrel{\text{def}}{=} B_1 - 2\eta B_2 + \eta^2 B_3. \end{aligned}$$

Following the same approach as in Lemma 3.1, we get

$$B_2 \geq (\gamma - \varepsilon\beta) \|\theta_k - \theta_{\text{PS}}\|_2^2.$$

The bound on B_3 is slightly different, as we no longer make assumptions on the second moment of the gradients; we use $z^{(\theta_{\text{PS}})}$ to denote a sample from $\mathcal{D}(\theta_{\text{PS}})$ and proceed as follows:

$$\begin{aligned}
\|\mathbb{E} \nabla \ell(z^{(k)}; \theta_k)\|_2^2 &= \|\mathbb{E} \nabla \ell(z^{(k)}; \theta_k) - \mathbb{E} \nabla \ell(z^{(\theta_{\text{PS}})}; \theta_{\text{PS}})\|_2^2 \\
&\leq \|\mathbb{E} \nabla \ell(z^{(k)}; \theta_k) - \mathbb{E} \nabla \ell(z^{(k)}; \theta_{\text{PS}}) + \mathbb{E} \nabla \ell(z^{(k)}; \theta_{\text{PS}}) - \mathbb{E} \nabla \ell(z^{(\theta_{\text{PS}})}; \theta_{\text{PS}})\|_2^2 \\
&\leq 2\|\mathbb{E} \nabla \ell(z^{(k)}; \theta_k) - \mathbb{E} \nabla \ell(z^{(k)}; \theta_{\text{PS}})\|_2^2 \\
&\quad + 2\|\mathbb{E} \nabla \ell(z^{(k)}; \theta_{\text{PS}}) - \mathbb{E} \nabla \ell(z^{(\theta_{\text{PS}})}; \theta_{\text{PS}})\|_2^2 \\
&\leq 2\beta^2 \|\theta_k - \theta_{\text{PS}}\|_2^2 + 2\beta^2 \varepsilon^2 \|\theta_k - \theta_{\text{PS}}\|_2^2 \\
&\leq 2\beta^2 (1 + \varepsilon^2) \|\theta_k - \theta_{\text{PS}}\|_2^2,
\end{aligned}$$

where in the third inequality we apply the fact that the loss is β -jointly smooth, together with Lemma B.2. Putting everything together, this implies

$$\|\theta_{k+1} - \theta_{\text{PS}}\|_2^2 \leq (1 - 2\eta(\gamma - \varepsilon\beta) + 2\eta^2\beta^2(1 + \varepsilon^2)) \|\theta_k - \theta_{\text{PS}}\|_2^2.$$

Using the fact that $\sqrt{1-x} \leq 1 - \frac{x}{2}$ for $x \in [0, 1]$, we get

$$\|\theta_{k+1} - \theta_{\text{PS}}\|_2 \leq (1 - \eta(\gamma - \varepsilon\beta) + \eta^2\beta^2(1 + \varepsilon^2)) \|\theta_k - \theta_{\text{PS}}\|_2.$$

By setting $\eta = \frac{\gamma - \varepsilon\beta}{2(1 + \varepsilon^2)\beta^2}$, we can conclude

$$\|\theta_{k+1} - \theta_{\text{PS}}\|_2 \leq \left(1 - \frac{(\gamma - \varepsilon\beta)^2}{4(1 + \varepsilon^2)\beta^2}\right) \|\theta_k - \theta_{\text{PS}}\|_2.$$

Note that $\frac{(\gamma - \varepsilon\beta)^2}{4(1 + \varepsilon^2)\beta^2} < 1$ because $(\gamma - \varepsilon\beta)^2 \leq \gamma^2 + \varepsilon^2\beta^2 \leq (1 + \varepsilon^2)\beta^2$.

We can unroll the above recursion to get

$$\begin{aligned}
\|\theta_{k+1} - \theta_{\text{PS}}\|_2 &\leq \left(1 - \frac{(\gamma - \varepsilon\beta)^2}{4(1 + \varepsilon^2)\beta^2}\right)^k \|\theta_1 - \theta_{\text{PS}}\|_2 \\
&\leq \exp\left(-\frac{k(\gamma - \varepsilon\beta)^2}{4(1 + \varepsilon^2)\beta^2}\right) \|\theta_1 - \theta_{\text{PS}}\|_2.
\end{aligned}$$

Setting the right-hand side to δ and expressing k completes the proof.

