
Regret Minimization with Performative Feedback

Meena Jagadeesan¹ Tijana Zrnic¹ Celestine Mendl-Dünner²

Abstract

In performative prediction, the deployment of a predictive model triggers a shift in the data distribution. As these shifts are typically unknown ahead of time, the learner needs to deploy a model to get feedback about the distribution it induces. We study the problem of finding near-optimal models under performativity while maintaining low regret. On the surface, this problem might seem equivalent to a bandit problem. However, it exhibits a fundamentally richer feedback structure that we refer to as *performative feedback*: after every deployment, the learner receives samples from the shifted distribution rather than bandit feedback about the reward. Our main contribution is regret bounds that scale only with the complexity of the distribution shifts and not that of the reward function. The key algorithmic idea is careful exploration of the distribution shifts that informs a novel construction of confidence bounds on the risk of unexplored models. The construction only relies on smoothness of the shifts and does not assume convexity. More broadly, our work establishes a conceptual approach for leveraging tools from the bandits literature for the purpose of regret minimization with performative feedback.

1. Introduction

Predictive models deployed in social settings are often performative. This means that the model’s predictions—by means of being used to inform consequential decisions—influence the outcomes the model aims to predict in the first place. For example, travel time estimates influence routing decisions and thus realized travel times, stock price predictions influence trading activity and hence prices. Such

¹University of California, Berkeley ²Max Planck Institute for Intelligent Systems, Tübingen. Correspondence to: Meena Jagadeesan <mjagadeesan@berkeley.edu>, Tijana Zrnic <tijana.zrnic@berkeley.edu>, Celestine Mendl-Dünner <cmendler@tuebingen.mpg.de>.

feedback-loop behavior arises in a variety of domains, including public policy, trading, traffic predictions, and recommendation systems.

Perdomo et al. (2020) formalized this phenomenon under the name *performative prediction*. A key concept in this framework is the *distribution map*, which formalizes the dependence of the data distribution on the deployed predictive model. This object maps a model, encoded by a parameter vector θ , to a distribution $\mathcal{D}(\theta)$ over instances. Naturally, in a performative environment, a model’s performance is measured on the distribution that results from its deployment. That is, given a loss function $\ell(z; \theta)$, which measures the learner’s loss when they predict on instance z using model θ , we evaluate a model based on its *performative risk*:

$$\text{PR}(\theta) := \mathbb{E}_{z \sim \mathcal{D}(\theta)} \ell(z; \theta). \quad (1)$$

In contrast with the risk function studied in usual supervised learning, the performative risk takes an expectation over a model-dependent distribution. Importantly, this distribution is *unknown* ahead of time; for example, one can hardly anticipate the distribution of travel times induced by a traffic forecasting system without deploying the system first.

Due to this inherent uncertainty about $\mathcal{D}(\theta)$, it is not possible to find a model with low performative risk offline. The learner needs to interact with the environment and deploy models θ to explore the induced distributions $\mathcal{D}(\theta)$. Given the online nature of this task, we measure the loss incurred by deploying a sequence of models $\theta_1, \dots, \theta_T$ by evaluating the *performative regret*:

$$\text{Reg}(T) := \sum_{t=1}^T \left(\mathbb{E} \text{PR}(\theta_t) - \min_{\theta} \text{PR}(\theta) \right),$$

where the expectation is taken over the possible randomness in the choice of $\{\theta_t\}_{t=1}^T$. Performative regret measures the suboptimality of the deployed sequence of models relative to a *performative optimum* $\theta_{\text{PO}} \in \text{argmin}_{\theta} \text{PR}(\theta)$.

At first glance, performative regret minimization might seem equivalent to a classical bandit problem. Bandit solutions minimize regret while requiring only noisy zeroth-order access to the unknown reward function—in our case PR. The resulting regret bounds generally grow with some notion of complexity of the reward function.

However, a naive application of bandit baselines misses out on a crucial fact: performative regret minimization exhibits significantly richer feedback than bandit feedback. When deploying a model θ , the learner gains access to samples from the induced distribution $\mathcal{D}(\theta)$, rather than only a noisy estimate of the risk $\text{PR}(\theta)$. We call this feedback model *performative feedback*. Together with the fact that the learner knows the loss $\ell(z; \theta)$, performative feedback can be used to inform the reward of unexplored arms. For instance, it allows the computation of an unbiased estimate of $\mathbb{E}_{z \sim \mathcal{D}(\theta)} \ell(z; \theta')$ for any point θ' .

To illustrate the power of this feedback model, consider the limiting case in which the performative effects entirely vanish and the distribution map is constant, i.e. $\mathcal{D}(\theta) \equiv \mathcal{D}_*$ for some fixed distribution \mathcal{D}_* independent of θ . With zeroth-order feedback, the learner would still need to deploy different models to explore the landscape of PR and find a point with low risk. However, with performative feedback, a *single* deployment gives samples from \mathcal{D}_* , thus resolving all uncertainty in the objective (1) apart from finite-sample uncertainty. This raises the question: *with performative feedback, can one achieve regret bounds that scale only with the complexity of the distribution map, and not that of the performative risk?*

1.1. Our Contribution

We study the problem of performative regret minimization based on performative feedback. Our main contribution is performative regret bounds that scale primarily with the complexity of the distribution map. The key conceptual idea is to apply bandits tools to carefully explore the distribution map, and then propagate this knowledge to the objective (1) in order to minimize performative regret.

Performative confidence bounds algorithm. Our main focus is on a setting where the distribution map is Lipschitz in an appropriate sense. We propose a new algorithm that takes advantage of performative feedback in order to construct non-trivial confidence bounds on the performative risk in unexplored regions of the parameter space and thus guide exploration. A crucial implication of these bounds is that the algorithm can discard highly suboptimal regions of the parameter space without ever deploying a model nearby. We summarize the regret guarantee of our *performative confidence bounds* algorithm:

Theorem 1.1 (Informal). *Suppose that the distribution map $\mathcal{D}(\theta)$ is ϵ -Lipschitz and that the loss $\ell(z; \theta)$ is L_z -Lipschitz in z . Then, after T deployments, the performative confidence bounds algorithm achieves a regret bound of*

$$\text{Reg}(T) = \tilde{O}\left(\sqrt{T} + T^{\frac{d+1}{d+2}} (L_z \epsilon)^{\frac{d}{d+2}}\right),$$

where d denotes the zooming dimension of the problem.

We compare the bound in Theorem 1.1 to a baseline Lipschitz bandits regret bound. The concept of zooming dimension stems from the work of Kleinberg et al. (2008) and serves as an instance-dependent notion of dimensionality.

Kleinberg et al. showed that sublinear regret $\tilde{O}\left(T^{\frac{d'+1}{d'+2}} L^{\frac{d'}{d'+2}}\right)$ can be achieved if the reward function is L -Lipschitz, where d' is a zooming dimension. The performative risk can be guaranteed to be Lipschitz if the distribution map is Lipschitz and the loss $\ell(z; \theta)$ is Lipschitz in *both* arguments.

The primary benefit of Theorem 1.1 is that our regret bound scales only with the Lipschitz constant of the distribution map, rather than the Lipschitz constant of the performative risk. In particular, our result allows $\text{PR}(\theta)$ to be highly irregular as a function of θ , seeing that $\ell(z; \theta)$ as a function of θ is unconstrained. This difference becomes salient when $\epsilon \rightarrow 0$, meaning that the performative effects vanish: our regret bound grows as $\tilde{O}(\sqrt{T})$ in an essentially dimension-independent manner. More precisely, the dimension can only arise implicitly through a model class complexity term. On the other hand, the rate of classical Lipschitz bandits remains exponential in the dimension.

Another difference between our regret bound and that of Lipschitz bandits is in the zooming dimension. In particular, d' is a zooming dimension *no smaller than* the zooming dimension we obtain in Theorem 1.1. As we will elaborate on in later sections, the benefit we derive from the zooming dimension comes from the fact that it implicitly depends on the Lipschitz constant driving the objective, which is smaller when making full use of performative feedback.

Extension to location families. In addition, we study performative regret minimization for the special case where the distribution map has a *location family* form (Miller et al., 2021). We again prove regret bounds that scale only with the complexity of the distribution map, rather than the complexity of the performative risk. We adapt the LinUCB algorithm (Li et al., 2010) to learn the hidden parameters of the location family. This enables us to achieve a $\tilde{O}(\sqrt{T})$ regret without placing any strong convexity assumptions on the performative risk that are required in prior work (Miller et al., 2021). In particular, our result again allows $\text{PR}(\theta)$ to be highly irregular as a function of θ .

More broadly, our work establishes a connection between performative prediction and the bandits literature, which we believe is a worthwhile direction for future inquiry.

Consequences for finding performative optima. While we have contextualized our work within online regret minimization, our performative confidence bounds algorithm has the additional property that it converges to the set of performative optima. Thus, if run for sufficiently many time steps, it generates a model with near-minimal performative risk. From this perspective, our algorithm offers superior guaran-

tees over retraining methods (e.g., (Perdomo et al., 2020)), which are unlikely to reach minima of the performative risk.

1.2. Related Work

Performative prediction. Prior work on performative prediction has largely studied gradient-based optimization methods (Perdomo et al., 2020; Mendler-Dünner et al., 2020; Drusvyatskiy and Xiao, 2020; Brown et al., 2022; Miller et al., 2021; Izzo et al., 2021; Maheshwari et al., 2022; Li and Wai, 2022; Ray et al., 2022; Dong and Ratliff, 2021). Many of the studied procedures only converge to *performatively stable* points, that is, points θ that satisfy the fixed-point condition $\theta \in \operatorname{argmin}_{\theta'} \mathbb{E}_{z \sim \mathcal{D}(\theta)} \ell(z; \theta')$. In general, stable points are not minimizers of the performative risk (Perdomo et al., 2020; Miller et al., 2021), which implies that procedures converging to stable points do not achieve sublinear performative regret. There are exceptions in the literature that focus on finding performative optima (Miller et al., 2021; Izzo et al., 2021), but those algorithms rely on proving or assuming convexity of the performative risk; in this work we make no convexity assumptions. In fact, it is known that the performative risk can be nonconvex even when the loss $\ell(z; \theta)$ is convex and the performative effects are relatively weak (Perdomo et al., 2020; Miller et al., 2021). One other work that studies performative optimality, without imposing convexity, is that of Dong and Ratliff (2021), but they focus on optimization heuristics that are not guaranteed to minimize performative regret.

Learning in Stackelberg games. Performative prediction is closely related to learning in *Stackelberg games*: if $\mathcal{D}(\theta)$ is thought of as a best response to the deployment of θ according to some unspecified utility function, then performative optima can be thought of as Stackelberg equilibria. There have been many works on learning dynamics in Stackelberg games in recent years (Balcan et al., 2015; Jin et al., 2020; Fiez et al., 2020; Fiez and Ratliff, 2020). Notably, Balcan et al. (2015) also study the benefit of a richer feedback model: they assume the agent’s type is revealed after taking an action. When combined with a known agent-response model, this allows them to directly infer the loss of unexplored strategies. In contrast, performative feedback does not imply full-information feedback. One instance of performative prediction that has an explicit Stackelberg structure, meaning $\mathcal{D}(\theta)$ is defined as a best response, is *strategic classification* (Hardt et al., 2016). Several works have studied learning dynamics in strategic classification (Dong et al., 2018; Chen et al., 2020; Bechavod et al., 2021; Zrnic et al., 2021); notably, Dong et al. (2018) and Chen et al. (2020) provide solutions that minimize Stackelberg regret, of which performative regret is an analog in the performative prediction context. However, all of these works rely on strong structural assumptions, such as linearity of the predictor or convexity of the risk function, which significantly reduce

the amount of necessary exploration compared to the mild Lipschitzness conditions we impose in our work.

Continuum-armed bandits. Particularly inspiring for our work is the literature on *continuum-armed bandits* (Agrawal, 1995; Kleinberg, 2004; Auer et al., 2007; Kleinberg et al., 2008; Podimata and Slivkins, 2021). As we will elaborate on in Section 2, performative prediction can be cast as a Lipschitz continuum-armed bandit problem. However, while this means that one can use an off-the-shelf Lipschitz bandit algorithm to minimize performative regret, this would generally be a conservative solution. After “pulling an arm” θ in performative prediction the learner observes samples from $\mathcal{D}(\theta)$. As explained earlier, in combination with the structure of our objective, this feedback model is more powerful than classical bandit feedback, where a noisy version of the mean reward at θ is observed. Moreover, it is fundamentally different from partial-feedback and side-information models studied in the literature, e.g. (Mannor and Shamir, 2011; Kocák et al., 2014; Wu et al., 2015; Cohen et al., 2016).

1.3. Preliminaries

Performative prediction, set up as an online learning problem, can be formalized as follows. At every step the learner chooses a model θ in the parameter space $\Theta \subset \mathbb{R}^{d_\theta}$. We assume¹ $\max\{\|\theta\| : \theta \in \Theta\} \leq 1$ for simplicity. The expected loss of model θ is given by $\operatorname{PR}(\theta) = \mathbb{E}_{z \sim \mathcal{D}(\theta)} \ell(z; \theta)$. We assume that the objective function is bounded so that $\ell(z; \theta) \in [0, 1]$ for all z and θ .

At every time step t , the learner chooses a model θ_t and observes a constant number m_0 of i.i.d. samples,

$$\{z_t^{(i)}\}_{i \in [m_0]}, \text{ where } z_t^{(i)} \sim \mathcal{D}(\theta_t).$$

The regret incurred by choosing θ_t at time step t is $\Delta(\theta_t) := \operatorname{PR}(\theta_t) - \operatorname{PR}(\theta_{\text{PO}})$, where θ_{PO} is the performative optimum.

The constant m_0 quantifies how many samples the learner can collect in a time window determined by how often they incur regret. For example, at the beginning of each week the learner might update the model, and thus at the end of each week they incur regret for the model they chose to deploy. In that case, m_0 is the number of samples the learner collects per week. Note that a learner with larger m_0 collects an empirical distribution that more accurately reflects $\mathcal{D}(\theta_t)$ and thus naturally minimizes regret at a faster rate.

To formally disentangle the effects of the parameter vector θ on the performative risk through the distribution map and the loss function, we use the notion of the *decoupled performative risk* (Perdomo et al., 2020):

$$\operatorname{DPR}(\theta, \theta') := \mathbb{E}_{z \sim \mathcal{D}(\theta)} \ell(z; \theta').$$

¹Throughout we use $\|\cdot\|$ to denote the ℓ_2 -norm for vectors and the operator norm for matrices.

This object captures the risk incurred by a model θ' on the distribution $\mathcal{D}(\theta)$. Note that $\text{PR}(\theta) = \text{DPR}(\theta, \theta)$ by definition.

To measure the complexity of the distribution map we consider how much the distribution $\mathcal{D}(\theta)$ can change with changes in θ , as formalized by ϵ -sensitivity.

Assumption 1.2 (ϵ -sensitivity (Perdomo et al., 2020)). A distribution map $\mathcal{D}(\cdot)$ is ϵ -sensitive if for any pair $\theta, \theta' \in \Theta$ it holds that

$$\mathcal{W}(\mathcal{D}(\theta), \mathcal{D}(\theta')) \leq \epsilon \|\theta - \theta'\|,$$

where \mathcal{W} denotes the Wasserstein-1 distance.

In the context of a traffic forecasting app, ϵ can be thought of as being proportional to the size of the user base of the app. When $\mathcal{D}(\theta)$ arises from the aggregate behavior of strategic agents manipulating their features in response to a model θ , ϵ grows when features are more easily manipulable.

