

# Outside the Echo Chamber: Optimizing the Performative Risk

John Miller\* Juan C. Perdomo\* Tijana Zrnic\*  
{miller\_john, jcperdomo, tijana.zrnic}@berkeley.edu

University of California, Berkeley

June 17, 2021

## Abstract

In performative prediction, predictions guide decision-making and hence can influence the distribution of future data. To date, work on performative prediction has focused on finding performatively stable models, which are the fixed points of repeated retraining. However, stable solutions can be far from optimal when evaluated in terms of the performative risk, the loss experienced by the decision maker when deploying a model. In this paper, we shift attention beyond performative stability and focus on optimizing the performative risk directly. We identify a natural set of properties of the loss function and model-induced distribution shift under which the performative risk is convex, a property which does not follow from convexity of the loss alone. Furthermore, we develop algorithms that leverage our structural assumptions to optimize the performative risk with better sample efficiency than generic methods for derivative-free convex optimization.

## 1 Introduction

Predictions in social settings are rarely made in isolation, but rather to inform decision-making. This link between predictions and decisions causes predictive models to often be *performative*, meaning they can alter their environment once deployed. For example, election forecasts impact campaign spending and affect voter turnout, hence influencing the final election outcome [44]. Similarly, long-term climate forecasts shape policy decisions which can then affect future weather patterns.

Performative prediction is a recent framework introduced by Perdomo et al. [28] which formalizes the idea that predictive models can impact the data-generating process. So far, work in this area has focused on a particular equilibrium notion known as *performative stability* [5, 10, 22]. Stability is a local definition of optimality, by which a model minimizes the expected risk for the specific distribution that it induces. However, stability provides no general guarantees of performance beyond this equilibrium notion. In fact, stable models can have exceedingly poor *performative risk*, the central measure of performance in the performative prediction framework which captures the true risk incurred by the learner when deploying the model.

---

\* Equal contribution.

Reasoning by analogy, stable classifiers can be thought of as an *echo chamber* in an online platform. In an echo chamber, one is reassured of their ideas by voicing them, but it’s not clear whether they are reasonable outside of this niche community. Similarly, stable classifiers minimize risk on the distribution that they induce, but they provide no global guarantees of performance.

Therefore, to develop accurate predictions in performative settings, we shift attention past performative stability and study optimizing the performative risk directly. This task has so far remained elusive due to the complexities of model-induced distribution shift, i.e. performative effects. In particular, even in simple settings with convex losses, these distribution shifts can make the performative risk non-convex as noted in [28]. Furthermore, optimizing the performative risk requires a different algorithmic approach than what was previously studied in performative prediction. For instance, the learner needs to actively *anticipate* performative effects rather than myopically retrain until convergence, as the latter would only lead to stability. In short, repeated retraining is an inadequate method of overcoming performative distribution shifts.

## 1.1 Our Contributions

In this paper, we provide the first set of results describing when and how the performative risk may be optimized efficiently. We identify natural assumptions under which the performative risk is convex, even in settings where performative effects can be arbitrarily strong. Furthermore, we study optimization algorithms which explicitly model distribution shift and provably minimize the performative risk in an efficient manner.

To give an overview of our main results, we recall the relevant concepts from the performative prediction framework. Relative to supervised learning, where the learner observes data from a single *static* distribution, the key conceptual innovation in the performative prediction framework is the notion of a *distribution map*  $\mathcal{D}(\cdot)$ , which maps model parameters  $\theta \in \mathbb{R}^d$  to a distribution  $\mathcal{D}(\theta)$  over instances  $z$ . Given a loss  $\ell$ , the quality of a predictive model parameterized by  $\theta$  is measured according to its *performative risk*,

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

A classifier  $\theta_{\text{PO}}$  is *performatively optimal* if it minimizes the performative risk, i.e  $\theta_{\text{PO}} \in \text{argmin}_{\theta} \text{PR}(\theta)$ . On the other hand, a classifier  $\theta_{\text{PS}}$  is *performatively stable* if it satisfies the fixed-point condition,

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

In other words, stable classifiers are those which are optimal for the particular distribution they induce. However, stability has little bearing on whether a classifier has low performative risk. More specifically, the following observation motivates a large part of our later analysis:

*Stable classifiers can maximize the performative risk even when the loss is well-behaved and performative effects are small.*

Not only can stable points maximize the performative risk, but they can also have an arbitrarily large suboptimality gap,  $\text{PR}(\theta_{\text{PS}}) - \text{PR}(\theta_{\text{PO}})$ . The most natural first step towards optimizing the performative risk is to ensure that it is *convex*. Our first main result states that under an appropriate stochastic dominance condition which ensures the distribution map is well-behaved, there exists a critical threshold on the strength of performative effects that guarantees convexity:

**Theorem 1.1** (Informal). *Assume that the loss is  $\beta$ -smooth in  $z$  and  $\gamma$ -strongly convex in  $\theta$ . If the map  $\mathcal{D}(\cdot)$  is  $\varepsilon$ -Lipschitz and satisfies an appropriate stochastic dominance condition, then the performative risk is guaranteed to be convex if and only if  $\varepsilon \leq \frac{\gamma}{2\beta}$ .*

Interestingly, previous work has established that  $\varepsilon < \gamma/\beta$  is a threshold for repeated retraining to provably converge to a performatively stable point. We show that if we halve this quantity, we get another threshold which determines whether the performative risk is provably convex.

While [Theorem 1.1](#) suggests that performative effects need to be small in order to guarantee convexity, we prove that this need not be the case for the setting of *location-scale* families. These are natural classes of distribution maps in which performative effects enter through an additive or multiplicative factor that is linear in  $\theta$ . Many examples of distribution maps that have appeared in prior work are in fact location-scale families. For this setting, we generalize [Theorem 1.1](#) to prove the following structural result.

**Theorem 1.2** (Informal). *If the loss is smooth, strongly convex and the map  $\mathcal{D}(\cdot)$  is a location-scale family, then the performative risk can be convex irrespective of the Lipschitz constant of  $\mathcal{D}(\cdot)$ .*

Finally, having established these structural properties, we turn to algorithms for finding performative optima. Modulo weak regularity assumptions, convexity alone is sufficient to apply classical zeroth-order algorithms in order to find optima in polynomial time. That said, the convergence rate of these algorithms is typically quite slow.

To address this problem, we propose a *two-stage* approach, by which the learner first creates an explicit model of the distribution map  $\widehat{\mathcal{D}}$ , and then optimizes a proxy objective for the performative risk obtained by “plugging in”  $\widehat{\mathcal{D}}$  as if it were really the true distribution map. We instantiate this two-stage procedure in the context of location families, and prove that it optimizes the performative risk with significantly better sample efficiency than generic zeroth-order algorithms.

## 1.2 Related Work

We build on the recent line of work on performative prediction started by Perdomo et al. [28]. While previous papers in this area have focused on performative stability [5, 10, 22], we move past this solution concept and instead analyze conditions under which one can compute performatively optimal classifiers.

Given that strategic classification is formally a special case of performative prediction (see [Section 5](#) or discussion in [28] for further details), the study of performative optimality has been implicitly considered in the growing body of work on strategic classification [3, 6, 13–15, 23, 36, 41]. More specifically, performatively optimal classifiers correspond to Stackelberg equilibria in strategic classification. In contrast to papers within this literature, our analysis relies on identifying macro-level assumptions on the loss and the distribution shift which make the problem tractable, rather than specific micro-level assumptions on the costs or utilities of the agents. For example, Dong et al. [9] prove that the institution’s objective (performative risk) is convex by assuming that the agents are rational and compute best-responses according to particular utilities and cost functions. On the other hand, our conditions are on the distribution map and do not directly constrain behavior at the agent level.

Similarly, several papers in strategic classification [9, 26] and policy design [43] have recognized that one can apply zeroth-order algorithms [1, 11, 35] to find optima of the institution’s risk. The main challenge in applying zeroth-order optimization is the fact that, in general,

the performative risk might not satisfy any structural properties which would imply that its stationary points have low risk. One of the main contributions of this paper is precisely to identify under what conditions we can expect this behavior to hold.

Several works within the economics literature [12, 26] have also contrasted fixed points of retraining and institutional optima; these analyses resemble our comparisons of stability and optimality, albeit in a more specific setting. Furthermore, there are other settings beyond strategic classification that have similarly studied optimality in the face of performative effects, such as in the context of rankings or selection bias [17, 31, 39].

Lastly, our two-stage approach to minimizing the performative risk, whereby we first estimate a model of the distribution map and then optimize a proxy objective, is closely related to ideas in neighboring fields. At a high level, this general principle has appeared in semiparametric statistics [4, 16, 18, 27, 29] and more recently in double machine learning [7, 8, 19]. Furthermore, this idea has been extensively studied in the controls literature where it is referred to as certainty equivalence [20, 37, 38, 40], or as model-based planning in reinforcement learning [2].

### 1.3 Additional Preliminaries

As done by previous works in this area, we limit ourselves to considering predictive models parameterized by a finite-dimensional vector  $\theta \in \Theta \subseteq \mathbb{R}^d$ , where  $\Theta$  is a closed, convex set. The distribution map  $\mathcal{D}(\cdot)$  maps parameter vectors to data distributions over real-valued instances  $z \in \mathbb{R}^m$ . While each model  $\theta$  can induce a potentially distinct distribution  $\mathcal{D}(\theta)$ , we expect similar classifiers to induce similar distributions. This intuition is captured by the notion of  $\varepsilon$ -sensitivity, which is essentially a Lipschitz condition on the distribution map  $\mathcal{D}(\cdot)$ . We state that  $\mathcal{D}(\cdot)$  is  $\varepsilon$ -sensitive for some  $\varepsilon \geq 0$  if for all  $\theta, \theta' \in \Theta$ ,

$$W_1(\mathcal{D}(\theta), \mathcal{D}(\theta')) \leq \varepsilon \|\theta - \theta'\|_2. \quad (\text{A1})$$

Here,  $W_1$  denotes the Wasserstein-1 or earth mover’s distance between two distributions.

## 2 Contrasting Optimality and Stability

Up until now, all works within the performative prediction literature have focused on analyzing when different algorithms converge to stable points. While the primary motivation for stability was eliminating the need for retraining, it was observed as a useful byproduct that stable points can approximately minimize the performative risk.

More specifically, Perdomo et al. [28] prove that all stable points and performative optima lie within  $\ell_2$ -distance at most  $2L_z\varepsilon/\gamma$  of each other, where  $\varepsilon$  is the sensitivity of the distribution map,  $\gamma$  denotes the strong convexity parameter of the loss, and  $L_z$  denotes the Lipschitz constant of the loss in  $z$ . At first glance, this result implicitly suggests that stable points also have good predictive performance. While this is sometimes the case, in many settings  $L_z$  is large enough to make the bound vacuous. For example, there exist cases where the performative risk is strongly convex, but stable points actually *maximize* the performative risk.

**Proposition 2.1.** *For any  $\gamma, \Delta > 0$ , there exists a performative prediction problem where the loss is  $\gamma$ -strongly convex in  $\theta$ , yet the unique stable point  $\theta_{\text{PS}}$  maximizes the performative risk and  $\text{PR}(\theta_{\text{PS}}) - \min_{\theta} \text{PR}(\theta) \geq \Delta$ .*

*Proof.* We prove the proposition by constructing an example. Let  $z \sim \mathcal{D}(\theta)$  be a point mass at  $\varepsilon\theta$ , and define the loss to be:

$$\ell(z; \theta) = -\beta \cdot \theta^\top z + \frac{\gamma}{2} \|\theta\|_2^2,$$

for some  $\beta \geq 0$ . This loss is  $\gamma$ -strongly convex and the distribution map is  $\varepsilon$ -sensitive. A short calculation shows that the performative risk simplifies to

$$\text{PR}(\theta) = \left( \frac{\gamma}{2} - \varepsilon\beta \right) \cdot \|\theta\|_2^2. \quad (1)$$

For  $\varepsilon \neq \gamma/\beta$ , there is a unique performatively stable point at the origin, and if  $\varepsilon > \frac{\gamma}{2\beta}$  this point is the unique maximizer of the performative risk. Moreover, for  $\varepsilon > \frac{\gamma}{2\beta}$ ,  $\min_{\theta} \text{PR}(\theta) = (\gamma/2 - \varepsilon\beta) \cdot \max_{\theta \in \Theta} \|\theta\|_2^2$ . Therefore, depending on the radius of  $\Theta$ , the suboptimality gap of  $\theta_{\text{PS}}$  can be arbitrarily large. ■

In the above example,  $\nabla_{\theta} \ell(z; \theta)$  is  $\beta$ -Lipschitz in  $z$ , a condition commonly referred to as *smoothness* in prior work on performativity. The previous proposition thus shows that stable points can have an arbitrary suboptimality gap when  $\varepsilon > \frac{\gamma}{2\beta}$ . This is important since  $\varepsilon < \frac{\gamma}{\beta}$  is the regime where previously studied algorithms for optimizing under performativity—such as repeated risk minimization or different variants of gradient descent [22, 28]—converge to stability. Applying these methods when  $\varepsilon \in (\gamma/(2\beta), \gamma/\beta)$  would hence maximize the performative risk on this problem.

Moreover, we remark that the Lipschitz constant  $L_z$  is equal to  $\beta \cdot \max_{\theta \in \Theta} \|\theta\|_2$ . Therefore, the results of [28] imply that stable points and optima are at distance at most  $\frac{2L_z \varepsilon}{\gamma} = \frac{2\beta \varepsilon}{\gamma} \max_{\theta \in \Theta} \|\theta\|_2$ . When  $\varepsilon > \frac{\gamma}{2\beta}$ , as assumed in the proof of Proposition 2.1, this bound on the distance becomes vacuous:  $\|\theta_{\text{PS}} - \theta_{\text{PO}}\|_2 \leq \max_{\theta \in \Theta} \|\theta\|_2$ .

Lastly, we point out that  $\varepsilon = \frac{\gamma}{2\beta}$  is a sharp threshold for convexity of the performative risk in this example, as can be seen in Equation (1). In the following section, we show that this threshold behavior is not an artifact of this particular setting, but rather a phenomenon that holds more generally.

