Alternative Microfoundations for Strategic Classification

Meena Jagadeesan Celestine Mendler-Dünner Moritz Hardt {mjagadeesan, mendler, hardt}@berkeley.edu

University of California, Berkeley

Abstract

When reasoning about strategic behavior in a machine learning context it is tempting to combine *standard microfoundations* of rational agents with the statistical decision theory underlying classification. In this work, we argue that a direct combination of these ingredients leads to brittle solution concepts of limited descriptive and prescriptive value. First, we show that rational agents with perfect information produce discontinuities in the aggregate response to a decision rule that we often do not observe empirically. Second, when any positive fraction of agents is not perfectly strategic, desirable stable points—where the classifier is optimal for the data it entails—no longer exist. Third, optimal decision rules under standard microfoundations maximize a measure of negative externality known as *social burden* within a broad class of possible assumptions about agent behavior.

Recognizing these limitations we explore alternatives to standard microfoundations for binary classification. We start by describing a set of desiderata that help navigate the space of possible assumptions about agent responses. In particular, we analyze a natural constraint on feature manipulations, and discuss properties that are sufficient to guarantee the robust existence of stable points and to upper bound social burden. Building on these insights, we then propose the *noisy response* model. Inspired by smoothed analysis and empirical observations, noisy response incorporates imperfection in the agent responses, which we show mitigates the limitations of standard microfoundations. Our model retains analytical tractability, leads to more robust insights about stable points, and imposes a lower social burden at optimality.

1 Introduction

Consequential decisions compel individuals to react in response to the specifics of the decision rule. This individual-level response in aggregate can disrupt both statistical patterns and social facts that motivated the decision rule, leading to unforeseen consequences. A similar conundrum in the context of macroeconomic policy making fueled the *microfoundations* program following the influential critique of macroeconomics by Lucas in the 1970s [Lucas Jr, 1976]. Microfoundations refers to a vast theoretical project that aims to ground theories of aggregate outcomes and population forecasts in microeconomic assumptions about individual behavior. Oversimplifying a broad endeavor, the hope was that if economic policy were *microfounded*, it would anticipate more accurately the response that the policy induces.

Predominant in neoclassical economic theory is the assumption of an agent that exhaustively maximizes a utility function on the basis of perfectly accurate information. This modeling assumption about agent behavior underwrites many celebrated results on markets, mechanisms, and games. Although called into question by behavioral economics and related fields (e.g. see [Camerer

et al., 2004]), the assumption remains central to economic theory and has become standard in computer science, as well.

When reasoning about incentives and strategic behavior in the context of classification tasks, it is tempting to combine the predominant modeling assumptions from microeconomic theory with the statistical decision theory underlying classification. In the resulting model, agents have perfect information about the decision rule and compute the best response according to their utility function with the goal to achieve a more favorable classification outcome. We refer to this agent model as *standard microfoundations*. Building on the assumption that agents follow standard microfoundations, the decision maker then chooses the decision rule that maximizes their own optimization objective in anticipation of the resulting agent response. This is the conceptual route taken in the area of *strategic classification*, but similar observations may apply more broadly to the intersection of economics and learning.

1.1 Our work

We argue that standard microfoundations are a poor basis for studying strategic behavior in binary classification problems. We make this point through three observations that illustrate the limited descriptive power of the standard model, and the problematic solution concepts it implies. In response, we explore the space of alternative agent models for strategic classification, and we identify desirable properties that when satisfied by a model of microfoundations lead to more realistic and robust insights. Guided by these desiderata, we propose *noisy response* as a promising alternative to the standard model.

A) Limitations of standard microfoundations

In strategic classification, agents respond strategically to the deployment of a binary decision rule f_{θ} specified by model parameters θ . The assumption of the decision maker is that agents follow the standard microfoundations: they have full information about f_{θ} and change their features so as to maximize their utility function. The utility function captures the benefit of a positive classification outcome, as well as the cost of feature change. Consequently, an agent does not invest in changing its features if the cost of change exceeds the benefit of positive classification.

We show that these standard microfoundations are not descriptive of reality, and undesirable under natural model misspecifications of agent responses.

Our first observation concerns the *aggregate response*—the distribution $\mathcal{D}(\theta)$ over feature, label pairs induced by a classifier f_{θ} . We show that in the standard model, the aggregate response necessarily exhibits discontinuities that we often do not observe in empirical settings. The problem persists even if we assume an approximate best response and allow for heterogeneous cost functions.

Our second observation reveals that, apart from lacking descriptive power, the standard model also leads to brittle conclusions about the solution concept of *performative stability*. Performative stability refers to decision rules that are optimal on the particular distribution they entail and thus represent fixed points of retraining methods. We show that the existence of performatively stable classifiers breaks down when *any* fraction of agents in the population are non-strategic. This brittleness of the existence of fixed points suggests that the standard model does not constitute a reliable model for investigating dynamics of retraining algorithms.

Our last observation concerns the solution concept of *performative optimality*. Performative optimality expresses a global property of a decision rule and is achieved by a classifier that overall exhibits the highest accuracy on the distribution it induces. Such a classifier is not necessarily performatively stable, and finding performatively optimal points requires the decision maker to

understand and anticipate strategic feedback effects. We prove that relying on standard microfoundations to model strategic behavior leads to extreme decision rules that maximize a measure of negative externality called *social burden* within a broad class of alternative models. Social burden, proposed in recent work, quantifies the expected cost that positive instances of a classification problem have to incur in order to be accepted. Thus, given natural model uncertainty, standard microfoundations produce optimal solutions that are least favorable in terms of social burden.

B) Alternative microfoundations

Recognizing the limitations of standard microfoundations, we systematically explore alternatives to the standard model. We investigate a broad space of alternative assumptions on agent responses, encompassing general agent behavior that need not be fully informed, strategic, or utility maximizing. We formalize microfoundations as a randomized map $M: X \times Y \to T$ that assigns each agent to a response type $t \in T$ that is associated with a response function $\mathcal{R}_t \colon X \times \Theta \to X$ specifying how agents of type t change their features t in response to a decision rule given by the parameters t.

Letting \mathcal{D}_{XY} be the base distribution over features and labels before any strategic adaptation, the *aggregate response* to a classifier f_{θ} is given by the distribution $\mathcal{D}(\theta; M)$ over induced feature, label pairs $(\mathcal{R}_t(x,\theta),y)$ for a random draw $(x,y) \sim \mathcal{D}_{XY}$, and t=M(x,y). In this sense, the mapping M *microfounds* the distribution induced by a classifier f_{θ} , endowing it with structure that allows the decision maker to deduce the aggregate response from a model of individual behavior.

To guide our search for more appropriate microfoundations for binary classification, we introduce a collection of properties that are desirable for a model of agent responses to satisfy. The first condition, that we call *aggregate smoothness* rules out the discontinuities arising from standard microfoundations. Conceptually, it requires that varying the model parameters slightly must change the aggregate response smoothly. We find that this property alone is sufficient to guarantee the robust existence of stable points under mixtures with non-strategic agents.

The second condition, that we call the *expenditure constraint*, helps ensure that the model encodes realistic agent-level responses \mathcal{R}_t . At the high level it requires that agents do not spend more on gaming than the utility of a positive outcome. This natural constraint gives raise to a large set of potential models, for which we can show that, under weak assumptions, the social burden of the optimal classifier is no larger than the social burden of the optimal classifier deduced from the standard model. Moreover, under the same assumptions the optimal points are determined by local behavior. This frees the decision maker from fully understanding the aggregate response $\mathcal{D}(\theta)$ and makes the task of finding an approximately optimal classifier more tractable.

C) Noisy response—an alternative model

Using the properties described above as a compass through the space of alternative models, we identify *noisy response* as a compelling model of microfoundations that complies with the aforementioned desiderata. In this model, each agent best responds with respect to $\theta + \xi$, where ξ is an independent sample from a zero mean noise distribution. This model is inspired by smoothed analysis and encodes imperfection in the population's response to a classifier by perturbing individual agent's manipulation target.

We show that noisy response satisfies a number of desirable properties that make it a promising model of microfoundations for classification in strategic settings, both for theoretical analyses and from a practical standpoint. First, noisy response satisfies aggregate smoothness, and thus leads to the robust existence of stable points. Moreover, the model satisfies the expenditure constraint, and thus encodes natural agent-level responses which can be used to reason about metrics such

as social burden. When used to anticipate strategic feedback effects and compute optimal points, noisy response leads to strictly less pessimistic acceptance thresholds than those computed under standard microfoundations, given the same constraints on manipulation expenditure. In fact, we show via simulations that larger variance σ of the noise in the manipulation target leads to more conservative optimal thresholds, and for $\sigma \to 0$, we approximate the extreme case of standard microfoundations. Finally, from a practical perspective, we demonstrate that noisy response enjoys the desirable property that the distribution map $\mathcal{D}(\theta)$ can be estimated from individual experiments alone, without ever deploying a classifier.