2. A Black-Box Bandits Approach

Performative regret minimization can be set up as a continuum-armed bandits problem where an arm corresponds to a choice of model parameters θ . Performative feedback is sufficient to simulate noisy zeroth-order feedback about the reward function, as assumed in bandits. When we deploy θ_t , the samples from $\mathcal{D}(\theta_t)$ enable us to compute an unbiased estimate

$$\widehat{\text{PR}}(\theta_t) = \frac{1}{m_0} \sum_{i=1}^{m_0} \ell(z_t^{(i)}; \theta_t)$$

of the risk $\text{PR}(\theta_t)$. Moreover, since we assume the loss function is bounded, the noise in the estimate $\widehat{\text{PR}}(\theta_t)$ is subgaussian, as typically required in bandits.

A standard condition that makes continuum-armed bandit problems tractable is a bound on how fast the reward can change when moving from one arm to a nearby arm. Formally, this regularity is ensured by assuming Lipschitzness of the reward function—in our case, Lipschitzness of the performative risk.

The dependence of $\text{PR}(\theta)$ on θ is twofold, as seen in Equation (1). Thus, the most natural way to ensure that $\text{PR}(\theta)$ is Lipschitz is to ensure that each of these two dependencies is Lipschitz. This yields the following bound:

Lemma 2.1 (Lipschitzness of PR). *If the loss $\ell(z; \theta)$ is L_z -Lipschitz in z and L_θ -Lipschitz in θ and the distribution map is ϵ -sensitive, then the performative risk is $(L_\theta + \epsilon L_z)$ -Lipschitz.*

The intuition behind Lemma 2.1 is that $\text{PR}(\theta)$ is guaranteed to be Lipschitz if $\text{DPR}(\theta, \theta')$ is Lipschitz in each argument

individually. Lipschitzness in the second argument follows from requiring that the loss be Lipschitz in θ . Lipschitzness in the first argument follows from combining Lipschitzness of the loss in z and ϵ -sensitivity of the distribution map.

2.1. Adaptive Zooming

Once we have established Lipschitzness of the performative risk, we can apply techniques from the Lipschitz bandits literature. Kleinberg et al. (2008) proposed a bandit algorithm that adaptively discretizes promising regions of the space of arms, using Lipschitzness of the reward function to bound the additional loss due to discretization. Their method, called the *zooming algorithm*, will serve as a baseline for our problem. The algorithm enjoys an instance-dependent regret that takes advantage of nice problem instances, while maintaining tight guarantees in the worst case. The rate depends on the *zooming dimension*, which is upper bounded in the worst case by the dimension of the full space d_Θ .

Proposition 2.2 (Zooming algorithm (Kleinberg et al., 2008)). *Suppose $m_0 = o(\log T)$. Then, after T deployments, the zooming algorithm achieves a regret bound of*

$$\text{Reg}(T) = \mathcal{O} \left(T^{\frac{d+1}{d+2}} \left(\frac{\log T}{m_0} \right)^{\frac{1}{d+2}} (L_\theta + \epsilon L_z)^{\frac{d}{d+2}} \right),$$

where d denotes the $(L_\theta + \epsilon L_z)$ -zooming dimension.

The zooming dimension quantifies the niceness of a problem instance by measuring the size of a covering of near-optimal arms, instead of the entire parameter space. Roughly speaking, if the reward function is very “flat” in that there are many near-optimal points, then the zooming dimension is close to the dimension d_Θ of the parameter space. However, if the reward has sufficient curvature, then the zooming dimension can be much smaller than d_Θ . The zooming dimension is defined formally as follows:

Definition 2.3 (α -zooming dimension). A performative prediction problem instance has α -zooming dimension equal to d if any minimal s -cover of any subset of $\{\theta : \Delta(\theta) \leq 16\alpha s\}$ includes at most a constant multiple of $(3/s)^d$ elements from $\{\theta : 16\alpha r \leq \Delta(\theta) < 32\alpha r\}$, for all $0 < r \leq s \leq 1$.

For well-behaved instances, the definition intuitively requires every minimal s -cover of $\{\theta : 16\alpha r \leq \Delta(\theta) < 32\alpha r\}$ to have size at most of order $(3/s)^d$. Definition 2.3 slightly differs from the definition presented in (Kleinberg et al., 2008) and makes the dependence on the Lipschitz constant explicit; we use Definition 2.3 to later ease the comparison to our new algorithm. The differences between the two definitions are minor technicalities that we do not expect to alter the zooming dimension in a meaningful way, neither formally nor conceptually. See Appendix E for a discussion.




Figure 1. Confidence bounds after deploying θ_1 and θ_2 . (left) Confidence bounds via Lipschitzness, Equation (2). (right) Performative confidence bounds, Equation (3). The performative feedback model used for this illustration can be found in Appendix G.




Figure 2. Performative feedback allows discarding unexplored suboptimal models even in regions that have not been explored. A model θ is discarded if $\text{PR}_{\text{LB}}(\theta) > \text{PR}_{\text{min}}$. The loss function and feedback model are the same as in Figure 1.

3. Making Use of Performative Feedback

In this section, we illustrate how we can take advantage of performative feedback beyond computing a point estimate of the deployed model’s risk. For now, we ignore finite-sample considerations and assume access to the entire distribution $\mathcal{D}(\theta)$ after deploying a model θ . We will address finite-sample uncertainty when presenting our main algorithm in the next section.

3.1. Constructing Performative Confidence Bounds

First, we demonstrate how performative feedback allows constructing tighter confidence bounds on the performative risk of unexplored models, compared to only relying on Lipschitzness of the risk function $\text{PR}(\theta)$.

Suppose we deploy a set of models $\mathcal{S} \subseteq \Theta$ and for each $\theta \in \mathcal{S}$ we observe $\mathcal{D}(\theta)$. Then, under the regularity conditions of Lemma 2.1, we can bound the risk of any $\theta' \in \Theta$ as

$$\begin{aligned} \max_{\theta \in \mathcal{S}} \text{PR}(\theta) - (L_\theta + L_z \epsilon) \|\theta - \theta'\| &\leq \text{PR}(\theta') \\ &\leq \min_{\theta \in \mathcal{S}} \text{PR}(\theta) + (L_\theta + L_z \epsilon) \|\theta - \theta'\|. \end{aligned} \quad (2)$$

These confidence bounds only use $\mathcal{D}(\theta)$ for the purpose of computing $\text{PR}(\theta)$ and rely on Lipschitzness to construct confidence sets around the risk of unexplored models. However, in light of the structure of our objective function (1), the bounds in Equation (2) do not make full use of performative feedback; in particular, access to $\mathcal{D}(\theta)$ actually allows us to evaluate $\text{DPR}(\theta, \theta')$ for any θ' . Importantly, this information can further reduce our uncertainty about $\text{PR}(\theta')$, and we can bound:

$$\begin{aligned} \text{PR}(\theta') &= \text{DPR}(\theta, \theta') + (\text{DPR}(\theta', \theta') - \text{DPR}(\theta, \theta')) \\ &\leq \text{DPR}(\theta, \theta') + L_z \epsilon \|\theta - \theta'\|. \end{aligned}$$

Thus we can get tighter bounds on the performative risk at

an unexplored parameter θ' :

$$\begin{aligned} \max_{\theta \in \mathcal{S}} \text{DPR}(\theta, \theta') - L_z \epsilon \|\theta - \theta'\| &\leq \text{PR}(\theta') \\ &\leq \min_{\theta \in \mathcal{S}} \text{DPR}(\theta, \theta') + L_z \epsilon \|\theta - \theta'\|. \end{aligned} \quad (3)$$

We call the confidence bounds computed in (3) *performative confidence bounds*. In Figure 1, we visualize and contrast these confidence bounds with the confidence bounds obtained via Lipschitzness. We observe that by computing DPR we can significantly tighten the confidence regions.

The tightness of the confidence bounds depends on the set \mathcal{S} of deployed models. By choosing a cover of the parameter space, we can get an estimate of the performative risk that has low approximation error on the whole parameter space.

Proposition 3.1. *Let \mathcal{S}_γ be a γ -cover of Θ and suppose we deploy all models $\theta \in \mathcal{S}_\gamma$. Then, using performative feedback we can compute an estimate $\widehat{\text{PR}}(\theta)$ such that for any $\theta \in \Theta$ it holds that $|\text{PR}(\theta) - \widehat{\text{PR}}(\theta)| \leq \gamma L_z \epsilon$.*

Proposition 3.1 implies that after exploring the cover \mathcal{S}_γ , we can find a model whose suboptimality is at most $\mathcal{O}(\gamma L_z \epsilon)$. To contextualize the bound in Proposition 3.1, consider an approach that uses the same cover \mathcal{S}_γ but only relies on zeroth-order feedback, that is, $\{\text{PR}(\theta) : \theta \in \mathcal{S}_\gamma\}$. Then, the only feasible estimate of PR over the whole space is $\widehat{\text{PR}}(\theta) = \text{PR}(\Pi_{\mathcal{S}_\gamma}(\theta))$, where $\Pi_{\mathcal{S}_\gamma}(\theta) = \arg\min_{\theta' \in \mathcal{S}_\gamma} \|\theta - \theta'\|$ is the projection onto the cover \mathcal{S}_γ . This zeroth-order approach only guarantees an accuracy of $|\text{PR}(\theta) - \widehat{\text{PR}}(\theta)| \leq (L_z \epsilon + L_\theta) \gamma$, a strictly weaker approximation than the one in Proposition 3.1.

3.2. Sequential Elimination of Suboptimal Models

Now we show how performative confidence bounds can guide exploration. Specifically, we show that every deployment informs the risk of unexplored models, which allows us to sequentially discard suboptimal regions of the space.

To develop a formal procedure for discarding points, let

$\text{PR}_{\text{LB}}(\theta)$ denote a lower confidence bound on $\text{PR}(\theta)$ and PR_{min} denote an upper confidence bound on $\text{PR}(\theta_{\text{PO}})$ based on the information from the models deployed so far:

$$\text{PR}_{\text{LB}}(\theta) = \max_{\theta' \text{ already deployed}} (\text{DPR}(\theta', \theta) - L_z \epsilon \|\theta - \theta'\|),$$

$$\text{PR}_{\text{min}} = \min_{\theta \in \Theta} \min_{\theta' \text{ already deployed}} (\text{DPR}(\theta', \theta) + L_z \epsilon \|\theta' - \theta\|).$$

It is not difficult to see that the following lower bound on the suboptimality of model θ holds:

Proposition 3.2. *We have $\Delta(\theta) \geq \text{PR}_{\text{LB}}(\theta) - \text{PR}_{\text{min}}, \forall \theta$.*

In particular, models θ with $\text{PR}_{\text{LB}}(\theta) > \text{PR}_{\text{min}}$ cannot be optimal. We recall our toy example from Figure 1 and illustrate in Figure 2 the parameter configurations we can discard after the deployment of two models, θ_1 and θ_2 . We can see that access to DPR allows us to discard a large portion of the parameter space, and, in contrast to the baseline black-box approach, it is possible to discard regions of the space that have not been explored.

4. Performative Confidence Bounds Algorithm

We introduce our main algorithm that builds on the two insights from the previous section. We furthermore provide a rigorous, finite-sample analysis of its guarantees.

4.1. Algorithm Overview

Our *performative confidence bounds* algorithm, formally stated in Algorithm 1, takes advantage of performative feedback by assessing the risk of unexplored models and thus guiding exploration. We give an overview of the main steps.

Inspired by the successive elimination algorithm (Even-Dar et al., 2002), the algorithm keeps track of and refines an *active* set of models $\mathcal{A} \subseteq \Theta$. Roughly speaking, active models are those that are estimated to have low risk and only they are admissible to deploy. To deal with finite-sample uncertainty, the algorithm proceeds in phases which progressively refine the precision of the finite-sample risk estimates. More precisely, in phase p the algorithm chooses an error tolerance γ_p and deploys a model for n_p steps. In each step m_0 samples induced by the deployed model are collected, and n_p is chosen so that the inferred estimates of DPR are γ_p -accurate. Formally, if θ is deployed in phase p , we collect an empirical distribution $\widehat{\mathcal{D}}(\theta)$ of $n_p m_0$ samples so that $|\widehat{\text{DPR}}(\theta, \theta') - \text{DPR}(\theta, \theta')| \leq \gamma_p$ for all θ' with high probability, where

$$\widehat{\text{DPR}}(\theta, \theta') := \mathbb{E}_{z \sim \widehat{\mathcal{D}}(\theta)} \ell(z; \theta').$$

These estimates of DPR are used to construct performative confidence bounds and refine \mathcal{A} .

Algorithm 1 Performative Confidence Bounds Algorithm

Require: time horizon T , number of samples m_0 , sensitivity ϵ , Lipschitz constant L_z , complexity bound \mathfrak{C}

- 1: Initialize $\mathcal{A} = \Theta$
- 2: **for** phase $p = 0, 1, \dots$ **do**
- 3: Set error tolerance $\gamma_p = 2^{-p}$ and net radius $r_p = \frac{\gamma_p}{L_z \epsilon}$
- 4: Let $n_p = \left\lceil \left(2\mathfrak{C} + 3\sqrt{\log T} \right)^2 / (\gamma_p^2 m_0) \right\rceil$
- 5: Initialize $\mathcal{S}_p \leftarrow \mathcal{N}_{r_p}(\mathcal{A})$, $\mathcal{P}_p \leftarrow \emptyset$
- 6: **while** $\mathcal{S}_p \neq \emptyset$ **do**
- 7: Draw $\theta_{\text{net}} \in \mathcal{S}_p$ uniformly at random
- 8: Deploy θ_{net} for n_p steps to form $\widehat{\text{DPR}}(\theta_{\text{net}}, \cdot)$
- 9: Update $\mathcal{S}_p \leftarrow \mathcal{S}_p \setminus \theta_{\text{net}}$ and $\mathcal{P}_p \leftarrow \mathcal{P}_p \cup \theta_{\text{net}}$
- 10: Update estimate of upper bound on $\text{PR}(\theta_{\text{PO}})$:

$$\text{PR}_{\text{min}} = \min_{\theta \in \Theta} \min_{\theta' \in \mathcal{P}_p} \left(\widehat{\text{DPR}}(\theta', \theta) + L_z \epsilon \|\theta' - \theta\| \right)$$

- 11: Update lower bound for all models $\theta \in \mathcal{A}$:

$$\text{PR}_{\text{LB}}(\theta) = \max_{\theta' \in \mathcal{P}_p} \left(\widehat{\text{DPR}}(\theta', \theta) - L_z \epsilon \|\theta' - \theta\| \right)$$

- 12: $\mathcal{A} \leftarrow \mathcal{A} \setminus \{ \theta \in \mathcal{A} : \text{PR}_{\text{LB}}(\theta) > \text{PR}_{\text{min}} + 2\gamma_p \}$
 - 13: $\mathcal{S}_p \leftarrow \mathcal{S}_p \setminus \{ \theta \in \mathcal{S}_p : \text{Ball}_{r_p}(\theta) \cap \mathcal{A} = \emptyset \}$
 - 14: **end while**
 - 15: **end for**
-

Each phase begins by constructing a net of the current active set \mathcal{A} . The points in the net are sequentially deployed in the phase, unless they are deemed to be suboptimal based on previous deployments in that phase and are in that case eliminated. During phase p , we denote by \mathcal{P}_p the running set of deployed points and by \mathcal{S}_p the running set of net points that have not been discarded. We initialize \mathcal{S}_p to a minimal r_p -net of the current set of active points \mathcal{A} , denoted $\mathcal{N}_{r_p}(\mathcal{A})$, where r_p is proportional to γ_p . A net point θ gets eliminated from \mathcal{S}_p if no point in $\text{Ball}_{r_p}(\theta) := \{ \theta' \in \Theta : \|\theta' - \theta\| \leq r_p \}$ is active. This means that we may deploy suboptimal points in the net if they help inform active points nearby.