D Greedy deploy: proofs

D.1 Proof of Lemma 3.1

Throughout the proof, we will use $z^{(\theta_{\text{PS}})}$ to denote a sample from $\mathcal{D}(\theta_{\text{PS}})$ which is independent from the whole trajectory of greedy deploy (e.g. $\{\theta_j, z^{(j)}\}_j$, etc.).

Since Θ is closed and convex, we know

$$\|\theta_{k+1} - \theta_{\text{PS}}\|_2^2 = \|\Pi_{\Theta}(\theta_k - \eta_k \nabla \ell(z^{(k)}; \theta_k)) - \theta_{\text{PS}}\|_2^2 \leq \|\theta_k - \eta_k \nabla \ell(z^{(k)}; \theta_k) - \theta_{\text{PS}}\|_2^2.$$

Squaring the right-hand side and expanding out the square,

$$\begin{aligned}
& \mathbb{E} \left[\|\theta_k - \eta_k \nabla \ell(z^{(k)}; \theta_k) - \theta_{\text{PS}}\|_2^2 \right] \\
&= \mathbb{E} \left[\|\theta_k - \theta_{\text{PS}}\|_2^2 \right] - 2\eta_k \mathbb{E} \left[\nabla \ell(z^{(k)}; \theta_k)^\top (\theta_k - \theta_{\text{PS}}) \right] + \eta_k^2 \mathbb{E} \left[\|\nabla \ell(z^{(k)}; \theta_k)\|_2^2 \right] \\
&\stackrel{\text{def}}{=} B_1 - 2\eta_k B_2 + \eta_k^2 B_3.
\end{aligned}$$

We begin by lower bounding B_2 . Since θ_{PS} is optimal for the distribution it induces, by Lemma B.3 we have $\mathbb{E} \left[\nabla \ell(z^{(\theta_{\text{PS}})}; \theta_{\text{PS}})^\top (\theta_k - \theta_{\text{PS}}) \right] \geq 0$. This allows us to bound B_2 as:

$$\begin{aligned}
B_2 &\geq \mathbb{E} \left[(\nabla \ell(z^{(k)}; \theta_k) - \nabla \ell(z^{(\theta_{\text{PS}})}; \theta_k) + \nabla \ell(z^{(\theta_{\text{PS}})}; \theta_k) - \nabla \ell(z^{(\theta_{\text{PS}})}; \theta_{\text{PS}}))^\top (\theta_k - \theta_{\text{PS}}) \right] \\
&= \mathbb{E} \left[(\nabla \ell(z^{(k)}; \theta_k) - \nabla \ell(z^{(\theta_{\text{PS}})}; \theta_k))^\top (\theta_k - \theta_{\text{PS}}) \right] \\
&\quad + \mathbb{E} \left[(\nabla \ell(z^{(\theta_{\text{PS}})}; \theta_k) - \nabla \ell(z^{(\theta_{\text{PS}})}; \theta_{\text{PS}}))^\top (\theta_k - \theta_{\text{PS}}) \right].
\end{aligned}$$

For the first term, we have that

$$\begin{aligned}
& \mathbb{E} \left[(\nabla \ell(z^{(k)}; \theta_k) - \nabla \ell(z^{(\theta_{\text{PS}})}; \theta_k))^\top (\theta_k - \theta_{\text{PS}}) \right] \\
&= \mathbb{E} \left[\mathbb{E} \left[(\nabla \ell(z^{(k)}; \theta_k) - \nabla \ell(z^{(\theta_{\text{PS}})}; \theta_k))^\top (\theta_k - \theta_{\text{PS}}) \mid \theta_k \right] \right] \\
&\geq -\varepsilon \beta \mathbb{E} \left[\|\theta_k - \theta_{\text{PS}}\|_2^2 \right].
\end{aligned}$$

Having applied the law of iterated expectation, the above inequality follows from the fact that, conditional on θ_k , the function $\nabla \ell(z; \theta_k)^\top (\theta_k - \theta_{\text{PS}})$ is $\beta \|\theta_k - \theta_{\text{PS}}\|_2$ -Lipschitz in z . To verify this claim, we can apply the Cauchy-Schwarz inequality followed by the fact that the gradient is β -jointly smooth. Then, we apply Lemma B.1 and the fact that $\mathcal{D}(\cdot)$ is ε -sensitive to get the final bound.

