The Stability of Online Algorithms in Performative Prediction

Gabriele Farina
MIT
gfarina@mit.edu

Juan Carlos Perdomo
NYU
j.perdomo.silva@nyu.edu

Abstract

The use of algorithmic predictions in decision-making leads to a feedback loop where the models we deploy actively influence the data distributions we see, and later use to retrain on. This dynamic was formalized by Perdomo et al. in their work on performative prediction. Our main result is an unconditional reduction showing that any no-regret algorithm deployed in performative settings converges to a (mixed) performatively stable equilibrium: a solution in which models actively shape data distributions in ways that their own predictions look optimal in hindsight. Prior to our work, all positive results in this area imposed strong restrictions on how models influenced distributions. By using a martingale argument and allowing randomization, we avoid any assumption on how populations respond to predictions and sidestep recent hardness results showing that deterministic stable models are in general PPAD-hard to compute. Lastly, on a more conceptual note, our connection sheds light on why common algorithms, like gradient descent, are naturally stabilizing and prevent runaway feedback loops. We hope our work enables future technical transfer of ideas between online optimization and performativity.

1 Introduction

Social prediction is an inherently dynamic affair. The predictions we make inform decisions that actively shape the data we see. These influenced outcomes are then fed back into the datasets used to train models, leading to a continual feedback loop between prediction systems and their encompassing social contexts. Once recognized, we start to see this impact, or *performativity*, of social prediction everywhere. Predicting individuals' health outcomes influences their behavior and hence their realized outcomes. A bank's choice of credit risk model influences the features that people strategically report to the system. Link prediction algorithms in professional networks determine who we know, and hence what links will end up forming in the future. The reader can certainly come up with more examples.

Perdomo, Zrnic, Mendler-Dünner, and Hardt [Per+20] introduced a learning-theoretic framework called *performative prediction* to reason about learning in these dynamic environments. In their work, they identify a natural solution concept termed *performative stability*. A model is performatively stable if it actively shapes data in a way that its own predictions look optimal in hindsight.

The motivation behind stability is that it prevents a runaway feedback loop where learners repeatedly update their models because they are not risk-minimizing on the observed data. In more detail, performativity often manifests as an unexpected distribution shift. Decision-makers often respond to this shift by *retraining* their models on the most recent examples. Performatively stable models have the desirable property that they are *fixed points* of retraining: if you deploy a performatively stable model and then retrain on its induced distributions, you get the same model back. There is no incentive to deviate, see Figure 1 for a visual illustration.

It seems intuitive that to guarantee stability, one needs tight control over whether the retraining has any fixed points. And indeed, there is now a rich literature on performativity showing that if $a)$ the learner's loss function is pointwise strongly convex and $b)$ populations are only weakly influenced by predictions, then retraining is a contraction (*i.e.*, the feedback loop has Lipschitz constant < 1 and hence converges to stability). While both of these conditions severely limit the regimes where we can

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

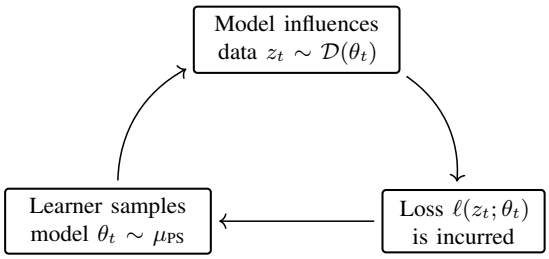

Informally, the *mixture* μ_{PS} is ϵ-performatively stable if for any other model θ', in expectation,

$$\frac{1}{n}\sum_{t=1}^{n}\ell(z_t;\theta_t) \le \frac{1}{n}\sum_{t=1}^{n}\ell(z_t;\theta')$$

That is, no other model has lower loss over the data induced by μ.

Figure 1: Visualizing the performative stability guarantee. A mixture μ over models θ is stable if it shapes the data in a way that its own predictions are risk minimizing. Our notion, Definition 1, generalizes the prior notion where μ is a point mass on a single model.

guarantee stability, they are, in a precise sense, unavoidable. If losses are not strongly convex, or if predictions are strongly performative, then finding a single model that is performatively stable can be impossible (they may not exist) or PPAD-hard to compute [Ana+26].

However, in a fortunate twist, these difficulties can be easily avoided by simply allowing the learner to randomize across a small number of models. Our main contribution is an unconditional reduction from regret minimization to performative stability. We show that any learner that uniformly randomizes across the iterates produced by any no-regret algorithm is provably stable. This holds even when making *no assumptions* about the distributional response: predictions can influence both features and outcomes in arbitrary ways. Some of the consequences of this result are that we establish the first positive results for finding stable points for losses such as squared loss or log-loss that are convex but not pointwise strongly convex, as well as the first polynomial sample complexity guarantees for repeated risk minimization. Furthermore, all of these results also carry over to the multiplayer setting (see Section 4), where the observed data is a function of many agents making predictions simultaneously, and to the stateful setting (see Section 5), where the observed data is a function of the entire history of deployed models.

Taken together, our work significantly expands the space of problems for which one can efficiently achieve performative stability and highlights how classical supervised learning algorithms are naturally equilibrating when deployed in these more dynamic domains. Looking forward, we hope our results will inspire future work at the intersection of performativity and online learning.

1.1 Contributions

To discuss our results, we give a brief overview of the performative prediction framework.

Performative prediction extends supervised learning to account for the fact that deploying predictive models actively shapes the distribution of future data. More formally, if the learner makes predictions using a model f_θ where $\theta \in \mathbb{R}^d$ denotes a vector of parameters, then they see examples $z \sim \mathcal{D}(\theta)$ drawn i.i.d. from the distribution $\mathcal{D}(\theta)$ induced by the model. $\mathcal{D}(\cdot)$ is the distribution map, a function from model parameters θ to distributions over points in a set $\mathcal{Z}$. Typically, we let $z = (x, y)$ consist of features $x \in \mathcal{X}$ and outcomes $y \in \mathcal{Y}$. In this case, the distribution map $\mathcal{D}(\cdot)$ captures how deploying models can influence both the marginal distribution over features and the conditional distribution over outcomes. Performative prediction therefore strictly subsumes the setting known as strategic classification [Har+16]. Given the dynamic effects of prediction, the following stability notion is one of the key desiderata within performativity:

Definition 1. *Given a loss function $\ell(z;\theta)$, a model θ_{PS} is ε-performatively stable if $\mathbb{E}_{z\sim\mathcal{D}(\theta_{\text{PS}})}\ell(z;\theta_{\text{PS}}) \le \min_\theta \mathbb{E}_{z\sim\mathcal{D}(\theta_{\text{PS}})}\ell(z;\theta) + \varepsilon$. More generally, we say that a mixture μ_{PS} over models θ is ε-performatively stable if*

$$\mathbb{E}_{\theta\sim\mu_{\text{PS}}}\mathbb{E}_{z\sim\mathcal{D}(\theta)}\ell(z;\theta) \le \min_{\theta'\in\Theta}\mathbb{E}_{\theta\sim\mu_{\text{PS}}}\mathbb{E}_{z\sim\mathcal{D}(\theta)}\ell(z;\theta') + \varepsilon.$$

Performatively stable models are the fixed points of the following retraining procedure,

$$\theta_{t+1} = \arg\min_\theta \mathbb{E}_{z\sim\mathcal{D}(\theta_t)}\ell(z;\theta). \tag{1}$$

This is a natural dynamic that arises in recommender systems, strategic classification, and even traffic prediction, where new models are trained on data that was influenced by previous models. Our main

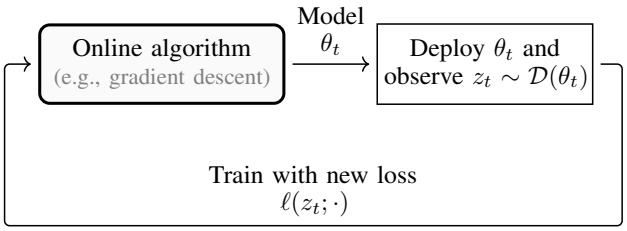

Main Theorem (Informal):

For any $T \geq 0$ and $\mathcal{D}(\cdot)$, the uniform distribution over the models $\theta_1, \ldots, \theta_T$ produced by the online algorithm is $\mathrm{Regret}(T)/T$-stable.

Figure 2: Our main result shows that any no-regret algorithm yields a sequence of models whose uniform mixture is performatively stable. The idea is simple and surprisingly strong. It sidesteps recent hardness results and establishes the first proofs of stability that do not require any continuity assumptions on $\mathcal{D}(\cdot)$ or curvature assumptions on ℓ.

result shows that any online algorithm that guarantees sublinear external regret against the stochastic losses $\ell(z_t; \theta_t)$ will, in fact, converge to performative stability. This result holds regardless of the underlying distribution map $\mathcal{D}(\cdot)$. See Figure 2 for a visual illustration.