4.2. Comparison with Adaptive Zooming Algorithm

While we borrow the idea of an instance-dependent zooming dimension from Kleinberg et al. (2008), Algorithm 1 and its analysis are substantially different from prior work. In particular, Kleinberg et al. (2008) study an adaptive zooming algorithm which combines a UCB-based approach with an arm activation step. Adapting this method to our setting encounters several obstacles that we describe below.

First, a naive application of the adaptive zooming algorithm proposed by Kleinberg et al. (2008) does not lead to sub-linear regret in our setting, unless we assume Lipschitzness of PR. Their rule for activating new arms requires that the

reward of arms within a given radius in Euclidean distance of the pulled arm is similar. However, without Lipschitzness of PR, there is no radius that would ensure this property.

Given the shortcomings of this exploration strategy, one might imagine that selecting a better distance between arms, e.g. one based on performative confidence bounds, would result in a better algorithm. A natural distance function would be $d(\theta, \theta')$ taken as (an empirical estimate of) $\text{PR}(\theta) - \text{DPR}(\theta, \theta') + L_z \epsilon \|\theta - \theta'\|$. The challenge is that the analysis in Kleinberg et al. (2008) explicitly requires symmetry of the distance function, which $d(\theta, \theta')$ violates.

Therefore, to single out the $L_z \epsilon$ dependence, it is necessary to disentangle learning the structure of the distribution map from the elimination of arms based on reward, which is in stark contrast with UCB-style adaptive zooming algorithms. Algorithm 1 achieves this by relying on a novel adaptation of successive elimination.

4.3. Regret Bound

Before we state the regret bound for Algorithm 1, let us comment on an important component in the analysis. Recall that throughout the algorithm we operate with finite-sample estimates of the decoupled performative risk to bound the risk of unexplored models. Specifically, for any deployed θ , we make use of $\widehat{\text{DPR}}(\theta, \theta')$ for all θ' . Since we need these estimates to be valid simultaneously for all θ' , we rely on uniform convergence. As such, the Rademacher complexity of the loss function class naturally enters the bound.

Definition 4.1 (Rademacher complexity). Given a loss function $\ell(z; \theta)$, we define $\mathfrak{C}^*(\ell)$ to be:

$$\mathfrak{C}^*(\ell) = \sup_{\theta \in \Theta} \sup_{n \in \mathbb{N}} \sqrt{n} \cdot \mathbb{E}_{\epsilon, z^\theta} \left(\sup_{\theta' \in \Theta} \left| \frac{1}{n} \sum_{j=1}^n \epsilon_j \ell(z_j^\theta; \theta') \right| \right),$$

where $\epsilon_j \sim \text{Rademacher}$ and $z_j^\theta \sim \mathcal{D}(\theta)$, $\forall j \in [n]$, which are all independent of each other.

Now we can state our regret guarantee for Algorithm 1.

Theorem 4.2 (Main regret bound). *Assume the loss $\ell(z; \theta)$ is L_z -Lipschitz in z and let ϵ denote the sensitivity of the distribution map. Suppose that \mathfrak{C} is any value such that $\mathfrak{C}^*(\ell) \leq \mathfrak{C}$ and $m_0 = o(\mathcal{B}_{\log T, \mathfrak{C}}^2)$, where $\mathcal{B}_{\log T, \mathfrak{C}} := \sqrt{\log T} + \mathfrak{C}$. Then, after T time steps, Algorithm 1 achieves a regret bound of*

$$\text{Reg}(T) = \mathcal{O} \left(T^{\frac{d+1}{d+2}} \left(\frac{(L_z \epsilon)^d \mathcal{B}_{\log T, \mathfrak{C}}^2}{m_0} \right)^{\frac{1}{d+2}} + \sqrt{T} \frac{\mathcal{B}_{\log T, \mathfrak{C}}}{\sqrt{m_0}} \right),$$

where d is the $(L_z \epsilon)$ -sequential zooming dimension.

The sequential zooming dimension, formally defined in Definition 4.4, accounts for the sequential elimination of models

within each phase. It will become clear in the next section that the sequential zooming dimension is upper bounded by the usual zooming dimension from Definition 2.3. In Appendix F, we provide an example where the sequential zooming is *strictly* smaller than the zooming dimension.

Proposition 4.3. *For all $\alpha > 0$, the α -zooming dimension is at least as large as the α -sequential zooming dimension.*

The primary advantage of Theorem 4.2 over the Lipschitz bandit baseline can be seen by examining the first term in the regret bound. This term resembles the black-box regret bound from Section 2; however, the key difference is that that the bound of Theorem 4.2 depends on the complexity of the distribution map rather than that of the performative risk. In particular, the Lipschitz constant is $L_z \epsilon$ and not $L_\theta + L_z \epsilon$. The advantage is pronounced when $\epsilon \rightarrow 0$, making the first term of the bound in Theorem 4.2 vanish so only the $\mathcal{O}(\sqrt{T})$ term remains. On the other hand, the bound in Proposition 4.3 maintains an exponential dimension dependence.

Taking the limit as $\epsilon \rightarrow 0$ also reveals why the second term in the bound emerges. Even if the distribution map is constant, there is regret arising from finite-sample error. This is a key conceptual difference in the meaning of Lipschitzness of the distribution map versus that of the performative risk: $L_\theta + L_z \epsilon$ being 0 implies that PR is flat and thus all models are optimal, while performative regret minimization is non-trivial even if $L_z \epsilon = 0$. Unlike the first term, the second term due to finite samples is dimension-independent apart from any dependence implicit in the Rademacher complexity.

We note that the presence of the Rademacher complexity term $\mathfrak{C}^*(\ell)$ makes a direct comparison of the bound in Theorem 4.2 and the bound in Proposition 4.3 subtle. When the Rademacher complexity is very high, the regret bound in Theorem 4.2 may be worse. Nonetheless, for many natural function classes, the Rademacher complexity is polynomial in the dimension; in these cases, Theorem 4.2 can substantially outperform the regret bound in Proposition 4.3.

Another key feature of the regret bound in Theorem 4.2 worth highlighting is the zooming dimension. Definition 2.3 allows us to directly compare the dimension in Theorem 4.2 with the dimension in Proposition 4.3: the $(L_z \epsilon)$ -zooming dimension of Algorithm 1 is no larger than, and most likely smaller than, the $(L_\theta + L_z \epsilon)$ -zooming dimension in the black-box approach. Moreover, the sequential variant of zooming dimension in Theorem 4.2 can further reduce the dimension.

Finally, the main assumption underpinning the bound in Theorem 4.2 is that DPR is $(L_z \epsilon)$ -Lipschitz in its first argument. Assumption 1.2 coupled with Lipschitzness of the loss in the data achieves this. However, this property can hold with different regularity assumptions on the distribution map and loss function; e.g., if the loss is bounded and the distribution map is Lipschitz in total variation distance.




Figure 3. Sequential deployment of models allows Algorithm 1 to eliminate points from \mathcal{S}_p , reducing the number of deployments. The deployment of $\theta_{\text{net},1}$ and $\theta_{\text{net},2}$ allows one to eliminate $\theta_{\text{net},3}$.

4.4. Sequential Zooming Dimension

The zooming dimension of Definition 2.3 does not take into account that, using performative feedback, our algorithm can eliminate unexplored models *within* a phase. We illustrate the benefits of this sequential exploration strategy in Figure 3, where the deployment of two models is sufficient to eliminate the remaining model in the cover. This motivates a sequential definition of zooming dimension that captures the benefits of sequential exploration.

To set up this definition, we need to introduce some notation. For a set of points \mathcal{S} , enumeration $\pi : \mathcal{S} \rightarrow \{1, \dots, |\mathcal{S}|\}$ that specifies an ordering on \mathcal{S} , and number $k \in \{1, \dots, |\mathcal{S}|\}$, let

$$\begin{aligned} \text{PR}_{\text{LB}}(\theta; k) &:= \max_{\theta' \in \mathcal{S}: \pi(\theta') < k} (\text{DPR}(\theta', \theta) - L_z \epsilon \|\theta - \theta'\|), \\ \text{PR}_{\text{LB}}^s(k) &:= \min_{\theta \in \text{Ball}_s(\pi^{-1}(k))} \text{PR}_{\text{LB}}(\theta; k), \\ \text{PR}_{\text{min}}(k) &:= \min_{\theta} \min_{\theta' \in \mathcal{S}: \pi(\theta') < k} (\text{DPR}(\theta', \theta) + L_z \epsilon \|\theta' - \theta\|). \end{aligned}$$

Here, $\text{PR}_{\text{LB}}(\theta; k)$ is a lower bound on $\text{PR}(\theta)$ arising from the first $k-1$ deployments of the phase. Similarly, $\text{PR}_{\text{LB}}^s(k)$ captures the minimal lower confidence bound on PR for any point in an s -ball around the k -th deployed model, $\pi^{-1}(k)$. Finally, $\text{PR}_{\text{min}}(k)$ captures an upper bound on $\text{PR}(\theta_{\text{PO}})$, estimated from the first $k-1$ deployments.

Using the above terms, we see that $\text{PR}_{\text{LB}}^s(k) \leq \text{PR}_{\text{min}}(k) + 4\alpha s$ is the population version of the condition that a model in the cover does not get discarded. The sequential zooming dimension captures the maximal number of models in each suboptimality band that can be deployed.

Definition 4.4 (Sequential zooming dimension). A performative prediction problem instance has α -zooming dimension equal to d if for any minimal s -cover \mathcal{S} of any subset of $\{\theta : \Delta(\theta) \leq 16\alpha s\}$ and all $0 < r \leq s \leq 1$, the expected number of models $\theta \in \mathcal{S} \cap \{\theta : 16\alpha r \leq \Delta(\theta) < 32\alpha r\}$ with

$$\text{PR}_{\text{LB}}^s(\pi(\theta)) \leq \text{PR}_{\text{min}}(\pi(\theta)) + 4\alpha s \quad (4)$$

is at most a constant multiple of $(3/s)^d$, where the expectation is taken over a uniformly sampled enumeration $\pi : \mathcal{S} \rightarrow \{1, \dots, |\mathcal{S}|\}$.

The claim of Proposition 4.3—that the sequential zooming dimension is bounded by the zooming dimension in Definition 2.3—follows by definition. To see this, let d be the α -zooming dimension. This means that \mathcal{S} includes at most a constant multiple of $(3/s)^d$ elements from $\{\theta : 16\alpha r \leq \Delta(\theta) < 32\alpha r\}$, for all $0 < r \leq s \leq 1$. This immediately guarantees that the subset of \mathcal{S} characterized by (4) is at most a multiple of $(3/s)^d$, as desired.

5. Regret Minimization for Location Families

In this section, we show how further knowledge about the structure of the distribution map can help reduce the complexity of performative regret minimization, without necessarily implying favorable structure of the performative risk. Once again, we apply our guiding principle of focusing exploration on learning the distribution map. Since the loss function is known, we can extrapolate knowledge about the distribution map to estimate the performative risk.

We focus on the setting of *location families* (Miller et al., 2021), which are distribution maps that depend on θ via a linear shift. More precisely, location families are distribution maps of the form $z \sim \mathcal{D}(\theta) \Leftrightarrow z \stackrel{d}{=} z_0 + \mu_*^\top \theta$, where $\mu_* \in \mathbb{R}^{d_\Theta \times m}$ is an unknown matrix and $z_0 \in \mathbb{R}^m$ is a zero-mean subgaussian sample from a base distribution \mathcal{D}_0 . In Appendix C, we provide an example of a standard location family in strategic classification.

At a high level, our algorithm can be described as follows: at every step t , the learner deploys a model θ_t and collects m_0 samples from $\mathcal{D}(\theta_t)$. We will write $\bar{z}_t := \frac{1}{m_0} \sum_{i=1}^{m_0} z_t^{(i)}$ for the corresponding sample average at time t . Then, based on all samples collected so far, the algorithm computes the least-squares estimate of μ_* along with a confidence region for μ_* . In the next step the algorithm picks the model that minimizes a lower confidence bound $\text{PR}_{\text{LB}}(\theta)$. See Algorithm 2 for details.

Algorithm 2 Regret Minimization for Location Families

Require: time horizon T , number of samples m_0 , base distribution \mathcal{D}_0 , bound M_* such that $\|\mu_*\| \leq M_*$

- 1: Initialize confidence set $\mathcal{C}_1 \leftarrow \{\mu : \|\mu\| \leq M_*\}$
- 2: **for** step $t = 1, 2, \dots$ **do**
- 3: $\text{PR}_{\text{LB}}(\theta) \leftarrow \min_{\mu \in \mathcal{C}_t} \mathbb{E}_{z_0 \sim \mathcal{D}_0} \ell(z_0 + \mu\theta; \theta) \forall \theta \in \Theta$
- 4: Deploy $\theta_t = \arg\min_{\theta} \text{PR}_{\text{LB}}(\theta)$, compute \bar{z}_t
- 5: Let $\Sigma_t \leftarrow \sum_{i=1}^t \theta_i \theta_i^\top + \frac{1}{m_0} I$
- 6: Update estimate of μ_* , $\hat{\mu}_t = \Sigma_t^{-1} \left(\sum_{i=1}^t \theta_i \bar{z}_i^\top \right)$
- 7: Update set $\mathcal{C}_{t+1} \leftarrow \left\{ \mu : \left\| \Sigma_t^{1/2} (\hat{\mu}_t - \mu) \right\| < \xi \right\}$, where

$$\xi := \frac{M_* + \sqrt{8m_0 + 8 \log T + 2d_\Theta \log \left(1 + \frac{Tm_0}{d_\Theta} \right)}}{\sqrt{m_0}}$$

- 8: **end for**

This algorithm is inspired by LinUCB (Li et al., 2010), a standard bandits algorithm for linear rewards whose regret scales as $\tilde{O}(d\sqrt{T})$, where d is the dimension of the linear map. Importantly, unlike in the LinUCB analysis, our objective function $\text{PR}(\theta)$ is *not* linear in θ . Still, the nature of performative feedback allows us to learn the hidden linear structure in the distribution map and apply this knowledge to obtain confidence bounds on the performative risk. Below we state our algorithm for performative regret minimization for location families together with its regret guarantees.

Theorem 5.1. *Suppose that $\ell(z; \theta)$ is L_z -Lipschitz in z , \mathcal{D}_0 is 1-subgaussian, and $m_0 = o(\log T)$. Then, after T time steps, Algorithm 2 achieves a regret bound of*

$$\text{Reg}(T) = \tilde{O}\left(\frac{1}{\sqrt{m_0}} \max\{L_z, 1\} \sqrt{T} \max\{d_\Theta, \sqrt{d_\Theta m}\}\right).$$

Remark 5.2. For simplicity, we assume that \mathcal{D}_0 is known in Algorithm 2. This assumption is justified, for example, when we have plenty of historical data about a population, before any model deployment. We note that Theorem 5.1 can be extended to the case where we only have a finite data set from \mathcal{D}_0 , by relying on a uniform convergence argument.

Theorem 5.1 shows that by leveraging the hidden linear structure of the distribution map, Algorithm 2 inherits the $\tilde{O}(\sqrt{T})$ rate of LinUCB. This bears resemblance to the regret bound in Theorem 4.2 that also scaled primarily with the complexity of the distribution map. Furthermore, similarly to Algorithm 1, we see that the regret bound for Algorithm 2 holds while allowing the loss to have arbitrary dependence on θ . For example, the loss need not be convex and, as a result, the performative risk need not be convex either.