Now, we use strong convexity to bound the second term,

$$\begin{aligned}
& \mathbb{E} \left[(\nabla \ell(z^{(\theta_{\text{PS}})}; \theta_k) - \nabla \ell(z^{(\theta_{\text{PS}})}; \theta_{\text{PS}}))^\top (\theta_k - \theta_{\text{PS}}) \right] \\
&= \mathbb{E} \left[\mathbb{E} \left[(\nabla \ell(z^{(\theta_{\text{PS}})}; \theta_k) - \nabla \ell(z^{(\theta_{\text{PS}})}; \theta_{\text{PS}}))^\top (\theta_k - \theta_{\text{PS}}) \mid \theta_k \right] \right] \\
&\geq \gamma \mathbb{E} \left[\|\theta_k - \theta_{\text{PS}}\|_2^2 \right].
\end{aligned}$$

Therefore, we get that

$$B_2 \geq (\gamma - \varepsilon \beta) \mathbb{E} \left[\|\theta_k - \theta_{\text{PS}}\|_2^2 \right].$$

Now we move on to bounding B_3 . Using our assumption on the variance on the gradients yields the following bound, we get

$$\begin{aligned}
\mathbb{E} \left[\|\nabla \ell(z^{(k)}; \theta_k)\|_2^2 \right] &\leq \sigma^2 + L^2 \mathbb{E} \left[\|\theta_k - G(\theta_k)\|_2^2 \right] \\
&= \sigma^2 + L^2 \mathbb{E} \left[\|\theta_k - \theta_{\text{PS}} + \theta_{\text{PS}} - G(\theta_k)\|_2^2 \right] \\
&\leq \sigma^2 + L^2 \left(\mathbb{E} \left[(\|\theta_k - \theta_{\text{PS}}\|_2 + \|\theta_{\text{PS}} - G(\theta_k)\|_2)^2 \right] \right) \\
&\leq \sigma^2 + L^2 \left(1 + \varepsilon \frac{\beta}{\gamma} \right)^2 \mathbb{E} \left[\|\theta_k - \theta_{\text{PS}}\|_2^2 \right],
\end{aligned}$$

where in the last step we use Lemma B.4, which implies $\|\theta_{\text{PS}} - G(\theta_k)\|_2 \leq \varepsilon \frac{\beta}{\gamma} \|\theta_k - \theta_{\text{PS}}\|_2$.

Putting all the steps together completes the proof.

D.2 Proof of Theorem 3.2

From Lemma 3.1, we have that the following recursion holds:

$$\mathbb{E}[\|\theta_{k+1} - \theta_{\text{PS}}\|_2^2] \leq \left(1 - 2\eta_k(\gamma - \varepsilon\beta) + \eta_k^2 L^2 \left(1 + \varepsilon \frac{\beta}{\gamma}\right)^2\right) \mathbb{E}[\|\theta_k - \theta_{\text{PS}}\|_2^2] + \eta_k^2 \sigma^2.$$

Using the fact that $\varepsilon < \frac{\gamma}{\beta}$, we get that,

$$\mathbb{E}[\|\theta_{k+1} - \theta_{\text{PS}}\|_2^2] \leq \left(1 - 2\eta_k(\gamma - \varepsilon\beta) + 4\eta_k^2 L^2\right) \mathbb{E}[\|\theta_k - \theta_{\text{PS}}\|_2^2] + \eta_k^2 \sigma^2.$$

We proceed by using induction. As in the theorem statement, we let $\eta_k = \frac{1}{(\gamma - \varepsilon\beta)(k + k_0)}$, where we denote $k_0 = \frac{8L^2}{(\gamma - \varepsilon\beta)^2}$. The base case, $k = 0$, is trivially true by construction of the bound and choice of k_0 . Now, we adopt the inductive hypothesis that

$$\mathbb{E}[\|\theta_{k+1} - \theta_{\text{PS}}\|_2^2] \leq \frac{\max\{2\sigma^2, 8L^2\|\theta_1 - \theta_{\text{PS}}\|_2^2\}}{(\gamma - \varepsilon\beta)^2(k + k_0)}.$$