Theorem 1 (Informal)**.** *Let $\theta_1, \ldots, \theta_T$ be models produced by an online algorithm that achieves no regret on the sequence of losses $\ell_t(\theta) = \ell(z_t; \theta)$ where $z_t \sim \mathcal{D}(\theta_t)$ for $t = 1$ to T. And, define μ to be the uniform distribution over $\theta_1, \ldots, \theta_T$. Then, μ is $\mathrm{Regret}(T)/T$-performatively stable.*

It is by now a celebrated fact that online algorithms have good generalization properties in supervised learning [CCG04]. Here, we show that when data is no longer drawn i.i.d. from a fixed distribution, but is instead *performative* and influenced by the chosen model, no-regret dynamics have the desirable property that they converge to a natural notion of stability.

As we will discuss, this connection significantly expands the kinds of problems for which one can guarantee performative stability. As easy corollaries of this main result, we can show that by taking a mixture over their iterates, common no-regret algorithms such as gradient descent and retraining (follow the leader) converge to performative stability under substantially weaker assumptions on the loss function ℓ and the distribution map $\mathcal{D}(\cdot)$ than those previously used in the performative prediction literature. Readers already familiar with prior work can jump to Table 1 for a summarized comparison.

Crucially, earlier analyses required that predictions have only a very weak, smoothly varying influence on the induced distribution. Formally, they assumed that $\mathcal{D}(\cdot)$ is Lipschitz (see Definition 3). More than just a technical condition, this assumption ruled out stability in important domains such as medicine or education, where people take actions that influence outcomes by thresholding model predictions (see, e.g., [Per+25] and Section 2 for an example).

Moreover, recent work by [Ana+26] showed that finding a single stable model even for Lipschitz $\mathcal{D}(\cdot)$ and strongly convex and smooth ℓ is PPAD-complete. By working with mixtures over models, we sidestep these complexity barriers, greatly expand the settings in which stability can be achieved efficiently, and provide a simple theoretical justification for why widely used learning practices prevent runaway feedback loops in social prediction.

1.2 Related Work

Performative prediction. [Per+20] introduced the performative prediction framework and established that retraining, as described in Equation (1), converged to performative stability if losses were smooth, strongly convex, and if the distribution map was sufficiently Lipschitz. Subsequent work [Men+20; DX23] proved finite-sample results showing that various stochastic optimization algorithms, such as gradient descent, also converged to stability under similar assumptions on the loss and distribution map. Later, [TH23] and [MMG23] analyzed convergence to stability under alternative Lipschitz conditions on $\mathcal{D}(\cdot)$.

In addition to these, there has been work analyzing performative stability in stateful settings [BHK22; LW22; Ray+22; WBD21] where the induced distribution is a function of the entire history of deployed models, as well as in multi-agent settings where data depends on the predictions made by multiple decision-makers [LYW22; PY23; Nar+23]. All of these results also imposed strong Lipschitzness conditions on the distribution map $\mathcal{D}(\cdot)$. Lastly, there is a parallel line of work in this area that designs algorithms which converge to a different solution concept called *performative optimality* [MPZ21;

Solution	$\rho < 1$	$\rho > 1 + \Omega(1)$	Convex losses, arbitrary $\mathcal{D}(\cdot)$	Squared or log loss, arbitrary $\mathcal{D}(\cdot)$
Stable Model θ_{PS}	$\mathcal{O}(1/T)$ †	PPAD-Complete ‡	Not guaranteed to exist	Not guaranteed to exist
Stable Mixture μ_{PS} (this work)	$\tilde{\mathcal{O}}(1/T)$	$\tilde{\mathcal{O}}(1/T)$	$\mathcal{O}(1/\sqrt{T})$	$\tilde{\mathcal{O}}(1/T)$

Table 1: Comparison to prior work. We let $\rho = \varepsilon\beta/\gamma$ denote the "standard" setting where ℓ is γ-strongly convex and β-smooth in θ, and $\mathcal{D}(\cdot)$ is ε-Lipschitz. Our reduction establishes the first known stability guarantees for arbitrary distribution maps $\mathcal{D}(\cdot)$, and weakly convex losses, while matching the rates in prior work up to log factors. †: [Men+20; DX23]. ‡: [Ana+26].

IYZ21; JZM22; KP23; LZ24]. We discuss the differences between these two concepts in the next section and point the reader to [HM25] for a broader survey of this area.

No-regret learning. Our reduction places performative stability within a broader literature connecting regret minimization to other areas. In game theory, regret-minimizing dynamics are closely related to equilibrium concepts such as coarse correlated and correlated equilibria [HM00; Dag+24; PR24], Nash equilibria in zero-sum games, and more general Φ-equilibria [GGM08]. These connections have been studied in a range of settings, including normal-form games, extensive-form games [Far+22; FP23], harmonic games [Leg+24], and concave games [Das+25]. In a similar vein, calibrated forecasting can also be famously reduced to no-regret learning [FV97; FV98]. Blackwell's approachability theorem has been shown to be algorithmically equivalent to online linear optimization, yielding a tight reduction between approachability and no-regret learning [ABH11].

In nonlinear optimization, recent reductions from nonsmooth, nonconvex optimization (in the sense of Goldstein stationarity) to online learning have led to optimal stochastic algorithms [CMO23]. In statistics and probability theory, sequential testing and anytime-valid inference can be reduced to no-regret algorithms, giving rise to e-processes and confidence sequences [OP16; Ram+23]. Deterministic regret inequalities have been shown to be equivalent to sharp martingale tail bounds and associated concentration inequalities [RS17]. Finally, multiplicative-weights and related online learning algorithms yield conceptually simple proofs of von Neumann's minimax theorem [PR08; FP24; Far26], spectral sparsification in graph theory [ALO15], and certain results in derandomization [Tre19]. Our result contributes to this line of work by showing that performative stability can also be understood through the lens of no-regret learning.

2 Preliminaries and Technical Background

Before presenting our main results, we provide some further context on definitions and prior results in performative prediction. In addition to performative stability, there is a second main solution concept in performative prediction known as *performative optimality*.

Definition 2. *θ_{PO} is performatively optimal if* $\mathbb{E}_{z\sim\mathcal{D}(\theta_{\text{PO}})}\ell(z;\theta_{\text{PO}}) \leq \min_\theta \mathbb{E}_{z\sim\mathcal{D}(\theta)}\ell(z;\theta)$.

Here, $\text{PR}(\theta) = \mathbb{E}_{z\sim\mathcal{D}(\theta)}\ell(z;\theta)$ is known as the *performative risk* of the model specified by θ. Optimal models minimize the performative risk. If we denote by $\text{DPR}(\theta, \mathcal{D}(\theta')) = \mathbb{E}_{z\sim\mathcal{D}(\theta')}\ell(z;\theta)$ the *decoupled* performative risk, we can relate optimality and stability as follows,

$$\underline{\text{Stability:}} \quad \text{PR}(\theta_{\text{PS}}) \leq \min_\theta \text{DPR}(\theta, \mathcal{D}(\theta_{\text{PS}})); \qquad \underline{\text{Optimality:}} \quad \text{PR}(\theta_{\text{PO}}) \leq \min_\theta \text{DPR}(\theta, \mathcal{D}(\theta)).$$

Both solution concepts are distinct. [Per+20] showed how stable models need not be performatively optimal and optimal models need not be stable. Given its connection to the retraining feedback loops we discussed in the introduction, the question of when and why various algorithms converge to stability has been a main focus of the performative prediction literature. The bulk of prior work analyzed convergence to performative stability under the following Lipschitz condition on the distribution map:

Definition 3. *A distribution map $\mathcal{D}(\cdot)$ is ε-sensitive (or Lipschitz) if for all θ, θ'*

$$\mathcal{W}_1\big(\mathcal{D}(\theta), \mathcal{D}(\theta')\big) \leq \varepsilon \cdot \|\theta - \theta'\|_2.$$

Here, $\mathcal{W}_1$ is the earth mover's or Wasserstein-1 distance between the distributions $\mathcal{D}(\theta)$ and $\mathcal{D}(\theta')$.

Intuitively, the sensitivity condition formalizes the notion that small changes in predictions lead to small changes in the induced distribution. A canonical example where this assumption holds is in models of strategic classification [Har+16], where people manipulate their features to game a classifier. In particular, some works assume that if a linear classifier with parameters θ is deployed, individuals modify their features from x to $x - \varepsilon\theta$ in an effort to move to the positive side of the decision boundary. Using this sensitivity condition (Definition 3), prior work in this area established the following result.