We conclude by comparing Theorem 5.1 to prior work (Miller et al., 2021), which provided an algorithm for finding performative optima for location families in the special case when the performative risk is *strongly convex*. Converting their optimization error into a regret bound yields a bound of $\mathcal{O}(\sqrt{T}(d_\Theta + m))$. While this bears resemblance to Theorem 5.1, the rates are not directly comparable. The algorithm from Miller et al. (2021) does not assume knowledge of the base distribution \mathcal{D}_0 , but rather deploys the model $\theta = 0$ in initial steps to collect samples from \mathcal{D}_0 (see Remark 5.2 for how to combine this strategy with our algorithm). In any case, the main benefit of Theorem 5.1 is that it applies to a more general setting, placing significantly fewer restrictions on the loss function and the performative risk.

6. Future Directions

Having illuminated the connection between performative prediction and bandit problems, our work opens the door for interesting further investigations. We highlight several directions we consider promising.

Structural knowledge of the distribution map. Domain knowledge about performative distribution shifts is sometimes available: for example, a parametric approximation to the aggregate response (Miller et al., 2021; Izzo et al., 2021), a microfoundations model for individual behavior (Hardt et al., 2016; Jagadeesan et al., 2021), or basic constraints on the agents’ action set (Chen et al., 2020). For linear shifts, we demonstrated how such structural knowledge about the distribution map can help guide exploration. We expect this principle to apply to other structures of $\mathcal{D}(\theta)$.

Consequences of exploration. An important limitation of exploration in performative environments are social welfare concerns. Performative shifts can rarely be analyzed offline and every model deployment is consequential for the population the model acts upon. The ability to discard highly suboptimal regions of the parameter space without having to deploy a model within is highly appealing from a welfare perspective. Beyond this, we believe that incorporating constraints on what constitutes safe exploration (Wu et al., 2016; Turchetta et al., 2019; Kazerouni et al., 2017) is crucial for performative optimization in practice.

Costs of a new deployment. Our notion of regret quantifies the statistical complexity of regret minimization, but it does not differentiate between collecting more samples induced by the currently deployed model and deploying a new model. This difference has previously been studied by Mendler-Dünnler et al. (2020). Due to the costs associated with a new deployment, collecting more samples from the same model typically comes at a reduced cost for the learner, and there may be a better notion of regret that reflects this.

Adapting to unknown sensitivity. Our algorithm relies on knowing $L_z\epsilon$. While the Lipschitzness of a classifier in the data has been studied in the context of adversarial robustness (Szegedy et al., 2014; Cisse et al., 2017; Hein and Andriushchenko, 2017; Yang et al., 2020), which could help inform L_z , the sensitivity ϵ of an environment is generally unknown. Adapting the tools by Bubeck et al. (2011) could help relax the requirement of a known sensitivity.

“Best of both worlds” algorithm. When the Rademacher complexity of the function class is high, the Lipschitz bandit baseline may provide a better regret bound than Algorithm 1. It would be interesting to design an algorithm that intersects the confidence sets of both algorithms and inherits the better of the two regret bounds.

Acknowledgements

The authors would like to thank Moritz Hardt for helpful conversations during the course of this project, Nilesch Tripuraneni for pointers to relevant literature, and Jacob Steinhardt, Alex Wei, and Clara Wong-Fannjiang for valuable feedback on the manuscript.

References

- Rajeev Agrawal. The continuum-armed bandit problem. *SIAM Journal on Control and Optimization*, 33(6):1926–1951, 1995.
- Peter Auer, Ronald Ortner, and Csaba Szepesvári. Improved rates for the stochastic continuum-armed bandit problem. In *International Conference on Computational Learning Theory*, pages 454–468. Springer, 2007.
- Maria-Florina Balcan, Avrim Blum, Nika Haghtalab, and Ariel D Procaccia. Commitment without regrets: Online learning in Stackelberg security games. In *Proceedings of the 16th ACM Conference on Economics and Computation*, pages 61–78, 2015.
- Yahav Bechavod, Katrina Ligett, Steven Wu, and Juba Ziani. Gaming helps! learning from strategic interactions in natural dynamics. In *International Conference on Artificial Intelligence and Statistics*, pages 1234–1242, 2021.
- Gavin Brown, Shlomi Hod, and Iden Kalemaj. Performative prediction in a stateful world. In *International Conference on Artificial Intelligence and Statistics*, pages 6045–6061, 2022.
- Sébastien Bubeck, Gilles Stoltz, and Jia Yuan Yu. Lipschitz bandits without the Lipschitz constant. In *International Conference on Algorithmic Learning Theory*, pages 144–158, 2011.
- Yiling Chen, Yang Liu, and Chara Podimata. Learning strategy-aware linear classifiers. *Advances in Neural Information Processing Systems*, 33:15265–15276, 2020.
- Moustapha Cisse, Piotr Bojanowski, Edouard Grave, Yann Dauphin, and Nicolas Usunier. Parseval networks: Improving robustness to adversarial examples. In *Proceedings of the 34th International Conference on Machine Learning*, page 854–863, 2017.
- Alon Cohen, Tamir Hazan, and Tomer Koren. Online learning with feedback graphs without the graphs. In *International Conference on Machine Learning*, pages 811–819, 2016.
- Jinshuo Dong, Aaron Roth, Zachary Schutzman, Bo Waggoner, and Zhiwei Steven Wu. Strategic classification from revealed preferences. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 55–70, 2018.
- Roy Dong and Lillian J Ratliff. Approximate regions of attraction in learning with decision-dependent distributions. *arXiv preprint arXiv:2107.00055*, 2021.
- Dmitriy Drusvyatskiy and Lin Xiao. Stochastic optimization with decision-dependent distributions. *arXiv preprint arXiv:2011.11173*, 2020.
- Eyal Even-Dar, Shie Mannor, and Yishay Mansour. PAC bounds for multi-armed bandit and Markov decision processes. In *International Conference on Computational Learning Theory*, pages 255–270, 2002.
- Tanner Fiez and Lillian J Ratliff. Local convergence analysis of gradient descent ascent with finite timescale separation. In *International Conference on Learning Representations*, 2020.
- Tanner Fiez, Benjamin Chasnov, and Lillian Ratliff. Implicit learning dynamics in Stackelberg games: Equilibria characterization, convergence analysis, and empirical study. In *International Conference on Machine Learning*, pages 3133–3144, 2020.
- Moritz Hardt, Nimrod Megiddo, Christos Papadimitriou, and Mary Wootters. Strategic classification. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, pages 111–122, 2016.
- Matthias Hein and Maksym Andriushchenko. Formal guarantees on the robustness of a classifier against adversarial manipulation. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, page 2263–2273, 2017.
- Zachary Izzo, Lexing Ying, and James Zou. How to learn when data reacts to your model: Performative gradient descent. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139, pages 4641–4650, 2021.
- Meena Jagadeesan, Celestine Mendler-Dünner, and Moritz Hardt. Alternative microfoundations for strategic classification. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139, pages 4687–4697. PMLR, 2021.
- Chi Jin, Praneeth Netrapalli, and Michael Jordan. What is local optimality in nonconvex-nonconcave minimax optimization? In *International Conference on Machine Learning*, pages 4880–4889, 2020.
- Abbas Kazerouni, Mohammad Ghavamzadeh, Yasin Abbasi Yadkori, and Benjamin Van Roy. Conservative contextual linear bandits. *Advances in Neural Information Processing Systems*, 30, 2017.
- Robert Kleinberg. Nearly tight bounds for the continuum-armed bandit problem. *Advances in Neural Information Processing Systems*, 17:697–704, 2004.

- Robert Kleinberg, Aleksandrs Slivkins, and Eli Upfal. Multi-armed bandits in metric spaces. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 681–690, 2008.
- Tomáš Kocák, Gergely Neu, Michal Valko, and Remi Munos. Efficient learning by implicit exploration in bandit problems with side observations. In *Advances in Neural Information Processing Systems*, volume 27, 2014.
- Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- Lihong Li, Wei Chu, John Langford, and Robert E Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th International Conference on World Wide Web*, pages 661–670, 2010.
- Qiang Li and Hoi-To Wai. State dependent performative prediction with stochastic approximation. In *International Conference on Artificial Intelligence and Statistics*, pages 3164–3186, 2022.
- Chinmay Maheshwari, Chih-Yuan Chiu, Eric Mazumdar, Shankar Sastry, and Lillian Ratliff. Zeroth-order methods for convex-concave min-max problems: Applications to decision-dependent risk minimization. In *International Conference on Artificial Intelligence and Statistics*, pages 6702–6734, 2022.
- Shie Mannor and Ohad Shamir. From bandits to experts: On the value of side-observations. In *Proceedings of the 24th International Conference on Neural Information Processing Systems*, page 684–692, 2011.
- Celestine Mender-Dünner, Juan Perdomo, Tijana Zrnic, and Moritz Hardt. Stochastic optimization for performative prediction. In *Advances in Neural Information Processing Systems*, volume 33, pages 4929–4939, 2020.
- John P Miller, Juan C Perdomo, and Tijana Zrnic. Outside the echo chamber: Optimizing the performative risk. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139, pages 7710–7720, 2021.
- Juan Perdomo, Tijana Zrnic, Celestine Mender-Dünner, and Moritz Hardt. Performative prediction. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119, pages 7599–7609, 2020.
- Chara Podimata and Alex Slivkins. Adaptive discretization for adversarial Lipschitz bandits. In *Conference on Learning Theory*, pages 3788–3805, 2021.
- Mitas Ray, Dmitriy Drusvyatskiy, Maryam Fazel, and Lillian J Ratliff. Decision-dependent risk minimization in geometrically decaying dynamic environments. In *Proceedings of the Association for the Advancement of Artificial Intelligence Conference on AI (AAAI)*, 2022.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *2nd International Conference on Learning Representations*, 2014.
- Matteo Turchetta, Felix Berkenkamp, and Andreas Krause. Safe exploration for interactive machine learning. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- Yifan Wu, András György, and Csaba Szepesvári. Online learning with Gaussian payoffs and side observations. In *Proceedings of the 28th International Conference on Neural Information Processing Systems*, page 1360–1368, 2015.
- Yifan Wu, Roshan Shariff, Tor Lattimore, and Csaba Szepesvári. Conservative bandits. In *International Conference on Machine Learning*, pages 1254–1262, 2016.
- Yao-Yuan Yang, Cyrus Rashtchian, Hongyang Zhang, Russ R Salakhutdinov, and Kamalika Chaudhuri. A closer look at accuracy vs. robustness. In *Advances in Neural Information Processing Systems*, volume 33, pages 8588–8601, 2020.
- Tijana Zrnic, Eric Mazumdar, Shankar Sastry, and Michael Jordan. Who leads and who follows in strategic classification? *Advances in Neural Information Processing Systems*, 34:15257–15269, 2021.

A. Proofs from Section 2 and Section 3

A.1. Proof of Lemma 2.1

Notice that $\text{PR}(\theta) - \text{PR}(\theta') = (\text{DPR}(\theta, \theta) - \text{DPR}(\theta, \theta')) + (\text{DPR}(\theta, \theta') - \text{DPR}(\theta', \theta'))$. We bound the first difference using Lipschitzness of ℓ in θ as $|\text{DPR}(\theta, \theta) - \text{DPR}(\theta, \theta')| = |\mathbb{E}_{z \sim \mathcal{D}(\theta)}[\ell(z; \theta) - \ell(z; \theta')]| \leq L_\theta \|\theta - \theta'\|$. For the second term we combine Assumption 1.2 and Lipschitzness of ℓ in z via the Kantorovich-Rubinstein duality theorem. In particular, we get $|\text{DPR}(\theta, \theta') - \text{DPR}(\theta', \theta')| = |\mathbb{E}_{z \sim \mathcal{D}(\theta)}\ell(z; \theta') - \mathbb{E}_{z \sim \mathcal{D}(\theta')} \ell(z; \theta')| \leq \epsilon L_z \|\theta - \theta'\|$. Putting both bounds together, we obtain the claimed Lipschitz bound.

A.2. Proof of Proposition 3.1

We construct a γ -cover of the parameter space, denoted \mathcal{S}_γ , and deploy all models in this cover. This gives us access to the distributions $\{\mathcal{D}(\theta) : \theta \in \mathcal{S}_\gamma\}$. Using this information, for any $\theta \in \Theta$ we can compute

$$\widehat{\text{PR}}(\theta) = \text{DPR}(\Pi_{\mathcal{S}_\gamma}(\theta), \theta) = \mathbb{E}_{z \sim \mathcal{D}(\Pi_{\mathcal{S}_\gamma}(\theta))} \ell(z; \theta),$$

where $\Pi_{\mathcal{S}_\gamma}(\theta) := \operatorname{argmin}_{\theta' \in \mathcal{S}_\gamma} \|\theta' - \theta\|$ is the projection onto \mathcal{S}_γ . Note that $\|\theta - \Pi_{\mathcal{S}_\gamma}(\theta)\| \leq \gamma$ all $\theta \in \Theta$ since \mathcal{S}_γ is a cover. Therefore, for any $\theta \in \Theta$, we can bound $\text{PR}(\theta)$ as

$$\begin{aligned} \text{PR}(\theta) &\leq \text{DPR}(\Pi_{\mathcal{S}_\gamma}(\theta), \theta) + L_z \epsilon \|\Pi_{\mathcal{S}_\gamma}(\theta) - \theta\| \\ &\leq \text{DPR}(\Pi_{\mathcal{S}_\gamma}(\theta), \theta) + L_z \epsilon \gamma \\ &= \widehat{\text{PR}}(\theta) + L_z \epsilon \gamma. \end{aligned}$$

Similarly we obtain $\text{PR}(\theta) \geq \widehat{\text{PR}}(\theta) - L_z \epsilon \gamma$, which completes the proof.

A.3. Proof of Proposition 3.2

We will show that $\text{PR}_{\text{LB}}(\theta) \leq \text{PR}(\theta)$ and $\text{PR}_{\text{min}} \geq \text{PR}(\theta_{\text{PO}})$; these two facts immediately imply $\Delta(\theta) := \text{PR}(\theta) - \text{PR}(\theta_{\text{PO}}) \geq \text{PR}_{\text{LB}}(\theta) - \text{PR}_{\text{min}}$.

The first bound follows because $\text{PR}(\theta) = \text{DPR}(\theta, \theta) \geq \text{DPR}(\theta', \theta) - L_z \epsilon \|\theta' - \theta\|$ for all θ' , where we use $(L_z \epsilon)$ -Lipschitzness of DPR in the first argument. Similarly, the second bound follows because

$$\text{PR}(\theta_{\text{PO}}) = \min_{\theta} \text{DPR}(\theta, \theta) \leq \min_{\theta} (\text{DPR}(\theta', \theta) + L_z \epsilon \|\theta - \theta'\|),$$

for all θ' .

B. Regret Analysis of Algorithm 1

In this section, we prove a regret bound for Algorithm 1. At a high-level, Theorem 4.2 combines bounds specific to performative prediction with ingredients from the analysis of successive elimination (Even-Dar et al., 2002). First, using a finite-sample analogue of Proposition 3.2, we show that after phase p all models $\theta \in \mathcal{A}$ have suboptimality $\Delta(\theta) \leq 8\gamma_p$. We then upper bound the number of models in each suboptimality band $\{\theta : 16L_z \epsilon r \leq \Delta(\theta) < 32L_z \epsilon r\}$, for fixed r , that are deployed in each phase, by leveraging the definition of sequential zooming dimension. The remainder of the proof separately analyzes the regret incurred from the first $\log_2(1/(L_z \epsilon))$ phases, in which the finite-sample error dominates the discretization error, and the regret from the later phases, in which the finite-sample error and the discretization error are of the same order.

We use $\text{Reg}_{\text{ph}}(p_1 : p_2)$ to denote the regret incurred from phase p_1 to phase p_2 :

$$\text{Reg}_{\text{ph}}(p_1 : p_2) = \mathbb{E} \sum_{p=p_1}^{p_2} \Delta(\theta_p).$$

We let $\text{Reg}_{\text{ph}}(0 : p) \equiv \text{Reg}_{\text{ph}}(p)$. For phases p that happen after the time horizon T , we assume that the incurred regret is 0; for example, if phases $p_1 \leq p_2$ happen after T , then $\text{Reg}_{\text{ph}}(p_1 : p_2) = 0$.