1.2 Related work

Existing work on strategic classification in machine learning has mostly followed standard microfoundations for modeling agent behavior in response to a decision rule, e.g., [Dalvi et al., 2004; Brückner and Scheffer, 2011; Hardt et al., 2016a; Khajehnejad et al.; Tsirtsis and Gomez-Rodriguez, 2020] to name a few. This includes works that focus on minimizing Stackelberg regret [Dong et al., 2018; Chen et al., 2020], quantify the price of transparency [Akyol et al., 2016], and investigate the benefits of randomization in the decision rule [Braverman and Garg, 2020]. Also investigations on externalities such as social cost [Milli et al., 2019; Hu et al., 2019] and whether classifiers incentivize improvement as opposed to gaming [Kleinberg and Raghavan, 2019; Miller et al., 2020; Shavit et al., 2020; Haghtalab et al., 2020] have mostly built on the standard assumption of best-responding agents with perfect information. Recent work by Levanon and Rosenfeld [2021] studied practical implications of standard microfoundations and how to make the resulting optimization problem more amenable for optimization.

A handful of works have suggested potential limitations of the standard strategic classification framework. Brückner et al. [2012] recognized that the standard model leads to very conservative Stackelberg solutions, and proposed to resort to Nash equilibria as an alternative solution concept. We instead take a different route and advocate for rethinking standard microfoundations that lead to these conservative acceptance thresholds altogether. Concurrent and independent work questioned the perfect information assumption in the standard model and studied strategic classification when the classifier is not fully revealed to the agents. In particular, Ghalme et al. [2021] find that hiding information about the classifier harms the decision-maker's accuracy, and Bechavod et al. [2021] emphasize disparate impact on subgroups, when agents can only obtain additional side-information from individuals within their communities. While we also question the assumptions of the standard model, we argue that the agents often do not perfectly respond to the classifier even when the decision rule is fully transparent. Therefore, we advocate for incorporating imperfection into the model of microfoundations in order to anticipate natural deviations from the standard model.

Related work in economics also investigates strategic responses to decision rules. This line of work, initiated by Spence [1973], has shown that information about individuals can become *muddled* as a result of heterogeneous gaming behavior [Frankel and Kartik, 2019], investigated the role of commitment power of the decision maker [Frankel and Kartik, 2020], considered the impact of an intermediary who aggregates the agents' multi-dimensional features [Ball, 2020], and considered the performance of different training approaches in strategic environments Hennessy and Goodhart [2020]. A notable work by Björkegren et al. [2020] investigates strategic behavior through a field experiment in the micro-lending domain, with a focus on evaluating approaches for designing strategy-robust classifiers. An important distinction is that these works tend to study regression, while we focus on classification. These settings appear to be qualitatively different in the context of strategic feedback effects (e.g. see note in [Hennessy and Goodhart, 2020]).

Our work is conceptually related to recent work in economics that has recognized mismatches between the predictions of standard models and empirical reality, for example in macroeconomic policy [Stiglitz, 2018; Kaplan and Violante, 2018; Coibion et al., 2018] and in mechanism design [Li, 2017]. These works, and many others, have explored incorporating richer behavioral and informational assumptions into the standard model in economic settings. Although our work also explores alternatives to standard microfoundations, we focus on algorithmic decision-making, where the limitations of the standard model had not been previously identified. We believe that our approach of navigating the entire space of potential models using a collection of properties could be of broader interest when developing alternative microfoundations.

1.3 Setup and basic notation

Let $X \subseteq \mathbb{R}^m$ denote the feature space, and let $Y = \{0,1\}$ be the space of binary outcomes. The base distribution \mathcal{D}_{XY} is a joint distribution over $X \times Y$ describing the population prior to any strategic adaption. Throughout the paper we assume that \mathcal{D}_{XY} is continuous and has zero mass on the boundary of X. We focus on binary classification where the model $f_{\theta}: X \to \{0,1\}$ is parameterized by $\theta \in \mathbb{R}^d$, and the decision-maker selects classifier weights θ from $\Theta \subseteq \mathbb{R}^d$ which is a compact, convex set. We assume that for every $\theta \in \Theta$, the set $\{x \in X \mid f_{\theta}(x) = 1\}$ is closed, and the decision boundary is measure 0. We adopt the notion of a distribution map $\mathcal{D}(\theta)$ from [Perdomo et al., 2020] to describe the distribution over $X \times Y$ induced by strategic manipulation of agents in the base distribution in response to the classifier f_{θ} .

2 Limitations of standard microfoundations

In the strategic classification literature, the typical agent model is a *rational agent with perfect information*. At the core of this model lies the assumption that agents have perfect knowledge of the classifier and maximize their utility given the classifier weights. The utility composes of two terms: a reward for obtaining a positive classification, and a cost of manipulating features. The reward is denoted $\gamma > 0$ and the manipulation cost is represented by a function $c: X \times X \to \mathbb{R}$ where c(x,x') reflects how much agents need to expend to change their features from x to x'. A valid cost function satisfies a natural monotonicity requirement as stated in Assumption 1. Given a feature vector x and a classifier f_{θ} , agents solve the following utility maximization problem:

$$\underset{x' \in X}{\operatorname{arg\,max}} \left[\gamma f_{\theta}(x') - c(x, x') \right]. \tag{1}$$

Assumption 1. A cost function $c: X \times X \to \mathbb{R}$ is *valid*, if it is continuous in both arguments, it holds that c(x, x') = 0 for x = x', and c increases with distance in the sense that $c(x, \bar{x}) < c(x, x')$ and $c(\bar{x}, x) < c(x', x)$ for every $\bar{x} \in X$ that lies on the line segment connecting the two points $x, x' \in X$.

We will refer to this response model as the *standard microfoundations*.

2.1 Discontinuities in the aggregate response

A striking property of distribution shifts induced by standard microfoundations in response to a classifier f_{θ} is that they are necessarily either trivial or discontinuous. The underlying cause is that

¹We model non-zero cost for all modifications to features, regardless of whether they result in beneficial classification or not. Generalizing beyond standard microfoundations this accounts for how agents may erroneously expend effort on changing their features in an incorrect direction, as empirically demonstrated by Björkegren et al. [2020].

agents behaving according to standard microfoundations either change their features exactly up to the decision boundary, or they do not change their features at all.

Proposition 1. Given a base distribution \mathcal{D}_{XY} , let $\mathcal{D}(\theta)$ be the distribution induced by a classifier f_{θ} . Then, if $\mathcal{D}(\theta)$ is continuous and $\mathcal{D}(\theta) \neq \mathcal{D}_{XY}$, there does not exist a valid cost function c such that $\mathcal{D}(\theta)$ is an aggregate of agents following standard microfoundations.

In addition to point masses, standard microfoundations face additional degeneracies in the aggregate response. Namely, a similar argument shows that any non-trivial distribution arising from the standard model must have a region of zero density below the decision boundary. These properties are highly unnatural in practice. To illustrate this, consider banking lending decisions and the corresponding distribution over credit scores in Example 1. If lenders decisions are based on the FICO score, then under standard microfoundations the distribution over credit scores should exhibit a discontinuity at the threshold. However this is not what we observe empirically.²

Example 1 (Lending decision). Bank's lending decisions typically use FICO credit scores as an indicator of creditworthiness. Previous work [Hardt et al., 2016b] studied a FICO dataset from 2003, where credit scores range from 300 to 850, and a cutoff of 620 was commonly used for prime-rate loans. The observed distribution over credit scores appears continuous and is supported across the full range of scores.

It is important to note that the degeneracies of standard microfoundations arise from the fact that classification decisions are binary and based on a *hard* decision. Agents who are not classified positively receive no reward: it does not matter how close to the decision boundary the agent is. This discontinuity in the utility is specific to classification and does not arise in regression problems that are predominantly studied in the economics literature. However, in machine learning and statistical decision theory, binary classification is ubiquitous, and the degeneracies that we have identified are relevant for modeling choices in a wide variety of applications.

The reader might imagine that common variations and generalizations of standard microfoundations can mitigate these issues. Unfortunately, the two variations of standard microfoundations that are typically considered—heterogeneous cost functions [Hu et al., 2019], and approximate best response [Miller et al., 2020]—result in similar degeneracies. Heterogeneity in the cost (or utility) function can only change whether or not an agent decides to change their features, but it does not change their target of manipulation. If agents approximately best-respond, and thus move to features x' that maximize their utility to within a constant factor approximation of the optimum, the model no longer leads to point masses at the decision boundary, but agents will never undershoot the decision boundary. This means that any nontrivial aggregate distribution must have a region of zero density below the decision boundary to comply with standard microfoundations and any of these variants.

In fact, agent behavior that is not captured by standard microfoundations or its variants has been directly observed in field experiments. In particular, agents have been observed to both overshoot and undershoot the decision boundary as well as generally exhibit noisy responses, even if the classifier is fully transparent.

Example 2 (Field Experiment [Björkegren et al., 2020]). The authors developed an app that mimicked aspects of "digital credit" applications, and deployed it in Kenya in order to empirically investigate strategic behavior. Participants were rewarded if the app guessed that they were a

²There could be a number of reasons why we do not observe discontinuities in practice: different lenders having different cutoffs, interest rates being non-binary, and the use of features beyond credit scores. In any case, the classical model for strategic classification does not describe this setting accurately.

high-income earner. When the participants were given access to the coefficients of the decision rule, they tended to change their features in the right direction, but a high variance in their responses was observed—see Table 5 in their work. The noise in the response was even more pronounced when participants were only given opaque access to the decision rule. In this case, agents often did not even change their features in the right direction.

In the following section we demonstrate that the model of standard microfoundations is not only lacking in descriptive power for strategic behavior in the context of binary classification, but also exhibits unfavorable properties under natural deviations from the modeling assumptions.