Theorem 2 ([Per+20; Men+20])**.** *Assume that $\ell(z;\theta)$ is β-smooth, and γ-strongly convex in θ for every z. Furthermore, assume that the distribution map $\mathcal{D}(\cdot)$ is ε-sensitive with $\rho = \varepsilon\beta/\gamma < 1$. Then, retraining, $\theta_{t+1} = \arg\min_\theta \mathbb{E}_{z\sim\mathcal{D}(\theta_t)}\ell(z;\theta)$ converges to a unique stable point θ_{PS} at a linear rate, $\|\theta_T - \theta_{\mathrm{PS}}\| \leq \mathcal{O}(\exp(-T))$. Additionally, if $\nabla_\theta\ell(z;\theta)$ is β-Lipschitz as a function of z for every θ, then stochastic gradient descent, $\theta_{t+1} = \theta_t - \eta_t\nabla_\theta\ell(z_t;\theta_t), \quad z_t \sim \mathcal{D}(\theta_t)$, converges to a unique stable point at a polynomial rate, $\mathbb{E}\|\theta_T - \theta_{\mathrm{PS}}\|^2 \leq \mathcal{O}(1/T)$.*

In other words, if the loss function is well-conditioned (*pointwise* smooth and strongly convex) and the distribution map is Lipschitz with a small Lipschitz constant, then repeatedly retraining overcomes performativity and strategic behavior. While promising, the regime in which these results hold is limited. Prior work provides examples showing that retraining and gradient descent can diverge if losses are only convex but not strongly convex, or non-smooth, or if $\rho = \varepsilon\beta/\gamma > 1$.[1] More importantly, they illustrate how things can destabilize if the distribution map $\mathcal{D}(\cdot)$ is not continuous.

Discontinuous $\mathcal{D}(\cdot)$ commonly arise in practice whenever decision-makers choose actions by thresholding a prediction. To give a concrete example, public schools often assign interventions to students whose predicted probability of graduation, $f_\theta(x)$, falls below a cutoff. [Per+25] empirically show how the conditional likelihood of graduation for public school students in Wisconsin has a discontinuous jump if $f_\theta(x)$ is below this cutoff: $\Pr[y = 1 \mid f_\theta(x) > \tau] \ll \Pr[y = 1 \mid f_\theta(x) \leq \tau]$. These discrete jumps are not just common in practice; they also pose a technical barrier for convergence analyses.

To illustrate this, we recall the following example from [Per+20].

Example 1 ([Per+20])**.** *Fix $\ell(z;\theta)$ to be the squared loss $\ell(z;\theta) = (z-\theta)^2$ and assume that $z \sim \mathcal{D}(\theta)$ is a point mass distribution where $z = 0$ if $\theta > 1/2$ and $z = 1$ if $\theta \leq 1/2$ for $\theta \in [0,1]$. The performatively optimal solution for this problem is $\theta_{\mathrm{PO}} = 1/2$. However, there is no performatively stable point since the following function has no fixed points $g(\theta) = \arg\min_{\theta'\in[0,1]} \mathbb{E}_{z\sim\mathcal{D}(\theta)}(z-\theta')^2$. Furthermore, retraining as per Equation (1) will produce the infinite alternating sequence $0, 1, 0, 1, \ldots$*

From this example, we see that performatively stable models are not guaranteed to exist in general. Moreover, recent work by [Ana+26] illustrates how even if we do assume that stable models exist, it is PPAD-hard to compute a deterministic performatively stable point even if ℓ is γ-strongly convex, β-smooth, and $\mathcal{D}(\cdot)$ is ε-sensitive with $\rho = \beta\varepsilon/\gamma > 1$. In short, it is provably impossible to find deterministic stable points in full generality. In this work, we sidestep these negative results by allowing the learner to *randomize* across models. This enables us to guarantee stability without requiring continuity assumptions on $\mathcal{D}(\cdot)$. We summarize these comparisons to prior work in Table 1.

Mixture models are not just necessary to guarantee stability; they are also practically well-motivated. To see this, it helps to go back to the Wisconsin dropout example. Assume that the department uses a mixture of models μ. When deployed, students respond to the specific prediction $f_\theta(x)$ made on them, not the mixture μ ($\theta \sim \mu$ is the random model deployed at the time of their prediction). The mixture is stable if, on average across all students, no model would have done any better, and the department has no incentive to retrain. That is,

$$\mathbb{E}_{\theta\sim\mu_{\mathrm{PS}}}\mathbb{E}_{z\sim\mathcal{D}(\theta)}\ell(z;\theta) \leq \min_{\theta'} \mathbb{E}_{\theta\sim\mu_{\mathrm{PS}}}\mathbb{E}_{z\sim\mathcal{D}(\theta)}\ell(z;\theta').$$

This is the idea illustrated visually in Figure 1. We close this section by showing that, while not true in general, for certain problems performatively stable mixtures can also be performatively optimal. Furthermore, these mixtures are the natural limit of no-regret algorithms.

Example 1 (Continued)**.** *Let $\theta_1, \ldots, \theta_T$ be the sequence of iterates produced by repeatedly retraining, that is, running the follow-the-leader algorithm, $\theta_t = \arg\min_\theta \sum_{\tau=1}^{t-1} \ell(z_\tau;\theta)$ where $z_\tau \sim \mathcal{D}(\theta_\tau)$.*

[1] Note that the squared loss $(y - \theta^\top x)^2$ is not pointwise strongly convex.

For this problem, we can check that $\theta_t = \frac{1}{t-1}\sum_{\tau=1}^{t-1} z_\tau = 1/2 - 1\{t \text{ is odd}\}/[2(t-1)]$. *Moreover, the uniform mixture* μ_{PS} *over* $\theta_1, \dots, \theta_T$ *is* $\widetilde{\mathcal{O}}(1/T)$ *performatively stable.*

And, for any $\alpha > 0$, $\Pr_{\theta\sim\mu_{\mathrm{PS}}}[\theta \in (1/2 - \alpha, 1/2]] \to 1$ *as* $T \to \infty$. *Hence, the limiting distribution is also performatively optimal since* $\theta_{\mathrm{PO}} = 1/2$.

3 Performative Stability of Online Algorithms

The convergence to stability in the previous example is not an accident, but rather a reflection of a more general result that we will now present. Before proving our main theorem, we recall the definition of a no-regret learning algorithm.

Definition 4. *Consider the online learning protocol, where at every time step* t, *the Learner selects a vector* $\theta_t \in \Theta$, *and then, having seen* θ_t, *Nature selects a loss* $\ell_t : \Theta \to \mathbb{R}$. *An algorithm guarantees no regret if at every round* t *it deterministically selects* $\theta_t \in \Theta$ *such that no matter the choice of function* ℓ_t, $\sum_{t=1}^{T} \ell_t(\theta_t) - \min_{\theta\in\Theta} \sum_{t=1}^{T} \ell_t(\theta) \le \mathrm{Regret}_T \le o(T)$.

3.1 Main Result

As mentioned in Section 1.2, online algorithms are popular not just because they have guarantees in adversarial settings, they also serve as a powerful template for algorithm design in more structured domains. In that vein, our main result shows that if the ℓ_t are equal to the stochastic losses $\ell(z_t;\cdot)$ then the uniform mixture over the iterates of any no-regret algorithm is performatively stable.

Theorem 3. *Let* $\theta_1, \dots, \theta_T$ *be the sequence of iterates produced over* T *rounds by a deterministic online learning algorithm that guarantees sublinear regret on the sequence of losses* $\ell_t(\theta) = \ell(z_t;\theta)$ *where* $z_t \sim \mathcal{D}(\theta_t)$ *and* ℓ *is bounded. Let* $S = \{z_t\}_{t=1}^{T}$ *be the dataset of* T *draws from* $\mathcal{D}(\cdot)$ *and* μ *be the uniform distribution over* $\theta_1, \dots, \theta_T$. *Then, the mixture* μ *is* Regret_T/T*-performatively stable:*

$$\mathop{\mathbb{E}}_{S,\theta\sim\mu,z\sim\mathcal{D}(\theta)} \ell(z;\theta) \le \min_{\theta'\in\Theta} \mathop{\mathbb{E}}_{S,\theta\sim\mu,z\sim\mathcal{D}(\theta)} \ell(z;\theta') + \frac{\mathrm{Regret}_T}{T}.$$

Furthermore, if ℓ *is* L*-Lipschitz such that for any* z, $|\ell(z;\theta) - \ell(z;\theta')| \le L \cdot \|\theta - \theta'\|_2$, *and uniformly bounded by* B *such that* $\sup_{z,\theta} |\ell(z;\theta)| \le B$, *then with probability* $1-\delta$ *over* S,

$$\mathop{\mathbb{E}}_{\theta\sim\mu,z\sim\mathcal{D}(\theta)} \ell(z;\theta) \le \min_{\theta'\in\Theta} \mathop{\mathbb{E}}_{\theta\sim\mu,z\sim\mathcal{D}(\theta)} \ell(z;\theta') + (1+L)\frac{\mathrm{Regret}_T}{T} + 2B\sqrt{\frac{2\big(\log|\mathcal{N}_\eta| + \log(1/\delta)\big)}{T}}$$

where $\mathcal{N}_\eta \subseteq \Theta$ *is an* η*-Net, the smallest subset of* Θ *such that* $\sup_{\theta\in\Theta} \min_{\theta'\in\mathcal{N}_\eta} \|\theta - \theta'\|_2 \le \eta$ *with* $\eta = \mathrm{Regret}_T/T$.

On a technical level, the proof is an online-to-batch conversion that relies on a martingale argument. The main difference relative to standard online-to-batch analyses is that at every time step, we draw samples z_t from $\mathcal{D}(\theta_t)$ instead of a fixed distribution as in supervised learning. This makes critical use of the fact that online algorithms can guarantee low regret against adaptive adversaries that choose ℓ_t as a function of θ_t. This allows us to set $\ell_t(\theta) = \ell(z_t;\theta)$ where $z_t \sim \mathcal{D}(\theta_t)$ is drawn from the distribution induced by the iterate at time t. We also made use of the fact that online algorithms deterministically guarantee low regret against any comparator over the realized sequence of losses.