Then, by Lemma 3.1, it is true that

$$\begin{aligned} \mathbb{E}[\|\theta_{k+2} - \theta_{\text{PS}}\|_2^2] &\leq \left(1 - 2\eta_k(\gamma - \varepsilon\beta) + 4\eta_k^2 L^2\right) \mathbb{E}[\|\theta_{k+1} - \theta_{\text{PS}}\|_2^2] + \eta_k^2 \sigma^2 \\ &\leq \frac{1}{(\gamma - \varepsilon\beta)^2} \left(\frac{k + k_0 - 2 + \frac{4L^2}{(\gamma - \varepsilon\beta)^2 k_0}}{(k + k_0)^2} \max\{2\sigma^2, 8L^2\|\theta_1 - \theta_{\text{PS}}\|_2^2\} + \frac{\sigma^2}{(k + k_0)^2} \right) \\ &\leq \frac{1}{(\gamma - \varepsilon\beta)^2} \left(\frac{k + k_0 - 1.5}{(k + k_0)^2} \max\{2\sigma^2, 8L^2\|\theta_1 - \theta_{\text{PS}}\|_2^2\} + \frac{\sigma^2}{(k + k_0)^2} \right) \\ &\leq \frac{1}{(\gamma - \varepsilon\beta)^2} \left(\frac{k + k_0 - 1}{(k + k_0)^2} \max\{2\sigma^2, 8L^2\|\theta_1 - \theta_{\text{PS}}\|_2^2\} - \frac{0.5 \cdot 2\sigma^2 - \sigma^2}{(k + k_0)^2} \right) \\ &= \frac{1}{(\gamma - \varepsilon\beta)^2} \cdot \frac{k + k_0 - 1}{(k + k_0)^2} \max\{2\sigma^2, 8L^2\|\theta_1 - \theta_{\text{PS}}\|_2^2\} \\ &\leq \frac{1}{(\gamma - \varepsilon\beta)^2} \cdot \frac{1}{k + 1 + k_0} \max\{2\sigma^2, 8L^2\|\theta_1 - \theta_{\text{PS}}\|_2^2\}, \end{aligned}$$

where the last step follows because $(k + k_0)^2 > (k + k_0)^2 - 1 = (k + k_0 + 1)(k + k_0 - 1)$. Therefore, we have shown $\mathbb{E}[\|\theta_{k+2} - \theta_{\text{PS}}\|_2^2] \leq \frac{M_{\text{greedy}}}{(\gamma - \varepsilon\beta)^2(k + 1 + k_0)}$, which completes the proof by induction.

E Lazy deploy: proofs

To prove Theorem 3.3, we use the following classical result about convergence of SGD on a static distribution (see, e.g., [15]). The step size is chosen such that it matches the step size of Theorem 3.2 when $\varepsilon = 0$. We include the proof for completeness.

Lemma E.1. *Under assumptions (A1), (A2), and (A3), lazy deploy satisfies the following:*

$$\mathbb{E}[\|\varphi_{k,j+1} - G(\theta_k)\|_2^2] \leq \left(1 - 2\eta_{k,j}\gamma + \eta_{k,j}^2 L^2\right) \mathbb{E}[\|\varphi_{k,j} - G(\theta_k)\|_2^2] + \eta_{k,j}^2 \sigma^2.$$

If, additionally, $\eta_{k,j} = \frac{1}{\gamma j + 8L^2/\gamma}$, then for all $k \geq 1, j \geq 0$, the following is true

$$\mathbb{E} \left[\|\varphi_{k,j+1} - G(\theta_k)\|_2^2 \right] \leq \frac{M_{\text{lazy}}}{\gamma^2 j + L^2},$$

where $M_{\text{lazy}} \stackrel{\text{def}}{=} \max \left\{ 1.2\sigma^2, 8L^2 \mathbb{E}[\|\theta_k - G(\theta_k)\|_2^2] \right\}$.

Proof. First we prove the recursion. Since Θ is closed and convex, we know

$$\begin{aligned} & \mathbb{E} \left[\|\varphi_{k,j+1} - G(\theta_k)\|_2^2 \right] \\ &= \mathbb{E} \left[\left\| \Pi_{\Theta} \left(\varphi_{k,j} - \eta_{k,j} \nabla \ell(z_j^{(k)}; \varphi_{k,j}) \right) - G(\theta_k) \right\|_2^2 \right] \\ &\leq \mathbb{E} \left[\left\| \varphi_{k,j} - \eta_{k,j} \nabla \ell(z_j^{(k)}; \varphi_{k,j}) - G(\theta_k) \right\|_2^2 \right] \\ &= \mathbb{E} \left[\left\| \varphi_{k,j} - G(\theta_k) \right\|_2^2 \right] - 2\eta_{k,j} \mathbb{E} \left[\nabla \ell(z_j^{(k)}; \varphi_{k,j})^\top (\varphi_{k,j} - G(\theta_k)) \right] + \eta_{k,j}^2 \mathbb{E} \left[\left\| \nabla \ell(z_j^{(k)}; \varphi_{k,j}) \right\|_2^2 \right]. \end{aligned}$$