### 3 Convexity of the Performative Risk

We now introduce our main structural results illustrating how the performative risk can be convex in various natural settings, and hence amenable to direct optimization. Throughout our presentation, we adopt the following convention. We state that the performative risk is  $\lambda$ -convex, for some  $\lambda \in \mathbb{R}$ , if the objective,

$$\text{PR}(\theta) - \frac{\lambda}{2} \|\theta\|_2^2$$

is convex. In other words, if  $\lambda$  is positive, then  $\text{PR}(\theta)$  is  $\lambda$ -strongly convex. If  $\lambda$  is negative, then adding the analogous regularizer  $\frac{\lambda}{2} \|\theta\|_2^2$  ensures  $\text{PR}(\theta)$  is convex. Furthermore, in addition to  $\varepsilon$ -sensitivity, we will make repeated use of the following assumptions throughout the remainder of the paper. To facilitate readability, we let  $\mathcal{Z} \stackrel{\text{def}}{=} \cup_{\theta \in \Theta} \text{supp}(\mathcal{D}(\theta))$ . We say that a loss function  $\ell(z; \theta)$  is  $\beta$ -smooth in  $z$  if for all  $\theta \in \Theta$  and  $z, z' \in \mathcal{Z}$ ,

$$\|\nabla_{\theta} \ell(z; \theta) - \nabla_{\theta} \ell(z'; \theta)\|_2 \leq \beta \|z - z'\|_2. \quad (\text{A2})$$

Furthermore, a loss function  $\ell(z; \theta)$  is  $\gamma$ -strongly convex in  $\theta$  if for all  $\theta, \theta', \theta_0 \in \Theta$ ,

$$\mathbb{E}_{z \sim \mathcal{D}(\theta_0)} \ell(z; \theta) \geq \mathbb{E}_{z \sim \mathcal{D}(\theta_0)} \ell(z; \theta') + \mathbb{E}_{z \sim \mathcal{D}(\theta_0)} \nabla_{\theta} \ell(z; \theta')^{\top} (\theta - \theta') + \frac{\gamma}{2} \|\theta - \theta'\|_2^2. \quad (\text{A3a})$$

If  $\gamma = 0$ , this assumption is equivalent to convexity. Similarly, we say that the loss is  $\gamma_z$ -strongly convex in  $z$  if for all  $\theta \in \Theta$  and  $z, z' \in \mathcal{Z}$ ,

$$\ell(z; \theta) \geq \ell(z'; \theta) + \nabla_z \ell(z'; \theta)^{\top} (z - z') + \frac{\gamma_z}{2} \|z - z'\|_2^2. \quad (\text{A3b})$$

Lastly, we state that a distribution map, loss pair  $(\mathcal{D}(\cdot), \ell)$  satisfies *mixture dominance* if the following condition holds for all  $\theta, \theta', \theta_0 \in \Theta$  and  $\alpha \in (0, 1)$ :

$$\mathbb{E}_{z \sim \mathcal{D}(\alpha\theta + (1-\alpha)\theta')} \ell(z; \theta_0) \leq \mathbb{E}_{z \sim \alpha\mathcal{D}(\theta) + (1-\alpha)\mathcal{D}(\theta')} \ell(z; \theta_0). \quad (\text{A4})$$

Smoothness and strong convexity are standard and have appeared previously in the context of performative prediction. The mixture dominance condition is novel and plays a central role in our analysis of when the performative risk is convex. To provide some intuition for this condition, we recall the definition of the *decoupled performative risk*:

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

Notice that asserting convexity of the performative risk is equivalent to showing convexity of  $\text{DPR}(\theta, \theta)$  when both arguments are forced to be the same. While convexity (A3a) guarantees that DPR is convex in the second argument, mixture dominance (A4) essentially posits convexity of DPR in the first argument. Importantly, assuming convexity in each argument separately does *not* directly imply that the performative risk is convex.

On a more intuitive level, this assumption (A4) is essentially a stochastic dominance statement: the mixture distribution  $\alpha\mathcal{D}(\theta) + (1 - \alpha)\mathcal{D}(\theta')$  “dominates”  $\mathcal{D}(\alpha\theta + (1 - \alpha)\theta')$  under a certain loss function. Similar conditions have been extensively studied within the literature on stochastic orders [33], which we further discuss in Appendix A. Part of our analysis relies on incorporating tools from this literature, and we believe that further exploring technical connections between this field and performative prediction could be valuable. For example, using results from stochastic orders we can show that (A4) holds when the loss is convex in  $z$  and the distribution map  $\mathcal{D}(\cdot)$  forms a *location-scale family* of the form:

$$z_{\theta} \sim \mathcal{D}(\theta) \Leftrightarrow z_{\theta} \stackrel{d}{=} (\Sigma_0 + \Sigma(\theta))z_0 + \mu_0 + \mu\theta, \quad (2)$$

where  $z_0 \sim \mathcal{D}_0$  is a sample from a fixed zero-mean distribution  $\mathcal{D}_0$ , and  $\Sigma(\theta), \mu$  are linear maps (see Proposition A.4 for a formal proof). Distribution maps of this sort are ubiquitous throughout the performative prediction literature and hence satisfy mixture dominance if the loss  $\ell$  is convex. For instance, the distribution map for the strategic classification simulator in [28] is a location family. Other examples of location families can be found in previous work on strategic classification [12, 13]. Mixture dominance can also hold in discrete settings, e.g.  $\mathcal{D}(\theta) = \text{Bernoulli}(a^{\top}\theta + b)$  satisfies this condition for any loss. Having provided some context on the mixture dominance condition, we can now state the main result of this section:

**Theorem 3.1.** *Suppose that the loss function  $\ell(z; \theta)$  is  $\gamma$ -strongly convex in  $\theta$  (A3a),  $\beta$ -smooth in  $z$  (A2), and that  $\mathcal{D}(\cdot)$  is  $\varepsilon$ -sensitive (A1). If mixture dominance (A4) holds, then the performative risk is  $\lambda$ -convex for  $\lambda = \gamma - 2\varepsilon\beta$ .*



Together with the example from the proof of [Proposition 2.1](#), this theorem shows that  $\frac{\gamma}{2\beta}$  is a sharp threshold for convexity of the performative risk. If  $\varepsilon$  is strictly less than this threshold, then under mixture dominance and appropriate conditions on the loss, the performative risk is strongly convex by [Theorem 3.1](#). On the other hand, if  $\varepsilon$  is above this threshold, the example from [Proposition 2.1](#) shows that there exists a performative prediction instance which satisfies the remaining assumptions, yet is non-convex; in particular, for  $\varepsilon > \frac{\gamma}{2\beta}$  the performative risk is strictly concave in that example. This threshold was also implicitly observed by Perdomo et al. [28] in the proof of Proposition 4.2 as byproduct of showing that the performative risk can be non-convex for  $\varepsilon \leq \frac{\gamma}{\beta}$ . However, they provide no general analysis of when the performative risk is convex. Note that all of the above examples satisfy mixture dominance.

While the threshold  $\varepsilon = \gamma/(2\beta)$  is in general tight as argued above, for certain families of distribution maps the conclusion of [Theorem 3.1](#) can be made considerably stronger. Indeed, in some cases the performative risk is convex *regardless* of the magnitude of performative effects, as observed for the following location family.

**Example 3.2.** Consider the following stylized model of predicting the final vote margin in an election contest. Features  $x$ , such as past polling averages, are drawn i.i.d. from a static distribution,  $x \sim \mathcal{D}_x$ . Since predicting a large margin in either direction can dissuade people from voting, we consider outcomes drawn from the conditional distribution:  $y|x \sim g(x) + \mu^\top \theta + \xi$ , where  $g: \mathbb{R}^d \rightarrow \mathbb{R}$  is an arbitrary map,  $\mu \in \mathbb{R}^d$  is a fixed vector, and  $\xi$  is a zero-mean noise variable. If  $\ell$  is the squared loss,  $\ell((x, y); \theta) = \frac{1}{2}(y - x^\top \theta)^2$ , or the absolute loss,  $\ell((x, y); \theta) = |y - x^\top \theta|$ , then the performative risk is convex for any  $g$  and  $\mu$ .

The proof follows by simply observing that in both cases, the performative risk can be written as a linear function in  $\theta$  composed with a convex function. Another interesting property of this example is that the distribution map is  $\varepsilon$ -sensitive with  $\varepsilon = \|\mu\|_2$ , yet the sensitivity parameter plays no role in the characterization of convexity. Motivated by this observation, we specialize the analysis in [Theorem 3.1](#) to the particular case of location-scale families, and obtain a result that is at least as tight as the previous theorem.

**Theorem 3.3.** *Suppose that  $\ell(z; \theta)$  is  $\gamma$ -strongly convex in  $\theta$  ([A3a](#)),  $\beta$ -smooth ([A2](#)), and  $\gamma_z$ -strongly convex in  $z$  ([A3b](#)). Furthermore, suppose that  $\mathcal{D}(\theta)$  forms a location-scale family ([2](#)) with  $\varepsilon$  as its sensitivity parameter<sup>1</sup>. Define  $\Sigma_{z_0}$  to be the covariance matrix of  $z_0 \sim \mathcal{D}_0$ , and let*

$$\sigma_{\min}(\mu) = \min_{\|\theta\|_2=1} \|\mu\theta\|_2, \sigma_{\min}(\Sigma) = \min_{\|\theta\|_2=1} \|\Sigma^{1/2}\Sigma(\theta)^\top\|_F.$$

*Then, the performative risk is  $\lambda$ -convex for  $\lambda$  equal to:*

$$\max\{\gamma - \beta^2/\gamma_z, \gamma - 2\varepsilon\beta + \gamma_z(\sigma_{\min}^2(\mu) + \sigma_{\min}^2(\Sigma))\}.$$

This tighter bound leverages the fact that some losses are strongly convex in the performative variables, such as the squared loss when only the outcome variable exhibits performative effects. In general, one can achieve a tighter analysis of when the performative risk is convex by distinguishing between variables which are *static*, whose distribution is the same under  $\mathcal{D}(\theta)$  for all  $\theta$ , and performative variables which are influenced by the deployed classifier. For the most part we avoid this distinction in the main body for the sake of readability, however, we elaborate

<sup>1</sup>The sensitivity parameter  $\varepsilon$  for location-scale families can be explicitly bounded in terms of the parameters  $\mu$  and  $\Sigma(\theta)$ ; see [Remark C.3](#) in the Appendix.

on how the analysis can be strengthened in [Appendix B](#). We now illustrate an application of [Theorem 3.3](#) on a scale family example.

**Example 3.4.** Suppose that  $x > 0$  is a one-dimensional feature drawn from a fixed distribution  $\mathcal{D}_x$ , and let  $y|x \sim \theta x \cdot \text{Exp}(1)$  be distributed as an exponential random variable with mean  $\theta x$ . Let the loss be the squared loss,  $\ell((x, y); \theta) = \frac{1}{2}(y - \theta \cdot x)^2$  and let  $\Theta = \mathbb{R}_+$ . Note that this example exhibits a self-fulfilling prophecy property whereby all solutions are performatively stable. On the other hand,  $\text{PR}(\theta) = \theta^2 \mathbb{E} x^2$ , and the unique performative optimum is  $\theta_{\text{PO}} = 0$ . Again, we see how stability has no bearing on whether a solution has low performative risk.

However, we note that the loss is 1-strongly convex in  $y$ . Furthermore, by averaging over the static features, we observe that  $\text{PR}(\theta)$  is  $\mathbb{E} x^2$ -strongly convex in  $\theta$  and  $\mathbb{E} x$ -smooth in  $y$ . Therefore, according to [Theorem 3.3](#), the performative risk is convex and hence tractable to optimize, since  $\gamma - \beta^2/\gamma_z = \mathbb{E} x^2 - (\mathbb{E} x)^2 \geq 0$  by Jensen's inequality.

While this example, like most others in this section, is intended as a toy problem to provide the reader with some intuition regarding the intricacies of performativity, many instances of performative prediction in the real world do exhibit a self-fulfilling prophecy aspect whereby predicting a particular outcome increases the likelihood that it occurs. For instance, predicting that a student is unlikely to do well on a standardized exam may discourage them from studying in the first place and hence lower their final grade. Settings like these where stability is a vacuous guarantee of performance remind us how developing reliable predictive models requires going outside the stability echo chamber.

As a final note, to prove the results in this section, we have imposed additional assumptions such as mixture dominance, or analyzed the special case of location-scale families. The reader might naturally ask whether these settings are so restrictive that one can optimize the performative risk using previous optimization methods for performative prediction which find stable points. Or in particular, whether stable points and performative optima now identify.

It turns out that both solutions can still have qualitatively different behavior, regardless of the strength of performative effects. First, notice that the example in the proof of [Proposition 2.1](#) is a location family, and as such it satisfies mixture dominance. In that example, when  $\varepsilon \in (\frac{\gamma}{2\beta}, \frac{\gamma}{\beta})$ , methods for finding stable points converge to a maximizer of the performative risk; however, this is outside the regime where the performative risk is convex. In what follows, by relying on [Theorem 3.3](#), we provide another scale family example where the performative risk is convex regardless of  $\varepsilon$ , yet stable points can be arbitrarily suboptimal.

**Example 3.5.** Suppose that  $\mathcal{D}(\theta) = \mathcal{N}(\mu, \varepsilon^2 \theta^2)$  for some  $\mu \in \mathbb{R}$  and  $\varepsilon > 0$ . This distribution map is  $\varepsilon$ -sensitive. Furthermore, if  $\ell$  is the squared loss,  $\ell(z; \theta) = \frac{1}{2}(z - \theta)^2$ , then there is a unique stable point  $\theta_{\text{PS}} = \mu$ . On the other hand,  $\theta_{\text{PO}} = \mu/(1 + \varepsilon^2)$ .

Notice how, contrary to the performative optimum  $\theta_{\text{PO}}$ , the stable point  $\theta_{\text{PS}}$  is independent of  $\varepsilon$  and hence oblivious to the performative effects. Depending on  $\mu$ , the stable point can be arbitrarily suboptimal, since  $\text{PR}(\theta_{\text{PS}}) - \text{PR}(\theta_{\text{PO}}) = \Omega(\mu^2)$ . Note also that, according to [Theorem 3.3](#), the performative risk is  $\gamma - 2\varepsilon\beta + \gamma_z \sigma_{\min}^2(\Sigma) = 1 - 2\varepsilon + \varepsilon^2$ -convex. Since  $1 - 2\varepsilon + \varepsilon^2 = (\varepsilon - 1)^2 \geq 0$ , the performative risk is always convex and hence tractable to optimize.

## 4 Optimization Algorithms

Having identified conditions under which the performative risk is convex, we now consider methods for efficiently optimizing it. One of the main challenges of carrying out this task



is that, even in convex settings, the learner can only access the objective via noisy function evaluations corresponding to classifier deployments. Without knowledge of the underlying distribution map, it is infeasible to compute gradients of the performative risk. A naive solution is to apply a zeroth-order method, however, these algorithms are in general hard to tune, and their performance scales poorly with the problem dimension.

Our main algorithmic contribution is to show how one can address these issues by creating an explicit *model* of the distribution map and then optimizing a proxy objective for the performative risk offline. We refer to this as the two-stage procedure for optimizing the performative risk and show it is provably efficient for the case of location families.

To develop further intuition, consider the following simple example. Let  $z \sim \mathcal{N}(\varepsilon\theta, 1)$  be a one-dimensional Gaussian and let  $\ell(z; \theta) = \frac{1}{2}(z - \theta)^2$  be the squared loss. Then, the performative risk,  $\text{PR}(\theta) = \frac{1}{2}(\varepsilon - 1)^2\theta^2$ , is a simple, convex function for all values of  $\varepsilon$  (as indeed confirmed by [Theorem 3.3](#), since  $\gamma - 2\varepsilon\beta + \gamma_z\sigma_{\min}^2(\mu) = 1 - 2\varepsilon + \varepsilon^2 \geq 0$ ). However, gradients are unavailable since they depend on the density of  $\mathcal{D}(\theta)$ , denoted  $p_\theta$ , which is typically unknown:

$$\begin{aligned} \nabla_\theta \text{PR}(\theta) &= \mathbb{E}_{z \sim \mathcal{D}(\theta)} \nabla_\theta \ell(z; \theta) + \mathbb{E}_{z \sim \mathcal{D}(\theta)} \ell(z; \theta) \nabla_\theta \log p_\theta(z) \\ &= \mathbb{E}_{z \sim \mathcal{D}(\theta)} -(z - \theta) + \varepsilon(\varepsilon - 1)\theta. \end{aligned}$$

Despite the simplicity of this example, earlier approaches to optimization in performative prediction, such as repeated retraining [28], fail on this problem. The reason is that they essentially ignore the second term in the gradient computation which requires explicitly anticipating performative effects. For example, retraining computes the sequence of updates  $\theta_{t+1} = \arg \min_\theta \mathbb{E}_{z \sim \mathcal{D}(\theta_t)} \frac{1}{2}(z - \theta)^2 = \varepsilon\theta_t$ , which diverges for  $|\varepsilon| > 1$ .