The statement is non-asymptotic (it holds for all $T \ge 1$). It also establishes stability both in expectation and with high probability over the samples S drawn from $\mathcal{D}(\cdot)$. To the best of our knowledge, this is the first high-probability result for finding stable points.

Comparison to Prior Work. Prior work in the performative prediction literature has also explored connections to online learning. [DX23] show that under the assumptions of Theorem 2, namely that losses are (pointwise) β-smooth, γ-strongly-convex, and $\mathcal{D}(\cdot)$ is ε-sensitive with $\varepsilon < \gamma/\beta$, then a single model equal to the average of the iterates $\bar{\theta} = \frac{1}{T}\sum_{t=1}^{T} \theta_t$ produced by an online algorithm is performatively stable. We on the other hand show that the mixture over iterates is performatively stable, but do not require smoothness, strong convexity, or that $\mathcal{D}(\cdot)$ is continuous (much less Lipschitz).

Perhaps the most closely related work to ours is [Per25]. He shows that in the *outcome* performative setting where features x are sampled from a fixed distribution $\mathcal{D}_x$, and a binary outcome y is sampled

from a (possibly discontinuous) performative distribution $\mathcal{D}_y(x, f_\theta(x))$ that depends on the learner's forecast $f_\theta(x)$, any algorithm that guarantees online multicalibration with respect to a particular set of functions can be converted to a (randomized) performatively stable model if ℓ is equal to the squared loss. By anchoring performativity on the model θ instead of the specific prediction $f_\theta(x)$, we achieve stability for general distribution maps $\mathcal{D}(\cdot)$ where θ influences the *joint* distribution over $z = (x, y)$ instead of just the distribution over y. Furthermore, our result holds for any loss where one can guarantee low regret and general outcomes y.

3.2 Implications

Theorem 3 is a general reduction showing we can extract a performatively stable mixture from the iterates produced by any no-regret learning algorithm. Here, we instantiate this reduction for specific choices of loss functions and algorithms, illustrating how Theorem 3 can be used to show that popular procedures are indeed stable in previously unknown settings. While prior work showed convergence to stability for different algorithms using tailored, bespoke arguments, our work provides a simple, unifying analysis that leverages known regret bounds to show convergence to stability. Most importantly, our results hold for arbitrary distribution maps $\mathcal{D}(\cdot)$, significantly expanding the domains where we can guarantee stability.

Throughout this section, we will assume that losses have gradients that are uniformly bounded by G and that the compact, convex domain Θ has diameter bounded by D:

$$\sup_{z\in\mathcal{Z},\theta\in\Theta} \|\nabla_\theta \ell(z;\theta)\|_2 \le G, \qquad \sup_{\theta,\theta'\in\Theta} \|\theta - \theta'\|_2 \le D.$$

Finite-Sample Guarantees for Retraining. Retraining, as defined in Equation (1), is the most commonly studied approach to finding a performatively stable solution. All analyses of this procedure (a) worked in the standard model with $\rho < 1$, and (b) considered only an idealized procedure where the learner finds the risk-minimizing solution at the population level at every time step. The finite sample results of [Per+20] are exponential in the problem dimension.

However, retraining is exactly equal to the classical follow-the-leader algorithm. Using standard regret bounds, along with Theorem 3, we can establish the first proof that smoothed retraining (where the learner retrains and then randomizes over the iterates) efficiently converges to stability in finite samples without making any assumptions on $\mathcal{D}(\cdot)$. We recall the following lemma:

Lemma 1 ([SK08; HAK07; Abe+08])**.** *Assume that $\ell_1, \dots, \ell_T$ is a sequence of functions that are all γ-strongly convex. If θ_t are chosen according to the follow-the-leader algorithm, $\theta_t \in \arg\min_\theta \sum_{j=1}^{t-1} \ell_j(\theta)$. Then, the iterates satisfy: $\sum_{t=1}^T \ell_t(\theta_t) - \min_\theta \sum_{t=1}^T \ell_t(\theta) \le \frac{G^2}{2\gamma}(\log(T)+1)$.*

Using this result, we can establish the following convergence for repeated retraining:

Corollary 1. *Let ℓ be a γ-strongly convex loss with gradients bounded by G over the domain Θ. Assume that the learner runs the following repeated retraining procedure, $\theta_t \in \arg\min_\theta \sum_{j=1}^{t-1} \ell(z_j;\theta)$ where $z_j \sim \mathcal{D}(\theta_j)$. Then, in expectation over the samples $S = \{z_t\}_{t=1}^T$ drawn from $\mathcal{D}(\cdot)$, the uniform mixture over $\theta_1, \dots, \theta_T$ is performatively stable:*

$$\mathop{\mathbb{E}}_{S,\theta\sim\mu,z\sim\mathcal{D}(\theta)} \ell(z;\theta) \le \min_{\theta'\in\Theta} \mathop{\mathbb{E}}_{S,\theta\sim\mu,z\sim\mathcal{D}(\theta)} \ell(z;\theta') + \frac{G^2}{2\gamma}\cdot\frac{(\log(T)+1)}{T}.$$

This result makes two adjustments relative to the standard definition of retraining as per Equation (1). First, Equation (1) defines θ_t to be the population risk minimizer on $\mathcal{D}(\theta_{t-1})$. Here, at every timestep, the learner finds the risk-minimizing solution over the *entire* history of samples z_t where each z_t is drawn from a *different* distribution. Furthermore, at the end, the learner randomizes across models instead of deploying the final iterate θ_T. As discussed in Section 2, this modification is necessary both to avoid the hardness (or non-existence) of deterministic performatively stable models and to guarantee stability while making no assumptions on $\mathcal{D}(\cdot)$.

Improved Guarantees for Gradient Descent. Apart from retraining, prior work also studied whether the stochastic gradient method also converges to stability. As outlined previously, [Men+20] and [DX23] proved that gradient descent converges to stability if the loss is again smooth, strongly convex, and $\mathcal{D}(\cdot)$ is Lipschitz with $\rho < 1$. Using our reduction, we can remove any assumptions on $\mathcal{D}(\cdot)$ and establish the first performative stability guarantees of any algorithm for losses that are non-smooth or just convex but not strongly convex. To do this, we recall the standard regret bounds:

Lemma 2 ([Zin03; HAK07]). *Assume that $\ell_1, \dots, \ell_T$ is a sequence of convex functions. If θ_t are chosen using online gradient descent $\theta_{t+1} = \mathrm{Proj}_\Theta\left[\theta_t - \eta_t \nabla \ell_t(\theta_t)\right]$ where $\eta_t = D/(G\sqrt{t})$, then the iterates satisfy the following regret bound: $\sum_{t=1}^T \ell_t(\theta_t) - \min_\theta \sum_{t=1}^T \ell_t(\theta) \le 1.5GD\sqrt{T}$. If losses ℓ_t are additionally γ-strongly convex and we set $\eta_t = 1/(\gamma t)$ then,*

$$\sum_{t=1}^T \ell_t(\theta_t) - \min_\theta \sum_{t=1}^T \ell_t(\theta) \le \frac{G^2}{2\gamma}(1 + \log T).$$

Tying this lemma together with Theorem 3, we get the following corollary:

Corollary 2. *Assume that the learner runs the following (greedy) gradient descent algorithm, $\theta_{t+1} = \mathrm{Proj}_\Theta\big[\theta_t - \eta_t \nabla_\theta \ell(z_t; \theta_t)\big]$, where $\eta_t = D/(G\sqrt{t})$, $z_t \sim \mathcal{D}(\theta_t)$, and that the loss ℓ has gradients bounded by G over the domain Θ. If the losses are convex, $\eta_t = D/(G\sqrt{t})$, then the uniform distribution over $\theta_1, \dots, \theta_T$ satisfies,*

$$\mathop{\mathbb{E}}_{S, \theta\sim\mu, z\sim\mathcal{D}(\theta)} \ell(z;\theta) \le \min_{\theta'\in\Theta} \mathop{\mathbb{E}}_{S, \theta\sim\mu, z\sim\mathcal{D}(\theta)} \ell(z;\theta') + \frac{3}{2}\cdot\frac{GD}{\sqrt{T}},$$

where $S = \{z_t\}_{t=1}^T$. If the losses are γ-strongly convex, and $\eta_t = 1/(\gamma t)$, then

$$\mathop{\mathbb{E}}_{S, \theta\sim\mu, z\sim\mathcal{D}(\theta)} \ell(z;\theta) \le \min_{\theta'\in\Theta} \mathop{\mathbb{E}}_{S, \theta\sim\mu, z\sim\mathcal{D}(\theta)} \ell(z;\theta') + \frac{G^2}{2\gamma}\cdot\frac{1+\log(T)}{T}.$$

Ignoring quantities that depend on the diameter of the space Θ, the bounds we get for gradient descent in the strongly-convex case match those from [Men+20] and [DX23] up to a $\log(T)$ factor. However, we achieve this guarantee for general distribution maps $\mathcal{D}(\cdot)$, rather than just for strongly Lipschitz $\mathcal{D}(\cdot)$. This result holds for what [Men+20] refer to as *greedy* gradient descent where the learner performs a model update and redeploys at every time step. It is an interesting direction for future work to understand if by leveraging other online algorithms (e.g., the primal-dual analysis in [SK08]) one could also prove that *lazy* versions of gradient descent where the learner performs many updates before redeployment also lead to stability, while reducing the total number of deployed models. Another interesting question is whether our online-to-batch analysis can be sharpened to yield a $\widetilde{\mathcal{O}}(1/T)$ rate with high probability.