2.2 Brittleness under natural model misspecifications

We describe two scenarios, where the behavior of standard microfoundations is highly undesirable under natural model misspecifications. In particular, we show that the existence of stable points breaks down when only a small fraction of agents are non-strategic, and the optimal solutions associated with the standard model cause unnecessarily high social burden.

A) Stability as a fragile solution concept

Our first result demonstrates that the standard model does not lead to robust insights about the solution concept of performative stability. In particular, performatively stable points that are guaranteed to exist under standard microfoundations, as proven by Milli et al. [2019], no longer exist if *any* fraction of individuals is non-strategic. Since stable points represent fixed points of retraining, this suggests that repeated risk minimization (RRM) methods that provably converge under standard microfoundations diverge in settings where only a tiny fraction of agents do not comply with the model.

For our analysis, we consider a local relaxation of the notion of performative stability that corresponds to fixed points of *repeated gradient descent* [Perdomo et al., 2020]. We say θ_{PS} is *locally stable* if θ_{PS} is a local minima or a stationary point of the following optimization problem:

$$\min_{\theta \in \Theta} \mathbb{E}_{(x,y) \sim \mathcal{D}(\theta_{PS})} \mathbb{1} \{ y \neq f_{\theta}(x) \}. \tag{2}$$

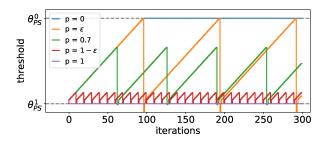
Local stability is closely related to the concept of a pure strategy (local) *Nash equilibrium* in a simultaneous game between the strategic agents that respond to the classifier f_{θ} and the decision maker who responds to the observed distribution $\mathcal{D}(\theta)$.

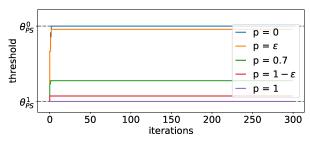
To showcase that the existence of locally stable points (or Nash equilibria) under standard microfoundations is very sensitive to deviations in agent behavior, we focus on the following simple 1-dimensional setting.

Setup 1 (1-dimensional). Let $X \subseteq \mathbb{R}$ and consider a threshold functions $f_{\theta}(\cdot) = \mathbb{1}\{\cdot \geq \theta\}$ with $\theta \in \Theta \subseteq \mathbb{R}$. Let $\mu(x)$ be the conditional probability over \mathcal{D}_{XY} of the true label being 1 given features x. Suppose that $\mu(x)$ is strictly increasing in x and there is an $\theta \in \text{Int}(\Theta)$ such that $\mu(\theta) = 0.5$.

Proposition 2. Consider Setup 1. Suppose that a p fraction of agents drawn from \mathcal{D}_{XY} do not ever change their features, and a 1-p fraction of agents drawn independently from \mathcal{D}_{XY} follow standard microfoundations with a valid cost function c. Then, we have the following properties:

- a) For $p \in \{0, 1\}$, locally stable points exist.
- b) For $p \in (0,1)$, locally stable points do not exist.





- (a) SM mixed with non-strategic agents
- (b) NR mixed with non-strategic agents

Figure 1: Convergence of retraining algorithm in a 1d-setting for different values of p with $\epsilon = 10^{-2}$. The population consists of 10^5 individuals. Half of the individuals are sampled from $x \sim \mathcal{N}(1,0.33)$ with true label 1 and the other half is sampled from $x \sim \mathcal{N}(0,0.33)$ with true label 0. The parameter of the noisy responses (NR) in (b) is taken to be $\sigma^2 = 0.1$. θ_{PS}^0 and θ_{PS}^1 are defined as in Proposition 2 for standard microfoundations (and similarly for noisy response).

c) For $p \in (0,1)$, RRM will oscillate as follows: Let θ_{PS}^q be the smallest locally performatively stable point for $q \in \{0,1\}$. Then, RRM oscillates between θ_{PS}^1 and a threshold $\tau(p) \in (\theta_{PS}^1, \theta_{PS}^0)$, where $\tau(p)$ is decreasing in p, approaching θ_{PS}^1 as $p \to 1$ and approaching θ_{PS}^0 as $p \to 0$.

Proof Sketch. For p=0 stability is trivial and for p=1 the claim follows from an argument similar to Lemma 3.2 in [Milli et al., 2019]. For $p \in (0,1)$, the core observation is that for any θ the distribution $\mathcal{D}(\theta)$ contains no *strategic* agents in the interval $\mathsf{Gap}(\theta) := [\theta - \Delta, \theta]$ for some $\Delta > 0$. Furthermore, for any $\theta > \theta_{PS}^1$ the misclassification rate on *non-strategic* agents could be improved by reducing the threshold to θ_{PS}^1 . Thus, it is not hard to see that or any $\theta > \theta_{PS}^1$, the threshold $\mathsf{max}(\theta - \Delta, \theta_{PS}^1)$ achieves smaller loss than θ , and thus θ can not be stable. We formalize this argument and the case for $\theta \leq \theta_{PS}^1$ in Appendix B.2. □

Proposition 2 implies that not only does the existence of local Nash equilibria break down if any fraction $p \in (0,1)$ of agents are non-strategic, but also repeated risk minimization oscillates between two extreme points. To illustrate this, we have implemented a simple instantiation of Setup 1 and visualize the trajectories of RRM for different values of p in Figure 1(a). The main insight is that retraining methods start oscillating substantially even when p is very close to 0 (only an ϵ fraction of agents are not following standard microfoundations). This sensitivity of the trajectory to natural deviations from the modeling assumptions suggests that standard microfoundations do not constitute a reliable model to study algorithm dynamics.

B) Maximal negative externalities at optimality

Our next result shows that it is undesirable for the decision-maker to rely on standard microfoundations to compute performatively optimal points when there is natural uncertainty about how exactly agents respond to a classifier. In particular, performatively optimal classifiers computed under standard microfoundations lead to highest negative externalities within a broad class of alternative models for agent responses.

Recall that a performatively optimal points correspond to the best classifier for the decision maker from a global perspective, but they are not necessarily stable under retraining. Formally, a classifier θ_{PO} is *performatively optimal* [Perdomo et al., 2020] if it minimizes the performative risk:

$$\theta_{PO} := \operatorname{argmin}_{\theta \in \Theta} \mathbb{E}_{(x,y) \sim \mathcal{D}(\theta)} \mathbb{1} \{ y \neq f_{\theta}(x) \}. \tag{3}$$

Performative optimality is closely related to the concept of a *Stackelberg equilibrium* in a leader-follower game. The key challenge of computing such a Stackelberg equilibrium, or performative optima, is that optimizing (3) requires the decision maker to anticipate the population's response $\mathcal{D}(\theta)$ to any classifier f_{θ} . A natural approach to model this response is to build on microfoundations and deduce properties of the distribution map from individual agent behavior. Different models for agent behavior can lead to solutions with qualitatively different properties.

While the decision-maker is unlikely to have a fully specified model for agent behavior at hand, we outline a few natural criteria that agent responses could reasonably satisfy. To formalize these criteria, we again focus on the 1-dimensional setting in Setup 1.

Property 1 (Expenditure monotonicity). For every agent with true features x, the manipulated features $x' = \mathcal{R}(x;\theta)$ in response to a classifier θ satisfy the following properties:

- a) $c(x, \mathcal{R}(x; \theta)) < \gamma$ for every $\theta \in \Theta$.
- b) if $f_{\theta}(\mathcal{R}(x;\theta)) = 1$, then $f_{\theta'}(\mathcal{R}(x;\theta')) = 1$ for all $\theta' \leq \theta$.

Property 1 describes agents that act rationally in the sense that a) they do not expend more on gaming than their utility from a positive outcome, and b) an agent's outcome does not worsen if the threshold is lowered. However, agents complying with Property 1 do not necessarily behave according to standard microfoundations, and instead may be imperfect and overestimate or underestimate the consequences of a feature manipulation.

We now show that within the broad class of microfoundations that exhibit Property 1, the standard model leads to an extreme acceptance threshold. For the formal statement, see Appendix B.3.

Proposition 3 (Informal). Consider Setup 1. Let \mathscr{D} be the class of distribution maps $\mathcal{D}: \Theta \to X \times Y$ that can be represented by a population of agents who all satisfy Property 1. Then under mild assumptions, for every $\mathcal{D} \in \mathscr{D}$, it holds that

$$\theta_{PO}(\mathcal{D}_{SM}) \ge \theta_{PO}(\mathcal{D})$$

where $\theta_{PO}(\mathcal{D})$ denotes the performative optima associated with the distribution map \mathcal{D} , and \mathcal{D}_{SM} is the distribution map induced by standard microfoundations.

A problematic implication of Proposition 3 is that standard microfoundations also maximize the negative externality called *social burden* [Miller et al., 2020]:

$$\mathsf{Burden}(\theta) := \mathbb{E}_{(x,y) \in \mathcal{D}_{XY}}[\min\left\{c(x,x') \mid f_{\theta}(x') = 1\right\} \mid y = 1].$$

Social burden quantifies the average cost that a positively labeled agent has to expend in order to be positively classified by f_{θ} . While previous work introduced and studied social burden within standard microfoundations, and showed that Nash equilibria lead to smaller social burden than Stackelberg equilibria, we use it to study implications of different modeling assumptions on agent behavior. In particular, the following corollary demonstrates that standard microfoundations lead to worst possible social burden across all microfoundations that satisfy Property 1.