B.1. Clean Event

First, we define a clean event that guarantees that the estimates $\widehat{\text{DPR}}(\theta, \theta')$ are close to the true values $\text{DPR}(\theta, \theta')$ at all phases. The clean event essentially guarantees uniform convergence over $\widehat{\text{DPR}}(\theta, \cdot)$ for every $\theta \in \mathcal{P}_p$.

Definition B.1 (Clean event). Denote the ‘‘clean event’’ by

$$E_{\text{clean}} = \left\{ \forall p : \sup_{\theta \in \mathcal{P}_p} \sup_{\theta' \in \Theta} \left| \widehat{\text{DPR}}(\theta, \theta') - \text{DPR}(\theta, \theta') \right| \leq \frac{2\mathfrak{C}^*(\ell) + 3\sqrt{\log(T)}}{\sqrt{n_p m_0}} \right\}, \quad (5)$$

where \mathcal{P}_p is the set of all models deployed in phase p during time horizon T .

We show that the clean event occurs with high probability.

Lemma B.2. *The clean event holds with high probability,*

$$\mathbb{P}\{E_{\text{clean}}\} \geq 1 - T^{-2}.$$

Proof. We consider each interval of length n_p in phase p , during which the same model is deployed, separately, and then take a union bound over these intervals across all phases. Therefore, we will say interval s in phase p to refer to steps $(s-1)n_p + 1, \dots, sn_p$ in phase p . For the sake of this proof, we consider a ‘‘counterfactual’’ set of samples for each model θ that augments the set of actually observed samples. In particular, for interval s in phase p , we let $\{z_j^{\theta, s}\}_{j=1}^{n_p m_0}$ denote i.i.d. samples from $\mathcal{D}(\theta)$. The samples for different time intervals and different phases are independent. When model θ is deployed, we observe the samples corresponding to the interval in which θ is deployed.

For each phase p and each time interval s within phase p , let $E_{\text{end}}^{s,p}$ denote the event that phase p terminates strictly before interval s is reached. Let $E_{\text{clean}}^{s,p}$ denote the event that one of the following two holds:

(E1) $E_{\text{end}}^{s,p}$ occurs;

(E2) $E_{\text{end}}^{s,p}$ does not occur, and for the model θ_s deployed in time interval s it holds that:

$$\sup_{\theta' \in \Theta} \left| \widehat{\text{DPR}}(\theta_s, \theta') - \text{DPR}(\theta_s, \theta') \right| \leq \frac{2\mathfrak{C} + 3\sqrt{\log(T)}}{\sqrt{n_p m_0}},$$

where θ_s is a random variable.

The probability that $E_{\text{clean}}^{s,p}$ does not occur is at most:

$$\begin{aligned} & \mathbb{P} \left[\neg E_{\text{end}}^{s,p} \ \& \ \sup_{\theta' \in \Theta} \left| \widehat{\text{DPR}}(\theta_s, \theta') - \text{DPR}(\theta_s, \theta') \right| > \frac{2\mathfrak{C} + 3\sqrt{\log(T)}}{\sqrt{n_p m_0}} \right] \\ &= \mathbb{P} \left[\neg E_{\text{end}}^{s,p} \right] \cdot \mathbb{P} \left[\sup_{\theta' \in \Theta} \left| \widehat{\text{DPR}}(\theta_s, \theta') - \text{DPR}(\theta_s, \theta') \right| > \frac{2\mathfrak{C} + 3\sqrt{\log(T)}}{\sqrt{n_p m_0}} \mid \neg E_{\text{end}}^{s,p} \right] \\ &\leq \mathbb{P} \left[\sup_{\theta' \in \Theta} \left| \widehat{\text{DPR}}(\theta_s, \theta') - \text{DPR}(\theta_s, \theta') \right| > \frac{2\mathfrak{C} + 3\sqrt{\log(T)}}{\sqrt{n_p m_0}} \mid \neg E_{\text{end}}^{s,p} \right]. \end{aligned}$$

We can equivalently write this as

$$\begin{aligned} & \mathbb{P} \left[\sup_{\theta' \in \Theta} \left| \frac{1}{n_p m_0} \sum_{j=1}^{n_p m_0} \ell(z_j^{\theta_s, s}; \theta') - \text{DPR}(\theta_s, \theta') \right| > \frac{2\mathfrak{C} + 3\sqrt{\log(T)}}{\sqrt{n_p m_0}} \mid \neg E_{\text{end}}^{s,p} \right] \\ &= \mathbb{E}_{\theta \sim \theta_s} \left[\mathbb{P} \left[\sup_{\theta' \in \Theta} \left| \frac{1}{n_p m_0} \sum_{j=1}^{n_p m_0} \ell(z_j^{\theta, s}; \theta') - \text{DPR}(\theta, \theta') \right| > \frac{2\mathfrak{C} + 3\sqrt{\log(T)}}{\sqrt{n_p m_0}} \mid \neg E_{\text{end}}^{s,p}, \theta_s = \theta \right] \right]. \end{aligned}$$

To upper bound this expression, it suffices to show an upper bound on

$$\mathbb{P} \left[\sup_{\theta' \in \Theta} \left| \frac{1}{n_p m_0} \sum_{j=1}^{n_p m_0} \ell(z_j^{\theta, s}; \theta') - \text{DPR}(\theta, \theta') \right| > \frac{2\mathfrak{C} + 3\sqrt{\log(T)}}{\sqrt{n_p m_0}} \middle| \neg E_{\text{end}}^{s,p}, \theta_s = \theta \right]$$

that holds for every θ . The first observation is that for any θ , the samples $\{z_j^{\theta, s}\}_{j=1}^{n_p m_0}$ are independent of the event $\{\theta_s = \theta, \neg E_{\text{end}}^{s,p}\}$, since the event depends only on the samples collected in previous time intervals and phases. This means that the above probability is equal to:

$$\mathbb{P} \left[\sup_{\theta' \in \Theta} \left| \frac{1}{n_p m_0} \sum_{j=1}^{n_p m_0} \ell(z_j^{\theta, s}; \theta') - \text{DPR}(\theta, \theta') \right| > \frac{2\mathfrak{C} + 3\sqrt{\log(T)}}{\sqrt{n_p m_0}} \right].$$

Let ϵ_j denote i.i.d. Rademacher random variables. Then, we can observe that with probability $1 - T^{-3}$, it holds that:

$$\begin{aligned} \sup_{\theta' \in \Theta} \left| \frac{1}{n_p m_0} \sum_{j=1}^{n_p m_0} \ell(z_j^{\theta, s}; \theta') - \text{DPR}(\theta, \theta') \right| &\leq \mathbb{E} \left[\sup_{\theta' \in \Theta} \left| \frac{1}{n_p m_0} \sum_{j=1}^{n_p m_0} \ell(z_j^{\theta, s}; \theta') - \text{DPR}(\theta, \theta') \right| \right] + \sqrt{\frac{6 \log(T)}{n_p m_0}} \\ &\leq 2 \cdot \mathbb{E} \left[\sup_{\theta' \in \Theta} \left| \frac{1}{n_p m_0} \sum_{j=1}^{n_p m_0} \ell(z_j^{\theta, s}; \theta') \cdot \epsilon_j \right| \right] + \sqrt{\frac{6 \log(T)}{n_p m_0}} \\ &\leq \frac{2}{\sqrt{n_p m_0}} \cdot \sup_{n \geq 1} \sqrt{n} \mathbb{E} \left[\sup_{\theta' \in \Theta} \left| \frac{1}{n} \sum_{j=1}^n \ell(z_j^{\theta}; \theta') \cdot \epsilon_j \right| \right] + \sqrt{\frac{6 \log(T)}{n_p m_0}} \\ &\leq \frac{2\mathfrak{C}^*(\ell) + 3\sqrt{\log(T)}}{\sqrt{n_p m_0}}, \end{aligned}$$

where the first step follows from the bounded differences inequality and the second step follows from a classical symmetrization argument. In the penultimate step we let $\{z_j^\theta\}_{j \in \mathbb{N}}$ denote an infinite sequence of samples from $\mathcal{D}(\theta)$. Putting this all together, we have that:

$$1 - \mathbb{P}[E_{\text{clean}}^{s,p}] \leq T^{-3}.$$

Finally, using that there are at most T intervals before time horizon T (across all phases), by a union bound we see that:

$$1 - \mathbb{P}[E_{\text{clean}}] \leq T^{-2},$$

as desired. □

B.2. Suboptimality of the Active Set

We show that the elimination strategy in Algorithm 1 will never eliminate any performatively optimal point.

Lemma B.3. *On the clean event (5), any performatively optimal point $\theta_{\text{PO}} \in \text{argmin}_\theta \text{PR}(\theta)$ will always remain in \mathcal{A} .*

Proof. It suffices to show that θ_{PO} cannot be eliminated in Step 14 of Algorithm 1. Fix any phase p and denote by \mathcal{P}_p the

running set of deployed points at any point during phase p . Then, we have:

$$\begin{aligned}
 \text{PR}_{\text{LB}}(\theta_{\text{PO}}) &= \max_{\theta' \in \mathcal{P}_p} \left(\widehat{\text{DPR}}(\theta', \theta_{\text{PO}}) - L_z \epsilon \|\theta_{\text{PO}} - \theta'\| \right) \\
 &\leq \max_{\theta' \in \mathcal{P}_p} \left(\text{DPR}(\theta', \theta_{\text{PO}}) - L_z \epsilon \|\theta_{\text{PO}} - \theta'\| \right) + \gamma_p \\
 &\leq \text{PR}(\theta_{\text{PO}}) + \gamma_p \\
 &= \min_{\theta} \text{DPR}(\theta, \theta) + \gamma_p \\
 &\leq \min_{\theta} \min_{\theta' \in \mathcal{P}_p} \text{DPR}(\theta', \theta) + L_z \epsilon \|\theta - \theta'\| + \gamma_p \\
 &\leq \min_{\theta} \min_{\theta' \in \mathcal{P}_p} \widehat{\text{DPR}}(\theta', \theta) + L_z \epsilon \|\theta' - \theta\| + 2\gamma_p \\
 &= \text{PR}_{\min} + 2\gamma_p.
 \end{aligned}$$

Therefore, $\text{PR}_{\text{LB}}(\theta_{\text{PO}}) \leq \text{PR}_{\min} + 2\gamma_p$, implying that θ_{PO} cannot be removed from \mathcal{A} during phase p . Since this is true for any phase p , that completes the proof of the lemma. \square

We next show that the elimination strategy is sufficiently effective that all models that remain active after a given phase p have suboptimality at most $8\gamma_p$.

Lemma B.4. *On the clean event (5), after phase p all models $\theta \in \mathcal{A}$ satisfy $\Delta(\theta) \leq 8\gamma_p$.*

Proof. Fix a phase p . We will analyze \mathcal{P}_p at the end of phase p . The proof relies on two key facts:

- (F1) If θ is active after phase p , then $\|\theta - \Pi_{\mathcal{P}_p}(\theta)\| \leq r_p$, where $\Pi_{\mathcal{P}_p}(\theta) = \arg\min_{\theta' \in \mathcal{P}_p} \|\theta - \theta'\|$.
- (F2) θ_{PO} is active after phase p .

The first fact follows since during phase p net points cannot be eliminated from \mathcal{S}_p in Step 13 while some parameter within an r_p -neighborhood is active. The second fact is proved in Lemma B.3. Note that from fact (F1) it further follows that there is always a model in \mathcal{P}_p within the r_p -neighborhood of θ_{PO} .

Now suppose that θ is active after phase p . Then, we have:

$$\begin{aligned}
 \text{PR}(\theta) &\leq \text{DPR}(\Pi_{\mathcal{P}_p}(\theta), \theta) + L_z \epsilon \|\Pi_{\mathcal{P}_p}(\theta) - \theta\| \\
 &\leq \widehat{\text{DPR}}(\Pi_{\mathcal{P}_p}(\theta), \theta) + L_z \epsilon \|\Pi_{\mathcal{P}_p}(\theta) - \theta\| + \gamma_p \\
 &\leq \min_{\theta'} \left(\widehat{\text{DPR}}(\Pi_{\mathcal{P}_p}(\theta'), \theta') + L_z \epsilon \|\Pi_{\mathcal{P}_p}(\theta') - \theta'\| \right) + 2L_z \epsilon \|\Pi_{\mathcal{P}_p}(\theta) - \theta\| + 3\gamma_p,
 \end{aligned}$$

where we used the definitions of PR_{\min} and $\text{PR}_{\text{LB}}(\theta)$, together with the fact that $\text{PR}_{\text{LB}}(\theta) \leq \text{PR}_{\min} + 2\gamma_p$ for active models. Now choosing $\theta' = \theta_{\text{PO}}$, applying (F1), (F2), and accounting for finite-sample uncertainty we find

$$\begin{aligned}
 \text{PR}(\theta) &\leq \widehat{\text{DPR}}(\Pi_{\mathcal{P}_p}(\theta_{\text{PO}}), \theta_{\text{PO}}) + L_z \epsilon \|\Pi_{\mathcal{P}_p}(\theta_{\text{PO}}) - \theta_{\text{PO}}\| + 2L_z \epsilon \|\Pi_{\mathcal{P}_p}(\theta) - \theta\| + 3\gamma_p \\
 &\leq \widehat{\text{DPR}}(\Pi_{\mathcal{P}_p}(\theta_{\text{PO}}), \theta_{\text{PO}}) + 3L_z \epsilon r_p + 3\gamma_p \\
 &\leq \text{DPR}(\theta_{\text{PO}}, \theta_{\text{PO}}) + L_z \epsilon \|\Pi_{\mathcal{P}_p}(\theta_{\text{PO}}) - \theta_{\text{PO}}\| + 3L_z \epsilon r_p + 4\gamma_p \\
 &\leq \text{PR}(\theta_{\text{PO}}) + 4(L_z \epsilon r_p + \gamma_p) \\
 &= \text{PR}(\theta_{\text{PO}}) + 8\gamma_p,
 \end{aligned}$$

where we use the fact that $r_p = \frac{\gamma_p}{L_z \epsilon}$. Rearranging the terms we obtain $\Delta(\theta) = \text{PR}(\theta) - \text{PR}(\theta_{\text{PO}}) \leq 8\gamma_p$ as claimed in Lemma B.4. \square

B.3. Bounding the Number of Suboptimal Deployments

For $i \geq 1$, we consider the suboptimality bands

$$\mathcal{E}_i = \{\theta : \Delta(\theta) \in [8 \cdot 2^{-i} L_z \epsilon, 16 \cdot 2^{-i} L_z \epsilon]\}.$$

In the following lemma, we bound the number of times that models in \mathcal{E}_i can be deployed in a given phase.

Lemma B.5. *Suppose that the clean event (5) holds. For $i \geq 1$, in phase $\log_2(1/(L_z \epsilon)) \leq p \leq \log_2(1/(L_z \epsilon)) + i + 1$, the number of models in \mathcal{E}_i that are deployed is at most $O\left(\frac{3}{r_p}\right)^d$ in expectation, where d is the $(L_z \epsilon)$ -sequential zooming dimension.*

To provide intuition for Lemma B.5, it is informative to consider a weaker version of the lemma where d is taken to be the $(L_z \epsilon)$ -zooming dimension rather than the $(L_z \epsilon)$ -sequential zooming dimension. To see why this weaker version of the lemma is true, notice that at the beginning of phase p , the set of active models \mathcal{A} is a subset of $\{\theta : \Delta(\theta) \leq 8\gamma_{p-1}\} = \{\theta : \Delta(\theta) \leq 16\gamma_p\} = \{\theta : \Delta(\theta) \leq 16L_z \epsilon r_p\}$. The set of models deployed in phase p is contained in a minimal r_p -net of \mathcal{A} .