Guaranteeing Performative Stability in New Regimes. As a final corollary, we highlight how the reduction can be used to guarantee stability in new domains (or in old domains at faster rates). In particular, to get a $\widetilde{\mathcal{O}}(1/T)$ rate for stability, prior work required that the losses ℓ were *pointwise* strongly convex. That is, for every z, $\ell(z;\theta)$ is a strongly convex function in θ. If we don't add a regularizer, this fails for the two most common losses in machine learning: log loss $\log(1+\exp(-y\cdot\theta^\top x))$ and squared loss $(y-\theta^\top x)^2$. These losses are, however, exp-concave and one can achieve $\log(T)$ regret for these problems using second-order methods. Using this connection, we can establish the first positive results for performative stability in these common settings (and with fast rates).

Corollary 3. *Let $\ell_1, \dots, \ell_T$ be a sequence of losses that are α-exp-concave for $\alpha > 0$. That is, for every $\theta \in \mathbb{R}^d$, $\nabla^2 \ell_t(\theta) - \alpha \nabla \ell_t(\theta) \nabla \ell_t(\theta)^\top$ is a positive semi-definite matrix. There exists an efficient deterministic online algorithm, the online Newton step algorithm, that guarantees that for $T \ge 4$:*

$$\sum_{t=1}^T \ell_t(\theta_t) - \min_\theta \sum_{t=1}^T \ell_t(\theta) \le 2\left(\frac{1}{\alpha} + GD\right)\cdot d\cdot\log(T).$$

Consequently, if the learner produces iterates θ_t by running online Newton step on the losses $\ell(z_t;\theta)$ for $z_t \sim \mathcal{D}(\theta_t)$, the uniform distribution over iterates $\theta_1, \dots, \theta_T$ is performatively stable:

$$\mathop{\mathbb{E}}_{S, \theta\sim\mu, z\sim\mathcal{D}(\theta)} \ell(z;\theta) \le \min_{\theta'\in\Theta} \mathop{\mathbb{E}}_{S, \theta\sim\mu, z\sim\mathcal{D}(\theta)} \ell(z;\theta') + 2\left(\frac{1}{\alpha} + GD\right)\cdot\frac{d\cdot\log T}{T}.$$

where $S = \{z_t\}_{t=1}^T$ denotes the dataset drawn from $\mathcal{D}$.

We point the reader to the excellent monograph [Haz16] for details on the online Newton step algorithm. The corollary follows again by tying together known regret bounds [HAK07] with Theorem 3. Lastly, we note that we do not intend this section to be an exhaustive list of corollaries.

Online optimization is by now a very rich and mature field with sophisticated algorithms that in many ways improve on the ones we present here (in particular on their dimension, and G, D dependence). We chose these due to their simplicity and to highlight some of their consequences for performativity. Furthermore, in these corollaries we take expectation over the samples in S for simplicity, but we could have stated the high probability version by using the second part of Theorem 3.

4 Applications to Multiplayer Performative Prediction

Our results so far considered the standard performative prediction setting with a single learner. However, many prediction systems operate in shared environments with multiple decision-makers, each deploying their own models. In traffic prediction, for instance, commute times are jointly shaped by the routing recommendations of several navigation apps. In digital advertising, user behavior is influenced by the combined choices of platforms competing for attention. This dynamic was formalized by [Nar+23] under the name of *multiplayer* performative prediction. As we now show, our reduction extends to this setting with essentially the same proof.

In the multiplayer problem, there are n players. Each player $i \in [n]$ deploys a model $\theta^i \in \Theta^i \subseteq \mathbb{R}^{d^i}$, and the data observed by player i is drawn from a distribution that depends on *all* of the deployed models. That is, $z^i \sim \mathcal{D}^i(\theta^1, \dots, \theta^n)$ where $\mathcal{D}^i(\cdot)$ is the distribution map for player i which depends on the choice of all n players. Furthermore, each player has their own loss $\ell^i(z^i; \theta^i)$. The natural generalization of performative stability to this setting is the following.

Definition 5 (Multiplayer performative stability)**.** *A joint distribution μ over profiles $(\theta^1, \dots, \theta^n) \in \Theta^1 \times \cdots \times \Theta^n$ is ε-performatively stable if for every player $i \in [n]$,*

$$\mathbb{E}_{(\theta^1,\dots,\theta^n)\sim\mu}\mathbb{E}_{z^i\sim\mathcal{D}^i(\theta^1,\dots,\theta^n)}\ell^i(z^i;\theta^i) \le \min_{\hat{\theta}^i\in\Theta^i} \mathbb{E}_{(\theta^1,\dots,\theta^n)\sim\mu}\mathbb{E}_{z^i\sim\mathcal{D}^i(\theta^1,\dots,\theta^n)}\ell^i(z^i;\hat{\theta}^i) + \varepsilon.$$

That is, no player can unilaterally reduce their expected loss under μ by switching to a different model. Note that μ is a distribution over joint profiles, not a product of independent marginals. This is what allows us to capture the correlation between players' iterates that naturally arises from their shared interaction history. Our main result extends to this multiplayer setting with essentially the same proof: if every player runs a no-regret algorithm on their observed losses, then the uniform distribution over the joint trajectory is performatively stable.

Theorem 4. *Suppose each player $i \in [n]$ produces iterates $\theta^i_1, \dots, \theta^i_T$ using a deterministic online algorithm that achieves sublinear regret $\mathrm{Regret}^i(T)$ against the sequence of losses $\ell^i_t(\theta^i) = \ell^i(z^i_t; \theta^i)$, where $z^i_t \sim \mathcal{D}^i(\theta^1_t, \dots, \theta^n_t)$ and each ℓ^i is bounded. Let μ be the uniform distribution over the joint profiles $\{(\theta^1_t, \dots, \theta^n_t)\}_{t=1}^T$. Then, in expectation over the draws from each $\mathcal{D}^i(\cdot)$, the mixture μ is $(\max_i \mathrm{Regret}^i(T)/T)$-performatively stable.*

Similarly to Theorem 3, the result can be extended to hold path-wise in high probability by using a concentration inequality and taking a union bound over all comparators and players.

Comparison to prior work. [Nar+23] established the first convergence results to multiplayer performative stability, showing that repeated retraining and gradient descent converge if every player's loss ℓ^i is β_i-smooth and γ_i-strongly convex, every distribution map $\mathcal{D}^i(\cdot)$ is ε_i-sensitive in the sense of Definition 3, and $\sqrt{\sum_{i=1}^n (\beta_i\varepsilon_i/\gamma_i)^2} < 1$. That is, the *norm* across n players of the individual parameters $\beta_i\varepsilon_i/\gamma_i$ needs to be small, a condition which becomes harder to satisfy as the number of players n increases.

Theorem 4 sidesteps this entirely. We require no smoothness or strong convexity on any ℓ^i, no sensitivity assumption on any $\mathcal{D}^i(\cdot)$, and in particular no joint condition relating constants across players. As long as each player runs a no-regret algorithm on their own observed losses, the uniform mixture over the joint trajectory is performatively stable, regardless of how many players there are or how strongly their predictions influence the shared environment.

5 Applications to Stateful Performative Prediction

The same analysis template also applies to the *stateful* setting, where data is a function of the entire history of deployed models rather than just the current one. This setting was introduced by [BHK22] and captures settings such as recommender systems, where recommendations shape a user's future behavior, and that behavior depends on the entire history of recommendations rather than just the most recent one.

We consider the following generalized version of the stateful problem. The learner predicts sequentially over rounds $t = 1, \ldots, T$ using the model θ_t. Unlike the standard setting where there is a fixed distribution map $\mathcal{D}(\cdot)$ and $z_t \sim \mathcal{D}(\theta_t)$, we now assume that $z_t \sim \mathcal{D}_t$, where $\mathcal{D}_t$ is an arbitrary distribution over $\mathcal{Z}$ that is potentially a function not just of θ_t, but of the entire history of models $\theta_1, \ldots, \theta_t$. Unlike prior work, we *make no assumptions* whatsoever on how $\mathcal{D}_t$ depends on this history. For instance, [BHK22] assumed that distributions were produced according to the geometric decay dynamic $\mathcal{D}_t = (1-\alpha)\mathcal{D}_{t-1} + \alpha\mathcal{D}(\theta_t)$, where $\mathcal{D} : \Theta \to \Delta(\mathcal{Z})$ is a fixed, Lipschitz distribution map and $\alpha \in (0, 1)$ is a problem parameter. In recent work, [LZ26] considered a different "partially performative" setup where $\mathcal{D}_t = (1-\alpha)\mathcal{D}(\theta_t) + \alpha\mathcal{P}_t$ for an exogenous sequence of distributions $\mathcal{P}_t$ that is unaffected by performativity, with $\mathcal{D}(\cdot)$ again a fixed, Lipschitz distribution map. Both are special, restricted cases of our setup. Without any kind of restriction, the sequence of distributions $\mathcal{D}_t$ may not have well-defined limits no matter the sequence of models. We therefore consider the following stability guarantee.