Next, we examine the cross-term. By the first-order optimality conditions for convex functions (Lemma B.3), we know that $\mathbb{E} \left[\nabla \ell(z_j^{(k)}; G(\theta_k))^\top (\varphi_{k,j} - G(\theta_k)) \right] \geq 0$. Using this lemma along with strong convexity, we can lower bound this term as follows,

$$\begin{aligned} \mathbb{E} \left[\nabla \ell(z_j^{(k)}; \varphi_{k,j})^\top (\varphi_{k,j} - G(\theta_k)) \right] &\geq \mathbb{E} \left[(\nabla \ell(z_j^{(k)}; \varphi_{k,j}) - \nabla \ell(z_j^{(k)}; G(\theta_k)))^\top (\varphi_{k,j} - G(\theta_k)) \right] \\ &\geq \gamma \mathbb{E} \left[\|\varphi_{k,j} - G(\theta_k)\|_2^2 \right]. \end{aligned}$$

For the final term, we use our assumption on the second moment of the gradients,

$$\mathbb{E} \left[\left\| \nabla \ell(z_j^{(k)}; \varphi_{k,j}) \right\|_2^2 \right] \leq \sigma^2 + L^2 \mathbb{E} \left[\|\varphi_{k,j} - G(\theta_k)\|_2^2 \right].$$

Putting everything together, we get the desired recursion,

$$\mathbb{E} \left[\|\varphi_{k,j+1} - G(\theta_k)\|_2^2 \right] \leq (1 - 2\eta_{k,j}\gamma + \eta_{k,j}^2 L^2) \mathbb{E} \left[\|\varphi_{k,j} - G(\theta_k)\|_2^2 \right] + \eta_{k,j}^2 \sigma^2.$$

Now we turn to proving the second part of the lemma. Similarly to Theorem 3.2, we prove the result using induction. As in the theorem statement, we let $\eta_{k,j} = \frac{1}{\gamma(j+k_0)}$, where we denote $k_0 = \frac{8L^2}{\gamma^2}$. The base case, $j = 0$, is trivially true by construction of the bound and choice of k_0 . Now, we adopt the inductive hypothesis that

$$\mathbb{E} \left[\|\varphi_{k,j+1} - G(\theta_k)\|_2^2 \right] \leq \frac{\max \left\{ 1.2\sigma^2, 8L^2 \mathbb{E} \left[\|\theta_k - G(\theta_k)\|_2^2 \right] \right\}}{\gamma^2(j+k_0)}.$$

Then, by part (a) of this lemma, it is true that

$$\begin{aligned}
\mathbb{E}\left[\|\varphi_{k,j+2} - G(\theta_k)\|_2^2\right] &\leq \left(1 - 2\eta_{k,j}\gamma + \eta_{k,j}^2 L^2\right) \mathbb{E}\left[\|\varphi_{k,j+1} - G(\theta_k)\|_2^2\right] + \eta_{k,j}^2 \sigma^2 \\
&\leq \frac{1}{\gamma^2} \left(\frac{j+k_0-2 + \frac{L^2}{\gamma^2 k_0}}{(j+k_0)^2} \max\{1.2\sigma^2, 8L^2 \mathbb{E}\left[\|\theta_k - G(\theta_k)\|_2^2\right]\} + \frac{\sigma^2}{(j+k_0)^2} \right) \\
&\leq \frac{1}{\gamma^2} \left(\frac{j+k_0-15/8}{(j+k_0)^2} \max\{1.2\sigma^2, 8L^2 \mathbb{E}\left[\|\theta_k - G(\theta_k)\|_2^2\right]\} + \frac{\sigma^2}{(j+k_0)^2} \right) \\
&\leq \frac{1}{\gamma^2} \left(\frac{j+k_0-1}{(j+k_0)^2} \max\{1.2\sigma^2, 8L^2 \mathbb{E}\left[\|\theta_k - G(\theta_k)\|_2^2\right]\} - \frac{7/8 \cdot 1.2\sigma^2 + \sigma^2}{(j+k_0)^2} \right) \\
&= \frac{1}{\gamma^2} \cdot \frac{j+k_0-1}{(j+k_0)^2} \max\{1.2\sigma^2, 8L^2 \mathbb{E}\left[\|\theta_k - G(\theta_k)\|_2^2\right]\} \\
&\leq \frac{1}{\gamma^2} \cdot \frac{1}{j+1+k_0} \max\{1.2\sigma^2, 8L^2 \mathbb{E}\left[\|\theta_k - G(\theta_k)\|_2^2\right]\},
\end{aligned}$$

where the last step follows because $(j+k_0)^2 > (j+k_0)^2 - 1 = (j+k_0+1)(j+k_0-1)$. Therefore, we have shown $\mathbb{E}\left[\|\varphi_{k,j+2} - G(\theta_k)\|_2^2\right] \leq \frac{M_{\text{lazy}}}{\gamma^2(j+1+k_0)}$, which completes the proof by induction. \blacksquare