#### 4.1 Generic Derivative-Free Methods

Having observed the difficulty of computing gradients, the most natural starting point for optimizing the performative risk is to consider derivative-free methods for convex optimization [1, 11, 35]. These methods work by constructing a noisy estimate of the gradient by querying the objective function at a randomly perturbed point around the current iterate. For instance, Flaxman et al. [11] sample a vector  $u \sim \text{Unif}(\mathcal{S}^{d-1})$  to get a slightly biased gradient estimator,

$$\nabla_\theta \text{PR}(\theta) \approx \frac{d}{\delta} \mathbb{E}[\text{PR}(\theta + \delta u)u],$$

for some small  $\delta > 0$ . Generic derivative-free algorithms for convex optimization require few assumptions beyond those given in the previous section to ensure convexity. Moreover, they guarantee convergence to a performative optimum given sufficiently many samples. However, their rate of convergence can be slow and scales poorly with the problem dimension. In general, zeroth-order methods require  $\widetilde{O}(d^2/\Delta^2)$  samples to obtain a  $\Delta$ -suboptimal point [1, 35], which can be prohibitively expensive if samples are hard to come by.

#### 4.2 Two-Stage Approach

In cases where we have further structure, an alternative solution to derivative-free methods is to utilize a *two-stage* approach to optimizing the performative risk. In the first stage, we estimate a coarse model of the distribution map,  $\widehat{\mathcal{D}}(\cdot)$  via experiment design. Then, in the second stage, the

---

**Algorithm 1** Two-Stage Algorithm for Location Families

---

**Stage 1:** Construct a model of the distribution map

// Estimate location parameter  $\mu$  with experiment design

**for**  $i = 1$  **to**  $n$  **do**

-Sample and deploy classifier  $\theta_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_d)$ .

-Observe  $z_i \sim \mathcal{D}(\theta_i)$ .

**end for**

-Estimate  $\mu$  via ordinary least squares,  $\widehat{\mu} \in \arg \min_{\mu} \sum_{i=1}^n \|z_i - \mu \theta_i\|_2^2$ .

// Gather samples from the base distribution

**for**  $j = n + 1$  **to**  $2n$  **do**

-Deploy classifier  $\theta_j = 0$ , and observe  $z_j \sim \mathcal{D}(0)$ .

**end for**

**Stage 2:** Minimize a finite-sample approximation of the performative risk,  $\arg \min_{\theta \in \Theta} \frac{1}{n} \sum_{j=n+1}^{2n} \ell(z_j + \widehat{\mu} \theta; \theta)$ .

---

algorithm optimizes a proxy to the performative risk treating the estimated  $\widehat{\mathcal{D}}$  as if it were the true distribution map:

$$\widehat{\theta}_{\text{PO}} \in \arg \min_{\theta} \widehat{\text{PR}}(\theta) \stackrel{\text{def}}{=} \mathbb{E}_{z \sim \widehat{\mathcal{D}}(\theta)} \ell(z; \theta).$$

The exact implementation of this idea depends on the problem setting at hand; to make things concrete, we instantiate the approach in the context of location families and prove that it optimizes the performative risk with significantly better sample complexity than generic zeroth-order methods. For the remainder of this section, we assume the distribution map  $\mathcal{D}$  is parameterized by a location family

$$z_{\theta} \sim \mathcal{D}(\theta) \Leftrightarrow z_{\theta} \stackrel{d}{=} z_0 + \mu \theta,$$

where the matrix  $\mu \in \mathbb{R}^{m \times d}$  is an unknown parameter, and  $z_0 \sim \mathcal{D}_0$  is a zero-mean random variable.<sup>2</sup>

As discussed previously, location-scale families encompass many formal examples discussed in prior work. They capture the intuition that in performative settings, the data points are composed of a *base* component  $z_0$ , representing the natural data distribution in the absence of performativity, and an additive performative term.

In the first stage of our two-stage procedure we build a model of the distribution map  $\widehat{\mathcal{D}}$  that in effect allows us to draw samples  $z \sim \widehat{\mathcal{D}}(\theta) \approx \mathcal{D}(\theta)$ . To do this, we perform experiment design to recover the unknown parameter  $\mu$  which captures the performative effects. In particular, we sample and deploy  $n$  classifiers  $\theta_i$ ,  $i \in [n]$ , observe data  $z_i \sim \mathcal{D}(\theta_i)$ , and then construct an estimate  $\widehat{\mu}$  of the location map  $\mu$  using ordinary least squares. We then gather samples from the base distribution  $\mathcal{D}_0$  by repeatedly deploying the zero classifier. In the location-family model, deploying the zero classifier ensures we observe data points  $z_0$ , without performative effects. With both of these components, given any  $\theta'$ , we can simulate  $z \sim \widehat{\mathcal{D}}(\theta')$  by taking  $z = z_0 + \widehat{\mu} \theta'$ .

---

<sup>2</sup>The variable  $z_0$  being zero-mean is only to simplify the exposition; the same analysis carries over when there is an additional intercept term. Similarly, the choice of Gaussian noise in the experiment design phase of Algorithm 1 is made for convenience. In general, any subgaussian distribution with full rank covariance would suffice.

In the second stage, we use the estimated model to construct a proxy objective. Define the perturbed performative risk:

$$\widehat{\text{PR}}(\theta) = \mathbb{E}_{z \sim \widehat{\mathcal{D}}(\theta)} \ell(z; \theta) = \mathbb{E}_{z_0 \sim \mathcal{D}_0} \ell(z_0 + \widehat{\mu}\theta; \theta).$$

Note that  $\text{PR}(\theta) = \mathbb{E}_{z_0 \sim \mathcal{D}_0} \ell(z_0 + \mu\theta; \theta)$ . Using the estimated parameter  $\widehat{\mu}$  and samples  $z_i \sim \mathcal{D}_0$ , we can construct a finite-sample approximation to the perturbed performative risk and find the following optimizer:

$$\widehat{\theta}_n \in \arg \min_{\theta \in \Theta} \widehat{\text{PR}}_n(\theta) \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=n+1}^{2n} \ell(z_i + \widehat{\mu}\theta; \theta).$$

The main technical result in this section shows that, under appropriate regularity assumptions on the loss, Algorithm 1 efficiently approximates the performative optimum. In particular, when the data dimensionality  $m$  is comparable to the model dimensionality  $d$ , i.e.  $m = O(d)$ , then computing a  $\Delta$ -suboptimal classifier requires  $O(d/\Delta)$  samples. In contrast, the derivative-free methods considered previously require  $\widetilde{O}(d^2/\Delta^2)$  samples to compute a classifier of similar quality. The formal statement and proof of this result is deferred to [Appendix C.2](#).

**Theorem 4.1** (Informal). *Under appropriate smoothness and strong convexity assumptions on the loss  $\ell$ , if the distribution of  $z_0$  is subgaussian, and if the number of samples  $n \geq \Omega(d + m + \log(1/\delta))$ , then, with probability  $1 - \delta$ , Algorithm 1 returns a point  $\widehat{\theta}_n$  such that*

$$\text{PR}(\widehat{\theta}_n) - \text{PR}(\theta_{\text{PO}}) \leq O\left(\frac{d + m + \log(1/\delta)}{n} + \frac{1}{\delta n}\right).$$

While we analyze this two-stage procedure in the context of location families, the principles behind the approach can be extended to more general settings. Whenever the distribution map has enough structure to efficiently estimate a model  $\widehat{\mathcal{D}}$  that supports sampling new data, we can always use the “plug-in” approach above and construct and optimize a perturbed version of the performative risk.

## 5 Experiments

We complement our theoretical findings with an empirical evaluation of different methods on two tasks: the strategic classification simulator from [28], and a synthetic linear regression example.

We pay particular attention to understanding the differences in empirical performance between algorithms which converge to performative optima, such as the two-stage procedure or derivative-free methods from [Section 4.1](#), versus existing optimization algorithms for finding stable points, in particular greedy and lazy SGD due to Mendler-Dünner et al. [22]. In addition, we focus on highlighting the differences in the sample efficiency of the different algorithms and examine their sensitivity to the relevant structural assumptions outlined in [Section 3](#). To evaluate derivative-free methods, we implement the “gradient descent without a gradient” algorithm from [11], which we refer to from here on out as the “DFO algorithm.” For each of the following experiments, we run each algorithm 50 times and display 95% bootstrap confidence intervals. We provide a formal description of all the procedures, as well as a detailed description of the experimental setup in [Appendix D](#).




Figure 1: Suboptimality gap versus number of samples collected for the two-stage algorithm, DFO algorithm, greedy SGD, and lazy SGD, for  $\epsilon = 0.01$  (left) and  $\epsilon = 100$  (right). Each experiment is repeated 50 times, and we display 95% bootstrap confidence intervals.

**Linear regression experiments.** We begin by evaluating how increasing the strength of performative effects affects the behavior of the different optimization procedures in settings where the performative risk is convex. We recall the setup from [Example 3.2](#), where the learner attempts to solve a linear regression with performative labels. Given a parameter  $\theta$ , data are drawn from  $\mathcal{D}(\theta)$  according to:

$$x \sim \mathcal{N}(0, \Sigma_x), U_y \sim \mathcal{N}(0, \sigma_y^2), y = \beta^\top x + \mu^\top \theta + U_y.$$

This distribution map is a location family, and is  $\epsilon$ -sensitive with  $\epsilon = \|\mu\|_2$ . Performance is measured according to the squared loss,  $\ell((x, y); \theta) = \frac{1}{2}(y - \theta^\top x)^2$ . Furthermore, the performative risk is convex for all choices of  $\mu$ .

For small  $\epsilon$ , we see that greedy and lazy SGD converge to a stable point that approximately minimizes the performative risk (see left panel in [Figure 1](#)). However, as we increase the strength of performative effects, these methods fail to make any progress, and are outperformed by both the DFO algorithm and the two-stage approach by a considerable margin (see right panel in [Figure 1](#)). The two-stage procedure efficiently converges after a small number of samples and its behavior is largely unaffected as we increase the value of  $\epsilon$ , while the DFO algorithm becomes considerably slower when  $\epsilon$  is large.

**Strategic classification simulator.** We next consider experiments on the credit scoring simulator from [\[28\]](#), which has been employed as an empirical benchmark for performative prediction in several works [\[5, 10, 22\]](#). The simulator models a strategic classification problem between a bank and individual agents seeking a loan. The bank deploys a logistic regression classifier  $f_\theta$  to determine the individuals' default probabilities, while agents strategically manipulate their features to achieve a more favorable classification.

More specifically, individuals correspond to feature, label pairs  $(x, y)$  drawn i.i.d. from a base distribution  $\mathcal{D}_0$ . Given a classifier  $f_\theta$ , agents compute a best-response set of features  $x_{BR}$  by solving an optimization problem. The bank then observes the manipulated data points




Figure 2: Classification accuracy versus number of samples collected for the two-stage algorithm, DFO algorithm, greedy SGD, and lazy SGD, for  $\epsilon = 0.0001 \leq \frac{\gamma}{2\beta}$  (left) and  $\epsilon = 100 \gg \frac{\gamma}{2\beta}$  (right). Each experiment is repeated 50 times, and we display 95% bootstrap confidence intervals.

$(x_{\text{BR}}, y) \sim \mathcal{D}(\theta)$ . For an appropriate choice of the agents' objective function, the distribution map forms a location family,  $x_{\text{BR}} = x + \varepsilon\theta$ , where  $\varepsilon$  is a parameter of the agents' objective. It also serves as a measure of performativity, since this distribution map is  $\varepsilon$ -sensitive. As a final remark, we add  $\ell_2$ -regularization to the logistic loss to ensure strong convexity. See discussion in [28] and Appendix D for full details.

Since the logistic loss is not strongly convex in the features, we only have a certificate of convexity when  $\varepsilon$  is small enough (namely,  $\varepsilon \leq \frac{\gamma}{2\beta}$ ). We consider two values of  $\varepsilon$ : one which is below this critical threshold, and one large value for which we do not have theoretical guarantees. When  $\varepsilon$  is small, both the DFO algorithm and the two-stage method yield significantly higher accuracy solutions compared to the two variants of SGD (see left panel of Figure 2). Together with the linear regression experiments, this observation serves as further evidence that stable points have significantly worse performative risk relative to performative optima, even in regimes where  $\varepsilon < \gamma/(2\beta)$ . Note also that, although both the DFO algorithm and the two-stage algorithm improve upon methods for repeated retraining, the two-stage algorithm converges with significantly fewer samples and significantly lower variance. Indeed, a few thousand samples suffice for convergence of the two-stage method, whereas the DFO algorithm has still not fully converged after a million samples.

Lastly, on the top right plot, we evaluate these methods for  $\varepsilon \gg \gamma/(2\beta)$  which is outside the regime of our theoretical analysis. Consequently, we have no convergence guarantees for any of the four algorithms. Despite the lack of guarantees and the increased strength of performative effects, we see that the two-stage procedure achieves only a slightly lower accuracy than in the previous setting. On the other hand, as described in our echo chamber analogy, greedy and lazy SGD rapidly converge to a local minimum and do not significantly improve predictive performance after the 10k sample mark. Despite extensive tuning, we were unable to improve the performance of the DFO algorithm and achieve nontrivial accuracy with this method.

## 6 Discussion and Future Work

Given the stark difference between performative stability and optimality, the goal of our work is to identify the first set of conditions and algorithmic procedures by which one might be able to provably optimize the performative risk. To this end, we focus on analyzing the problem at a broad level of generality, identifying simple, structural conditions under which the optimization problem becomes tractable.

However, when applying these ideas in practice, there are a number of important considerations determined by the relevant social context that are not explicitly addressed by our theoretical analysis and which we believe are an important direction for future work. In social settings, such as credit scoring or election forecasts, the choice of loss function must balance predictive accuracy with any externalities that arise from the classifier’s impact on the observed distribution. For example, in lending we may wish to find a classifier  $f_\theta$  that accurately predicts individual defaults, but that also induces a distribution  $\mathcal{D}(\theta)$  over which the mean probability of default is low (or which satisfies some other socially desirable property).

In order to balance between predictive accuracy and other concerns about the observed distribution, one possibility is to directly incorporate a penalty on  $\mathcal{D}(\theta)$ . For instance, if we would like  $\mathcal{D}(\theta)$  to satisfy  $\mathbb{E}_{z \sim \mathcal{D}(\theta)}[z] \approx z_\star$ , e.g., to control the mean probability of default, we can modify our loss by including a regularization term of the form  $(z - z_\star)^\top Q(z - z_\star)$ , where  $Q$  is some PSD matrix. Importantly, incorporating such a regularization term does not alter the convexity of the performative risk. More formally, given any loss  $\ell$  and distribution map  $\mathcal{D}(\cdot)$ , which satisfy the conditions from [Theorem 3.3](#), and hence for which  $\text{PR}(\theta) = \mathbb{E}_{z \sim \mathcal{D}(\theta)} \ell(z; \theta)$  is convex, our analysis proves that the regularized objective,

$$\text{PR}'(\theta) \stackrel{\text{def}}{=} \mathbb{E}_{z \sim \mathcal{D}(\theta)} [\ell(z; \theta) + (z - z_\star)^\top Q(z - z_\star)], \quad (3)$$

is also convex. Furthermore, this phenomenon does not just hold for quadratics, [Theorem 3.3](#) shows that the performative risk remains convex after we incorporate any convex function  $f(z)$  independent of  $\theta$ .