Definition 6. *In the stateful setup above, we say that a sequence of models $\{\theta_t\}_{t=1}^T$ is a performatively stable sequence if*

$$\frac{1}{T}\sum_{t=1}^{T} \mathbb{E}_{z\sim\mathcal{D}_t}[\ell(z;\theta_t)] \leq \min_{\theta\in\Theta} \frac{1}{T}\sum_{t=1}^{T} \mathbb{E}_{z\sim\mathcal{D}_t}[\ell(z;\theta)] + o(1),$$

where $o(1)$ denotes a term that goes to 0 as $T \to \infty$.

Performatively stable sequences generalize the same desiderata from the stateless case to the stateful setting: from the perspective of retraining, there is no incentive to deviate to any alternative model θ since no fixed model has lower loss over the induced distributions. The same reduction as before shows that no-regret algorithms produce performatively stable sequences, again without any assumptions on the dynamics. The proof (Section C) is identical to that of Theorem 3: since $\mathcal{D}_t$ is a function of $\theta_1, \ldots, \theta_t$, it is $\mathcal{F}_{t-1}$-measurable, so the same martingale argument goes through. We state only the in-expectation version; the high-probability version follows identically.

Theorem 5. *Let $\theta_1, \ldots, \theta_T$ be the sequence of iterates produced over T rounds by a deterministic online learning algorithm that guarantees sublinear regret on the sequence of losses $\ell_t(\theta) = \ell(z_t; \theta)$, where $z_t \sim \mathcal{D}_t$, each $\mathcal{D}_t$ is an arbitrary function of $\theta_1, \ldots, \theta_t$, and ℓ is bounded. Then, $\{\theta_t\}_{t=1}^T$ is a performatively stable sequence; with $S = \{z_t\}_{t=1}^T$,*

$$\mathbb{E}_S\left[\frac{1}{T}\sum_{t=1}^{T} \mathbb{E}_{z\sim\mathcal{D}_t}\ell(z;\theta_t)\right] \leq \min_{\theta\in\Theta} \mathbb{E}_S\left[\frac{1}{T}\sum_{t=1}^{T} \mathbb{E}_{z\sim\mathcal{D}_t}\ell(z;\theta)\right] + \frac{\mathrm{Regret}_T}{T}.$$

Comparison to prior work. The theorem above establishes that, under no assumptions on the distribution map, online algorithms yield a generalized notion of stability. To the best of our knowledge, this is the first assumption-free notion of stability in the stateful case. Both [BHK22] and [LZ26] work with the classical assumptions that losses are well-conditioned and that at least part of the underlying data generating process is sufficiently Lipschitz. Our work illustrates how one can achieve a similarly motivated notion of stability while removing all assumptions around how people respond to predictions.

6 Discussion and Future Work

In this work, we establish a reduction between no-regret learning in online optimization and performative stability. As a consequence of this connection, we expand the regimes of problems in which one can achieve this natural solution concept and provide an explanation for why common learning procedures do not produce runaway feedback loops in dynamic, social environments. We believe that there are a number of exciting directions to pursue in future work.

One especially important direction pertains to connections between stability and optimality. As we discussed in Section 2, performatively stable models are fixed points of retraining, but they do not necessarily minimize the performative risk. Prior work by [MPZ21] and [Per25] shows that there exist worst-case instances where stable solutions in fact *maximize* the performative risk. This begs a natural question: what structural conditions can we impose on $\mathcal{D}(\cdot)$ to ensure that stable points are also optimal? We feel this question is particularly interesting in light of our results since we now have a robust set of algorithmic techniques to achieve stability. Finding natural settings where stability and optimality coincide would enable us to apply this algorithmic toolkit to achieve a thus far more challenging notion.

Acknowledgments and Disclosure of Funding

We are both graciously supported by AI2050 Early Career Fellowships.

References

[Abe+08] Jacob Abernethy, Peter L Bartlett, Alexander Rakhlin, and Ambuj Tewari. "Optimal strategies and minimax lower bounds for online convex games". In: (2008).

[ABH11] Jacob Abernethy, Peter L Bartlett, and Elad Hazan. "Blackwell approachability and no-regret learning are equivalent". In: *Conference on Learning Theory*. 2011.

[ALO15] Zeyuan Allen-Zhu, Zhenyu Liao, and Lorenzo Orecchia. "Spectral sparsification and regret minimization beyond matrix multiplicative updates". In: *ACM Symposium on Theory of Computing*. 2015.

[Ana+26] Ioannis Anagnostides, Rohan Chauhan, Ioannis Panageas, Tuomas Sandholm, and Jingming Yan. "On the Computational Complexity of Performative Prediction". In: *arXiv preprint arXiv:2601.20180* (2026).

[BHK22] Gavin Brown, Shlomi Hod, and Iden Kalemaj. "Performative prediction in a stateful world". In: *International Conference on Artificial Intelligence and Statistics*. 2022.

[CCG04] Nicolo Cesa-Bianchi, Alex Conconi, and Claudio Gentile. "On the generalization ability of on-line learning algorithms". In: *IEEE Transactions on Information Theory* (2004).

[CMO23] Ashok Cutkosky, Harsh Mehta, and Francesco Orabona. "Optimal stochastic non-smooth non-convex optimization through online-to-non-convex conversion". In: *International Conference on Machine Learning*. 2023.

[Dag+24] Yuval Dagan, Constantinos Daskalakis, Maxwell Fishelson, and Noah Golowich. "From external to swap regret 2.0: An efficient reduction and oblivious adversary for large action spaces". In: *ACM Symposium on Theory of Computing*. 2024.

[Das+25] Constantinos Daskalakis, Gabriele Farina, Maxwell Fishelson, Charilaos Pipis, and Jon Schneider. "Efficient learning and computation of linear correlated equilibrium in general convex games". In: *ACM Symposium on Theory of Computing*. 2025.

[DX23] Dmitriy Drusvyatskiy and Lin Xiao. "Stochastic optimization with decision-dependent distributions". In: *Mathematics of Operations Research* (2023).

[Far+22] Gabriele Farina, Andrea Celli, Alberto Marchesi, and Nicola Gatti. "Simple Uncoupled No-regret Learning Dynamics for Extensive-form Correlated Equilibrium". In: *Journal of the ACM* (2022).

[Far26] Gabriele Farina. "Turning defense into offense in $O(\log 1/\epsilon)$ steps: Efficient constructive proof of the minimax theorem". In: *ACM SIGecom Exchanges* (2026).

[FP23] Gabriele Farina and Charilaos Pipis. "Polynomial-time linear-swap regret minimization in imperfect-information sequential games". In: *Advances in Neural Information Processing Systems* (2023).

[FP24] Gabriele Farina and Charilaos Pipis. "Polynomial-Time Computation of Exact Φ-Equilibria in Polyhedral Games". In: *Advances in Neural Information Processing Systems* (2024).

[FV97] Dean P Foster and Rakesh V Vohra. "Calibrated learning and correlated equilibrium". In: *Games and Economic Behavior* (1997).

[FV98] Dean P Foster and Rakesh V Vohra. "Asymptotic calibration". In: *Biometrika* (1998).

[GGM08] Geoffrey J Gordon, Amy Greenwald, and Casey Marks. "No-regret learning in convex games". In: *International Conference on Machine learning*. 2008.

[HAK07] Elad Hazan, Amit Agarwal, and Satyen Kale. "Logarithmic regret algorithms for online convex optimization". In: *Machine Learning* (2007).

[Har+16] Moritz Hardt, Nimrod Megiddo, Christos Papadimitriou, and Mary Wootters. "Strategic classification". In: *Innovations in Theoretical Computer Science*. 2016.

[Haz16] Elad Hazan. "Introduction to online convex optimization". In: *Foundations and Trends in Optimization* (2016).

[HM00] Sergiu Hart and Andreu Mas-Colell. "A simple adaptive procedure leading to correlated equilibrium". In: *Econometrica* (2000).

[HM25] Moritz Hardt and Celestine Mendler-Dünner. “Performative prediction: Past and Future”. In: *Statistical Science* (2025).

[IYZ21] Zachary Izzo, Lexing Ying, and James Zou. “How to learn when data reacts to your model: performative gradient descent”. In: *International Conference on Machine Learning*. 2021.

[JZM22] Meena Jagadeesan, Tijana Zrnic, and Celestine Mendler-Dünner. “Regret minimization with performative feedback”. In: *International Conference on Machine Learning*. 2022.

[KP23] Michael P Kim and Juan C Perdomo. “Making decisions under outcome performativity”. In: *Innovations in Theoretical Computer Science* (2023).