E.1 Proof of Theorem 3.3

First we state two deterministic identities used in the proof, which follow from Lemma B.4:

$$\|G(\theta) - \theta_{\text{PS}}\|_2 \leq \varepsilon \frac{\beta}{\gamma} \|\theta - \theta_{\text{PS}}\|_2, \quad (1)$$

$$\|\theta - G(\theta)\|_2 \leq \|\theta - \theta_{\text{PS}}\|_2 + \|\theta_{\text{PS}} - G(\theta)\|_2 \leq \left(1 + \varepsilon \frac{\gamma}{\beta}\right) \|\theta - \theta_{\text{PS}}\|_2. \quad (2)$$

Note that identity (2) implies $\|\theta - G(\theta)\|_2 < 2\|\theta - \theta_{\text{PS}}\|_2$ if $\varepsilon < \frac{\gamma}{\beta}$.

By triangle inequality, we have

$$\begin{aligned}
&\mathbb{E}\left[\|\theta_{k+1} - \theta_{\text{PS}}\|_2^2\right] \\
&= \mathbb{E}\left[\|\theta_{k+1} - G(\theta_k) + G(\theta_k) - \theta_{\text{PS}}\|_2^2\right] \\
&\leq \mathbb{E}\left[\|\theta_{k+1} - G(\theta_k)\|_2^2\right] + 2\mathbb{E}\left[\|\theta_{k+1} - G(\theta_k)\|_2 \|G(\theta_k) - \theta_{\text{PS}}\|_2\right] + \mathbb{E}\left[\|G(\theta_k) - \theta_{\text{PS}}\|_2^2\right]. \quad (3)
\end{aligned}$$

Denoting $k_0 = \frac{8L^2}{\gamma^2}$, Lemma E.1 bounds the first term by

$$\begin{aligned}
\mathbb{E}\left[\|\theta_{k+1} - G(\theta_k)\|_2^2\right] &= \mathbb{E}\left[\mathbb{E}\left[\|\theta_{k+1} - G(\theta_k)\|_2^2 \mid \theta_k\right]\right] \\
&\leq \frac{1.2\sigma^2 + 8L^2 \mathbb{E}\left[\|\theta_k - G(\theta_k)\|_2^2\right]}{\gamma^2(n(k) + k_0)} \\
&\leq \frac{1.2\sigma^2 + 32L^2 \mathbb{E}\left[\|\theta_k - \theta_{\text{PS}}\|_2^2\right]}{\gamma^2(n(k) + k_0)},
\end{aligned}$$

where in the last step we apply identity (2). Note also that by Jensen's inequality, we know

$$\mathbb{E}\left[\|\theta_{k+1} - G(\theta_k)\|_2\right] \leq \frac{1.1\sigma + 6L \mathbb{E}\left[\|\theta_k - G(\theta_k)\|_2\right]}{\gamma \sqrt{n(k) + k_0}}.$$

We can use this inequality, together with identities (1) and (2), to bound the cross-term in equation (3) as follows:

$$\begin{aligned}
& 2 \mathbb{E} [\|\theta_{k+1} - G(\theta_k)\|_2 \|G(\theta_k) - \theta_{\text{PS}}\|_2] \\
& \leq 2\varepsilon \frac{\beta}{\gamma} \mathbb{E} [\|\theta_{k+1} - G(\theta_k)\|_2 \|\theta_k - \theta_{\text{PS}}\|_2] \\
& \leq \frac{2\varepsilon \frac{\beta}{\gamma}}{\sqrt{n(k) + k_0}} \mathbb{E} \left[\left(\frac{6L}{\gamma} \|\theta_k - G(\theta_k)\|_2 + \frac{1.1\sigma}{\gamma} \right) \|\theta_k - \theta_{\text{PS}}\|_2 \right] \\
& \leq \frac{2\varepsilon \frac{\beta}{\gamma}}{\sqrt{n(k) + k_0}} \mathbb{E} \left[\left(\frac{6L}{\gamma} \left(1 + \varepsilon \frac{\beta}{\gamma} \right) \|\theta_k - \theta_{\text{PS}}\|_2 + \frac{1.1\sigma}{\gamma} \right) \|\theta_k - \theta_{\text{PS}}\|_2 \right] \\
& \leq \frac{24\varepsilon\beta L}{\gamma^2 \sqrt{n(k) + k_0}} \mathbb{E} [\|\theta_k - \theta_{\text{PS}}\|_2^2] + \frac{2.2\sigma\varepsilon\beta}{\gamma^2 \sqrt{n(k) + k_0}} \mathbb{E} [\|\theta_k - \theta_{\text{PS}}\|_2].
\end{aligned}$$