Notice that $r_p \geq 2^{-(i+1)}$. By the definition of zooming dimension, we know that at most a multiple of $\left(\frac{3}{r_p}\right)^d$ elements from the set $\{\theta : \Delta(\theta) \in [8 \cdot 2^{-i} L_z \epsilon, 16 \cdot 2^{-i} L_z \epsilon]\} = \{\theta : \Delta(\theta) \in [16 \cdot 2^{-(i+1)} L_z \epsilon, 32 \cdot 2^{-(i+1)} L_z \epsilon]\}$ are deployed, as desired.

The proof of Lemma B.5 boils down to refining this proof sketch to account for the sequential elimination aspect of Algorithm 1.

Proof. For the purposes of this analysis, we condition on the clean event.

Fix a phase $\log_2(1/(L_z \epsilon)) \leq p \leq \log_2(1/(L_z \epsilon)) + i + 1$. Let \mathcal{S}_p^0 be the covering of \mathcal{A} chosen at the beginning of phase p , and let π be an ordering of \mathcal{S}_p^0 chosen uniformly at random. It is not difficult to see that Algorithm 1 is equivalent to drawing π at the beginning of the phase, and deploying models in the order given by π (naturally, skipping those that get eliminated). For technical convenience, we analyze this reformulation of the algorithm.

Condition on a realization π , and let $\mathcal{P}_p \subseteq \mathcal{S}_p^0$ be the set of models that are ultimately get deployed. Note that \mathcal{P}_p depends on the randomness arising from finite-sample noise at each step of the phase. We will show a bound on $|\mathcal{P}_p|$ that deterministically holds on the clean event. In particular, consider the models $\theta \in \mathcal{S}_p^0 \cap \{\theta : 8L_z \epsilon r_i \leq \Delta(\theta) < 16L_z \epsilon r_i\}$ such that:

$$\text{PR}_{\text{LB}}^{r_p}(\pi(\theta)) \leq \text{PR}_{\min}(\pi(\theta)) + 4L_z \epsilon r_p = \text{PR}_{\min}(\pi(\theta)) + 4\gamma_p. \quad (6)$$

We will show that \mathcal{P}_p is a subset of such models.

Suppose that $\theta_{\text{net}} \in \mathcal{S}_p^0$ is deployed in phase p . Then, that means that there exists $\theta'' \in \text{Ball}_{r_p}(\theta_{\text{net}})$ that remains active after the first $\pi(\theta_{\text{net}}) - 1$ deployments; that is:

$$\begin{aligned} \max_{\theta' : \pi(\theta') < \pi(\theta_{\text{net}})} (\widehat{\text{DPR}}(\theta', \theta'') - L_z \epsilon \|\theta' - \theta''\|) &= \text{PR}_{\text{LB}}(\theta'') \\ &\leq \text{PR}_{\min} + 2\gamma_p \\ &= \min_{\theta} \min_{\theta' : \pi(\theta') < \pi(\theta_{\text{net}})} (\widehat{\text{DPR}}(\theta', \theta) + L_z \epsilon \|\theta' - \theta\|) + 2\gamma_p. \end{aligned}$$

Since the clean event holds, we know that:

$$\max_{\theta' : \pi(\theta') < \pi(\theta_{\text{net}})} (\text{DPR}(\theta', \theta'') - L_z \epsilon \|\theta' - \theta''\|) - \gamma_p \leq \min_{\theta} \min_{\theta' : \pi(\theta') < \pi(\theta_{\text{net}})} (\text{DPR}(\theta', \theta) + L_z \epsilon \|\theta' - \theta\|) + 3\gamma_p.$$

Rearranging, this means that:

$$\begin{aligned} \max_{\theta' : \pi(\theta') < \pi(\theta_{\text{net}})} (\text{DPR}(\theta', \theta'') - L_z \epsilon \|\theta' - \theta''\|) \\ \leq \min_{\theta} \min_{\theta' : \pi(\theta') < \pi(\theta_{\text{net}})} (\text{DPR}(\theta', \theta) + L_z \epsilon \|\theta' - \theta\|) + 4\gamma_p = \text{PR}_{\min}(\pi(\theta_{\text{net}})) + 4\gamma_p. \end{aligned}$$

This further implies that:

$$\text{PR}_{\text{LB}}^{r_p}(\pi(\theta_{\text{net}})) = \min_{\theta'' \in \text{Ball}_{r_p}(\theta_{\text{net}})} \max_{\theta': \pi(\theta') < \pi(\theta_{\text{net}})} (\text{DPR}(\theta', \theta'') - L_z \epsilon \|\theta' - \theta''\|) \leq \text{PR}_{\min}(\pi(\theta_{\text{net}})) + 4\gamma_p.$$

We see that any $\theta_{\text{net}} \in \mathcal{P}_p$ must satisfy condition (6). By the definition of sequential zooming dimension, we know that the expected number of models in \mathcal{E}_i that satisfy (6), where the expectation is taken over the randomness of π , is at most a multiple of $\left(\frac{3}{r_p}\right)^d$, hence $\mathbb{E}|\mathcal{P}_p \cap \mathcal{E}_i| \leq \mathcal{O}\left(\left(\frac{3}{r_p}\right)^d\right)$, as desired. \square

B.4. Regret Bound on the Clean Event

To bound the regret on the clean event, we break the analysis into two cases: (a) the first $\log_2(1/(L_z \epsilon))$ phases, and (b) all remaining phases.

Lemma B.6. *Suppose that the clean event (5) holds. In the first $\lfloor \log_2(1/(L_z \epsilon)) \rfloor$ phases, the algorithm has incurred regret at most*

$$\text{Reg}_{\text{ph}}(\lfloor \log_2(1/(L_z \epsilon)) \rfloor) = \mathcal{O}\left(\sqrt{\frac{T}{m_0}} (\sqrt{\log T} + \mathfrak{C})\right).$$

Proof. During phases $p \leq \log_2(1/(L_z \epsilon))$, we deploy a single model since $r_p \geq 1$ and Θ is assumed to have radius 1.

We break the first $\lfloor \log_2(1/(L_z \epsilon)) \rfloor$ phases into two cases. For a value of $N \geq 0$ specified later, we consider cases $p < N$ and $p \geq N$ separately.

Case 1: phases $N \leq p \leq \lfloor \log_2(1/(L_z \epsilon)) \rfloor$. By Lemma B.4, we see that the model deployed in phase N must have suboptimality at most $8 \cdot 2^{-N+1} = 2^{-N+4}$. Since the algorithm runs for at most T time steps, this means that the total regret incurred in these phases is at most $T \cdot 2^{-N+4}$.

Case 2: phases $0 \leq p < \min\{N, \lfloor \log_2(1/(L_z \epsilon)) \rfloor\}$.

By Lemma B.4, we know that the model deployed in phase p must have suboptimality at most $8 \cdot 2^{-p+1} = 2^{-p+4}$. Moreover, this model is deployed for $n_p = \left\lfloor \frac{(2\mathfrak{C} + 3\sqrt{\log T})^2}{\gamma_p^2 m_0} \right\rfloor$ steps. The regret incurred up to phase N can thus be bounded as:

$$\begin{aligned} \text{Reg}_{\text{ph}}(N) &\leq \sum_{p=0}^{N-1} n_p 2^{-p+4} \\ &\leq 16 \sum_{p=0}^{N-1} 2^{-p} \left\lfloor \frac{2^{2p} (2\mathfrak{C} + 3\sqrt{\log T})^2}{m_0} \right\rfloor. \end{aligned}$$

Since we assume $m_0 = o((\mathfrak{C} + \sqrt{\log T})^2)$, for a large enough T we have $n_p \geq 1$ and thus $\lfloor n_p \rfloor \leq 2n_p$. Therefore,

$$\begin{aligned} \text{Reg}_{\text{ph}}(N) &\leq C \sum_{p=0}^{N-1} 2^{-p} \frac{2^{2p} (2\mathfrak{C} + 3\sqrt{\log T})^2}{m_0} \\ &\leq C \frac{(2\mathfrak{C} + 3\sqrt{\log T})^2}{m_0} \left(\sum_{p=0}^{N-1} 2^p \right) \\ &\leq C \cdot 2^N \frac{(2\mathfrak{C} + 3\sqrt{\log T})^2}{m_0}, \end{aligned}$$

for some large enough constant $C > 0$.

Putting the two cases together, on the clean event, the total regret incurred in phases $p = 0, \dots, \lfloor \log_2(1/(L_z \epsilon)) \rfloor$ can be upper bounded by

$$C \cdot 2^N \frac{(2\mathfrak{C} + 3\sqrt{\log T})^2}{m_0} + T \cdot 2^{-N+4}.$$

We can also trivially upper bound the regret by T , using the fact that the loss incurred at each step is at most 1. This means that we obtain a regret bound of:

$$\mathcal{O}\left(\min\left\{T, 2^N \frac{(2\mathfrak{C} + 3\sqrt{\log T})^2}{m_0} + T \cdot 2^{-N+4}\right\}\right).$$

We now choose N to minimize this bound. We let $\eta = 2^{-N}$ and optimize over $\eta \in (0, 1)$. Optimizing over η instead of an integral value of N changes the bound by constant factors at most. This means that we can upper bound the regret by:

$$\mathcal{O}\left(\min_{0 < \eta \leq 1} \min\left\{T, \eta^{-1} \frac{(2\mathfrak{C} + 3\sqrt{\log T})^2}{m_0} + T\eta\right\}\right).$$

If $\eta > 1$, then the minimum of the two terms would be T , which is at least as big as the above expression. Therefore, we can upper bound the above expression by:

$$\mathcal{O}\left(\min_{\eta > 0} \left(\eta^{-1} \frac{(2\mathfrak{C} + 3\sqrt{\log T})^2}{m_0} + T\eta\right)\right).$$

We set $\eta = \frac{3\sqrt{\log T} + 2\mathfrak{C}}{\sqrt{m_0 T}}$ and obtain a regret bound of:

$$\text{Reg}_{\text{ph}}(\lfloor \log_2(1/(L_z \epsilon)) \rfloor) = \mathcal{O}\left(\sqrt{\frac{T}{m_0}} (\sqrt{\log T} + \mathfrak{C})\right),$$

as desired. \square

Lemma B.7. *Suppose that the clean event (5) holds. Let $d \geq 0$ be such that for every $i \geq 0$ and every phase $p \in [\log_2(1/(L_z \epsilon)), \log_2(1/(L_z \epsilon)) + i + 1]$, the number of models in $\mathcal{E}_i = \{\theta : \Delta(\theta) \in [2^{-i+3}L_z \epsilon, 2^{-i+4}L_z \epsilon]\}$ that are deployed in phase p is upper bounded by $\mathcal{O}\left(\left(\frac{3}{r_p}\right)^d\right)$ in expectation. Then, the regret incurred in phases $p \geq \log_2(1/(L_z \epsilon))$, within time horizon T , can be upper bounded as*

$$\text{Reg}_{\text{ph}}(\lceil \log_2(1/(L_z \epsilon)) \rceil : \infty) \leq \mathcal{O}\left(T^{\frac{d+1}{d+2}} (L_z \epsilon)^{\frac{d}{d+2}} \left(\frac{(\sqrt{\log T} + \mathfrak{C})^2}{m_0}\right)^{\frac{1}{d+2}}\right).$$

Proof. By Lemma B.4, we see that all models θ that are active in phase $p = \lceil \log_2(1/(L_z \epsilon)) \rceil$ or later have $\Delta(\theta) \leq 8L_z \epsilon r_p \leq 8L_z \epsilon$. We split these models into suboptimality bands and define, for each $i \geq 1$, the set:

$$\mathcal{E}_i = \{\theta : \Delta(\theta) \in [8 \cdot 2^{-i}L_z \epsilon, 16 \cdot 2^{-i}L_z \epsilon]\}.$$

Note that all models deployed starting with phase $\lceil \log_2(1/(L_z \epsilon)) \rceil$ are in $\cup_{i \geq 1} \mathcal{E}_i$. For a value of N specified later, we break the analysis into two cases.

Case 1: models in $\cup_{i > N} \mathcal{E}_i$. Since the algorithm runs for at most T time steps, the total regret incurred due to deploying models in $\cup_{i > N} \mathcal{E}_i$ is at most

$$T \cdot 16 \cdot 2^{-N-1} L_z \epsilon \leq 8T 2^{-N} L_z \epsilon.$$

Case 2: models in $\cup_{1 \leq i \leq N} \mathcal{E}_i$. By Lemma B.4, we know that all models θ that are active in phases $p \geq N + \log_2(1/(L_z \epsilon))$ have $\Delta(\theta) \leq 82^{-p} = 8 \cdot 2^{-N} L_z \epsilon = 16 \cdot 2^{-N-1} L_z \epsilon$. This means that all models that are active after phase $N + \log_2(1/(L_z \epsilon))$ are in $\cup_{i > N} \mathcal{E}_i$. Thus, to bound the regret incurred by deploying models in $\cup_{1 \leq i \leq N} \mathcal{E}_i$ in phase $\lceil \log_2(1/(L_z \epsilon)) \rceil$ or later, we only need to consider phases $p = \lceil \log_2(1/(L_z \epsilon)) \rceil, \dots, N + \log_2(1/(L_z \epsilon))$.

For $1 \leq i \leq N$, consider \mathcal{E}_i . By Lemma B.4, we know that any $\theta \in \mathcal{E}_i$ can only be active during phases $p \leq \log_2(1/L_z\epsilon) + i + 1$. By assumption, in phase p , the number of points in \mathcal{E}_i that are deployed is at most of the order $\left(\frac{3}{r_p}\right)^d$ in expectation. Moreover, each point is deployed n_p times. Putting this all together, the expected number of points in \mathcal{E}_i deployed in phase p is at most:

$$\mathcal{O}\left(\left(\frac{3}{r_p}\right)^d n_p\right) = \mathcal{O}\left(\left(\frac{3}{r_p}\right)^d \frac{(2\mathfrak{C} + 3\sqrt{\log T})^2}{L_z^2 \epsilon^2 r_p^2 m_0}\right),$$

where we use the fact that, given the condition $m_0 = o((\mathfrak{C} + \sqrt{\log T})^2)$, $n_p \geq 1$ for large enough T and hence we can bound $\lceil n_p \rceil \leq 2n_p$. Take $p = j + \log_2(1/L_z\epsilon)$; then, $r_p = 2^{-j}$. We sum over phases $\log_2(1/(L_z\epsilon)) \leq p \leq \log_2(1/(L_z\epsilon)) + i + 1$ to obtain that in expectation, the total number of times that these models are deployed is at most:

$$\mathcal{O}\left(\frac{3^d (2\mathfrak{C} + 3\sqrt{\log T})^2}{L_z^2 \epsilon^2 m_0} \sum_{j=0}^{i+1} 2^{j(d+2)}\right) = \mathcal{O}\left(\frac{3^d (2\mathfrak{C} + 3\sqrt{\log T})^2}{L_z^2 \epsilon^2 m_0} 2^{(i+1)(d+2)}\right).$$

Using the fact that the models have suboptimality at most $16 \cdot 2^{-i} L_z \epsilon = 32 \cdot 2^{-(i+1)} L_z \epsilon$, we see that the regret incurred by deploying models in \mathcal{E}_i is upper bounded by:

$$\mathcal{O}\left(\frac{3^d (2\mathfrak{C} + 3\sqrt{\log T})^2}{L_z \epsilon m_0} 2^{(i+1)(d+1)}\right).$$