Corollary 4. Under the same assumptions as Proposition 3, for every $\mathcal{D} \in \mathcal{D}$, it holds that

$$Burden(\theta_{PO}(\mathcal{D}_{SM})) \geq Burden(\theta_{PO}(\mathcal{D})).$$

where $\theta_{PO}(\mathcal{D})$ denotes the performative optima associated with the distribution map \mathcal{D} , and \mathcal{D}_{SM} is the distribution map induced by standard microfoundations.

This result indicates that under natural modeling uncertainty it is hard for the decision-maker to justify the use of standard microfoundations. Implicit in our argument is the following moral stance: given a set of criteria for what defines a plausible model for microfoundations, the decision-maker should not select the one that maximizes negative externalities. In our setup, the set of criteria is given by agents satisfying expenditure monotonicity, and we showed that standard microfoundations maximize negative externalities within all models satisfying this property. What this means is that when agents deviate from the standard model, computing the performative optimal point using standard microfoundations results in suboptimal performative risk and unnecessarily high social burden, relative to the performative optimal point of the true agent responses.

3 Alternative microfoundations

The previous section revealed several issues with using standard microfoundations to analyze strategic feedback effects. In this section, we depart from this classical approach and search for models that are more appropriate for binary classification. First, we define an exhaustive space of alternative microfoundations. Then, we collect a set of useful properties that we show are desirable for a model of microfoundations to satisfy. These properties serve as a "compass" to guide our search for an alternative agent model for strategic classification.

3.1 Defining the space of alternatives

The principle behind microfoundations for strategic classification is to equip the distribution map with structure by viewing the distribution shifts induced by a decision rule as an aggregate of the responses of individual agents. We consider a space of alternative microfoundations that capture agent responses in full generality. To do so, we introduce a family of response types \mathcal{T} that represents the space of all possible ways that agents can perceive and react to the classifier f_{θ} . Each agent is associated with a response type $t \in \mathcal{T}$, true features $x \in X$, and a true label $y \in Y$. The response type fully determines an agent's behavior through the *agent response function* $\mathcal{R}_t: X \times \Theta \to X$, where $x' = \mathcal{R}_t(x, \theta)$ is the feature vector to which an agent with true features x and response type t games when the classifier t0 is deployed.

Remark. Using the language of agent response functions, non-strategic agents correspond to a response type $t_{\rm NS}$ such that $\mathcal{R}_{t_{\rm NS}}(x,\theta)=x$ for all $\theta\in\Theta$, and standard microfoundations correspond to a response type $t_{\rm SM}$ where $\mathcal{R}_{t_{\rm SM}}(x,\theta)$ is given by (1) for all $\theta\in\Theta$. Note that a population of agents could by heterogeneous and exhibit a mixture of different types, or even be described by a *continuum* of response types.

Within this framework, we formalize microfoundations through a *mapping* $M: X \times Y \to \mathcal{T}$ from agents to response types. We denote the set of possible mappings M by the collection \mathcal{M} that consists of all³ possible randomized functions $X \times Y \to \mathcal{T}$. Conceptually, the mapping $M \in \mathcal{M}$ sets up the rules of agent behavior, and the fact that these rules are independent of the deployed classifier, makes our framework immune to Lucas's critique [Lucas Jr, 1976]. One aspect that distinguishes our framework from typical approaches to microfoundations in the economics literature is that it directly specifies agent responses, rather than specifying an underlying behavioral mechanism. An advantage of this approach is that responses can be observed, whereas the behavioral mechanism is harder to infer.

³These mappings are subject to weak measurability constraints that we describe in detail in Appendix A.2

Importantly, the mapping M coupled with the base distribution \mathcal{D}_{XY} provides all the necessary information to specify the population's response to a classifier f_{θ} and reason about metric such as the performative risk or social burden. In particular, for each $\theta \in \Theta$, the *aggregate response* $\mathcal{D}(\theta;M)$ is the distribution over $(\mathcal{R}_t(x,\theta),y)$ where $(x,y) \sim \mathcal{D}_{XY}$ and t = M(x,y). We use the notation $\mathcal{D}(\cdot;M):\Theta \to \Delta(X\times Y)$ to denote the aggregate response map induced by M.

Naturally, with such a flexible model, *any* distribution map can be microfounded, albeit with complex response types, as long as feature manipulations do not change the fraction of positively labeled agents in the population. We refer to Appendix C.1 for an explicit construction of M.

Proposition 5. Let \mathcal{D}_{XY} be a non-atomic distribution. Let $\mathcal{D}(\theta)$ be any distribution map that preserves the marginal distribution over Y of \mathcal{D}_{XY} . Then, there exists a $M \in \mathcal{M}$ such that $\mathcal{D}(\cdot;M)$ is equal to $\mathcal{D}(\cdot)$.

This result primarily serves as a proof of existence that our general framework for microfoundations can capture any aggregate distribution—and in particular, continuous distributions that are observed empirically (e.g. Example 1). In the following subsections, we specify two properties that we believe a natural model of microfoundations should satisfy.

3.2 Aggregate smoothness

The first property we describe is an aggregate-level property pertaining to the induced distribution and its interactions with the model class. It rules out unnatural discontinuities in the distribution map. We call this property aggregate smoothness, and formalize it in terms of the decoupled performative risk [Perdomo et al., 2020].

Property 2 (Aggregate smoothness). Define the decoupled performative risk induced by M to be $\mathsf{DPR}_M(\theta, \theta') := \mathbb{E}_{(x,y) \sim \mathcal{D}(\theta;M)}[\ \mathbb{I}\{y \neq f_{\theta'}(x)\}]$. For a given base distribution \mathcal{D}_{XY} , a mapping M satisfies *aggregate smoothness* if the derivative of the decoupled performative risk with respect to θ' exists and is continuous in θ and θ' across all of Θ .

Intuitively, the existence of the partial derivative of $DPR_M(\theta, \theta')$ with respect to θ' guarantees that each distribution $\mathcal{D}(\theta; M)$ is sufficiently continuous (and cannot have a point mass at the decision boundary), and assuming continuity of the derivative we guarantee that $\mathcal{D}(\theta; M)$ changes continuously in θ . This connection can be made explicit in the case of 1-dimensional features:

Proposition 6. Suppose that $X \subseteq \mathbb{R}$, and let $\Theta \subseteq \mathbb{R}$ be a model class of threshold functions. Then, if the distribution map $\mathcal{D}(\cdot; M)$ has the following properties, M satisfies aggregate smoothness w.r.t. Θ :

- 1. For each θ , the probability density $p_{\theta}(x,y)$ of $\mathcal{D}(\theta;M)$ exists everywhere and is continuous in x.
- 2. For each x, y, the probability density $p_{\theta}(x,y)$ is continuous in θ .

We believe that these two continuity properties are natural and likely to hold in practice, given the empirical evidence in Example 1 and Example 2. A striking consequence of aggregate smoothness is that it is sufficient to guarantee the existence of locally stable points, using Brouwer's fixed point theorem.

Theorem 7. Given a base distribution \mathcal{D}_{XY} and model class Θ , for any M that satisfies aggregate smoothness, there exists a locally stable point.

In fact, this result actually implies that stable points exist even under deviations from the model, as long as aggregate smoothness is preserved. To demonstrate that this holds for mixtures with non-strategic agents, we use the fact that aggregate smoothness is closed under mixtures, and a population of non-strategic agents satisfies aggregate smoothness under weak assumptions on

the base distribution. For ease of notation we formalize mixtures with non-strategic agents through the operator $\Phi_p(M)$, where for $p \in [0,1]$, we let $\Phi_p(M(x,y))$ be equal to t_{NS} with probability p and equal to t_{NS} otherwise. Our next result proves that aggregate smoothness of t_{NS} guarantees the *robust* existence of stable points for any t_{NS} .

Corollary 8. Suppose that the non-performative risk $R(\theta) := \mathbb{E}_{(x,y) \in \mathcal{D}_{XY}} \mathbb{1}\{f_{\theta}(x) = y\}$ is continuously differentiable $\forall \theta \in \Theta$. Then, for any M that satisfies aggregate smoothness and any $p \in [0,1]$, there exists a locally stable point for $\Phi_p(M)$.

Conceptually, our investigations in this section have been inspired by the line of work on performative prediction [Perdomo et al., 2020] that demonstrated that regularity assumptions on the aggregate response alone can be sufficient to guarantee the existence of stable points for smooth, strongly convex loss functions. Our results differ from these previous analyses of performative stability in that we instead focus on the 0-1 loss function. In Appendix E, we provide a discussion of why the Lipschitzness assumptions on the distribution map used in prior work are not sufficient to guarantee the existence of stable points in our classification setting.

3.3 Constraint on manipulation expenditure

While aggregate smoothness focused on the population-level properties of the induced distribution, a model for microfoundations must also be descriptive of realistic agent-level responses in order to yield useful qualitative insights about metrics such as social burden or accuracy on subgroups. A minimal assumption on agent responses is that an agent never expends more on manipulation than the utility of a positive outcome.

Property 3 (Expenditure constraint). Given a model class Θ and a cost function c, a mapping $M \in \mathcal{M}$ is *expenditure-constrained* if $c(\mathcal{R}_t(x,\theta),x) \leq \gamma$ for every $\theta \in \Theta$ and every $t \in \text{Image}(M)$.