We bound the latter term by applying the AM-GM inequality; in particular, for all $\alpha_0 \in (0, 1)$, it holds that

$$\frac{2.2\sigma\varepsilon\beta}{\gamma^2 \sqrt{n(k) + k_0}} \mathbb{E} [\|\theta_k - \theta_{\text{PS}}\|_2] \leq \frac{1.1\sigma\varepsilon\beta}{\gamma^2} \left(\frac{1}{(n(k) + k_0)^{\alpha_0}} + \frac{\mathbb{E} [\|\theta_k - \theta_{\text{PS}}\|_2^2]}{(n(k) + k_0)^{1-\alpha_0}} \right).$$

Thus, the final bound on the cross-term in equation (3) is

$$\begin{aligned}
2 \mathbb{E} [\|\theta_{k+1} - G(\theta_k)\|_2 \|G(\theta_k) - \theta_{\text{PS}}\|_2] & \leq \left(\frac{24\varepsilon\beta L}{\gamma^2 \sqrt{n(k) + k_0}} + \frac{1.1\sigma\varepsilon\beta}{\gamma^2 (n(k) + k_0)^{1-\alpha_0}} \right) \mathbb{E} [\|\theta_k - \theta_{\text{PS}}\|_2^2] \\
& \quad + \frac{1.1\sigma\varepsilon\beta}{\gamma^2 (n(k) + k_0)^{\alpha_0}}.
\end{aligned}$$

The final term in equation (3) can be bounded by identity (1):

$$\mathbb{E} [\|G(\theta_k) - \theta_{\text{PS}}\|_2^2] \leq \left(\varepsilon \frac{\beta}{\gamma} \right)^2 \mathbb{E} [\|\theta_k - \theta_{\text{PS}}\|_2^2].$$

Putting all the steps together, we have derived the following recursion, true for all $\alpha_0 \in (0, 1)$:

$$\begin{aligned}
\mathbb{E} [\|\theta_{k+1} - \theta_{\text{PS}}\|_2^2] & \leq \left(\frac{32L^2}{\gamma^2 (n(k) + k_0)} + \frac{24\varepsilon\beta L}{\gamma^2 \sqrt{n(k) + k_0}} + \frac{1.1\sigma\varepsilon\beta}{\gamma^2 (n(k) + k_0)^{1-\alpha_0}} + \left(\varepsilon \frac{\beta}{\gamma} \right)^2 \right) \mathbb{E} [\|\theta_k - \theta_{\text{PS}}\|_2^2] \\
& \quad + \frac{1.2\sigma^2}{\gamma^2 (n(k) + k_0)} + \frac{1.1\sigma\varepsilon\beta}{\gamma^2 (n(k) + k_0)^{\alpha_0}} \\
& \leq c \mathbb{E} [\|\theta_k - \theta_{\text{PS}}\|_2^2] + \frac{1.2\sigma^2}{\gamma^2 (n(k) + k_0)} + \frac{1.1\sigma\varepsilon\beta}{\gamma^2 (n(k) + k_0)^{\alpha_0}}, \tag{4}
\end{aligned}$$

where we define

$$c \stackrel{\text{def}}{=} \frac{32L^2}{\gamma^2 n_0} + \frac{24\varepsilon\beta L}{\gamma^2 \sqrt{n_0}} + \frac{1.1\sigma\varepsilon\beta}{\gamma^2 n_0^{1-\alpha_0}} + \left(\varepsilon \frac{\beta}{\gamma} \right)^2.$$

We pick n_0 large enough such that there exists $\alpha_0 > 0$ for which $c < 1$.