We sum over $1 \leq i \leq N$ to obtain the total regret incurred due to deploying models in $\cup_{1 \leq i \leq N} \mathcal{E}_i$:

$$\mathcal{O}\left(\frac{3^d (2\mathfrak{C} + 3\sqrt{\log T})^2}{L_z \epsilon m_0} 2^{(N+2)(d+1)}\right).$$

Putting together the two cases we obtain a total regret bound of

$$\mathcal{O}\left(\frac{3^d (2\mathfrak{C} + 3\sqrt{\log T})^2}{L_z \epsilon m_0} 2^{(N+2)(d+1)} + T 2^{-N} L_z \epsilon\right).$$

We also can upper bound the regret by $8T L_z \epsilon$, since all models active after phase $\lceil \log_2(1/(L_z\epsilon)) \rceil$ have $\Delta(\theta) \leq 8L_z \epsilon$ and there are at most T time steps in total. This means that we can bound the regret by:

$$\mathcal{O}\left(\min\left\{T L_z \epsilon, \frac{3^d (2\mathfrak{C} + 3\sqrt{\log T})^2}{L_z \epsilon m_0} 2^{(N+2)(d+1)} + T 2^{-N} L_z \epsilon\right\}\right).$$

We now choose N to minimize this bound. We let $\eta = 2^{-N}$ and choose some $\eta \in (0, 1)$. The error from optimizing over $\eta \in (0, 1)$ instead of an integral value of N contributes at most constant factors. This means that we can upper bound the regret by:

$$\mathcal{O}\left(\min\left\{T L_z \epsilon, \frac{12^d (2\mathfrak{C} + 3\sqrt{\log T})^2}{L_z \epsilon m_0} \eta^{-(d+1)} + T \eta L_z \epsilon\right\}\right),$$

for any $\eta \in (0, 1)$. Note that, if $\eta \geq 1$, the second term in the bound is at least as large as the first term, hence we can choose any $\eta > 0$. In particular, we can further upper bound the regret by

$$\mathcal{O}\left(\min_{\eta > 0} \left(\frac{12^d (2\mathfrak{C} + 3\sqrt{\log T})^2}{L_z \epsilon m_0} \eta^{-(d+1)} + T \eta L_z \epsilon\right)\right).$$

Now, we set

$$\eta = \left(\frac{12^d (3\sqrt{\log T} + 2\mathfrak{C})^2}{T L_z^2 \epsilon^2 m_0}\right)^{\frac{1}{d+2}}.$$

Thus, we finally get a regret bound of

$$\mathcal{O} \left(T^{\frac{d+1}{d+2}} (L_z \epsilon)^{\frac{d}{d+2}} \left(\frac{(\sqrt{\log T} + \mathcal{C})^2}{m_0} \right)^{\frac{1}{d+2}} \right),$$

as desired. \square

B.5. Proof of Theorem 4.2

Now, we are ready to prove Theorem 4.2.

First, we handle the case where the clean event defined in (5) does not hold and the concentration bound is violated. By Lemma B.2, this happens with probability at most T^{-2} . The regret incurred in each deployment is at most 1 and there are T deployments, so these events contribute a negligible factor T^{-1} to the expected regret.

For the case where the clean event holds we can build on Lemma B.5, Lemma B.6, and Lemma B.7. From Lemma B.6, we obtain a bound for the total regret incurred in phases up to $\lfloor \log_2(1/(L_z \epsilon)) \rfloor$. By Lemma B.5 we can set the parameter d in Lemma B.7 to be the $(L_z \epsilon)$ -sequential zooming dimension, and thus from Lemma B.7 we obtain a regret bound for all later phases.

Putting all this together yields the desired bound.

C. Location Family Example: Strategic Classification

Location families arise in strategic classification (Hardt et al., 2016), where agents strategically manipulate their features in response to a deployed model. Suppose the learner uses a linear predictor $f_\theta(x) = \theta^\top x$ and the agents incur quadratic cost for changing their original features x to manipulated features x' , $c(x, x') = \frac{1}{2}(x - x')^\top \Lambda (x - x')$. Then, the best response of an agent, typically modeled as $x_{\text{BR}}(\theta) = \operatorname{argmax}_{x'} f_\theta(x') - c(x, x')$, satisfies the location family structural assumption with z_0 being the agent's original features and $\mu^* = \Lambda^{-1}$.

D. Regret Analysis of Algorithm 2

The proof of Theorem 5.1 relies on two key lemmas. One proves that \mathcal{C}_t are valid confidence sets for μ_* at every step, and the other one proves a regret bound assuming that \mathcal{C}_t are valid confidence sets.

Throughout we denote by \mathcal{B}_m the unit ball in \mathbb{R}^m . For a vector x and matrix M , we will use the notation $\|x\|_M = \sqrt{x^\top M x}$.

An important object in the proofs will be $S_t := \sum_{i=1}^t \theta_i \bar{z}_{0,i}^\top$, where $\bar{z}_{0,i} = \frac{1}{m_0} \sum_{j=1}^{m_0} z_i^{(j)} - \mu_*^\top \theta_i$. Essentially $\bar{z}_{0,i}$ is the average over m_0 samples from \mathcal{D}_0 , collected at step i . We will also denote $V_t(\lambda) = (\lambda I + \sum_{i=1}^t \theta_i \theta_i^\top)$, for an arbitrary offset $\lambda > 0$, and $V_t \equiv V_t(0)$. Note that in the algorithm statement we use $\Sigma_t = V_t\left(\frac{1}{m_0}\right)$.

D.1. Clean Event

As for Algorithm 1, we introduce a clean event. In this case, the clean event will be defined as

$$E_{\text{clean}} = \{\forall t \in \mathbb{N} : \mu_* \in \mathcal{C}_t\}, \quad (7)$$

where \mathcal{C}_t are the confidence sets constructed in Algorithm 2.

The technical subtlety lies in the fact that the points θ_t are chosen adaptively, hence one cannot simply apply standard least-squares confidence intervals to argue that the sets \mathcal{C}_t are valid. The same difficulty is resolved in the analysis of the LinUCB algorithm and our proof builds on the proof technique of that analysis.

Before stating the main technical lemma, we start with an auxiliary result that we will use in the proof.

Lemma D.1. *Suppose that \mathcal{D}_0 is 1-subgaussian. Then, for all $x \in \mathcal{B}_m$ and $y \in \mathbb{R}^{d_\theta}$, the process*

$$M_t(x, y) = \exp \left(y^\top S_t x - \frac{1}{2m_0} \|y\|_{V_t}^2 \right)$$

is a supermartingale with respect to the natural filtration, with $M_0(x, y) = 1$.

Proof. Since $\bar{z}_{0,i}$ are $\frac{1}{\sqrt{m_0}}$ -subgaussian, we know that all one-dimensional projections are also $\frac{1}{\sqrt{m_0}}$ -subgaussian, hence $\bar{z}_{0,i}^\top x$ are independent $\frac{1}{\sqrt{m_0}}$ -subgaussian as well. Using this, we know

$$\mathbb{E} \left[\exp(y^\top \theta_t z_{0,t}^\top x) \mid \mathcal{F}_{t-1} \right] \leq \exp \left(\frac{(y^\top \theta_t)^2}{2m_0} \right) = \exp \left(\frac{\|y\|_{\theta_t \theta_t^\top}^2}{2m_0} \right)$$

almost surely. Hence,

$$\begin{aligned} \mathbb{E}[M_t(x, y) \mid \mathcal{F}_{t-1}] &= \mathbb{E} \left[\exp \left(y^\top S_t x - \frac{1}{2m_0} \|y\|_{V_t}^2 \right) \mid \mathcal{F}_{t-1} \right] \\ &= M_{t-1}(x, y) \mathbb{E} \left[\exp \left(y^\top \theta_t z_{0,t}^\top x - \frac{1}{2m_0} \|y\|_{\theta_t \theta_t^\top}^2 \right) \mid \mathcal{F}_{t-1} \right] \\ &\leq M_{t-1}(x, y) \end{aligned}$$

almost surely. Furthermore, $M_0(x, y) = 1$ is trivially true. \square

Now we are ready to state the main technical lemma about the validity of \mathcal{C}_t .

Lemma D.2. *We have that*

$$\mathbb{P}\{E_{\text{clean}}\} \geq 1 - T^{-2}.$$

Proof. First we will show that for any $\delta \in (0, 1)$,

$$\mathbb{P} \left\{ \exists t \in \mathbb{N} : \|V_t(\lambda)^{-1/2} S_t\|^2 \geq \frac{1}{m_0} \left(8m + 4 \log \left(\frac{1}{\delta} \right) + 2 \log \left(\frac{\det(V_t(\lambda))}{\lambda^{d_\Theta}} \right) \right) \right\} \leq \delta, \quad (8)$$

for all $\lambda > 0$.

Let $\Sigma = \frac{m_0}{\lambda} I \in \mathbb{R}^{d_\Theta \times d_\Theta}$ and let h be the density of $\mathcal{N}(0, \Sigma)$. Then, for any fixed $x \in \mathcal{B}_m$ and $M_t(x, y)$ as in Lemma D.1, define

$$\bar{M}_t(x) = \int_{\mathbb{R}^{d_\Theta}} M_t(x, y) h(y) = \frac{1}{\sqrt{(2\pi)^{d_\Theta} \det(\Sigma)}} \int_{\mathbb{R}^{d_\Theta}} \exp \left(y^\top S_t x - \frac{1}{2m_0} \|y\|_{V_t}^2 - \frac{1}{2} \|y\|_{\Sigma^{-1}}^2 \right) dy.$$

Notice that we can write

$$y^\top S_t x - \frac{1}{2m_0} \|y\|_{V_t}^2 - \frac{1}{2} \|y\|_{\Sigma^{-1}}^2 = \frac{1}{2} \|S_t x\|_{(\Sigma^{-1} + \frac{V_t}{m_0})^{-1}}^2 - \frac{1}{2} \left\| y - \left(\Sigma^{-1} + \frac{V_t}{m_0} \right)^{-1} S_t x \right\|_{\Sigma^{-1} + \frac{V_t}{m_0}}^2.$$

Thus, by integrating out the Gaussian density, we get

$$\begin{aligned} \bar{M}_t(x) &= \exp \left(\frac{1}{2} \|S_t x\|_{(\Sigma^{-1} + \frac{V_t}{m_0})^{-1}}^2 \right) \frac{1}{\sqrt{(2\pi)^{d_\Theta} \det(\Sigma)}} \int_{\mathbb{R}^{d_\Theta}} \exp \left(-\frac{1}{2} \left\| y - \left(\Sigma^{-1} + \frac{V_t}{m_0} \right)^{-1} S_t x \right\|_{\Sigma^{-1} + \frac{V_t}{m_0}}^2 \right) dy \\ &= \exp \left(\frac{1}{2} \|S_t x\|_{(\Sigma^{-1} + \frac{V_t}{m_0})^{-1}}^2 \right) \left(\frac{\det((\Sigma^{-1} + \frac{V_t}{m_0})^{-1})}{\det(\Sigma)} \right)^{1/2} \\ &= \exp \left(\frac{m_0}{2} \|V_t^{-1/2}(\lambda) S_t x\|^2 \right) \left(\frac{\lambda^{d_\Theta}}{\det(V_t(\lambda))} \right)^{1/2}. \end{aligned}$$

Now, by Lemma 20.3 in (Lattimore and Szepesvári, 2020), since $M_t(x, y)$ is a supermartingale then $\bar{M}_t(x)$ is a non-negative supermartingale with $\bar{M}_0(x) = 1$. Thus, we can apply the maximal inequality to get

$$\begin{aligned} & \mathbb{P}\left\{\exists t \in \mathbb{N} : \log \bar{M}_t(x) \geq \log(1/\delta)\right\} \\ &= \mathbb{P}\left\{\exists t \in \mathbb{N} : \frac{m_0}{2} \|V_t(\lambda)^{-1/2} S_t x\|^2 - \frac{1}{2} \log\left(\frac{\det(V_t(\lambda))}{\lambda^{d_\Theta}}\right) \geq \log(1/\delta)\right\} \leq \delta. \end{aligned} \quad (9)$$

Inequality (9) is valid for all fixed $x \in \mathcal{B}_m$; to prove inequality (8), we use a covering argument. Let $N_{\frac{1}{2}, m}$ denote a $\frac{1}{2}$ -net of \mathcal{B}_m , and note that we can make $|N_{\frac{1}{2}, m}| \leq 5^m$. Then,

$$\|V_t(\lambda)^{-1/2} S_t\| = \max_{x \in \mathcal{B}_m} \|V_t(\lambda)^{-1/2} S_t x\| \leq 2 \max_{x \in N_{\frac{1}{2}, m}} \|V_t(\lambda)^{-1/2} S_t x\|.$$

Therefore, we can apply a union bound to conclude that for all $s > 0$,

$$\begin{aligned} \mathbb{P}\left\{\exists t \in \mathbb{N} : \|V_t(\lambda)^{-1/2} S_t\|^2 \geq s\right\} &\leq \mathbb{P}\left\{\exists t \in \mathbb{N} : \max_{x \in N_{1/2, m}} \|V_t(\lambda)^{-1/2} S_t x\|_2^2 \geq \frac{s}{4}\right\} \\ &\leq \sum_{x \in N_{1/2, m}} \mathbb{P}\left\{\exists t \in \mathbb{N} : \|V_t(\lambda)^{-1/2} S_t x\|_2^2 \geq \frac{s}{4}\right\}. \end{aligned}$$

By picking $s = \frac{1}{m_0} (8m + 4 \log \frac{1}{\delta} + 2 \log(\frac{\det(V_t(\lambda))}{\lambda^{d_\Theta}})) \geq \frac{1}{m_0} (4 \log \frac{5^m}{\delta} + 2 \log(\frac{\det(V_t(\lambda))}{\lambda^{d_\Theta}}))$ and applying Equation (9), we get

$$\mathbb{P}\left\{\exists t \in \mathbb{N} : \|V_t(\lambda)^{-1/2} S_t\|^2 \geq \frac{1}{m_0} \left(8m + 4 \log\left(\frac{1}{\delta}\right) + 2 \log\left(\frac{\det(V_t(\lambda))}{\lambda^{d_\Theta}}\right)\right)\right\} \leq \sum_{x \in N_{1/2, m}} \frac{\delta}{5^m} \leq \delta.$$

This completes the proof of inequality (8).

It remains to relate this bound to the definition of \mathcal{C}_t . We can write

$$\hat{\mu}_t - \mu_* = V_t(\lambda)^{-1} S_t + V_t(\lambda)^{-1} V_t \mu_* - \mu_*,$$

and therefore

$$\begin{aligned} \|V_t(\lambda)^{1/2} (\hat{\mu}_t - \mu_*)\| &= \|V_t(\lambda)^{-1/2} S_t + V_t(\lambda)^{1/2} (V_t(\lambda)^{-1} V_t - I) \mu_*\| \\ &\leq \|V_t(\lambda)^{-1/2} S_t\| + \sqrt{\|\mu_*^\top (V_t(\lambda)^{-1} V_t - I) V_t(\lambda) (V_t(\lambda)^{-1} V_t - I) \mu_*\|} \\ &= \|V_t(\lambda)^{-1/2} S_t\| + \sqrt{\lambda} \sqrt{\|\mu_*^\top (I - V_t(\lambda)^{-1} V_t) \mu_*\|} \\ &= \|V_t(\lambda)^{-1/2} S_t\| + \sqrt{\lambda} \|\mu_*\|, \end{aligned}$$

where the second equality follows by writing $V_t = V_t(\lambda) - \lambda I$. Note additionally that by $\max\{\|\theta\| : \theta \in \Theta\} \leq 1$ and the AM-GM inequality,

$$\det(V_t(\lambda)) \leq \left(\frac{1}{d_\Theta} \text{trace } V_t(\lambda)\right)^{d_\Theta} \leq \left(\frac{d_\Theta \lambda + t}{d_\Theta}\right)^{d_\Theta}.$$

Applying Equation (8), setting $\delta = \frac{1}{T^2}$ and $\lambda = \frac{1}{m_0}$ completes the proof. \square

D.2. Regret Bound on the Clean Event

The place where the structure of the performative risk comes into play is the following lemma, where we relate the suboptimality of the deployed model θ_t to properties of the confidence set \mathcal{C}_t .