This constraint is implicitly encoded in standard microfoundations and many of its variants. Furthermore, if c is a valid cost function, then this property, together with a basic monotonicity requirement on agent's feature manipulations, implies *expenditure monotonicity*. As discussed in Section 2.2, expenditure monotonicity is satisfied by a large set of microfoundations models among which the standard model achieves extreme social burden at optimality. In Section 4 we will focus on one particular model for microfoundations within this set which results in a *strictly* lower social burden than the standard model.

Reducing empirical burden. Apart from defining a natural class of feasible microfoundations models, an additional advantage of Property 3 is that it naturally constrains each agent's range of manipulations. This can significantly reduce empirical burden for the decision maker who wants to estimate the distribution map to compute a strategy robust classifier offline.

Assume the decision maker follows a two-stage estimation procedure to estimate a performatively optimal point, similar to [Miller et al., 2021]. First, they compute an estimate \tilde{M} of the true mapping M and infer $\mathcal{D}(\cdot;\tilde{M})$ from the base distribution \mathcal{D}_{XY} . Second, they assume the model reflects the true decision dynamics and approximate optimal points as follows:

$$\theta_{PO}(\tilde{M}) := \operatorname{argmin}_{\theta \in \Theta} \mathbb{E}_{(x,y) \sim \mathcal{D}(\theta; \tilde{M})} [\mathbb{1}\{y \neq f_{\theta}(x)\}]. \tag{4}$$

Using a naive bound (see Lemma D.2) it is not difficult to see that it suffices to compute an estimate \tilde{M} of M, such that $\sup_{\theta} \text{TV}(\mathcal{D}(\theta; \tilde{M}), \mathcal{D}(\theta; M)) \leq \xi$ to guarantee that $\text{PR}(\theta_{\text{PO}}(M)) - \text{PR}(\theta_{\text{PO}}(\tilde{M})) \leq 2\xi$.

However, achieving this level of accuracy fundamentally requires a full specification of the response types for every agent in the population.

The expenditure constraint helps to make this task more tractable, in that the decision-maker only needs to estimate a small fraction of agent responses to achieve the same bound on the suboptimality of the obtained performative risk. To formalize this, let's assume the decision maker can define a set $\Theta_0 \subseteq \Theta$ that contains the performatively optimal classifier $\theta_{PO}(M)$. Then, given the implied restriction in the search space in (4), the expenditure constraint enables us to restrict the set of covariates that are relevant for the optimization problem to

$$S(\Theta_0, c) := \bigcup_{\theta \in \Theta_0} \{ x \in X : \exists x' \in X : f_\theta(x') \neq f_\theta(x) \land c(x, x') \leq \gamma \}. \tag{5}$$

The salient part $S(\Theta_0, c) \subseteq X$ captures all agents who are sufficiently close to the decision boundary for some $\theta \in \Theta_0$ so they are able to cross it without expending more than γ units of cost. The subset $S(\Theta_0, c)$ can be entirely specified by the cost function c and can be much smaller than X. We explicitly construct a pruned search space Θ_0 and $S(\Theta_0, c)$ in an example setting in Appendix F.1.

We now describe the implications of constraining to the salient part for a 1-dimensional setting where $X \subseteq \mathbb{R}$ and f_{θ} is a threshold function.⁴ Let us define an *agent response oracle* that given x and θ , outputs a draw x' from the response distribution $(\mathcal{R}_t(x,\theta),y)$ where $(x,y) \sim \mathcal{D}_{XY}$. We show with few calls to the oracle, the decision-maker can build an sufficiently precise estimate of M.

Lemma 9. Let $X \subseteq \mathbb{R}$, let $\Theta \subseteq \mathbb{R}$ be the model class of threshold functions. Suppose that M satisfies the expenditure constraint, the distribution map $\mathcal{D}(\cdot;M)$ is 1-Lipschitz with respect to TV distance, and $\Theta_0 \subseteq \Theta : \theta_{PO}(M) \in \Theta_0$. We further assume that an agent's type does not depend on their label, i.e., M(x,0) = M(x,1) for all $x \in X$. Then, with $O\left(\zeta^2 \frac{\ln(1/\epsilon)}{2\epsilon^3}\right)$ calls to the agent response oracle, where $\zeta := \mathbb{P}_{\mathcal{D}_{XY}}[x \in S(\Theta_0,c)]$, the decision maker can create an estimate \tilde{M} so that:

$$PR(\theta_{PO}(\tilde{M})) \le PR(\theta_{PO}(M)) + \epsilon.$$

with probability 0.9.

The number of necessary calls to the response function oracle for estimating M decays with $\zeta := \mathbb{P}_{\mathcal{D}_{XY}}[x \in S(\Theta_0, c)]$. Without any assumption on agent actions, the value of ζ is equal to 1. However, when the decision-maker is able to constrain $S(\Theta_0, c)$ to a small part of the input space by relying on the manipulation constraints, domain knowledge, or stronger assumptions on agent behavior, ζ and thus the number of oracle calls can be reduced significantly.

The concept of a salient part bears resemblance to the approaches by Zhang and Conitzer [2021]; Zhang et al. [2021], which directly specify the set of feature changes that an agent may make, rather than implicitly specifying agent actions through a cost function. While these models assume that agents best-respond, our key finding is that constraining agent behavior alone can lessen the empirical burden on the decision-maker.

4 Microfoundations based on imperfect agents

Using the properties established in the previous section as a guide, we propose an alternate model for microfoundations that naturally allows agents to undershoot or overshoot the decision boundary, while complying with aggregate smoothness and expenditure rationality. Furthermore, we show that this model, called *noisy response*, leads to strictly smaller social burden than the standard model while retaining analytical tractability.

⁴Lemma 9 directly extends to *posterior threshold functions* [Milli et al., 2019]. We also present results regarding the expenditure constraint in general settings in Appendix F.2.

4.1 Noisy response

Noisy response captures the idea of an *imperfect agent* who does not perfectly best-respond to the classifier weights. This imperfection can arise from many different sources—including interpretability issues, imperfect control over manipulations, or opaque access to the classifier. Inspired by *smoothed analysis* [Spielman and Teng, 2009], we do not directly specify the source of imperfection but instead capture imperfection in an agnostic manner, by adding small random perturbations to the classifier weights targeted by the agents. Since smoothed analysis has been successful in explaining convergence properties of algorithms in practical (instead of worst case) situations, we similarly hope to better capture empirically observed strategic phenomenas.

More formally, we assume that an agent of type t perceives θ according to $\theta + t$, and given a valid cost function c responds to a classifier f_{θ} as

$$R_t(x,\theta) := \underset{x' \in X'}{\operatorname{arg\,max}} \left[\gamma \cdot f_{\theta+t}(x') - c(x,x') \right], \tag{6}$$

where $\gamma > 0$ denotes the utility of a positive classification outcome, and $X' \subseteq \mathbb{R}^d$ is a compact, convex set that contains X.⁵ To model noise that is independent across agents, we define the mapping $M_{\sigma}: X \times Y \to T$ as follows: For every $x \in X$ and $y \in Y$, the random variable $M_{\sigma}(x,y)$ is distributed according to the multivariate Gaussian distribution $\mathcal{N}(\mathbf{0},\Sigma)$ with mean zero and diagonal covariance matrix $\Sigma = \operatorname{diag}(\sigma^2,\ldots,\sigma^2)$. This model results in a distribution over perceived values of θ across the entire population that follows a Gaussian distribution centered at θ . The noise level σ reflects the degree of imperfection in the population.

Conceptually, our model of noisy response bears similarities to models of *incomplete information* [Harsanyi, 1968] that are standard in game theory (but that have not been traditionally considered in the strategic classification literature). However, a crucial difference is that we advocate for modeling agents actions as imperfect even if the classifier is fully transparent, because we believe that imperfection can also arise from other sources. Evidence for this was found in the empirical study of Björkegren et al. [2020] discussed in Example 2 where agents act imperfectly even when the classifier weights are fully revealed.

We want to emphasize that we instantiate imperfection by adding noise to the perceptions, instead of directly adding noise to the *responses*. While both approaches would mitigate the discontinuities in the aggregate distribution, the approach of adding noise directly to the responses results in a less natural model for agent behavior that violates the expenditure constraint.

4.2 Aggregate-level properties of noisy response

Intuitively, the noise in the manipulation target of noisy responses smooths out the discontinuities of standard microfoundations, eliminating the point mass at the decision boundary and region of zero density below the decision boundary. We show this explicitly in a 1-dimensional setting.

Proposition 10. Let $\Theta \subseteq \mathbb{R}$ be a model class of threshold functions, and suppose also that $X \subseteq \mathbb{R}$. For any $\sigma \in (0, \infty)$, the distribution map $\mathcal{D}(\cdot; M_{\sigma})$ satisfies the continuity properties in Proposition 6, and thus M_{σ} satisfies aggregate smoothness.

Remark. Proposition 10 implies that noisy response inherits the robust existence of stable points from Theorem 7. Furthermore, we illustrate in Figure 1(b) how noisy response mitigates the large oscillations of repeated retraining that we observed for standard microfoundations.

⁵We assume that c is defined on all of $X' \times X'$, and for every $x \in X$, it holds that $c(x, x') > \gamma$ for all x' that are on the boundary of X'.