Reasoning about regularizers of this form illustrates another important difference between performatively stable solutions and performative optima. In particular, the set of stable points is the same for both  $\text{PR}(\theta)$  and its regularized version  $\text{PR}'(\theta)$  as defined in equation (3), since the regularization term is *independent* of  $\theta$ . Therefore, retraining algorithms such as RRM [28] or greedy/lazy SGD [22] essentially ignore any kind of signal provided by these regularizers and converge to the same point on both the regularized and unregularized objective. To find performatively optimal classifiers whose induced distributions satisfy socially desirable criteria, we must directly engage with and anticipate performative effects, rather than just passively retrain.

Lastly, to date, work on performative prediction has mostly studied the problem from a theoretical perspective. We believe that evaluating the ideas and different design choices raised by these papers in the context of specific applications and case studies on performative prediction would be of high value to the community and an exciting direction for future work.

## Acknowledgements

We thank Moritz Hardt and Celestine Mendler-Dünner for many helpful conversations during the course of this project as well as for providing detailed feedback on a draft of this manuscript.



We would also like to thank the anonymous reviewers whose comments helped improve the quality of our work.

This research was generously supported in part by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE 1752814.

## References

- [1] Alekh Agarwal and Ofer Dekel. Optimal algorithms for online convex optimization with multi-point bandit feedback. In *Conference on Learning Theory*, pages 28–40, 2010.
- [2] Alekh Agarwal, Sham Kakade, and Lin F Yang. Model-based reinforcement learning with a generative model is minimax optimal. In *Conference on Learning Theory*, pages 67–83, 2020.
- [3] 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.
- [4] Peter J Bickel. On adaptive estimation. *The Annals of Statistics*, pages 647–671, 1982.
- [5] Gavin Brown, Shlomi Hod, and Iden Kalemaj. Performative prediction in a stateful world. *arXiv preprint arXiv:2011.03885*, 2020.
- [6] Yiling Chen, Yang Liu, and Chara Podimata. Learning strategy-aware linear classifiers. In *Advances in Neural Information Processing Systems*, volume 33, pages 15265–15276, 2020.
- [7] Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, and Whitney Newey. Double/debiased/Neyman machine learning of treatment effects. *American Economic Review*, 107(5):261–65, 2017.
- [8] Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters, 2018.
- [9] 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. ACM, 2018.
- [10] Dmitriy Drusvyatskiy and Lin Xiao. Stochastic optimization with decision-dependent distributions. *arXiv preprint arXiv:2011.11173*, 2020.
- [11] Abraham D Flaxman, Adam Tauman Kalai, and H Brendan McMahan. Online convex optimization in the bandit setting: gradient descent without a gradient. In *Symposium on Discrete Algorithms*, pages 385–394, 2005.
- [12] Alex Frankel and Navin Kartik. Improving information from manipulable data. *Journal of the European Economic Association*, 2021.
- [13] Nika Haghtalab, Nicole Immorlica, Brendan Lucier, and Jack Z Wang. Maximizing welfare with incentive-aware evaluation mechanisms. In *International Joint Conference on Artificial Intelligence*, 2020.

- [14] Moritz Hardt, Nimrod Megiddo, Christos Papadimitriou, and Mary Wootters. Strategic classification. In *Proceedings of the ACM Conference on Innovations in Theoretical Computer Science*, pages 111–122, 2016.
- [15] Lily Hu, Nicole Immorlica, and Jennifer Wortman Vaughan. The disparate effects of strategic manipulation. In *Proceedings of the 2nd ACM Conference on Fairness, Accountability, and Transparency*, pages 259–268, 2019.
- [16] Il’dar Abdulovich Ibragimov and Rafail Zalmanovich Has’ Minskii. *Statistical estimation: asymptotic theory*, volume 16. Springer Science & Business Media, 2013.
- [17] Niki Kilbertus, Manuel Gomez Rodriguez, Bernhard Schölkopf, Krikamol Muandet, and Isabel Valera. Fair decisions despite imperfect predictions. In *International Conference on Artificial Intelligence and Statistics*, pages 277–287, 2020.
- [18] B Ya Levit. On the efficiency of a class of non-parametric estimates. *Theory of Probability & Its Applications*, 20(4):723–740, 1976.
- [19] Lester Mackey, Vasilis Syrgkanis, and Ilias Zadik. Orthogonal machine learning: Power and limitations. In *International Conference on Machine Learning*, pages 3375–3383, 2018.
- [20] Horia Mania, Stephen Tu, and Benjamin Recht. Certainty equivalence is efficient for linear quadratic control. In *Advances in Neural Information Processing Systems*, pages 10154–10164, 2019.
- [21] Nikolai Matni and Stephen Tu. A tutorial on concentration bounds for system identification. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 3741–3749. IEEE, 2019.
- [22] Celestine Mendler-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.
- [23] Smitha Milli, John Miller, Anca D Dragan, and Moritz Hardt. The social cost of strategic classification. In *Proceedings of the 2nd ACM Conference on Fairness, Accountability, and Transparency*, pages 230–239, 2019.
- [24] Alfred Müller and Ludger Rüschendorf. On the optimal stopping values induced by general dependence structures. *Journal of applied probability*, pages 672–684, 2001.
- [25] Alfred Müller and Dietrich Stoyan. *Comparison methods for stochastic models and risks*, volume 389. Wiley, 2002.
- [26] Evan Munro. Learning to personalize treatments when agents are strategic. *arXiv preprint arXiv:2011.06528*, 2020.
- [27] Whitney K Newey. Semiparametric efficiency bounds. *Journal of applied econometrics*, 5(2):99–135, 1990.
- [28] Juan Perdomo, Tijana Zrnic, Celestine Mendler-Dünner, and Moritz Hardt. Performative prediction. In *International Conference on Machine Learning*, pages 7599–7609, 2020.

- [29] Peter M Robinson. Root-n-consistent semiparametric regression. *Econometrica: Journal of the Econometric Society*, pages 931–954, 1988.
- [30] Horn Roger and R Johnson Charles. Topics in matrix analysis, 1994.
- [31] Nir Rosenfeld, Anna Hilgard, Sai Srivatsa Ravindranath, and David C Parkes. From predictions to decisions: Using lookahead regularization. In *Advances in Neural Information Processing Systems*, volume 33, pages 4115–4126, 2020.
- [32] Sheldon M Ross, John J Kelly, Roger J Sullivan, William James Perry, Donald Mercer, Ruth M Davis, Thomas Dell Washburn, Earl V Sager, Joseph B Boyce, and Vincent L Bristow. *Stochastic processes*, volume 2. Wiley New York, 1996.
- [33] Moshe Shaked and J George Shanthikumar. *Stochastic orders*. Springer Science & Business Media, 2007.
- [34] Shai Shalev-Shwartz, Ohad Shamir, Nathan Srebro, and Karthik Sridharan. Learnability, stability and uniform convergence. *The Journal of Machine Learning Research*, 11:2635–2670, 2010.
- [35] Ohad Shamir. On the complexity of bandit and derivative-free stochastic convex optimization. In *Conference on Learning Theory*, pages 3–24, 2013.
- [36] Yonadav Shavit, Benjamin Edelman, and Brian Axelrod. Causal strategic linear regression. In *International Conference on Machine Learning*, pages 8676–8686, 2020.
- [37] Max Simchowitz and Dylan Foster. Naive exploration is optimal for online LQR. In *International Conference on Machine Learning*, pages 8937–8948, 2020.
- [38] Herbert A Simon. Dynamic programming under uncertainty with a quadratic criterion function. *Econometrica, Journal of the Econometric Society*, pages 74–81, 1956.
- [39] Behzad Tabibian, Vicenç Gomez, Abir De, Bernhard Schölkopf, and Manuel Gomez Rodriguez. On the design of consequential ranking algorithms. In *Conference on Uncertainty in Artificial Intelligence*, pages 171–180, 2020.
- [40] Henri Theil. A note on certainty equivalence in dynamic planning. *Econometrica: Journal of the Econometric Society*, pages 346–349, 1957.
- [41] Stratis Tsirtsis and Manuel Gomez Rodriguez. Decisions, counterfactual explanations and strategic behavior. In *Advances in Neural Information Processing Systems*, volume 33, pages 16749–16760, 2020.
- [42] Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.
- [43] Stefan Wager and Kuang Xu. Experimenting in equilibrium. *Management Science*, 2021.
- [44] Sean Jeremy Westwood, Solomon Messing, and Yphtach Lelkes. Projecting confidence: How the probabilistic horse race confuses and demobilizes the public. *The Journal of Politics*, 82(4):1530–1544, 2020.

## A Background on Stochastic Orders

In this section we provide the necessary preliminaries from the literature on stochastic orders.

First, we recall the notion of the *convex order*: for two random vectors  $z, z' \in \mathbb{R}^m$ , we say that  $z$  is less than  $z'$  in the convex order, denoted  $z \leq_{cx} z'$ , if for all convex functions  $g: \mathbb{R}^m \rightarrow \mathbb{R}$ , it holds that

$$\mathbb{E} g(z) \leq \mathbb{E} g(z').$$

Using a slight abuse of notation, we will also write  $\mathcal{D}_1 \leq_{cx} \mathcal{D}_2$  for two distributions  $\mathcal{D}_1, \mathcal{D}_2$  when  $z \sim \mathcal{D}_1, z' \sim \mathcal{D}_2$  and  $z \leq_{cx} z'$ .

Therefore, an immediate way to satisfy condition (A4) is to assume that the loss function  $\ell(z; \theta)$  is convex in  $z$ , and to require  $\mathcal{D}(\alpha\theta + (1-\alpha)\theta') \leq_{cx} \alpha\mathcal{D}(\theta) + (1-\alpha)\mathcal{D}(\theta')$ . The latter condition has been long studied in classical statistical literature and many equivalent characterizations are known (see, e.g., [25, 32, 33]). This leads to the following corollary of Theorem 3.1.

**Corollary A.1.** *Suppose that the loss function is  $\gamma$ -strongly convex in  $\theta$  (A3a) and  $\beta$ -smooth in  $z$  (A2), and that the distribution map  $\mathcal{D}(\cdot)$  is  $\varepsilon$ -sensitive (A1). Further, assume that  $\ell(z; \theta)$  is convex in  $z$  and that  $\mathcal{D}(\alpha\theta + (1-\alpha)\theta') \leq_{cx} \alpha\mathcal{D}(\theta) + (1-\alpha)\mathcal{D}(\theta')$ . Then, the performative risk  $\text{PR}(\theta)$  is  $(\gamma - 2\varepsilon\beta)$ -convex.*

Now we discuss important families of distributions that satisfy the convex order condition  $\mathcal{D}(\alpha\theta + (1-\alpha)\theta') \leq_{cx} \alpha\mathcal{D}(\theta) + (1-\alpha)\mathcal{D}(\theta')$ .

**Example A.2.** An obvious example where  $\mathcal{D}(\alpha\theta + (1-\alpha)\theta') \leq_{cx} \alpha\mathcal{D}(\theta) + (1-\alpha)\mathcal{D}(\theta')$  is when  $\mathcal{D}(\alpha\theta + (1-\alpha)\theta') = \alpha\mathcal{D}(\theta) + (1-\alpha)\mathcal{D}(\theta')$ . An important setting which satisfies this linearity property is when the probability of a positive outcome of a binary variable is linear in  $\theta$ :  $z_\theta \sim \text{Bern}(a + w^\top \theta)$  defines  $z_\theta \sim \mathcal{D}(\theta)$ . In this case,  $\mathcal{D}(\alpha\theta + (1-\alpha)\theta') = \alpha\mathcal{D}(\theta) + (1-\alpha)\mathcal{D}(\theta')$ .

For further examples, we invoke a convenient characterization of the convex order condition.

**Lemma A.3** ([24]). *Two random vectors  $z$  and  $z'$  satisfy  $z \leq_{cx} z'$  if and only if there exists a coupling of  $z$  and  $z'$  such that  $\mathbb{E}[z'|z] = z$  a.s.*

By applying Lemma A.3, we show that the important case of *location-scale families* satisfies the convex order condition. Therefore, if the loss function is additionally convex in  $z$ , condition (A4) follows.

**Proposition A.4.** *Suppose that  $\mathcal{D}(\theta)$  forms a location-scale family (2) such that  $\Sigma_0 + \Sigma(\theta)$  has full rank for all  $\theta \in \Theta$ . Then,  $\mathcal{D}(\alpha\theta + (1-\alpha)\theta') \leq_{cx} \alpha\mathcal{D}(\theta) + (1-\alpha)\mathcal{D}(\theta')$  for all  $\theta, \theta' \in \Theta$ .*

*Proof.* We will construct a coupling  $(z, z')$  such that  $z \sim \mathcal{D}(\alpha\theta + (1-\alpha)\theta')$ ,  $z' \sim \alpha\mathcal{D}(\theta) + (1-\alpha)\mathcal{D}(\theta')$ , and  $\mathbb{E}[z'|z] = z$ . Let  $z \sim \mathcal{D}(\alpha\theta + (1-\alpha)\theta')$ ; then we define  $z'$  in terms of  $z$  as

$$z' = (\Sigma_0 + \Sigma(G))(\Sigma_0 + \Sigma(\alpha\theta + (1-\alpha)\theta'))^{-1} (z - \mu_0 - \mu(\alpha\theta + (1-\alpha)\theta')) + \mu_0 + \mu G, \quad (4)$$

where

$$G = \begin{cases} \theta, & \text{with probability } \alpha, \\ \theta', & \text{with probability } 1 - \alpha \end{cases}$$

is independent of  $z$ . Notice that

$$\begin{aligned}\mathbb{E}[z' | z] &= \mathbb{E}\left[(\Sigma_0 + \Sigma(G))(\Sigma_0 + \Sigma(\alpha\theta + (1-\alpha)\theta'))^{-1} (z - \mu_0 - \mu(\alpha\theta + (1-\alpha)\theta')) + \mu_0 + \mu G \mid z\right] \\ &= (\Sigma_0 + \mathbb{E}[\Sigma(G)])(\Sigma_0 + \Sigma(\alpha\theta + (1-\alpha)\theta'))^{-1} (z - \mu_0 - \mu(\alpha\theta + (1-\alpha)\theta')) + \mu_0 + \mathbb{E}[\mu G] \\ &= z,\end{aligned}$$

which follows by linearity of  $\mu$  and  $\Sigma(\cdot)$  and the fact that  $\mathbb{E}[G] = \alpha\theta + (1-\alpha)\theta'$ .