Lemma D.3. *Suppose that the clean event (7) holds. Then, we can bound the suboptimality of θ_t by*

$$\Delta(\theta_t) \leq \min \left\{ 1, L_z \sup_{\mu, \mu' \in \mathcal{C}_t} \|(\mu - \mu')^\top \theta_t\| \right\}.$$

Proof. In what follows, all expectations are taken only over a sample $z_0 \sim \mathcal{D}_0$ independent of everything else (i.e., all other random quantities are conditioned on).

Since the loss is bounded, we know $\Delta(\theta_t) \leq 1$. For the other bound, notice that

$$\Delta(\theta_t) = \mathbb{E}\ell(z_0 + \mu_*^\top \theta_t; \theta_t) - \mathbb{E}\ell(z_0 + \mu_*^\top \theta_{\text{PO}}; \theta_{\text{PO}}).$$

By the definition of the algorithm and the clean event, we can lower bound the second term $\mathbb{E}\ell(z_0 + \mu_*^\top \theta_{\text{PO}}; \theta_{\text{PO}})$ as follows:

$$\mathbb{E}\ell(z_0 + \mu_*^\top \theta_{\text{PO}}; \theta_{\text{PO}}) \geq \text{PR}_{\text{LB}}(\theta_{\text{PO}}) \geq \text{PR}_{\text{LB}}(\theta_t) = \mathbb{E}\ell(z_0 + \tilde{\mu}_t^\top \theta_t; \theta_t),$$

for some $\tilde{\mu}_t \in \mathcal{C}_t$. This means that:

$$\Delta(\theta_t) \leq \mathbb{E}\ell(z_0 + \mu_*^\top \theta_t; \theta_t) - \mathbb{E}\ell(z_0 + \tilde{\mu}_t^\top \theta_t; \theta_t).$$

To finish, we use Lipschitzness of the loss to upper bound this by $L_z \|(\mu_* - \tilde{\mu}_t)^\top \theta_t\|$. Using the clean event, we can further upper bound this by $L_z \sup_{\mu, \mu' \in \mathcal{C}_t} \|(\mu - \mu')^\top \theta_t\|$ as desired. \square

We now use this bound on the suboptimality of deployed models, along with the structure of the confidence sets, to bound the regret on the clean event.

Lemma D.4. *Let $1 \leq \beta_1 \leq \beta_2 \leq \dots \leq \beta_T$ and assume that the loss $\ell(z; \theta)$ is L_z -Lipschitz in z . Assume that the event*

$$\mu_* \in \mathcal{C}_t \subseteq \left\{ \mu \in \mathbb{R}^{d_\Theta \times m} : \left\| V_{t-1}^{1/2} \left(\frac{1}{m_0} \right) (\mu - \hat{\mu}_{t-1}) \right\|^2 \leq \beta_t \right\}$$

holds true, for all $2 \leq t \leq T$. Then, on this event, Algorithm 2 satisfies:

$$\sum_{t=1}^T \Delta(\theta_t) = \tilde{\mathcal{O}} \left(1 + \sqrt{d_\Theta T \beta_T \log \left(\frac{d_\Theta + T m_0}{d_\Theta} \right) \max\{L_z, 1\}} \right).$$

Proof. As in the proof of Lemma D.3, all expectations are taken only over a sample $z_0 \sim \mathcal{D}_0$ independent of everything else (i.e., all other random quantities are conditioned on).

First, we separately bound the regret of the first step as $\mathcal{O}(1)$, using the fact that the loss is bounded in $[0, 1]$.

For the remainder of the steps, we apply Lemma D.3 to upper bound $\Delta(\theta_t)$. Using this, coupled with structure of \mathcal{C}_t , we can obtain the following upper bound, for any $\lambda > 0$:

$$\begin{aligned} \Delta(\theta_t) &\leq \min \left\{ 1, L_z \sup_{\mu, \mu' \in \mathcal{C}_t} \|(\mu - \mu')^\top \theta_t\| \right\} \\ &\leq \min \left\{ 1, L_z \sup_{\mu, \mu' \in \mathcal{C}_t} \|(\mu - \mu')^\top V_{t-1}^{1/2}(\lambda)\| \cdot \|V_{t-1}^{-1/2}(\lambda) \theta_t\| \right\} \\ &\leq \min \left\{ 1, 2L_z \sqrt{\beta_t} \|V_{t-1}^{-1/2}(\lambda) \theta_t\| \right\} \\ &\leq 2\sqrt{\beta_T} \min \left\{ 1, L_z \|V_{t-1}^{-1/2}(\lambda) \theta_t\| \right\}, \end{aligned}$$

where the last line uses the fact that $\beta_T \geq \max\{1, \beta_t\}$.

By the Cauchy-Schwarz inequality,

$$\begin{aligned}
 \sum_{t=2}^T \Delta(\theta_t) &\leq \sqrt{T \sum_{t=2}^T \Delta(\theta_t)^2} \\
 &\leq 2 \sqrt{T \beta_T \sum_{t=2}^T \min\{1, L_z^2 \|V_{t-1}^{-1/2}(\lambda) \theta_t\|^2\}} \\
 &\leq 2 \sqrt{T \beta_T \sum_{t=2}^T \min\{1, \max\{1, L_z^2\} \|V_{t-1}^{-1/2}(\lambda) \theta_t\|^2\}} \\
 &\leq 2 \sqrt{T \max\{1, L_z^2\} \beta_T \sum_{t=2}^T \min\{1, \|V_{t-1}^{-1/2}(\lambda) \theta_t\|^2\}} \\
 &= 2 \max\{1, L_z\} \sqrt{T \beta_T \sum_{t=2}^T \min\{1, \|V_{t-1}^{-1/2}(\lambda) \theta_t\|^2\}}.
 \end{aligned}$$

Finally, we use Lemma 19.4 in (Lattimore and Szepesvári, 2020) that says

$$\sum_{t=2}^T \min\{1, \|V_{t-1}^{-1/2}(\lambda) \theta_t\|^2\} \leq 2d_\Theta \log \left(\frac{\text{trace } V_0(\lambda) + T}{d_\Theta \det(V_0(\lambda))^{1/d_\Theta}} \right) = 2d_\Theta \log \left(\frac{d_\Theta \lambda + T}{d_\Theta \lambda} \right).$$

Using this expression in the equation above and setting $\lambda = \frac{1}{m_0}$ yields the final result. \square

D.3. Proof of Theorem 5.1

We take $\sqrt{\beta_t} = \max \left\{ 1, \sqrt{\frac{1}{m_0}} M_* + \sqrt{\frac{8m+8 \log T + 2d_\Theta \log \left(\frac{d_\Theta + tm_0}{d_\Theta} \right)}{m_0}} \right\}$. By the constraint that $m_0 = o(\log T)$, we see that second

branch dominates over the first one and so, for large enough T , $\sqrt{\beta_t} = \sqrt{\frac{1}{m_0}} M_* + \sqrt{\frac{8m+8 \log T + 2d_\Theta \log \left(\frac{d_\Theta + tm_0}{d_\Theta} \right)}{m_0}}$. Lemma D.2 shows that:

$$\mu_* \in \mathcal{C}_t \subseteq \left\{ \mu \in \mathbb{R}^{d_\Theta \times m} : \left\| V_{t-1}^{1/2} \left(\frac{1}{m_0} \right) (\mu - \hat{\mu}_{t-1}) \right\|^2 \leq \beta_t \right\}.$$

Moreover, the contribution of the complement of the clean event to the overall regret is negligible. Plugging this choice of β_t into the bound of Lemma D.4 completes the proof of Theorem 5.1.

E. Zooming Dimension Definitions

We note that Definition 2.3 slightly differs from the definition presented in (Kleinberg et al., 2008). The statement of Definition 2.3 eases the comparison of the zooming algorithm of Kleinberg et al. to our new algorithm.

First, we introduce a multiplier α to emphasize that the zooming dimension implicitly depends on the Lipschitz constant of the problem (assumed to be fixed and equal to 1 by Kleinberg et al.), which can be smaller when we make full use of performative feedback.

Second, Definition 2.3 is slightly more conservative in two ways. One is that we intersect a cover of any subset of $\{\theta : \Delta(\theta) \leq 16\alpha s\}$ with $\{\theta : 16\alpha r < \Delta(\theta) \leq 32\alpha r\}$, rather than directly take a cover of the latter set. The other one is that we take a supremum over all covers with radius coarser than r , i.e. $s \in [r, 1]$, instead of only r . These differences are minor technicalities that we do not expect to alter the zooming dimension in a meaningful way, neither formally nor conceptually.

Lastly, rather than requiring the size of the relevant set of points to be at most of order $(1/s)^d$, we require the size to be at most of order $(3/s)^d$. In this regard, Definition 2.3 is less conservative than the zooming dimension in (Kleinberg et al.,

2008). We make this modification so that for the Euclidean ball of dimension d_Θ of radius 1, which contains Θ , the zooming dimension is guaranteed to be at most d_Θ . This would not be true without the factor of 3. We note that the analysis of adaptive zooming in (Kleinberg et al., 2008) can be modified in a straightforward way to allow for this change, only altering constant factors in the regret bounds.

F. Gains of Sequential Zooming Dimension

We provide an example where the sequential zooming dimension is strictly smaller than the zooming dimension.

Example F.1. Suppose that model parameters are 2-dimensional, $L_z\epsilon = 1/32$, and the distribution map is a fixed distribution: $\text{PR}(\theta) = \text{DPR}(\theta', \theta)$ for all θ, θ' . Let $\theta_0 = 0, \theta_1 = [1/2, 0], \theta_2 = [1/4, \sqrt{3}/4]$. Suppose that $\text{PR}(\theta_0) = 0, \text{PR}(\theta_1) = 1/8, \text{PR}(\theta_2) = 15/64$, and $\text{PR}(\theta) = 1$ otherwise.

Lemma F.2. In Example F.1, the $(L_z\epsilon)$ -zooming dimension is at least $d \geq 0.39$, and the $(L_z\epsilon)$ -sequential zooming dimension is at most $\bar{d} \leq \log_6(1.5) \approx 0.23$.

Proof. We begin by observing that for all $0 < s \leq 1$, it holds that: $\{\theta \mid \Delta(\theta) \leq 16L_z\epsilon s\} \subseteq \{\theta \mid \text{PR}(\theta) \leq 16L_z\epsilon\} = \{\theta \mid \text{PR}(\theta) \leq 1/2\} = \{\theta_0, \theta_1, \theta_2\}$. Note that θ_0 achieves the optimal performative risk and thus does not appear in any suboptimality band $\{\theta \mid 16L_z\epsilon r \leq \text{PR}(\theta) < 32L_z\epsilon r\}$.

First, we show that the zooming dimension is at least $\log_6(2) \approx 0.39$. Let $s = 1/2 - \epsilon$ for ϵ sufficiently small. Consider the set $\{\theta \mid \Delta(\theta) \leq 16L_z\epsilon s\} = \{\theta_0, \theta_1, \theta_2\}$. A minimal covering of the set will necessarily consist of all three points $\{\theta_0, \theta_1, \theta_2\}$. We see that the suboptimality band $\{\theta \mid 16L_z\epsilon r \leq \text{PR}(\theta) < 32L_z\epsilon r\}$ for $r = 1/4$ contains $\{\theta_1, \theta_2\}$. Taking $\epsilon \rightarrow 0$, we see that the zooming dimension is at least $\log_6(2)$.

Next, we show that the sequential zooming dimension is $\bar{d} \leq \log_6(1.5) \approx 0.23$. Let \mathcal{C} be a minimal covering of a subset of $\{\theta \mid \Delta(\theta) \leq 16L_z\epsilon s\}$. For $s \geq 1/2$, we see that $\{\theta \mid \Delta(\theta) \leq 16L_z\epsilon s\} = \{\theta_0, \theta_1, \theta_2\}$, and \mathcal{C} contains at most 1 point. If $s < 1/2$, then \mathcal{C} might contain up to 3 points. If \mathcal{C} does not contain both θ_1 and θ_2 , then any suboptimality band $\{\theta \mid 16L_z\epsilon r \leq \text{PR}(\theta) < 32L_z\epsilon r\}$ contains at most 1 point from \mathcal{C} . If \mathcal{C} contains both θ_1 and θ_2 , then we leverage the sequential properties of the sequential zooming dimension. We claim that pulling θ_1 first results in θ_2 being eliminated. Notice that if θ_1 is pulled first, then PR_{\min} will be equal to $1/8$. PR_{LB} will be equal to $\text{DPR}(\theta_1, \theta_2) + L_z\epsilon\|\theta_1 - \theta_2\| = \text{PR}(\theta_2) + L_z\epsilon\|\theta_1 - \theta_2\| = 15/64 - 1/64 = 14/64 = 7/32$. We see that $\text{PR}_{\min} + 4L_z\epsilon s \leq 1/8 + 1/16 = 3/16$. Since $7/32 > 3/16$, we see that θ_1 will be eliminated. This means that in expectation, at most 1.5 arms are pulled. This yields the desired bound. □

G. Details of Numerical Illustrations

For the purpose of the illustrations in Figure 1, Figure 2, and Figure 3 we use a one-dimensional example where $\theta \in \mathbb{R}$. The performative effects are modeled by a linear shift, i.e.,

$$\text{DPR}(\phi, \theta) = f(\theta) + \alpha\phi,$$

where f is a multi-modal function illustrated in the respective figures and specified as

$$f(\theta) = c_0 \cos(c_1\theta) + c_2 \sin(c_3(\theta - c_4)).$$

The shaded gray area in the figures illustrates the confidence sets computed as

$$\text{PR}_{\text{LB}}(\theta) = \max_{\theta' \in \mathcal{S}} \text{PR}(\theta') - L_{\text{PR}}\|\theta - \theta'\|, \quad \text{PR}_{\text{UB}}(\theta) = \min_{\theta' \in \mathcal{S}} \text{PR}(\theta') + L_{\text{PR}}\|\theta - \theta'\|$$

for the baseline approach, and as

$$\text{PR}_{\text{LB}}(\theta) = \max_{\theta' \in \mathcal{S}} \text{DPR}(\theta', \theta) - L_\phi\|\theta - \theta'\|, \quad \text{PR}_{\text{UB}}(\theta) = \min_{\theta' \in \mathcal{S}} \text{DPR}(\theta', \theta) + L_\phi\|\theta - \theta'\|$$

for the performative confidence bounds. We use $\mathcal{S} := \{\theta_1, \theta_2\}$ as shown in the figures. The Lipschitz constant L_{PR} of the performative risk $\text{PR}(\theta) = \text{DPR}(\theta, \theta)$ is evaluated numerically for each figure.

Regret Minimization with Performative Feedback

For Figure 1 and Figure 2 we use the following parameters: $c_0 = -1$, $c_1 = 0.7$, $c_2 = 0.3$, $c_3 = 3$, $c_4 = 0.5$, $\alpha = 1$, and a conservative Lipschitz bound $L_\phi = 1.6$ for the performative confidence bounds and $L_{PR} = 3.8$ for the performative risk.

For Figure 3 we use $c_0 = -3$, $c_1 = 1$, $c_2 = 0.9$, $c_3 = 3$, $c_4 = 0.5$, $\alpha = 0.5$, and a conservative Lipschitz bound $L_\phi = 1.3$ for illustrating the performative confidence bounds. If exact knowledge of the shifts were available these bounds could be made even tighter.