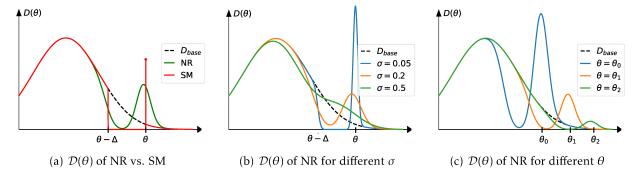


Figure 2: Probability density of the aggregate response $\mathcal{D}(\theta)$ in a 1d-setting, where the base distribution \mathcal{D}_X is a Gaussian with $x \sim \mathcal{N}(0, 0.5)$. We illustrate (a) $\mathcal{D}(\theta)$ for a population of agents that follow noisy response (NR) compared to standard microfoundations (SM), (b) how $\mathcal{D}(\theta)$ of NR changes for different θ , (c) variations in $\mathcal{D}(\theta)$ of NR for different values of σ .

To visualize the aggregate-level properties of noisy response and compare them to standard microfoundations we depict the respective density functions for a 1-dimensional setting with a Gaussian base distribution in Figure 2(a). We observe how the density around the threshold θ increases as a result of agents changing their features. The distribution $\mathcal{D}(\theta)$ can be bi-modal as agents closer to the threshold θ are more likely to change their features. The shape of the response distribution changes with σ as illustrated in Figure 2(b). As $\sigma \to 0$, the aggregate response of a population of noisy response agents approaches that of standard microfoundations. This means that noisy response maintains continuity while also being able to approximate the aggregate response of standard microfoundations to arbitrary accuracy. Finally, we note that the distribution map of noisy response changes *continuously* with θ , as visualized in Figure 2(c). In fact, the distribution map induced by noisy response is Lipschitz in total-variation distance, where the Lipschitz constant grows with $1/\sigma$.

Lemma 11. Given $\sigma \in (0, \infty)$, the distribution map $\mathcal{D}(\theta; M_{\sigma})$ is continuous in TV distance and supported on all of X. Moreover, the distribution map is Lipschitz in TV distance. That is, for any $\theta, \theta' \in \Theta$, we have that $\text{TV}(\mathcal{D}(\theta'; M_{\sigma}), \mathcal{D}(\theta; M_{\sigma})) \leq \frac{1}{\sqrt{2}\sigma} \|\theta - \theta'\|_2$.

This result displays a favorable property of noisy response compared to standard microfoundations, in that the performative risk changes smoothly with changes in the classifier weights.

Remark (Implications beyond 0-1 loss). Lemma 11 implies that noisy response induces a distribution map that is Lipschitz in Wasserstein distance *for any cost function*, where the Lipschitz constant depends on the diameter of the set X'. For smooth and strongly convex loss functions, this readily implies convergence of repeated retraining [Perdomo et al., 2020; Mendler-Dünner et al., 2020].

4.3 Trade-off between imperfection and social burden

Apart from satisfying desirable aggregate-level properties, noisy response also satisfies the expenditure monotonicity requirement in Property 1 (see Proposition D.1 for a proof). By Corollary 4, this implies that in Setup 1 the social burden of the optimal classifier computed under noisy perception is no larger than that of standard microfoundations. That is, $\operatorname{Burden}(\theta_{PO}(M_\sigma)) \leq \operatorname{Burden}(\theta_{PO}(\mathcal{D}_{SM}))$. In certain cases, we can obtain a stronger result and show that the social burden of noisy response is *strictly* lower than that of standard microfoundations.

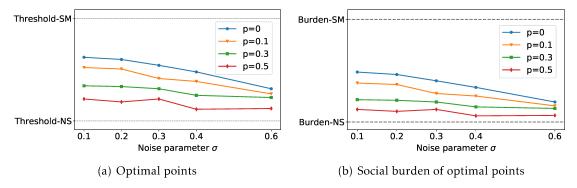


Figure 3: Optimal points in a 1d-setting for different values of σ (noise in the response model) and p (fraction of non-strategic agents). The population consists of 10^5 individuals. Half of the individuals are sampled from $x \sim \mathcal{N}\left(1,\frac{1}{3}\right)$ with true label 1 and the other half is sampled from $x \sim \mathcal{N}\left(0,\frac{1}{3}\right)$ with true label 0. We compute the optimal points and the social burden (see Section B)). [CM: not sure what this refers to]

Corollary 12. Consider Setup 1. Let M_{SM} be the map associated with standard microfoundations, let $\sigma \in (0, \infty)$, and let the cost function be of the form $c(x_1, x_2) = |x_1 - x_2|$. Suppose that $[\theta_{SL}, \theta_{SL} + 1] \in \Theta \cap X$, where θ_{SL} is defined so that $\mu(\theta_{SL}) = 0.5$. Then, it holds that:

Burden(
$$\theta_{PO}(M_{\sigma})$$
) < Burden($\theta_{PO}(M_{SM})$).

In fact, the social burden for fuzzy perception can be well below the social burden of standard microfoundations. To demonstrate this we visualize the threshold and social burden across a variety of different parameters of σ and p in Figure 3. The dashed lines indicate the respective reference values for standard microfoundations (SM) and for a population of non-strategic agents (NS). We observe that the threshold as well as the social burden decrease with the fraction p of non-strategic agents in the population. Furthermore, if every agent follows noisy response (p = 0), the threshold and the social burden are decreasing with σ (and hence the degree of imperfection in agents response). Overall, the acceptance threshold and the social burden of optimal classifiers derived under our new microfoundations is significantly lower than for standard microfoundations.

4.4 Using noisy response in practice

Microfoundations are useful in practice to model and anticipate how the population responds to a deployed classifier. This tool is particularly powerful if the decision maker can estimate agent behavior without needing to expose the population to potentially inaccurate and harmful classifiers to explore and learn about agents responses.

An appealing aspect of assuming a parameterized model of the agent responses is that the complex task of learning agent behavior is reduced to a parameter estimation problem. For noisy response, the aggregate response $\mathcal{D}(\theta)$ is parameterized by the variance σ of the perception noise. In practice such parameters of individual responses can often be estimated via *individual experiments*, i.e., by gathering information about individuals without ever deploying a classifier. For example, the decision maker can randomly sample agents in the population, ask survey questions to learn about their perceptions of the deployed classifier, and infer the variance σ from these samples. We refer to [Björkegren et al., 2020] for an actual field experiment that shows a example procedure for how to similarly obtain a reasonable estimates of the cost function c. An error in parameter estimation can then be translated into an error in the aggregate response.

Lemma 13. Given a population with noisy perceptions with parameter $\sigma \in (0, \infty)$, and an estimate $\tilde{\sigma}$ of σ , it holds that $\mathrm{TV} \Big(\mathcal{D}(\theta; M_{\tilde{\sigma}}), \mathcal{D}(\theta; M_{\tilde{\sigma}}) \Big) \leq \frac{3}{2} \sqrt{\frac{\epsilon m}{\min(\sigma^2, \tilde{\sigma}^2)}}.$

Combining Lemma 13 with a bound on the performative risk, we obtain the following robustness guarantee of the performative risk to estimation errors in noise parameter σ .

Corollary 14. Let the population be the aggregate of agents following noisy perceptions with parameter $\sigma \in (0, \infty)$, and let $\tilde{\sigma} \in (0, \infty)$ be an estimate of the perception parameter σ . Then the suboptimality of the estimated performative risk $\theta_{PO}(M_{\tilde{\sigma}})$ on the true population represented by M_{σ} is bounded by:

$$\Pr(\theta_{\text{PO}}(M_{\tilde{\sigma}})) - \Pr(\theta_{\text{PO}}(M_{\sigma}) \le 3\sqrt{\frac{\epsilon m}{\min(\sigma^2, \tilde{\sigma}^2)}}.$$

Hence, if the true distribution map is sufficiently close to some parameterization of noisy response, estimating the noise parameter σ provides a robust procedure to infer an estimate of the performative optima in practice.

Overall, noisy response offers a more descriptive and prescriptive model of agent behavior compared to the standard model, and still maintains analytical tractability. While we have focused on Gaussian noise in the perception function throughout this work, the outlined benefits of noisy response also apply to other *parameterized* noise distributions, as long as they are sufficiently smooth and continuous on all of \mathbb{R}^m . Hence, depending on the application, the decision-maker might prefer to pick a different noise model that can better capture the expected particularities of agents imperfections. The inference procedure via individual experimentation can then be adapted to obtain performative risk estimates that depend on the parameters of the noise distribution.

5 Discussion

Traditional approaches to decision making in strategic settings take either a purely individual-level approach like *strategic classification*, or a fully population-based approach like *performative prediction*. In this work, we maintain both perspectives. On one hand, we utilize microfoundations to endow the distribution shift induced by strategic behavior with natural structure, and on the other hand we also investigate population-level properties of the distribution they induce.

Taking this holistic view enabled us to identify degeneracies with standard microfoundations in the context of binary classification. In particular, it revealed that the aggregate response induced by the standard model is inherently discontinuous, which is the main cause for the brittleness of standard microfoundations identified in this work. By investigating the space of alternative microfoundations, we depart from this classical approach and focus on general properties of a model of microfoundations that are desirable. Finally, we proposed noisy response as an alternative microfoundation that is motivated by empirically observed imperfections in the agent response, satisfies the desired properties, and remains analytically tractable.

While we have focused on strategic classification in this work, we expect that considering the interplay between the individual-level and aggregate-level properties of strategic behavior could lead to novel insights in dynamic decision-making more generally. We believe that whenever distribution shifts result from individuals actions in social settings, a suitably chosen microfoundations model can provide a powerful tool for the decision-maker to reason about implications of a predictive model, and design effective and socially desirable algorithms, without exhaustive exploration.