We now only need to verify that  $z' \sim \alpha\mathcal{D}(\theta) + (1-\alpha)\mathcal{D}(\theta')$  in order to apply [Lemma A.3](#) and conclude that  $z' \leq_{cx} z$ . Indeed, with probability  $\alpha$  we have  $G = \theta$ , and on that event  $z' \stackrel{d}{=} (\Sigma_0 + \Sigma(\theta))z_0 + \mu_0 + \mu\theta$ ; a similar argument applies to  $\theta'$ . Therefore, putting everything together we conclude that  $z \leq_{cx} z'$ .  $\blacksquare$

[Proposition A.4](#) implies that for all convex functions  $g : \mathbb{R}^m \rightarrow \mathbb{R}$ ,

$$\mathbb{E}_{z \sim \mathcal{D}(\alpha\theta + (1-\alpha)\theta')} [g(z)] \leq \mathbb{E}_{z \sim \alpha\mathcal{D}(\theta) + (1-\alpha)\mathcal{D}(\theta')} [g(z)].$$

We now show that for *strongly* convex  $g$ , this conclusion can be made even stronger. This result allows for deriving a tighter version of [Theorem 3.1](#) for the important class of location-scale families, stated in [Theorem 3.3](#).

**Proposition A.5.** *Let  $g : \mathbb{R}^m \rightarrow \mathbb{R}$  be a  $\gamma_z$ -strongly convex function for some  $\gamma_z \geq 0$ , and let  $\mathcal{D}(\theta)$  form a location-scale family (2). Then,*

$$\mathbb{E}_{z \sim \mathcal{D}(\alpha\theta + (1-\alpha)\theta')} [g(z)] \leq \mathbb{E}_{z \sim \alpha\mathcal{D}(\theta) + (1-\alpha)\mathcal{D}(\theta')} [g(z)] - \frac{\alpha(1-\alpha)\gamma_z}{2} \mathbb{E} \|\Sigma(\theta - \theta')z_0 + \mu(\theta - \theta')\|_2^2.$$

*Proof.* Since  $g$  is strongly convex, we can write  $g(z) = g_0(z) + \frac{\gamma_z}{2} \|z\|_2^2$ , where  $g_0$  is a convex function. Thus, we want to prove

$$\begin{aligned}\mathbb{E}_{z \sim \mathcal{D}(\alpha\theta + (1-\alpha)\theta')} \left[ g_0(z) + \frac{\gamma_z}{2} \|z\|_2^2 \right] &\leq \mathbb{E}_{z' \sim \alpha\mathcal{D}(\theta) + (1-\alpha)\mathcal{D}(\theta')} \left[ g_0(z') + \frac{\gamma_z}{2} \|z'\|_2^2 \right] \\ &\quad - \frac{\alpha(1-\alpha)\gamma_z}{2} \mathbb{E} \|\Sigma(\theta - \theta')z_0 + \mu(\theta - \theta')\|_2^2.\end{aligned}$$

By [Proposition A.4](#), we know that

$$\mathbb{E}_{z \sim \mathcal{D}(\alpha\theta + (1-\alpha)\theta')} [g_0(z)] \leq \mathbb{E}_{z \sim \alpha\mathcal{D}(\theta) + (1-\alpha)\mathcal{D}(\theta')} [g_0(z)].$$

Therefore, we only need to argue that

$$\mathbb{E} \left[ \|z'\|_2^2 - \|z\|_2^2 \right] \geq \alpha(1-\alpha) \mathbb{E} \|\Sigma(\theta - \theta')z_0 + \mu(\theta - \theta')\|_2^2.$$

Without loss of generality, we take  $z, z'$  to be coupled as in equation (4). Then, we can write

$$\begin{aligned}\mathbb{E} \left[ \|z'\|_2^2 - \|z\|_2^2 \right] &= \mathbb{E} \left[ \|z' - z\|_2^2 + 2(z' - z)^\top z \right] \\ &= \mathbb{E} \left[ \|z' - z\|_2^2 \right] \\ &= \mathbb{E} \left[ \left\| \Sigma(G - (\alpha\theta + (1-\alpha)\theta'))z_0 + \mu(G - (\alpha\theta + (1-\alpha)\theta')) \right\|_2^2 \right],\end{aligned}$$

where the second steps follows by iterating expectations, because  $\mathbb{E}[z'|z] = z$ .

By further taking an expectation over  $G$ , we get:

$$\begin{aligned} & \mathbb{E} \left[ \left\| \Sigma(G - (\alpha\theta + (1-\alpha)\theta'))z_0 + \mu(G - (\alpha\theta + (1-\alpha)\theta')) \right\|_2^2 \right] \\ &= \alpha(1-\alpha)^2 \mathbb{E} \left\| \Sigma(\theta - \theta')z_0 + \mu(\theta - \theta') \right\|_2^2 + (1-\alpha)\alpha^2 \mathbb{E} \left\| \Sigma(\theta' - \theta)z_0 + \mu(\theta - \theta') \right\|_2^2 \\ &= \alpha(1-\alpha) \mathbb{E} \left\| \Sigma(\theta - \theta')z_0 + \mu(\theta - \theta') \right\|_2^2. \end{aligned}$$

■

## B Distinguishing between Static and Performative Variables

In many natural examples, the performative effects are only present in a subset of the variables that make up  $z$ . For example, in strategic classification, the performative effects are often only present in the strategically manipulated features, and not in the label. In [Example 3.2](#), on the other hand, the effects are only present in the label. For simplicity of exposition, we suppress this distinction between *performative* and *static* variables, that is, those whose distribution does not change for different  $\mathcal{D}(\theta)$ . However, the reader should think of all assumptions on  $z$ , such as strong convexity or various Lipschitz assumptions, as only having to apply to the performative variables, while the static ones can be averaged out. To give one example, suppose that  $z = (z_s, z_p)$ , where  $z_s$  denotes the static variables and  $z_p$  denotes the performative ones. Using this distinction, the step in equation (6) would proceed as follows:

$$\begin{aligned} & \mathbb{E}_{(z_s, z_p) \sim \mathcal{D}(\theta)} [\nabla_{\theta} \ell((z_s, z_p); \theta)]^{\top} (\theta' - \theta) - \mathbb{E}_{(z_s, z'_p) \sim \mathcal{D}(\theta')} [\nabla_{\theta} \ell((z_s, z'_p); \theta)]^{\top} (\theta' - \theta) \\ &= \mathbb{E}_{z_s} \left[ \left( \mathbb{E} [\nabla_{\theta} \ell((z_s, z_p); \theta) | z_s] - \mathbb{E} [\nabla_{\theta} \ell((z_s, z'_p); \theta) | z_s] \right)^{\top} (\theta' - \theta) \right] \\ &\leq \mathbb{E}_{z_s} [\beta(z_s) \varepsilon(z_s)] \|\theta - \theta'\|_2^2. \end{aligned}$$

Here,  $\beta(z_s)$  is the Lipschitz constant of  $\nabla_{\theta} \ell((z_s, \cdot); \theta)$ , and  $\varepsilon(z_s)$  is the sensitivity parameter of the distribution of  $z_p$ , conditional on  $z_s$ . As clear from the above example, stating all conditions and proofs while emphasizing this distinction is fairly cumbersome, so we opted for a simplified presentation. Similar calculations can be carried out for the rest of the proofs of the structural results.

## C Deferred Proofs

### C.1 Convexity of the Performative Risk

**Proof of [Theorem 3.1](#).** We begin by writing out the gradient of the performative risk:

$$\begin{aligned} \nabla_{\theta} \text{PR}(\theta) &= \nabla_{\theta} \left( \int \ell(z; \theta) p_{\theta}(z) dz \right) = \int \nabla_{\theta} \ell(z; \theta) p_{\theta}(z) dz + \int \ell(z; \theta) \nabla_{\theta} p_{\theta}(z) dz \\ &= \int \nabla_{\theta} \ell(z; \theta) p_{\theta}(z) dz + \int \ell(z; \theta) \nabla_{\theta} \log(p_{\theta}(z)) p_{\theta}(z) dz \\ &= \mathbb{E}_{z \sim \mathcal{D}(\theta)} [\nabla_{\theta} \ell(z; \theta)] + \mathbb{E}_{z \sim \mathcal{D}(\theta)} [\ell(z; \theta) \nabla_{\theta} \log(p_{\theta}(z))]. \end{aligned}$$



By the first-order condition for convexity, we know that  $\text{PR}(\theta)$  is  $(\gamma - 2\varepsilon\beta)$ -convex if and only if

$$\left( \mathbb{E}_{z \sim \mathcal{D}(\theta)} [\nabla_{\theta} \ell(z; \theta) + \ell(z; \theta) \nabla_{\theta} \log(p_{\theta}(z))] \right)^{\top} (\theta' - \theta) + \frac{\gamma - 2\varepsilon\beta}{2} \|\theta - \theta'\|_2^2 \leq \text{PR}(\theta') - \text{PR}(\theta), \quad (5)$$

for all  $\theta, \theta' \in \Theta$ . By assumption (A4), we know that for all  $\theta, \theta', \theta_0 \in \Theta$ ,

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

This assumption is equivalent to saying that  $g_{\theta_0}(\theta) = \mathbb{E}_{z \sim \mathcal{D}(\theta)} [\ell(z; \theta_0)]$  is a convex function of  $\theta$ , for all  $\theta_0$ . We can express this convexity condition using the equivalent first-order characterization:

$$\mathbb{E}_{z \sim \mathcal{D}(\theta)} [\ell(z; \theta_0) \nabla_{\theta} \log(p_{\theta}(z))]^{\top} (\theta' - \theta) \leq \mathbb{E}_{z \sim \mathcal{D}(\theta')} [\ell(z; \theta_0)] - \mathbb{E}_{z \sim \mathcal{D}(\theta)} [\ell(z; \theta_0)].$$

Since the mixture dominance condition holds for all  $\theta, \theta'$  and  $\theta_0$ , we can set  $\theta_0$  equal to  $\theta$  in the inequality above to conclude that

$$\mathbb{E}_{z \sim \mathcal{D}(\theta)} [\ell(z; \theta) \nabla_{\theta} \log(p_{\theta}(z))]^{\top} (\theta' - \theta) \leq \mathbb{E}_{z \sim \mathcal{D}(\theta')} [\ell(z; \theta)] - \mathbb{E}_{z \sim \mathcal{D}(\theta)} [\ell(z; \theta)].$$

Going back to equation (5), we see that a sufficient condition for  $(\gamma - 2\varepsilon\beta)$ -convexity of the performative risk is

$$\mathbb{E}_{z \sim \mathcal{D}(\theta)} [\nabla_{\theta} \ell(z; \theta)]^{\top} (\theta' - \theta) + \frac{\gamma - 2\varepsilon\beta}{2} \|\theta - \theta'\|_2^2 \leq \mathbb{E}_{z \sim \mathcal{D}(\theta')} \ell(z; \theta') - \mathbb{E}_{z \sim \mathcal{D}(\theta)} \ell(z; \theta).$$

By the assumption that the loss is  $\gamma$ -strongly convex in  $\theta$ , we know

$$\mathbb{E}_{z \sim \mathcal{D}(\theta')} \ell(z; \theta') - \mathbb{E}_{z \sim \mathcal{D}(\theta')} \ell(z; \theta) \geq \mathbb{E}_{z \sim \mathcal{D}(\theta')} [\nabla_{\theta} \ell(z; \theta)]^{\top} (\theta' - \theta) + \frac{\gamma}{2} \|\theta - \theta'\|_2^2,$$

and thus we have further simplified the sufficient condition to

$$\mathbb{E}_{z \sim \mathcal{D}(\theta)} [\nabla_{\theta} \ell(z; \theta)]^{\top} (\theta' - \theta) - \mathbb{E}_{z \sim \mathcal{D}(\theta')} [\nabla_{\theta} \ell(z; \theta)]^{\top} (\theta' - \theta) \leq \frac{2\varepsilon\beta}{2} \|\theta - \theta'\|_2^2.$$

Since the loss is  $\beta$ -smooth in  $z$ , we have that  $\nabla_{\theta} \ell(z; \theta)^{\top} (\theta' - \theta)$  is  $\beta \|\theta - \theta'\|_2$ -Lipschitz in  $z$ . Now, we can use the fact that the distribution map is  $\varepsilon$ -sensitive to upper bound the left-hand side by applying the Kantorovich-Rubinstein duality theorem:

$$\mathbb{E}_{z \sim \mathcal{D}(\theta)} [\nabla_{\theta} \ell(z; \theta)]^{\top} (\theta' - \theta) - \mathbb{E}_{z \sim \mathcal{D}(\theta')} [\nabla_{\theta} \ell(z; \theta)]^{\top} (\theta' - \theta) \leq \varepsilon \beta \|\theta - \theta'\|_2^2. \quad (6)$$

Therefore, we can conclude that the performative risk is  $(\gamma - 2\varepsilon\beta)$ -convex.

**Proof of Theorem 3.3.** Following the steps of Theorem 3.1, we know that  $\text{PR}(\theta)$  is  $\lambda$ -convex if and only if

$$\mathbb{E}_{z \sim \mathcal{D}(\theta)} [\nabla_{\theta} \ell(z; \theta)]^{\top} (\theta' - \theta) + \mathbb{E}_{z \sim \mathcal{D}(\theta)} [\ell(z; \theta) \nabla_{\theta} \log(p_{\theta}(z))]^{\top} (\theta' - \theta) + \frac{\lambda}{2} \|\theta - \theta'\|_2^2 \leq \text{PR}(\theta') - \text{PR}(\theta),$$

for all  $\theta, \theta' \in \Theta$ .

We now state a technical lemma which rephrases the conclusion of Proposition A.5 in an equivalent way, deferring its proof to the end of this section.

**Lemma C.1.** *Suppose that*

$$\mathbb{E}_{z \sim \mathcal{D}(\alpha\theta + (1-\alpha)\theta')} [g(z)] \leq \mathbb{E}_{z \sim \alpha\mathcal{D}(\theta) + (1-\alpha)\mathcal{D}(\theta')} [g(z)] - \frac{\alpha(1-\alpha)\gamma_z}{2} \mathbb{E} \|\Sigma(\theta - \theta')z_0 + \mu(\theta - \theta')\|_2^2.$$

Then,

$$\mathbb{E}_{z \sim \mathcal{D}(\theta')} [g(z)] \geq \mathbb{E}_{z \sim \mathcal{D}(\theta)} [g(z)] + (\nabla_{\theta} \mathbb{E}_{z \sim \mathcal{D}(\theta)} [g(z)])^\top (\theta' - \theta) + \frac{\gamma_z}{2} \mathbb{E} \|\Sigma(\theta - \theta')z_0 + \mu(\theta - \theta')\|_2^2.$$

Therefore, by [Proposition A.5](#) and [Lemma C.1](#), we know

$$\begin{aligned} \mathbb{E}_{z \sim \mathcal{D}(\theta)} [\ell(z; \theta) \nabla_{\theta} \log(p_{\theta}(z))]^\top (\theta' - \theta) &\leq \mathbb{E}_{z \sim \mathcal{D}(\theta')} [\ell(z; \theta')] - \mathbb{E}_{z \sim \mathcal{D}(\theta)} [\ell(z; \theta)] \\ &\quad - \frac{\gamma_z}{2} \mathbb{E} \|\Sigma(\theta - \theta')z_0 + \mu(\theta - \theta')\|_2^2, \end{aligned}$$

where we take  $g(z) = \ell(z; \theta)$ .