[Leg+24] Davide Legacci, Panayotis Mertikopoulos, Christos Papadimitriou, Georgios Piliouras, and Bary Pradelski. “No-regret learning in harmonic games: Extrapolation in the face of conflicting interests”. In: *Advances in Neural Information Processing Systems* (2024).

[LW22] Qiang Li and Hoi-To Wai. “State dependent performative prediction with stochastic approximation”. In: *International Conference on Artificial Intelligence and Statistics*. 2022.

[LYW22] Qiang Li, Chung-Yiu Yau, and Hoi-To Wai. “Multi-agent performative prediction with greedy deployment and consensus seeking agents”. In: *Advances in Neural Information Processing Systems* (2022).

[LZ24] Licong Lin and Tijana Zrnic. “Plug-in Performative Optimization”. In: *International Conference on Machine Learning*. 2024.

[LZ26] Jaewook Lee and Tijana Zrnic. *Partially Performative Prediction*. 2026.

[Men+20] Celestine Mendler-Dünner, Juan Perdomo, Tijana Zrnic, and Moritz Hardt. “Stochastic optimization for performative prediction”. In: *Advances in Neural Information Processing Systems* (2020).

[MMG23] Mehrnaz Mofakhami, Ioannis Mitliagkas, and Gauthier Gidel. “Performative prediction with neural networks”. In: *International Conference on Artificial Intelligence and Statistics*. 2023.

[MPZ21] John P Miller, Juan C Perdomo, and Tijana Zrnic. “Outside the echo chamber: Optimizing the performative risk”. In: *International Conference on Machine Learning*. 2021.

[Nar+23] Adhyyan Narang, Evan Faulkner, Dmitriy Drusvyatskiy, Maryam Fazel, and Lillian J Ratliff. “Multiplayer performative prediction: Learning in decision-dependent games”. In: *Journal of Machine Learning Research* (2023).

[OP16] Francesco Orabona and Dávid Pál. “Coin betting and parameter-free online learning”. In: *Advances in Neural Information Processing Systems* (2016).

[Per+20] Juan C. Perdomo, Tijana Zrnic, Celestine Mendler-Dünner, and Moritz Hardt. “Performative prediction”. In: *International Conference on Machine Learning* (2020).

[Per+25] Juan Carlos Perdomo, Tolani Britton, Moritz Hardt, and Rediet Abebe. “Difficult Lessons on Social Prediction from Wisconsin Public Schools”. In: *ACM Conference on Fairness, Accountability, and Transparency*. 2025.

[Per25] Juan C. Perdomo. “Revisiting the Predictability of Performative, Social Events”. In: *International Conference on Machine Learning* (2025).

[PR08] Christos H Papadimitriou and Tim Roughgarden. “Computing correlated equilibria in multi-player games”. In: *Journal of the ACM* (2008).

[PR24] Binghui Peng and Aviad Rubinstein. “Fast swap regret minimization and applications to approximate correlated equilibria”. In: *ACM Symposium on Theory of Computing*. 2024.

[PY23] Georgios Piliouras and Fang-Yi Yu. “Multi-agent performative prediction: From global stability and optimality to chaos”. In: *ACM Conference on Economics and Computation*. 2023.

[Ram+23] Aaditya Ramdas, Peter Grünwald, Vladimir Vovk, and Glenn Shafer. “Game-theoretic statistics and safe anytime-valid inference”. In: *Statistical Science* (2023).

[Ray+22] Mitas Ray, Lillian J Ratliff, Dmitriy Drusvyatskiy, and Maryam Fazel. “Decision-dependent risk minimization in geometrically decaying dynamic environments”. In: *AAAI Conference on Artificial Intelligence*. 2022.

[RS17] Alexander Rakhlin and Karthik Sridharan. “On equivalence of martingale tail bounds and deterministic regret inequalities”. In: *Conference on Learning Theory*. 2017.

[SK08] Shai Shalev-Shwartz and Sham M Kakade. "Mind the duality gap: Logarithmic regret algorithms for online optimization". In: *Advances in Neural Information Processing Systems* (2008).

[TH23] Rohan Taori and Tatsunori Hashimoto. "Data feedback loops: Model-driven amplification of dataset biases". In: *International Conference on Machine Learning*. 2023.

[Tre19] Luca Trevisan. *Online Optimization Post 2: Constructing Pseudorandom Sets*. Available online at `https://lucatrevisan.wordpress.com/2019/04/25/online-optimization-post-2-constructing-pseudorandom-sets`. [Online; accessed Feb. 2026]. May 2019.

[WBD21] Killian Wood, Gianluca Bianchin, and Emiliano Dall'Anese. "Online projected gradient descent for stochastic optimization with decision-dependent distributions". In: *IEEE Control Systems Letters* (2021).

[Zin03] Martin Zinkevich. "Online convex programming and generalized infinitesimal gradient ascent". In: *International Conference on Machine Learning*. 2003.

A Proof of Theorem 3

Proof. Let $\theta_\star$ be any vector in Θ. The goal is to upper bound,

$$\mathbb{E}_{\theta\sim\mu, z\sim\mathcal{D}(\theta)}\left[\ell(z;\theta) - \ell(z;\theta_\star)\right].$$

We start out by expanding the first expectation over μ. Since μ is chosen to be the uniform distribution over $\theta_1, \dots, \theta_T$, this expression then becomes:

$$\frac{1}{T}\sum_{t=1}^{T}\mathbb{E}_{z\sim\mathcal{D}(\theta_t)}\Big[\ell(z;\theta_t) - \ell(z;\theta_\star)\Big].$$

Next, we add and subtract $\ell(z_t;\theta_t) - \ell(z_t;\theta_\star)$ where $z_t \sim \mathcal{D}(\theta_t)$ for every t. Rearranging, the expression becomes

$$\underbrace{\frac{1}{T}\sum_{t=1}^{T}\mathbb{E}_{z\sim\mathcal{D}(\theta_t)}\Big[\ell(z;\theta_t) - \ell(z;\theta_\star) - (\ell(z_t;\theta_t) - \ell(z_t;\theta_\star))\Big]}_{\text{small by a martingale argument}} + \underbrace{\frac{1}{T}\sum_{t=1}^{T}\Big[\ell(z_t;\theta_t) - \ell(z_t;\theta_\star)\Big]}_{\text{small by regret}}.$$

In particular, if we define $V_t = \mathbb{E}_{z\sim\mathcal{D}(\theta_t)}\Big[\ell(z;\theta_t) - \ell(z;\theta_\star)\Big] - [\ell(z_t;\theta_t) - \ell(z_t;\theta_\star)]$, then

$$\frac{1}{T}\sum_{t=1}^{T}\mathbb{E}_{z\sim\mathcal{D}(\theta_t)}\left[\ell(z;\theta_t) - \ell(z;\theta_\star) - (\ell(z_t;\theta_t) - \ell(z_t;\theta_\star))\right] = \frac{1}{T}\sum_{t=1}^{T}V_t.$$

Here, $V_1, \dots, V_T$ is a martingale difference sequence with respect to the filtration $\{\mathcal{F}_t\}_{t\geq 0}$ defined by $\mathcal{F}_0 \coloneqq \sigma(\theta_1)$

$$\mathcal{F}_t \coloneqq \sigma(\theta_1, z_1, \theta_2, z_2, \dots, \theta_t, z_t, \theta_{t+1}) \quad \text{for } t \geq 1,$$

where $\sigma(\cdot)$ denotes the σ-algebra generated by the random variables. By construction, θ_t is measurable with respect to $\mathcal{F}_{t-1}$ (we assumed the online algorithm is deterministic), and conditional on $\mathcal{F}_{t-1}$ the sample z_t is an independent draw from $\mathcal{D}(\theta_t)$. Hence, V_t is $\mathcal{F}_t$-measurable and

$$\mathbb{E}\Big[V_t \mid \mathcal{F}_{t-1}\Big] = \mathbb{E}_{z\sim\mathcal{D}(\theta_t)}\Big[\ell(z;\theta_t) - \ell(z;\theta_\star)\Big] - \mathbb{E}\Big[\ell(z_t;\theta_t) - \ell(z_t;\theta_\star) \;\Big|\; \mathcal{F}_{t-1}\Big] = 0.$$

Since losses are bounded, $\mathbb{E}[|V_t|] < \infty$, and therefore $\mathbb{E}_S[\sum_{t=1}^{T} V_t] = 0$. Since $\theta_\star$ was arbitrary, this proves the first statement.

To show the second statement, we can just union bound over an η-net. In particular, let $\mathcal{N}_\eta$ be an η-net of Θ. That is, define $\mathcal{N}_\eta \subset \Theta$ to be the set consisting of the minimum number of points in Θ such that for every $\theta' \in \Theta$ there exists $\theta \in \mathcal{N}_\eta$ such that $\|\theta - \theta'\|_2 \leq \eta$.