Unrolling the recursion given by equation (4), we get

$$\mathbb{E}\left[\|\theta_{k+1} - \theta_{\text{PS}}\|_2^2\right] \leq c^k \|\theta_1 - \theta_{\text{PS}}\|_2^2 + \frac{1}{\gamma^2} \sum_{j=1}^k c^{k-j} \left(\frac{1.2\sigma^2}{n(j) + k_0} + \frac{1.1\sigma\varepsilon\beta}{(n(j) + k_0)^{\alpha_0}} \right).$$

Since $\alpha_0 < 1$, we can upper bound the second term as

$$\begin{aligned} & \frac{1}{\gamma^2} \sum_{j=1}^k c^{k-j} \left(\frac{1.2\sigma^2}{n(j) + k_0} + \frac{1.1\sigma\varepsilon\beta}{(n(j) + k_0)^{\alpha_0}} \right) \\ & \leq \frac{1.2\sigma^2}{\gamma^2} \sum_{j=1}^k c^{k-j} \frac{1}{n(j) + k_0} + \frac{1.1\sigma\varepsilon\beta}{\gamma^2} \sum_{j=1}^k c^{k-j} \frac{1}{(n(j) + k_0)^{\alpha_0}} \\ & \leq \frac{1}{\gamma^2(1-c)} \left(\frac{1.2\sigma^2}{n_0} (2k^{-\alpha} + c^{(1-2^{-1/\alpha})k}) + \frac{1.1\sigma\varepsilon\beta}{n_0^{\alpha_0}} (2k^{-\alpha\alpha_0} + c^{(1-2^{-1/(a\alpha_0)})k}) \right) \end{aligned}$$

where in the second inequality we apply Lemma B.5 after plugging in the choice of $n(k)$. Using the fact that $\alpha_0 \in (0, 1)$ and hence $c^{(1-2^{-1/(a\alpha_0)})k} < c^{(1-2^{-1/\alpha})k}$, as well as $\varepsilon < \frac{\gamma}{\beta}$ and $n_0 \geq 1$, gives

$$\begin{aligned} & \frac{1}{\gamma^2(1-c)} \left(\frac{1.2\sigma^2}{n_0} (2k^{-\alpha} + c^{(1-2^{-1/\alpha})k}) + \frac{1.1\sigma\varepsilon\beta}{n_0^{\alpha_0}} (2k^{-\alpha\alpha_0} + c^{(1-2^{-1/(a\alpha_0)})k}) \right) \\ & \leq \frac{1.2\sigma^2 + 1.1\sigma\gamma}{\gamma^2(1-c)} (4k^{-\alpha\alpha_0} + 2c^{(1-2^{-1/\alpha})k}) \\ & \leq \frac{3(\sigma + \gamma)^2}{\gamma^2(1-c)} (2k^{-\alpha\alpha_0} + c^{\Omega(k)}). \end{aligned}$$

It remains to set α_0 ; we set $\alpha_0 = \max\{\delta \in (0, 1) : c < 1\}$ (note that the existence of such α_0 is guaranteed by the choice of n_0). Clearly, $\alpha_0 \rightarrow 1$ as n_0 grows, and so putting everything together gives

$$\mathbb{E}\left[\|\theta_{k+1} - \theta_{\text{PS}}\|_2^2\right] \leq c^k \|\theta_1 - \theta_{\text{PS}}\|_2^2 + \frac{3(\sigma + \gamma)^2}{\gamma^2(1-c)} \left(\frac{2}{k^{\alpha \cdot (1-o(1))}} + c^{\Omega(k)} \right),$$

as desired.

F Proof of Corollary 3.4

From Theorem 3.2, we know that for greedy deploy, $\mathbb{E}\left[\|\theta_{k+1} - \theta_{\text{PS}}\|_2^2\right] = \mathcal{O}\left(\frac{1}{k}\right)$ where k indexes both the number of classifiers and the number of samples collected. By inverting this bound, we see that to ensure $\mathbb{E}\left[\|\theta_{k+1} - \theta_{\text{PS}}\|_2^2\right] \leq \delta$, it suffices to collect $\mathcal{O}\left(\frac{1}{\delta}\right)$ samples.

From our analogous convergence result for lazy deploy (Theorem 3.3), we know that after the k -th deployment, it holds that $\mathbb{E}\left[\|\theta_{k+1} - \theta_{\text{PS}}\|_2^2\right] = \mathcal{O}(1/k^{\alpha\omega})$, for some $\omega = 1 - o(1)$ which is independent of k and tends to 1 as n_0 grows. If we collect $\Theta(j^\alpha)$ samples for each deployment $j = 1 \dots k$, after k deployments the total number of samples N is $\Theta(k^{\alpha+1})$. Therefore,

$$\mathbb{E}\left[\|\theta_{k+1} - \theta_{\text{PS}}\|_2^2\right] = \mathcal{O}(1 / N^{\frac{\alpha\omega}{\alpha+1}}).$$

By inverting these bounds, we get our desired result for the asymptotics of lazy deploy.