Thus it suffices to show

$$\mathbb{E}_{z \sim \mathcal{D}(\theta)} [\nabla_{\theta} \ell(z; \theta)]^\top (\theta' - \theta) + \frac{\lambda}{2} \|\theta - \theta'\|_2^2 \leq \mathbb{E}_{z \sim \mathcal{D}(\theta')} \ell(z; \theta') - \mathbb{E}_{z \sim \mathcal{D}(\theta)} \ell(z; \theta) + \frac{\gamma_z}{2} \mathbb{E} \|\Sigma(\theta - \theta')z_0 + \mu(\theta - \theta')\|_2^2.$$

By the assumption that the loss is  $\gamma$ -strongly convex, we know

$$\mathbb{E}_{z \sim \mathcal{D}(\theta')} \ell(z; \theta') - \mathbb{E}_{z \sim \mathcal{D}(\theta)} \ell(z; \theta) \geq \mathbb{E}_{z \sim \mathcal{D}(\theta')} [\nabla_{\theta} \ell(z; \theta)]^\top (\theta' - \theta) + \frac{\gamma}{2} \|\theta - \theta'\|_2^2.$$

With this, we have simplified the sufficient condition for  $\gamma$ -convexity to

$$\left( \mathbb{E}_{z \sim \mathcal{D}(\theta)} [\nabla_{\theta} \ell(z; \theta)] - \mathbb{E}_{z \sim \mathcal{D}(\theta')} [\nabla_{\theta} \ell(z; \theta)] \right)^\top (\theta' - \theta) \leq \frac{\gamma - \lambda}{2} \|\theta - \theta'\|_2^2 + \frac{\gamma_z}{2} \mathbb{E} \|\Sigma(\theta - \theta')z_0 + \mu(\theta - \theta')\|_2^2. \quad (7)$$

We bound the left-hand side by applying smoothness of the loss together with the Kantorovich-Rubinstein duality theorem; for this, we need a bound on  $W(\mathcal{D}(\theta), \mathcal{D}(\theta'))$ . We will use the bound implied by  $\varepsilon$ -sensitivity, as well as the bound implied by the following lemma.

**Lemma C.2.** *Suppose that the distribution map  $\mathcal{D}(\theta)$  forms a location-scale family (2). Then,*

$$W(\mathcal{D}(\theta), \mathcal{D}(\theta')) \leq \mathbb{E} \|\Sigma(\theta - \theta')z_0 + \mu(\theta - \theta')\|_2.$$

*Proof of Lemma C.2.* By definition,  $W(\mathcal{D}(\theta), \mathcal{D}(\theta')) = \inf_{\Pi(\mathcal{D}(\theta), \mathcal{D}(\theta'))} \mathbb{E}_{(z_{\theta}, z_{\theta'}) \sim \Pi(\mathcal{D}(\theta), \mathcal{D}(\theta'))} [\|z_{\theta} - z_{\theta'}\|_2]$ , where  $\Pi(\mathcal{D}(\theta), \mathcal{D}(\theta'))$  denotes a coupling of  $\mathcal{D}(\theta)$  and  $\mathcal{D}(\theta')$ . The simplest way to couple  $\mathcal{D}(\theta)$  and  $\mathcal{D}(\theta')$ , or equivalently  $z_{\theta}$  and  $z_{\theta'}$ , is to sample  $z_0 \sim \mathcal{D}$ , and set  $z_{\theta} = (\Sigma_0 + \Sigma(\theta))z_0 + \mu_0 + \mu(\theta)$  and  $z_{\theta'} = (\Sigma_0 + \Sigma(\theta'))z_0 + \mu_0 + \mu(\theta')$ . With this choice,  $\|z_{\theta} - z_{\theta'}\|_2 = \|\Sigma(\theta - \theta')z_0 + \mu(\theta - \theta')\|_2$ , and hence  $W(\mathcal{D}(\theta), \mathcal{D}(\theta')) \leq \mathbb{E} \|\Sigma(\theta - \theta')z_0 + \mu(\theta - \theta')\|_2$ . ■

Therefore, the left-hand side in equation (7) can be bounded by

$$\mathbb{E}_{z \sim \mathcal{D}(\theta)} [\nabla_{\theta} \ell(z; \theta)]^\top (\theta' - \theta) - \mathbb{E}_{z \sim \mathcal{D}(\theta')} [\nabla_{\theta} \ell(z; \theta)]^\top (\theta' - \theta) \leq \beta \mathbb{E} \|\Sigma(\theta - \theta')z_0 + \mu(\theta - \theta')\|_2 \|\theta' - \theta\|_2,$$

but also by applying  $\varepsilon$ -sensitivity

$$\mathbb{E}_{z \sim \mathcal{D}(\theta)} [\nabla_{\theta} \ell(z; \theta)]^{\top} (\theta' - \theta) - \mathbb{E}_{z \sim \mathcal{D}(\theta')} [\nabla_{\theta} \ell(z; \theta)]^{\top} (\theta' - \theta) \leq \beta \varepsilon \|\theta' - \theta\|_2^2.$$

Finally, to show  $\lambda = \max\{\gamma - \beta^2/\gamma_z, \gamma + \gamma_z(\sigma_{\min}^2(\mu) + \sigma_{\min}^2(\Sigma)) - 2\beta\varepsilon\}$ -convexity it suffices to show both

$$\beta \mathbb{E} \|\Sigma(\theta - \theta')z_0 + \mu(\theta - \theta')\|_2 \|\theta' - \theta\|_2 \leq \frac{\beta^2/\gamma_z}{2} \|\theta - \theta'\|_2^2 + \frac{\gamma_z}{2} \mathbb{E} \|\Sigma(\theta - \theta')z_0 + \mu(\theta - \theta')\|_2^2 \quad (8)$$

and

$$\beta \varepsilon \|\theta' - \theta\|_2^2 \leq \frac{2\beta\varepsilon - \gamma_z(\sigma_{\min}^2(\mu) + \sigma_{\min}^2(\Sigma))}{2} \|\theta - \theta'\|_2^2 + \frac{\gamma_z}{2} \mathbb{E} \|\Sigma(\theta - \theta')z_0 + \mu(\theta - \theta')\|_2^2. \quad (9)$$

By the AM-GM inequality, we have

$$\beta \mathbb{E} \|\Sigma(\theta - \theta')z_0 + \mu(\theta - \theta')\|_2 \|\theta' - \theta\|_2 \leq \frac{1}{2} \frac{\beta^2}{\gamma_z} \|\theta' - \theta\|_2^2 + \frac{\gamma_z}{2} \mathbb{E} \|\Sigma(\theta - \theta')z_0 + \mu(\theta - \theta')\|_2^2,$$

and so condition (8) follows.

For condition (9), we observe that

$$\begin{aligned} \mathbb{E} \|\Sigma(\theta - \theta')z_0 + \mu(\theta - \theta')\|_2^2 &= \mathbb{E} \|\Sigma(\theta - \theta')z_0\|_2^2 + \|\mu(\theta - \theta')\|_2^2 \\ &= \text{Tr}\left(\Sigma(\theta - \theta')\Sigma_{z_0}\Sigma(\theta - \theta')^{\top}\right) + \|\mu(\theta - \theta')\|_2^2 \\ &= \|\Sigma_{z_0}^{1/2}\Sigma(\theta - \theta')^{\top}\|_F^2 + \|\mu(\theta - \theta')\|_2^2. \end{aligned}$$

Applying  $\sigma_{\min}(\Sigma)\|\theta - \theta'\|_2 \leq \|\Sigma_{z_0}^{1/2}\Sigma(\theta - \theta')^{\top}\|_F$  and  $\sigma_{\min}(\mu)\|\theta - \theta'\|_2 \leq \|\mu(\theta - \theta')\|_2$  completes the proof of the theorem.

*Proof of Lemma C.1.* The proof follows the standard argument for proving equivalent formulations of strong convexity.

First we show that  $\mathbb{E}_{z \sim \mathcal{D}(\theta)}[g(z)] - \frac{\gamma_z}{2} \mathbb{E} \|\Sigma(\theta)z_0 + \mu\theta\|_2^2$  is convex in  $\theta$ . This follows because:

$$\begin{aligned} &\mathbb{E}_{z \sim \mathcal{D}(\alpha\theta + (1-\alpha)\theta')} [g(z)] - \frac{\gamma_z}{2} \mathbb{E} \|\Sigma(\alpha\theta + (1-\alpha)\theta')z_0 + \mu(\alpha\theta + (1-\alpha)\theta')\|_2^2 \\ &\leq \mathbb{E}_{z \sim \alpha\mathcal{D}(\theta) + (1-\alpha)\mathcal{D}(\theta')} [g(z)] - \frac{\alpha(1-\alpha)\gamma_z}{2} \mathbb{E} \|\Sigma(\theta - \theta')z_0 + \mu(\theta - \theta')\|_2^2 \\ &\quad - \frac{\gamma_z}{2} \mathbb{E} \|\Sigma(\alpha\theta + (1-\alpha)\theta')z_0 + \mu(\alpha\theta + (1-\alpha)\theta')\|_2^2 \\ &= \mathbb{E}_{z \sim \alpha\mathcal{D}(\theta) + (1-\alpha)\mathcal{D}(\theta')} [g(z)] - \frac{\gamma_z}{2} \alpha^2 \mathbb{E} \|\Sigma(\theta)z_0 + \mu\theta\|_2^2 - \frac{\gamma_z}{2} (1-\alpha)^2 \mathbb{E} \|\Sigma(\theta')z_0 + \mu\theta'\|_2^2 \\ &\quad + \frac{\gamma_z}{2} 2\alpha(1-\alpha) \mathbb{E} (\Sigma(\theta) + \mu\theta)^{\top} (\Sigma(\theta') + \mu\theta') - \frac{\alpha(1-\alpha)\gamma_z}{2} \mathbb{E} \|\Sigma(\theta - \theta')z_0 + \mu(\theta - \theta')\|_2^2 \\ &= \mathbb{E}_{z \sim \alpha\mathcal{D}(\theta) + (1-\alpha)\mathcal{D}(\theta')} [g(z)] - \frac{\gamma_z}{2} \alpha \mathbb{E} \|\Sigma(\theta)z_0 + \mu\theta\|_2^2 - \frac{\gamma_z}{2} (1-\alpha) \mathbb{E} \|\Sigma(\theta')z_0 + \mu\theta'\|_2^2 \\ &= \alpha \left( \mathbb{E}_{z \sim \mathcal{D}(\theta)} [g(z)] - \frac{\gamma_z}{2} \mathbb{E} \|\Sigma(\theta)z_0 + \mu\theta\|_2^2 \right) - (1-\alpha) \left( \mathbb{E}_{z \sim \mathcal{D}(\theta')} [g(z)] - \frac{\gamma_z}{2} \mathbb{E} \|\Sigma(\theta')z_0 + \mu\theta'\|_2^2 \right). \end{aligned}$$

By the equivalent first-order characterization, this means that

$$\begin{aligned}
\mathbb{E}_{z \sim \mathcal{D}(\theta')} [g(z)] &\geq \frac{\gamma_z}{2} \mathbb{E} \|\Sigma(\theta')z_0 + \mu\theta'\|_2^2 + \mathbb{E}_{z \sim \mathcal{D}(\theta)} [g(z)] - \frac{\gamma_z}{2} \mathbb{E} \|\Sigma(\theta)z_0 + \mu\theta\|_2^2 \\
&\quad + (\nabla_{\theta} \mathbb{E}_{z \sim \mathcal{D}(\theta)} [g(z)])^\top (\theta' - \theta) - \frac{\gamma_z}{2} 2 \mathbb{E}(\Sigma(\theta)z_0 + \mu\theta)^\top (\nabla_{\theta}(\Sigma(\theta)z_0 + \mu\theta))^\top (\theta' - \theta) \\
&\geq \frac{\gamma_z}{2} \mathbb{E} \|\Sigma(\theta')z_0 + \mu\theta'\|_2^2 + \mathbb{E}_{z \sim \mathcal{D}(\theta)} [g(z)] - \frac{\gamma_z}{2} \mathbb{E} \|\Sigma(\theta)z_0 + \mu\theta\|_2^2 \\
&\quad + (\nabla_{\theta} \mathbb{E}_{z \sim \mathcal{D}(\theta)} [g(z)])^\top (\theta' - \theta) - \gamma_z \mathbb{E}(\Sigma(\theta)z_0 + \mu\theta)^\top (\Sigma(\theta' - \theta)z_0 + \mu(\theta' - \theta)) \\
&= \mathbb{E}_{z \sim \mathcal{D}(\theta)} [g(z)] + (\nabla_{\theta} \mathbb{E}_{z \sim \mathcal{D}(\theta)} [g(z)])^\top (\theta' - \theta) + \frac{\gamma_z}{2} \mathbb{E} \|\Sigma(\theta - \theta')z_0 + \mu(\theta - \theta')\|_2^2.
\end{aligned}$$

■

**Remark C.3.** We note that the sensitivity parameter  $\varepsilon$  can be bounded in terms of the location and scale parameters for location-scale families. In particular, in showing condition (9), we saw that

$$\mathbb{E} \|\Sigma(\theta - \theta')z_0 + \mu(\theta - \theta')\|_2^2 = \|\Sigma_{z_0}^{1/2} \Sigma(\theta - \theta')^\top\|_F^2 + \|\mu(\theta - \theta')\|_2^2.$$

If we then denote

$$\sigma_{\max}(\mu) = \max_{\|\theta\|_2=1} \|\mu\theta\|_2, \quad \sigma_{\max}(\Sigma) = \max_{\|\theta\|_2=1} \|\Sigma_{z_0}^{1/2} \Sigma(\theta)^\top\|_F,$$

we can see that  $\mathbb{E} \|\Sigma(\theta - \theta')z_0 + \mu(\theta - \theta')\|_2^2 \leq \sigma_{\max}^2(\mu) \|\theta - \theta'\|_2^2 + \sigma_{\max}^2(\Sigma) \|\theta - \theta'\|_2^2$ . Combining this result with Lemma C.2 and Jensen's inequality, we get that

$$W(\mathcal{D}(\theta), \mathcal{D}(\theta')) \leq \sqrt{\sigma_{\max}^2(\mu) + \sigma_{\max}^2(\Sigma)} \|\theta - \theta'\|_2,$$

and so  $\varepsilon \leq \sqrt{\sigma_{\max}^2(\mu) + \sigma_{\max}^2(\Sigma)}$ .

## C.2 Two-Stage Algorithm for Location Families

We carefully review the problem setup and introduce the remaining assumptions. The distribution map  $\mathcal{D}$  parameterizes a location family

$$z_\theta \sim \mathcal{D}(\theta) \Leftrightarrow z_\theta \stackrel{d}{=} z_0 + \mu\theta,$$

where  $z_0 \sim \mathcal{D}_0$ . We assume the base distribution  $\mathcal{D}_0$  is zero-mean and subgaussian with parameter  $K$ . The loss function  $\ell(z; \theta)$  is  $L_z$ -Lipschitz in  $z$ ,  $L$ -Lipschitz and in  $\theta$ , and  $\beta$ -smooth in  $(z, \theta)$  in the sense that  $\nabla \ell(z; \theta) \in \mathbb{R}^{m+d}$  is Lipschitz in  $(z, \theta)$ .

We also assume that  $\lambda = \max\{\gamma - \beta^2/\gamma_z, \gamma - 2\varepsilon\beta + \gamma_z \sigma_{\min}^2(\mu)\} > 0$ , where  $\gamma$  and  $\gamma_z$  are the strong convexity parameters of the loss in  $\theta$  and  $z$ , respectively. By Theorem 3.3, this implies that the performative risk is  $\lambda$ -strongly convex.