If we redo the same argument as before with respect to each $\theta \in \mathcal{N}_\eta$ and take a union bound over $\mathcal{N}_\eta$, we obtain that with probability at least $1-\delta$,

$$\sup_{\theta\in\mathcal{N}_\eta}\frac{1}{T}\sum_{t=1}^{T}V_t(\theta) \leq 2B\sqrt{\frac{2\big(\log|\mathcal{N}_\eta| + \log(1/\delta)\big)}{T}}. \tag{2}$$

To complete the argument, we need to relate the loss of an arbitrary θ to that of one in $\mathcal{N}_\eta$. By Lipschitzness of the loss, for any $\theta_\star$, there exists a $\bar{\theta} \in \mathcal{N}_\eta$ such that:

$$\begin{aligned}\frac{1}{T}\sum_{t=1}^{T}\mathbb{E}_{z\sim\mathcal{D}(\theta_t)}\left[\ell(z;\theta_t)-\ell(z;\theta_\star)\right] &\le \frac{1}{T}\sum_{t=1}^{T}\mathbb{E}_{z\sim\mathcal{D}(\theta_t)}\left[\ell(z;\theta_t)-\ell(z;\bar{\theta})\right] \\ &\quad + \frac{1}{T}\sum_{t=1}^{T}\mathbb{E}_{z\sim\mathcal{D}(\theta_t)}\left[\ell(z;\bar{\theta})-\ell(z;\theta_\star)\right] \\ &\le \frac{\mathrm{Regret}_T}{T} + 2B\sqrt{\frac{2\big(\log|\mathcal{N}_\eta|+\log(1/\delta)\big)}{T}} + \eta L.\end{aligned}$$

Setting η to Regret_T/T finishes the proof. □

B Proof of Theorem 4

Proof. Fix a player $i \in [n]$ and an arbitrary comparator $\theta^i_\star \in \Theta^i$. For brevity, let $\theta_t = (\theta^1_t, \dots, \theta^n_t)$ and $z_t = (z^1_t, \dots, z^n_t)$. The goal is to upper bound

$$\mathbb{E}_{(\theta^1,\dots,\theta^n)\sim\mu}\mathbb{E}_{z^i\sim\mathcal{D}^i(\theta^1,\dots,\theta^n)}\left[\ell^i(z^i;\theta^i)-\ell^i(z^i;\theta^i_\star)\right] = \frac{1}{T}\sum_{t=1}^{T}\mathbb{E}_{z^i\sim\mathcal{D}^i(\theta_t)}\left[\ell^i(z^i;\theta^i_t)-\ell^i(z^i;\theta^i_\star)\right],$$

where the equality follows because μ is the uniform distribution over the joint profiles. Following the proof of Theorem 3, we add and subtract $\ell^i(z^i_t;\theta^i_t) - \ell^i(z^i_t;\theta^i_\star)$ to obtain

$$\underbrace{\frac{1}{T}\sum_{t=1}^{T}\mathbb{E}_{z^i\sim\mathcal{D}^i(\theta_t)}\left[\ell^i(z^i;\theta^i_t)-\ell^i(z^i;\theta^i_\star)-\left(\ell^i(z^i_t;\theta^i_t)-\ell^i(z^i_t;\theta^i_\star)\right)\right]}_{\text{small by a martingale argument}} + \underbrace{\frac{1}{T}\sum_{t=1}^{T}\left[\ell^i(z^i_t;\theta^i_t)-\ell^i(z^i_t;\theta^i_\star)\right]}_{\text{small by regret}}.$$

Define $V^i_t = \mathbb{E}_{z^i\sim\mathcal{D}^i(\theta_t)}[\ell^i(z^i;\theta^i_t)-\ell^i(z^i;\theta^i_\star)] - [\ell^i(z^i_t;\theta^i_t)-\ell^i(z^i_t;\theta^i_\star)]$. We show that $V^i_1, \dots, V^i_T$ is a martingale difference sequence with respect to the filtration $\{\mathcal{F}_t\}_{t\ge0}$ defined by $\mathcal{F}_0 \coloneqq \sigma(\theta_1)$ and

$$\mathcal{F}_t \coloneqq \sigma(\theta_1, z_1, \theta_2, z_2, \dots, \theta_t, z_t, \theta_{t+1}) \quad \text{for } t \ge 1.$$

By construction, θ_t (and hence θ^i_t) is $\mathcal{F}_{t-1}$-measurable, and conditional on $\mathcal{F}_{t-1}$ the sample z^i_t is an independent draw from $\mathcal{D}^i(\theta_t)$. Hence V^i_t is $\mathcal{F}_t$-measurable and $\mathbb{E}[V^i_t \mid \mathcal{F}_{t-1}] = 0$. Since losses are bounded, $\mathbb{E}[|V^i_t|] < \infty$, and therefore $\mathbb{E}_S\big[\sum_{t=1}^T V^i_t\big] = 0$.

For the second term, the no-regret guarantee of player i's algorithm gives

$$\sum_{t=1}^{T}\ell^i(z^i_t;\theta^i_t) - \ell^i(z^i_t;\theta^i_\star) \le \mathrm{Regret}^i(T).$$

Combining the two bounds and noting that the argument holds for any $\theta^i_\star \in \Theta^i$ establishes the result for player i. Since i was arbitrary, the bound holds simultaneously for every player. □

C Proof of Theorem 5

Proof. Fix an arbitrary comparator $\theta_\star \in \Theta$. As in the proof of Theorem 3, we add and subtract $\ell(z_t;\theta_t) - \ell(z_t;\theta_\star)$ and write

$$\frac{1}{T}\sum_{t=1}^{T}\mathbb{E}_{z\sim\mathcal{D}_t}\Big[\ell(z;\theta_t)-\ell(z;\theta_\star)\Big] = \frac{1}{T}\sum_{t=1}^{T}V_t + \frac{1}{T}\sum_{t=1}^{T}\Big[\ell(z_t;\theta_t)-\ell(z_t;\theta_\star)\Big],$$

where $V_t \coloneqq \mathbb{E}_{z\sim\mathcal{D}_t}[\ell(z;\theta_t)-\ell(z;\theta_\star)] - [\ell(z_t;\theta_t)-\ell(z_t;\theta_\star)]$. The second term is at most Regret_T/T by the no-regret guarantee. For the first, consider the filtration $\{\mathcal{F}_t\}_{t\ge0}$ defined by $\mathcal{F}_0 \coloneqq \sigma(\theta_1)$ and $\mathcal{F}_t \coloneqq \sigma(\theta_1, z_1, \dots, \theta_t, z_t, \theta_{t+1})$ for $t \ge 1$. Since the algorithm is deterministic, $\theta_1, \dots, \theta_t$ are $\mathcal{F}_{t-1}$-measurable, and hence so is $\mathcal{D}_t$. Conditional on $\mathcal{F}_{t-1}$, z_t is a draw from $\mathcal{D}_t$, so $\mathbb{E}[V_t \mid \mathcal{F}_{t-1}] = 0$. Thus $V_1, \dots, V_T$ is a bounded martingale difference sequence and $\mathbb{E}_S[\sum_{t=1}^T V_t] = 0$. Since $\theta_\star$ was arbitrary, the result follows. □

NeurIPS Paper Checklist

1. **Claims**

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

   Answer: [Yes]

   Justification: Yes, all claims are substantiated through formal mathematical results.

   Guidelines:

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

2. **Limitations**

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

   Answer: [Yes]

   Justification: We explain the formal conditions necessary for our formal results and specify settings such as the stateful case where they do not directly apply.

   Guidelines:

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

3. **Theory assumptions and proofs**

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

   Answer: [Yes]

Justification: Yes, all theorems have correct proofs. All lemmas used from other works are cited.

Guidelines:

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

4. **Experimental result reproducibility**

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

   Answer: [N/A]

   Justification: This is a theory paper with no experimental results.

   Guidelines:

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

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

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

Answer: [N/A]

Justification: This is a theory paper with no experimental results.

Guidelines:

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

6. **Experimental setting/details**

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

   Answer: [N/A]

   Justification: This is a theory paper with no experimental results.

   Guidelines:

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

7. **Experiment statistical significance**

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

   Answer: [N/A]

   Justification: This is a theory paper with no experimental results.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
   - The assumptions made should be given (e.g., Normally distributed errors).
   - It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

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

   Answer: [N/A]

   Justification: This is a theory paper with no experimental results.

   Guidelines:

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

9. **Code of ethics**

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

   Answer: [Yes]

   Justification: Yes, it abides by the code of ethics.

   Guidelines:

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

10. **Broader impacts**

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

    Answer: [N/A]

    Justification: This paper focuses on theoretical results, we see this as foundational work that can help guide practice, but there are no specific societal impacts that merit discussion.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

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

    Answer: [N/A]

    Justification: This is a theory paper that does not propose a new model.

    Guidelines:

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

12. **Licenses for existing assets**

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

    Answer: [N/A]

    Justification: This is a theory paper.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

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

    Answer: [N/A]

    Justification: The assets of this paper are theoretical.

    Guidelines:

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

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

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

    Answer: [N/A]

    Justification: This is a theory paper.

    Guidelines:

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

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

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

    Answer: [N/A]

    Justification: No human subjects research was involved.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

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

Answer: [N/A]

Justification: We used tools only for spell checking and editing.

Guidelines:

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