Acknowledgments

We would like to thank Jacob Steinhardt for feedback. MJ acknowledges support from the Paul and Daisy Soros Fellowship. CMD acknowledges support from the Swiss National Science Foundation Postdoc.Mobility fellowship program.

References

- Emrah Akyol, Cedric Langbort, and Tamer Basar. Price of transparency in strategic machine learning. *Arxiv:1610.08210*, 2016.
- Ian Ball. Scoring strategic agents. ArXiv:1909.01888, 2020.
- Yahav Bechavod, Chara Podimata, Zhiwei Steven Wu, and Juba Ziani. Information discrepancy in strategic learning. *Arxiv:2103.01028*, 2021.
- Daniel Björkegren, Joshua E. Blumenstock, and Samsun Knight. Manipulation-proof machine learning. *Arxiv*:2004.03865, 2020.
- Mark Braverman and Sumegha Garg. The role of randomness and noise in strategic classification. In *Proc.* 1st FORC, 2020.
- Gavin Brown, Shlomi Hod, and Iden Kalemaj. Performative prediction in a stateful world. *Arxiv:2011.03885*, 2020.
- Michael Brückner and Tobias Scheffer. Stackelberg games for adversarial prediction problems. In *Proc.* 17th KDD, 2011.
- Michael Brückner, Christian Kanzow, and Tobias Scheffer. Static prediction games for adversarial learning problems. *JMLR*, 13(1):2617–2654, September 2012.
- Colin F. Camerer, George Loewenstein, and Matthew Rabin. *Advances in Behavioral Economics*. Princeton University Press, 2004.
- Yiling Chen, Yang Liu, and Chara Podimata. Learning strategy-aware linear classifiers. In *Proc.* 33rd NeurIPS, 2020.
- Olivier Coibion, Yuriy Gorodnichenko, and Rupal Kamdar. The formation of expectations, inflation, and the phillips curve. *Journal of Economic Literature*, 56(4):1447–1491, 2018.
- Nilesh N. Dalvi, Pedro M. Domingos, Mausam, Sumit K. Sanghai, and Deepak Verma. Adversarial classification. In *Proc.* 10th KDD, 2004.
- Ilias Diakonikolas, Gautam Kamath, Daniel Kane, Jerry Li, Ankur Moitra, and Alistair Stewart. Robust estimators in high-dimensions without the computational intractability. *SIAM J. Comput.*, 48(2):742–864, 2019.
- Jinshuo Dong, Aaron Roth, Zachary Schutzman, Bo Waggoner, and Zhiwei Steven Wu. Strategic classification from revealed preferences. In *Proc. EC*, 2018.
- Alex Frankel and Navin Kartik. Muddled Information. *Journal of Political Economy*, 127(4): 1739–1776, 2019.

- Alex Frankel and Navin Kartik. Improving Information via Manipulable Data. Working Paper, 2020.
- Dimitris Gatzouras. On images of borel measures under borel mappings. *Proceedings of the American Mathematical Society*, 130(9):2687–2699, 2002.
- Ganesh Ghalme, Vineet Nair, Itay Eilat, Inbal Talgam-Cohen, and Nir Rosenfeld. Strategic classification in the dark. *Arxiv*:2102.11592, 2021.
- Nika Haghtalab, Nicole Immorlica, Brendan Lucier, and Jack Z. Wang. Maximizing welfare with incentive-aware evaluation mechanisms. In *Proc.* 29th IJCAI, pages 160–166, 2020.
- Moritz Hardt, Nimrod Megiddo, Christos H. Papadimitriou, and Mary Wootters. Strategic classification. In *Proc. 7th ITCS*. ACM, 2016a.
- Moritz Hardt, Eric Price, and Nati Srebro. Equality of opportunity in supervised learning. In *Proc.* 29th NeurIPS, pages 3315–3323, 2016b.
- John C. Harsanyi. Games with incomplete information played by "bayesian" players, i-iii. part ii. bayesian equilibrium points. *Management Science*, 14(5):320–334, 1968.
- Christopher Hennessy and Charles Goodhart. Goodhart's law and machine learning. SSRN, 2020.
- Lily Hu, Nicole Immorlica, and Jennifer Wortman Vaughan. The disparate effects of strategic manipulation. In *Proc. FAccT*, 2019.
- Greg Kaplan and Giovanni L. Violante. Microeconomic Heterogeneity and Macroeconomic Shocks. *Journal of Economic Perspectives*, 32(3):167–194, 2018.
- Moein Khajehnejad, Behzad Tabibian, Bernhard Schölkopf, Adish Singla, and Manuel Gomez-Rodriguez. Optimal decision making under strategic behavior. *Arxiv:1905.09239*.
- Jon Kleinberg and Manish Raghavan. How do classifiers induce agents to invest effort strategically? In *Proc. EC*, 2019.
- Sagi Levanon and Nir Rosenfeld. Strategic classification made practical. Arxiv:2103.01826, 2021.
- Shengwu Li. Obviously strategy-proof mechanisms. American Economic Review, 107(11), 2017.
- Robert E Lucas Jr. Econometric policy evaluation: A critique. In *Carnegie-Rochester conference series on public policy*, volume 1, pages 19–46. North-Holland, 1976.
- Celestine Mendler-Dünner, Juan Perdomo, Tijana Zrnic, and Moritz Hardt. Stochastic optimization for performative prediction. *Proc.* 33rd NeurIPS, 2020.
- John Miller, Smitha Milli, and Moritz Hardt. Strategic classification is causal modeling in disguise. In *Proc.* 37th ICML, 2020.
- John Miller, Juan C. Perdomo, and Tijana Zrnic. Outside the echo chamber: Optimizing the performative risk. *Arxiv*:2102.08570, 2021.
- Smitha Milli, John Miller, Anca D. Dragan, and Moritz Hardt. The social cost of strategic classification. In *Proc. FAccT*, 2019.

- Juan C. Perdomo, Tijana Zrnic, Celestine Mendler-Dünner, and Moritz Hardt. Performative prediction. In *Proc.* 37th ICML, 2020.
- Yonadav Shavit, Benjamin L. Edelman, and Brian Axelrod. Causal strategic linear regression. In *Proc.* 37th ICML, 2020.
- Michael Spence. Job market signaling. The Quarterly Journal of Economics, 87(3):355–374, 1973.
- Daniel A. Spielman and Shang-Hua Teng. Smoothed analysis: An attempt to explain the behavior of algorithms in practice. *Commun. ACM*, 52(10):76–84, October 2009.
- Joseph E Stiglitz. Where modern macroeconomics went wrong. *Oxford Review of Economic Policy*, 34(1-2):70–106, 01 2018.
- Stratis Tsirtsis and Manuel Gomez-Rodriguez. Decisions, counterfactual explanations and strategic behavior. *Arxiv*:2002.04333, 2020.
- Hanrui Zhang and Vincent Conitzer. Incentive-aware pac learning. Proc. 35th AAAI, 2021.
- Hanrui Zhang, Yu Chen, and Vincent Conitzer. Automated mechanism design for classification with partial verification. *Proc.* 35th AAAI, 2021.

A Additional discussion of assumptions

A.1 Cost function

Let us discuss some context and implications of Assumption 1 defining a valid cost function. Unlike prior work [Milli et al., 2019; Miller et al., 2020; Braverman and Garg, 2020], we model a nonzero cost for *all* modifications to features, regardless of whether these modifications are in the right direction. In the spirit of generalizing beyond standard microfoundations, this accounts for how agents may erroneously expend effort on changing their features in an incorrect direction, as empirically demonstrated in Example 2. We further note that the definition of a valid cost function does not require symmetry in the arguments, which differentiates it from a metric.

A.2 Measurability requirements for alternative microfoundations

We now describe the measurability requirements that we need in order to define and work with maps $M \in \mathcal{M}$. If we ignore measurability requirements for a moment, then notice that each map $M \in \mathcal{M}$ can be associated with a distribution $\mathcal{D}_{\mathcal{T}XY} \in \Delta(\mathcal{T} \times X \times Y)$ given by (M(x,y),x,y). Since it is easier to define measurability requirements on $\mathcal{D}_{\mathcal{T}XY}$, we specify requirements on $\mathcal{D}_{\mathcal{T}XY}$, which gives an implicit specification of requirements on M. First, we define the probability space. Consider the sample space $\mathcal{T} \times X \times Y$. We can define a sigma algebra \mathcal{F} over Ω by viewing \mathcal{T} as the set of functions $X \times \Theta \to X$, and using that $X \subseteq \mathbb{R}^d$, $\Theta \subseteq \mathbb{R}^d$. The probability measure can then be given by $\mathcal{D}_{\mathcal{T}XY}$.

Since $\operatorname{image}(M) = \operatorname{supp}(\mathcal{D}_{\mathcal{T}XY})$ contains a very small fraction of the sample space $\mathcal{T} \times X \times Y$, we can work with a much smaller probability space in this context. This probability space is defined as follows: the sample space is $\operatorname{supp}(\mathcal{D}_{\mathcal{T}XY}) \in \mathcal{F}$ (i.e. a subset of $\mathcal{T} \times X \times Y$ in the sigma-algebra), and the sigma-algebra is intersections of every set in \mathcal{F} with $\operatorname{supp}(\mathcal{D}_{\mathcal{T}XY})$. The probability measure given by $\mathcal{D}_{\mathcal{T}XY}$ can be defined over this smaller probability space.