We assume that the performative optimum  $\theta_{\text{PO}}$  is contained in a ball of radius  $R$ , so in the second stage we can set the domain of optimization to be  $\Theta = \{\theta : \|\theta\|_2 \leq R\}$ . Finally, we assume that the minimizer of the perturbed performative risk at the population level,  $\widehat{\theta} \in \arg \min_{\theta \in \Theta} \widehat{\text{PR}}(\theta)$  is contained in the interior of  $\Theta$  with probability 1.

**Theorem C.4.** *Under the preceding assumptions, if  $n \geq \Omega(d + m + \log(1/\delta))$ , then, with probability  $1 - \delta$ , Algorithm 1 returns a point  $\widehat{\theta}_n$  such that*

$$\text{PR}(\widehat{\theta}_n) - \text{PR}(\theta_{\text{PO}}) \leq O\left(\frac{d + m + \log(1/\delta)}{n} + \frac{1}{\delta n}\right).$$

Before proceeding to the proof of this result, we first state four auxiliary lemmas, which constitute the bulk of our analysis. The proofs of the lemmas are included in [Appendix C.3](#). The first lemma is a standard result about ordinary least-squares estimation.

**Lemma C.5.** *If  $n \geq \Omega(d + m + \log(1/\delta))$ , then with probability  $1 - \delta$ ,*

$$\|\mu - \widehat{\mu}\| \leq O\left(\sqrt{\frac{(d + m) + \log(1/\delta)}{n}}\right).$$

The next lemma is a simple adaptation from Theorem 2 in [34] controlling the generalization gap of the empirical risk minimizer for strongly convex losses.

**Lemma C.6.** *Suppose  $\widehat{\text{PR}}_n$  is  $\widehat{\lambda}$ -strongly convex. Then, with probability at least  $1 - \delta$ ,*

$$\widehat{\text{PR}}(\widehat{\theta}_n) - \widehat{\text{PR}}(\widehat{\theta}) \leq \frac{4(L_z \|\widehat{\mu}\| + L)^2}{\delta \widehat{\lambda} n}.$$

The next lemma controls the difference in gradients between the true performative risk PR and the perturbed performative risk  $\widehat{\text{PR}}$ .

**Lemma C.7.** *For any  $\theta \in \Theta$ ,*

$$\|\nabla \text{PR}(\theta) - \nabla \widehat{\text{PR}}(\theta)\|_2^2 \leq O(\|\mu\|^2 \|\mu - \widehat{\mu}\|^2).$$

Finally, the last lemma shows that the smoothness assumptions on the loss ensure smoothness of the performative risk. Here, by  $\beta_\theta$ -smoothness we mean that  $\nabla_\theta \text{PR}(\theta)$  is  $\beta_\theta$ -Lipschitz.

**Lemma C.8.** *Under the preceding assumptions, the performative risk  $\text{PR}(\theta)$  is  $\beta_\theta = O(\|\mu\|^2)$ -smooth.*

With these lemmas in hand, we are now ready to prove [Theorem C.4](#).

*Proof of Theorem C.4.* By assumption, the performative risk  $\text{PR}(\theta)$  is  $\lambda$ -strongly convex, for some  $\lambda > 0$ . This implies

$$\text{PR}(\widehat{\theta}_n) - \text{PR}(\theta_{\text{PO}}) \leq \frac{1}{2\lambda} \|\nabla \text{PR}(\widehat{\theta}_n)\|_2^2.$$

Since  $\widehat{\theta}_{\text{PO}}$  is an interior minimizer of  $\widehat{\text{PR}}$ , we know  $\nabla \widehat{\text{PR}}(\widehat{\theta}_{\text{PO}}) = 0$ . Using  $\|a + b\|^2 \leq 2\|a\|^2 + 2\|b\|^2$ ,

$$\begin{aligned} \frac{1}{2\lambda} \|\nabla \text{PR}(\widehat{\theta}_n)\|_2^2 &= \frac{1}{2\lambda} \|\nabla \text{PR}(\widehat{\theta}_n) - \nabla \widehat{\text{PR}}(\widehat{\theta}_{\text{PO}})\|_2^2 \\ &= \frac{1}{2\lambda} \|\nabla \text{PR}(\widehat{\theta}_n) - \nabla \widehat{\text{PR}}(\widehat{\theta}_n) + \nabla \widehat{\text{PR}}(\widehat{\theta}_n) - \nabla \widehat{\text{PR}}(\widehat{\theta}_{\text{PO}})\|_2^2 \\ &\leq \frac{1}{\lambda} \|\nabla \text{PR}(\widehat{\theta}_n) - \nabla \widehat{\text{PR}}(\widehat{\theta}_n)\|_2^2 + \frac{1}{\lambda} \|\nabla \widehat{\text{PR}}(\widehat{\theta}_n) - \nabla \widehat{\text{PR}}(\widehat{\theta}_{\text{PO}})\|_2^2. \end{aligned} \quad (10)$$

We bound each of these terms separately. For the first term, by [Lemma C.7](#),

$$\|\nabla\text{PR}(\widehat{\theta}_n) - \nabla\widehat{\text{PR}}(\widehat{\theta}_n)\|_2^2 \leq O(\|\mu\|^2 \|\mu - \widehat{\mu}\|^2).$$

By [Lemma C.5](#), with probability  $1 - \delta$ , we can bound  $\|\mu - \widehat{\mu}\|^2 \leq O\left(\frac{d+m+\log(1/\delta)}{n}\right)$ , and thus

$$\|\nabla\text{PR}(\widehat{\theta}_n) - \nabla\widehat{\text{PR}}(\widehat{\theta}_n)\|_2^2 \leq O\left(\frac{d+m+\log(1/\delta)}{n}\right).$$

For the second term in equation (10), notice that  $\lambda = \max\{\gamma - \beta^2/\gamma_z, \gamma - 2\varepsilon\beta + \gamma_z\sigma_{\min}^2(\mu)\} > 0$  implies that  $\widehat{\text{PR}}$  is at least  $\widehat{\lambda} = \lambda - O(\frac{1}{\sqrt{n}})$ -strongly convex. This follows because  $|\sigma_{\min}(\mu) - \sigma_{\min}(\widehat{\mu})| \leq \|\mu - \widehat{\mu}\|$  by Weyl's inequality (see for example Theorem 3.3.16 in [30]), and  $\widehat{\text{PR}}$  is  $O(\|\widehat{\mu}\|)$ -sensitive, so by [Lemma C.5](#), each term depending on  $\varepsilon$  or  $\sigma_{\min}(\widehat{\mu})$  is within  $O(1/\sqrt{n})$  or  $O(1/n)$  of the corresponding values for the non-perturbed risk PR.

Hence, when  $n \geq \Omega(1/\lambda^2)$ , the strong convexity parameter of the perturbed performative risk,  $\widehat{\lambda}$ , is at least  $\lambda/2$ .

With this, we can apply the fact that  $\widehat{\theta}_{\text{PO}}$  is an interior minimizer of  $\widehat{\text{PR}}$  by assumption to conclude that when  $n \geq \Omega(1/\lambda^2)$ ,

$$\|\widehat{\theta}_n - \widehat{\theta}_{\text{PO}}\|_2^2 \leq \frac{4}{\lambda} (\widehat{\text{PR}}(\widehat{\theta}_n) - \widehat{\text{PR}}(\widehat{\theta}_{\text{PO}})).$$

Now, when  $\widehat{\text{PR}}$  is strongly convex, the finite-sample performative risk  $\widehat{\text{PR}}_n$  is also strongly convex because [Theorem 3.3](#) does not depend on the base distribution  $\mathcal{D}_0$ , and  $\widehat{\text{PR}}_n$  is simply  $\widehat{\text{PR}}$  when the base distribution  $\mathcal{D}_0$  is replaced with the uniform distribution on  $\{z_1, \dots, z_n\}$ . Consequently, by [Lemma C.6](#), with probability  $1 - \delta$ ,

$$\|\widehat{\theta}_n - \widehat{\theta}_{\text{PO}}\|_2^2 \leq O(\widehat{\text{PR}}(\widehat{\theta}_n) - \widehat{\text{PR}}(\widehat{\theta}_{\text{PO}})) \leq O\left(\frac{\|\widehat{\mu}\|^2}{\delta n}\right).$$

By [Lemma C.8](#),  $\widehat{\text{PR}}$  is  $O(\|\widehat{\mu}\|^2)$ -smooth. Applying the previous display then gives us,

$$\|\nabla\widehat{\text{PR}}(\widehat{\theta}_n) - \nabla\widehat{\text{PR}}(\widehat{\theta}_{\text{PO}})\|_2^2 \leq O\left(\|\widehat{\mu}\|^4 \|\widehat{\theta}_n - \widehat{\theta}_{\text{PO}}\|_2^2\right) \leq O\left(\frac{\|\widehat{\mu}\|^6}{\delta n}\right).$$

By the triangle inequality and repeated application of  $(a+b)^2 \leq 2a^2 + 2b^2$ ,  $\|\widehat{\mu}\|^6 \leq 128 \|\widehat{\mu} - \mu\|^6 + 128 \|\mu\|^6$ . Therefore, the above term is  $O(\|\mu\|^6/\delta n)$ . Putting everything together with a union bound, we have shown that with probability  $1 - \delta$ , if  $n \geq \Omega(d+m+\log(1/\delta))$ , it holds that

$$\text{PR}(\widehat{\theta}_n) - \text{PR}(\theta_{\text{PO}}) \leq O\left(\frac{d+m+\log(1/\delta)}{n} + \frac{1}{\delta n}\right),$$

as desired. ■



### C.3 Proofs of Lemmas for Two-Stage Algorithm Analysis

The proof of Lemma C.5 is essentially standard (see, e.g., [21]), but we include it for completeness.

*Proof of Lemma C.5.* Define  $Z \in \mathbb{R}^{n \times m}$  with rows  $z_i$  and  $\Theta \in \mathbb{R}^{n \times d}$  with rows  $\theta_i$ ,  $1 \leq i \leq n$ . Then,  $Z = \Theta \mu^\top + Z_0$ , where  $Z_0 \in \mathbb{R}^{n \times m}$  is a matrix with base samples from  $\mathcal{D}_0$  as rows. Temporarily assume that  $\Theta^\top \Theta$  is invertible; we will later condition on this event. Separately optimizing over each row of  $\mu$ , we can write the least-squares estimator as

$$\widehat{\mu}^\top = (\Theta^\top \Theta)^{-1} \Theta^\top Z.$$

Consequently, we can bound the estimation error as

$$\begin{aligned} \|\mu - \widehat{\mu}\| &= \|\mu^\top - \widehat{\mu}^\top\| = \left\| \mu^\top - (\Theta^\top \Theta)^{-1} \Theta^\top (\Theta \mu^\top + Z_0) \right\| \\ &= \left\| (\Theta^\top \Theta)^{-1} \Theta^\top Z_0 \right\| \\ &\leq \frac{1}{\lambda_{\min}(\Theta^\top \Theta)} \|\Theta^\top Z_0\|. \end{aligned}$$

Since  $\theta_i \sim \mathcal{N}(0, I)$ ,  $\Theta \in \mathbb{R}^{n \times d}$  has i.i.d.  $\mathcal{N}(0, 1)$  entries, and so  $\Theta^\top \Theta$  is a standard Wishart matrix. The standard bound on the minimum eigenvalue of a Wishart matrix (see Theorem 4.6.1 in [42]) gives, with probability  $1 - \delta$ ,

$$\sqrt{\lambda_{\min}(\Theta^\top \Theta)} \geq \Omega(\sqrt{n} - \sqrt{d} - \sqrt{\log(1/\delta)}).$$

Therefore, if  $n \geq \Omega(d + \log(2/\delta))$ , then, with probability  $1 - \delta/2$ ,

$$\sqrt{\lambda_{\min}(\Theta^\top \Theta)} \geq \Omega(\sqrt{n}/2). \quad (11)$$

Control of the second term,  $\|\Theta^\top Z_0\|$ , also follows from a standard covering argument followed by the Bernstein bound. Write  $\Theta^\top Z_0 = \sum_{i=1}^n \theta_i (z_0)_i^\top$ . Let  $\mathcal{B}^d$  and  $\mathcal{B}^m$  denote the unit balls in  $\mathbb{R}^d$  and  $\mathbb{R}^m$ , respectively. Then,

$$\|\Theta^\top Z_0\| = \sup_{x \in \mathcal{B}^d, y \in \mathcal{B}^m} x^\top \left( \sum_{i=1}^n \theta_i (z_0)_i^\top \right) y = \sup_{x \in \mathcal{B}^d, y \in \mathcal{B}^m} \sum_{i=1}^n (x^\top \theta_i) ((z_0)_i^\top y).$$

Let  $\mathcal{N}_\varepsilon$ , and  $\mathcal{M}_\varepsilon$  denote  $\varepsilon$ -coverings of  $\mathcal{B}^d$  and  $\mathcal{B}^m$ , respectively. A volumetric bound gives  $|\mathcal{N}_\varepsilon| \leq \left(1 + \frac{2}{\varepsilon}\right)^d$  and similarly  $|\mathcal{M}_\varepsilon| \leq \left(1 + \frac{2}{\varepsilon}\right)^m$  (see Corollary 4.2.13 in [42]). Taking  $\varepsilon = 1/4$ ,  $|\mathcal{N}_\varepsilon| \leq 9^d$  and  $|\mathcal{M}_\varepsilon| \leq 9^m$ . Approximating the supremum over the  $\varepsilon$ -nets gives

$$\|\Theta^\top Z_0\| \leq 2 \max_{x \in \mathcal{N}_\varepsilon, y \in \mathcal{M}_\varepsilon} \sum_{i=1}^n (x^\top \theta_i) ((z_0)_i^\top y).$$

Fix  $x, y \in \mathcal{N}_\varepsilon, \mathcal{M}_\varepsilon$ . Since  $\theta_i \sim \mathcal{N}(0, I)$  and  $\|x\|_2 = 1$ ,  $x^\top \theta_i \sim \mathcal{N}(0, 1)$ , which has subgaussian norm 1. Similarly, since  $(z_0)_i$  is subgaussian with parameter  $K$  and  $\|y\|_2 = 1$ , the marginal  $(z_0)_i^\top y$  is subgaussian with parameter  $K$ . Since  $z_0$  and  $\theta$  are independent and zero-mean, the

product  $(x^\top \theta_i)((z_0)_i^\top y)$  is zero-mean and subexponential with parameter  $K$ . Since each term is subexponential, by the Bernstein bound (see Theorem 2.8.1 in [42]), for any  $t > 0$ ,

$$\mathbb{P} \left\{ \sum_{i=1}^n (x^\top \theta_i)((z_0)_i^\top y) > t/2 \right\} \leq \exp \left( -c \min \left\{ \frac{t^2}{nK^2}, \frac{t}{K} \right\} \right),$$

for some universal constant  $c$ . Taking a union bound over the  $\varepsilon$ -nets,

$$\mathbb{P} \left\{ \|\Theta^\top Z_0\| > t \right\} \leq 9^{d+m} \exp \left( -c \min \left\{ \frac{t^2}{nK^2}, \frac{t}{K} \right\} \right).$$

If  $n \geq \Omega(d + m + \log(2/\delta))$ , then with probability at least  $1 - \delta/2$ ,

$$\|\Theta^\top Z_0\| \leq O(\sqrt{n((d+m) + \log(1/\delta))}). \quad (12)$$

Combining equations (11) and (12) with a union bound, if  $n \geq \Omega(d + m + \log(1/\delta))$ , then

$$\|\mu - \widehat{\mu}\| \leq O \left( \sqrt{\frac{(d+m) + \log(1/\delta)}{n}} \right).$$

■

*Proof of Lemma C.7.* Under the location-family parameterization, we can write

$$\text{PR}(\theta) = \mathbb{E}_{z \sim \mathcal{D}(\theta)} \ell(z; \theta) = \mathbb{E}_{z_0 \sim \mathcal{D}_0} \ell(z_0 + \mu\theta; \theta),$$

so the gradients are given by

$$\nabla \text{PR}(\theta) = \mathbb{E}_{z_0 \sim \mathcal{D}_0} \nabla \ell(z_0 + \mu\theta; \theta) \quad \text{and} \quad \nabla \widehat{\text{PR}}(\theta) = \mathbb{E}_{z_0 \sim \mathcal{D}_0} \nabla \ell(z_0 + \widehat{\mu}\theta; \theta).$$

This representation allows us to write

$$\|\nabla \text{PR}(\theta) - \nabla \widehat{\text{PR}}(\theta)\|_2^2 = \left\| \mathbb{E}_{z_0 \sim \mathcal{D}_0} [\nabla \ell(z_0 + \mu\theta; \theta) - \nabla \ell(z_0 + \widehat{\mu}\theta; \theta)] \right\|_2^2.$$

Applying the chain rule, together with the triangle-inequality, gives

$$\begin{aligned} \|\nabla \text{PR}(\theta) - \nabla \widehat{\text{PR}}(\theta)\|_2 &\leq \left\| \mathbb{E}_{z_0 \sim \mathcal{D}_0} [\nabla_\theta \ell(z_0 + \mu\theta; \theta) - \nabla_\theta \ell(z_0 + \widehat{\mu}\theta; \theta)] \right\|_2 \\ &\quad + \left\| \mathbb{E}_{z_0 \sim \mathcal{D}_0} [\mu^\top \nabla_z \ell(z_0 + \mu\theta; \theta) - \widehat{\mu}^\top \nabla_z \ell(z_0 + \widehat{\mu}\theta; \theta)] \right\|_2. \end{aligned}$$

We bound each of these terms separately. For the first term,  $\beta$ -smoothness in  $z$  immediately gives

$$\left\| \mathbb{E}_{z_0 \sim \mathcal{D}_0} [\nabla_\theta \ell(z_0 + \mu\theta; \theta) - \nabla_\theta \ell(z_0 + \widehat{\mu}\theta; \theta)] \right\|_2 \leq \beta \|\mu\theta - \widehat{\mu}\theta\|_2 \leq \beta \|\mu - \widehat{\mu}\| \|\theta\|_2.$$

For the second term, adding and subtracting  $\mu^\top \nabla_z \ell(z_0 + \widehat{\mu}\theta; \theta)$  and then using the triangle inequality,

$$\begin{aligned} & \left\| \mathbb{E}_{z_0 \sim \mathcal{D}_0} \left[ \mu^\top \nabla_z \ell(z_0 + \mu\theta; \theta) - \widehat{\mu}^\top \nabla_z \ell(z_0 + \widehat{\mu}\theta; \theta) \right] \right\|_2 \\ & \leq \|\mu\| \left\| \mathbb{E}_{z_0 \sim \mathcal{D}_0} [\nabla_z \ell(z_0 + \mu\theta; \theta) - \nabla_z \ell(z_0 + \widehat{\mu}\theta; \theta)] \right\|_2 + \|\mu - \widehat{\mu}\| \left\| \mathbb{E}_{z_0 \sim \mathcal{D}_0} [\nabla_z \ell(z_0 + \widehat{\mu}\theta; \theta)] \right\|_2 \\ & \leq \beta \|\mu\| \|\mu - \widehat{\mu}\| \|\theta\|_2 + L_z \|\mu - \widehat{\mu}\|, \end{aligned}$$

where the last line used  $\beta$ -smoothness in  $z$ . Combining both pieces, we have

$$\|\nabla \text{PR}(\theta) - \nabla \widehat{\text{PR}}(\theta)\|_2 \leq ((\beta + \beta \|\mu\|) \|\theta\|_2 + L_z) \|\mu - \widehat{\mu}\|.$$

Using the trivial bound  $\|\theta\|_2 \leq R$ , and then squaring both sides,

$$\|\nabla \text{PR}(\widehat{\theta}) - \nabla \widehat{\text{PR}}(\widehat{\theta})\|_2^2 \leq ((1 + \|\mu\|) \beta R + L_z)^2 \|\mu - \widehat{\mu}\|^2.$$

■

*Proof of Lemma C.8.* By applying the location family parameterization as in the proof of Lemma C.7, we get

$$\|\nabla \text{PR}(\theta) - \nabla \text{PR}(\theta')\|_2 = \left\| \mathbb{E}_{z_0 \sim \mathcal{D}_0} [\nabla \ell(z_0 + \mu\theta; \theta) - \nabla \ell(z_0 + \mu\theta'; \theta')] \right\|_2.$$

Using the chain rule and the triangle inequality,

$$\begin{aligned} \|\nabla \text{PR}(\theta) - \nabla \text{PR}(\theta')\|_2 & \leq \left\| \mathbb{E}_{z_0 \sim \mathcal{D}_0} \nabla_\theta \ell(z_0 + \mu\theta; \theta) - \nabla_\theta \ell(z_0 + \mu\theta'; \theta') \right\|_2 \\ & \quad + \left\| \mathbb{E}_{z_0 \sim \mathcal{D}_0} \mu^\top \nabla_z \ell(z_0 + \mu\theta; \theta) - \mu^\top \nabla_z \ell(z_0 + \mu\theta'; \theta') \right\|_2. \end{aligned} \quad (13)$$

For the first term in equation (13), adding and subtracting  $\nabla_\theta \ell(z_0 + \mu\theta'; \theta)$  and using the triangle inequality gives

$$\begin{aligned} \left\| \mathbb{E}_{z_0 \sim \mathcal{D}_0} [\nabla_\theta \ell(z_0 + \mu\theta; \theta) - \nabla_\theta \ell(z_0 + \mu\theta'; \theta')] \right\|_2 & \leq \left\| \mathbb{E}_{z_0 \sim \mathcal{D}_0} \nabla_\theta \ell(z_0 + \mu\theta; \theta) - \nabla_\theta \ell(z_0 + \mu\theta'; \theta) \right\|_2 \\ & \quad + \left\| \mathbb{E}_{z_0 \sim \mathcal{D}_0} \nabla_\theta \ell(z_0 + \mu\theta'; \theta) - \nabla_\theta \ell(z_0 + \mu\theta'; \theta') \right\|_2 \\ & \leq \beta \|\mu\| \|\theta - \theta'\|_2 + \beta \|\theta - \theta'\|_2, \end{aligned}$$

where we used Jensen's inequality and the assumption that  $\nabla_\theta \ell(z; \theta)$  is  $\beta$ -Lipschitz in  $z$  (for the first term) and  $\beta$ -Lipschitz in  $\theta$  (for the second term).

Now, for the second term in equation (13), similarly adding and subtracting  $\mu^\top \nabla_z \ell(z_0 + \mu\theta'; \theta)$  and using the triangle inequality gives

$$\begin{aligned} \left\| \mathbb{E}_{z_0 \sim \mathcal{D}_0} [\mu^\top \nabla_z \ell(z_0 + \mu\theta; \theta) - \mu^\top \nabla_z \ell(z_0 + \mu\theta'; \theta')] \right\|_2 & \leq \left\| \mathbb{E}_{z_0 \sim \mathcal{D}_0} \mu^\top \nabla_z \ell(z_0 + \mu\theta; \theta) - \mu^\top \nabla_z \ell(z_0 + \mu\theta'; \theta) \right\|_2 \\ & \quad + \left\| \mathbb{E}_{z_0 \sim \mathcal{D}_0} \mu^\top \nabla_z \ell(z_0 + \mu\theta'; \theta) - \mu^\top \nabla_z \ell(z_0 + \mu\theta'; \theta') \right\|_2 \\ & \leq \beta \|\mu\|^2 \|\theta - \theta'\|_2 + \beta \|\mu\|^2 \|\theta - \theta'\|_2, \end{aligned}$$

where we used  $\nabla_z \ell(z; \theta)$  is  $\beta$  Lipschitz in  $z$  (for the first term) and  $\beta$  Lipschitz in  $\theta$  (for the second term). This completes the proof. ■

## D Experimental Details

Lastly, we elaborate on the implementation details of the various simulators and algorithms evaluated in Section 5.

### D.1 Synthetic Linear Regression Example

**Data generating process.** Given a parameter vector  $\theta \in \mathbb{R}^d$ , as per Example 3.2, feature label pairs  $(x, y)$  are generated according to the following data generating process:

1.  $x \sim \mathcal{N}(0, \Sigma_x)$ .
2.  $y = \beta^\top x + \mu^\top \theta + U_y$  where  $U_y \sim \mathcal{N}(0, \sigma_y^2)$ .

In our experiments, we take  $d = 20$ , and set  $\sigma_y^2 = 0.01$ . For each trial, we sample  $\Sigma_x$  as a random symmetric positive-definite matrix with operator norm 0.01, sample  $\beta \sim \mathcal{N}(0, I_d)$ , and sample  $w$  uniformly on the sphere of radius  $\varepsilon$ , where  $\varepsilon$  is the sensitivity parameter of the distribution map. In our experiments, we choose  $\varepsilon \in \{0.01, 100\}$ .

For this example, the performative optimum  $\theta_{\text{PO}}$  can be computed in closed-form due to the squared-loss and the linearity of the performative effects. In particular,

$$\theta_{\text{PO}} = (\Sigma_x + \mu\mu^\top)^{-1} \Sigma_x \beta.$$

**Algorithms.** We compare four different algorithms. In all four cases, we set  $\Theta = \{\theta : \|\theta\|_2 \leq 10\}$ .

1. **The two-stage procedure.** For a budget of  $n$  samples, the two-stage procedure consists of first deploying  $n/2$  classifiers  $\theta_i \sim \mathcal{N}(0, I_d)$  and observing data  $(x_i, y_i)$ . We then compute an estimate  $\widehat{\mu}$  by solving a least-squares problem:

$$\widehat{\mu} \in \arg \min_{\mu, c} \sum_{i=1}^n (y_i - \mu^\top \theta_i - c)^2.$$

After computing  $\widehat{\mu}$ , the algorithm collects another  $n/2$  samples  $(x_i, y_i)$  by repeatedly deploying  $\theta_i = 0$  for  $i = n + 1, \dots, 2n$ , and computes  $\widehat{\theta}_n$  by solving another least-squares problem

$$\widehat{\theta}_n \in \arg \min_{\theta \in \Theta} \sum_{i=n+1}^{2n} (y_i - \theta^\top (x_i - \widehat{\mu}))^2.$$

2. **DFO.** We run the derivative-free optimization procedure from Flaxman et al. [11]. We initialize  $\theta_0 = \mathbf{1}$ , use step-size sequence  $c_0/t$ ,  $c_0 = 0.01$ , a batch size of 20 samples per-step, and take  $\delta = 10$ . These parameter were chosen via a small grid search over  $c_0 \in [1e-4, 1]$ , batch size in  $[1, 500]$ , and  $\delta \in [0.1, 100]$ . However, the algorithm still has variance across runs, especially in the small  $\varepsilon$  regime.
3. **Greedy SGD.** We use the greedy SGD variant introduced by Mendler-Dünner et al. [22] with initial point  $\theta_0 = \mathbf{1}$  and step-size sequence  $1/\sqrt{t}$ , which we found to slightly outperform the step-sequence  $1/t$  in our experiments. For the sake of brevity, we omit the full pseudocode of greedy/lazy SGD instead point the reader to Figure 1 in [22].
4. **Lazy SGD.** We use the lazy SGD algorithm [22] with initial point  $\theta_0 = \mathbf{1}$ , step-size sequence  $c/(k_0 + t)$  with parameters  $c = 1, k_0 = 1$ , and  $k^2$  collected samples in  $k$ -th update.

**Evaluation.** We ran each algorithm for 50 trials, and in Figure 1, we compare the suboptimality gap  $\text{PR}(\theta) - \text{PR}(\theta_{\text{PO}})$  of each algorithm as a function of the number of samples. For each sample size  $n$ , we bootstrap 95% confidence intervals over the 50 trials.

## D.2 Strategic Classification

**Data generation.** We use the same strategic classification simulator as [28]. For detailed information about the simulator, please refer to Appendix B.2 of [28].

The strategic responses are determined according to

$$x_{\text{BR}} = x + \varepsilon B\theta,$$

for some matrix  $B$  which determines the subset of features that are performative. Here,  $\theta \in \mathbb{R}^{11}$  parameterizes a logistic regression classifier. The logistic loss is regularized by an additional  $\ell_2$ -penalty, which makes it strongly convex. The computation of the smoothness parameter can be found in [28].

We consider two different values of the sensitivity parameter,  $\varepsilon \in \{0.0001, 100\}$ , and set the magnitude of the regularizer to be  $\lambda = 0.002$ . We restrict the radius of the optimization domain to be 10,  $\Theta = \{\theta : \|\theta\|_2 \leq 10\}$ . This choice of parameters ensures that  $\varepsilon = 0.0001$  is below the critical threshold  $\frac{\gamma}{2\beta}$ , while  $\varepsilon = 100$  is above the threshold.

**Algorithms.** We compare the same four algorithms as the previous section.

1. **Two-stage procedure.** In the first stage, we deploy random  $\theta_i \sim \mathcal{N}(0, I)$  and perform linear regression to estimate  $\mu$ ,

$$\widehat{\mu} = \arg \min_{\mu} \sum_{i=1}^n \|z_i - \mu \theta_i\|^2$$

Then, having collected samples from the base distribution, we solve the proxy logistic regression objective offline by running gradient descent with a line search procedure until a tolerance criterion is met. In particular, we solve,

$$\arg \min_{\theta \in \Theta} \frac{1}{n} \sum_{j=n+1}^{2n} \ell(z_j + \widehat{\mu}\theta; \theta),$$

where  $\ell(z; \theta)$ , is the regularized logistic regression objective, until the improvement between consecutive iterates is smaller than 1e-10.

2. **DFO.** We again run the derivative-free optimization procedure from Flaxman et al. [11]. We initialize  $\theta_0 = \mathbf{0}$ , use step-size sequence  $1/t$ , a batch size of 100 samples per-step, and set  $\delta = 1$ . We tried several other parameter configurations and found this one to perform best on this problem setting.
3. **Greedy SGD.** We run the greedy SGD variant with initial point  $\theta_0 = \mathbf{0}$  and step-size sequence as suggested by [22]. See Appendix A in [22] for details.
4. **Lazy SGD.** We use the lazy SGD algorithm with initial point  $\theta_0 = \mathbf{0}$  and  $k^2$  collected samples in  $k$ -th update. As for greedy SGD, we use the step-size sequence suggested by [22].

**Evaluation.** We ran each algorithm for 50 trials, and in [Figure 2](#), we compare the performative risk  $PR(\theta)$  of each algorithm as a function of the number of samples. For each sample size  $n$ , we bootstrap 95% confidence intervals over the 50 trials.