The distribution map \mathcal{D} can thus be viewed as random variables over this probability space. In particular, $\mathcal{D}(\theta)$ is the distribution of the random variable ($\mathcal{R}_t(x,\theta),y$). In order for this random variable to be well-defined, we place the following measurability assumption.

Assumption 2 (Measurability requirement on \mathcal{R}). We require that for each $\theta \in \Theta$ the function F_{θ} : supp $(\mathcal{D}_{TXY}) \to X \times Y$ given by $F_{\theta}(t, x, y) = (\mathcal{R}_{t}(x, \theta), y)$ is measurable.

A.3 Assumption on gaming behavior

In Section 3.3 we make the assumption that agent cannot have differing types solely on the basis of their true label. In other words, the map M cannot take into account the true label.

Assumption 3. For a map $M \in \mathcal{M}$, we require that M(x,0) = M(x,1) for all $x \in X$.

Assumption 3 means that agents with features *x* who have true label 0 versus true label 1 have identical distributions over response types in aggregate. We need this assumption to reason about performatively optimal points, because a decision maker has no access to the true labels beyond agents' reported features when anticipating strategic behavior.

A.4 Compactness of *X*

This assumption guarantees that the behavior of agents who follow standard microfoundations is well-defined.

Fact 1. Suppose that c is a valid cost function, and X is compact. Then $\sup_{x' \in X} (f_{\theta}(x') - c(x, x'))$ is attained on some $x^* \in X$ and the behavior of rational agents with perfect information is well-defined.

Proof. Let $X_{pos} := \{x \in X \mid f_{\theta}(x) = 1\} \subseteq X$ and $\Phi := \sup_{x' \in X} (f_{\theta}(x') - c(x, x'))$. Suppose that $\Phi = 0$. Then, by Assumption 1 the supremum is attained at x' = x. Now, suppose that $\Phi > 0$. Then $\Phi = 1 - \inf_{x' \in X_{pos}} \{c(x, x')\}$. By the fact that X_{pos} is a closed subset of a compact set (and thus compact), and c is continuous, we know that $\inf_{x' \in X_{pos}} \{c(x, x')\}$ is attained on some $x' \in X$.

B Proofs for Section 2

B.1 Proof of Proposition 1

In order to prove Proposition 1, we show that the gaming behavior of rational agents with perfect information can be characterized in the following way: Any rational agent with perfect information either will not change their features at all or will change their features exactly up to the decision boundary. We use the notation:

$$R_{t_{\text{SM}}}(x,\theta) := \underset{x' \in X}{\arg\max} (f_{\theta}(x') - c(x,x')) \tag{7}$$

to denote how an agent with features x who follows standard microfoundations will change their features in response to f_{θ} .

Proposition B.1. Suppose that c is a valid cost function. Then for any x the response (7) is either $R_{t_{SM}}(x,\theta) = x$ or $R_{t_{SM}}(x,\theta)$ is on the decision boundary of f_{θ} .

Proof of Proposition B.1. By Fact 1, we know that the quantity $\arg\max_{x' \in X} (f_{\theta}(x') - c(x, x'))$ is well defined. It suffices to show that if $R_{t_{\text{SM}}}(x,\theta) \neq x$, then $R_{t_{\text{SM}}}(x,\theta)$ is on the decision boundary of f_{θ} . If $R_{t_{\text{SM}}}(x,\theta) \neq x$, then we know that $c(x,R_{t_{\text{SM}}}(x,\theta)) > 0$. This means that $f_{\theta}(x) = 0$ and $R_{t_{\text{SM}}}(x,\theta) \in \arg\max_{x' \in X} (f_{\theta}(x') - c(x,x')) = \arg\min_{x' \in X_{\text{pos}}} c(x,x')$, where $X_{\text{pos}} := \{x \in X \mid f_{\theta}(x) = 1\}$. Assume for sake of contradiction that $R_{t_{\text{SM}}}(x,\theta)$ is not on the decision boundary. Then since $x \notin X_{\text{pos}}$ and $R_{t_{\text{SM}}}(x,\theta) \in X_{\text{pos}}$, there must exist x' on the line segment between x and $R_{t_{\text{SM}}}(x,\theta)$ such that x' is on boundary of X_{pos} , and thus the decision boundary of f_{θ} . Moreover, by Assumption 1, we know that $c(x,x') < c(x,R_{t_{\text{SM}}}(x,\theta))$. Since X_{pos} is closed, we see that $f_{\theta}(x') = 1$. Thus, $[f_{\theta}(x') - c(x,x')] < [f_{\theta}(R_{t_{\text{SM}}}(x,\theta)) - c(x,x')]$ which is a contradiction. □

Now, we use Proposition B.1 to prove Proposition 1.

Proof of Proposition 1. It suffices to show that $\mathcal{D}(\theta)$ is either equal to \mathcal{D}_{XY} or is a discontinuous distribution. Let $Q(\theta) \subseteq X$ be the set of agents who change their features at f_{θ} , i.e.

$$Q(\theta) := \left\{ x \in X \mid R_{t_{\text{SM}}}(x, \theta) \neq x \right\}.$$

If $\mathbb{P}_{(x,y)\in\mathcal{D}_{XY}}[x\in Q(\theta)]=0$, then $\mathcal{D}(\theta)=\mathcal{D}_{XY}$. Otherwise, suppose that $\mathbb{P}_{(x,y)\in\mathcal{D}_{XY}}[x\in Q(\theta)]>0$. By Proposition B.1, all of the agents in $Q(\theta)$ will game to somewhere on the decision boundary: that is, $R_{t_{\text{SM}}}(x,\theta)$ will be on the decision boundary for all $x\in Q(\theta)$. Thus, in $\mathcal{D}(\theta)$, there will be at least a $\mathbb{P}_{(x,y)\in\mathcal{D}_{XY}}[x\in Q(\theta)]$ probability mass of agents at the decision boundary, which is measure 0. This means that $\mathcal{D}(\theta)$ is not a continuous distribution.

B.2 Proof of Proposition 2

For convenience, we break down Proposition 2 into a series of propositions, roughly corresponding to part (a), part (b), and part (c), which we prove one-by-one.

First, let's consider the case where p=0. By the assumptions in Setup 1, we know that there exists a unique $\theta \in \Theta$ such that $p(\theta)=0.5$. We call this value θ_{SL} (and it is in the interior of Θ). We claim that this is the unique locally stable point when p=0.

Proposition B.2. Consider Setup 1, where a p = 1 fraction of agents are non-strategic. Then, θ_{SL} (defined above) is the unique locally stable point.

Proof. Since p = 1, the distribution map is given by $\mathcal{D}(\theta) = \mathcal{D}_{XY}$. A locally stable point θ must be a local minimum or a stationary point of the following optimization problem:

$$\min_{\theta \in \Theta} \mathbb{E}_{(x,y) \in \mathcal{D}_{XY}} [\mathbb{1}\{f_{\theta}(x) = y\}] = \min_{\theta \in \Theta} (\mathbb{E}_{(x,y) \in \mathcal{D}_{XY}} [\mathbb{1}\{x \geq \theta\}(1 - p(x))] + \mathbb{E}_{(x,y) \in \mathcal{D}_{XY}} [\mathbb{1}\{x < \theta\}p(x)]).$$

Notice that the unique such θ is θ_{SL} .

We introduce some basic properties and notation for agents who behave according to standard microfoundations. By Proposition B.1, we know if an agent games when the classifier f_{θ} is deployed, then they will game up to boundary, which in this case, is θ . We adopt similar notation to the proof of Proposition 1, and we denote the set of who game by:

$$Q(\theta) := \left\{ x \in X \mid R_{t_{\mathrm{SM}}}(x,\theta) \neq x \right\} = \left\{ x \in X \mid c(x,\theta) \leq 1, x < \theta \right\}.$$

(Technically, the agents $x \in Q(\theta)$ for whom $c(x,\theta) = 1$ are indifferent between not gaming and gaming to θ , but this is a measure 0 set by the assumption that \mathcal{D}_{XY} is continuous, and the assumption that c is valid (Assumption 1)). For $\theta \neq \min(\Theta)$, we see that for $\mathcal{D}(\theta)$, there will be a point mass at θ (from agents in $Q(\theta)$), the region $Q(\theta)$ will have zero probability density, and the rest of the distribution will remain identical to \mathcal{D}_{XY} .

We first characterize the set of stable points at p = 0. This follows a very similar argument to Lemma 3.2 in [Milli et al., 2019], but since our assumptions as well as our requirements for stability are slightly weaker, the characterization result looks slightly different. (In particular, points above the Stackelberg equilibrium can be locally stable points.)

Proposition B.3. Consider Setup 1, where a p=0 fraction of agents are non-strategic. Then, there exists a locally stable point, and moreover, the set of locally stable points forms an interval $[\theta_{min}, \max(\Theta)]$, where θ_{min} is the unique value such that:

$$\frac{\mathbb{E}_{(x,y)\in\mathcal{D}_{XY}}[p(x)I_{x\in Q(\theta_{min})}]}{\mathbb{E}_{(x,y)\in\mathcal{D}_{XY}}[I_{x\in Q(\theta_{min})}]} = 0.5.$$

(Moreover, it holds that $\theta_{min} > \theta_{SL}$, and $c(\theta_{SL}, \theta_{min}) < 1$.)

Proof. First, we show that $\theta^* = \min(\Theta)$ cannot be a stable point. Notice that $\mathcal{D}(\theta^*) = \mathcal{D}_{XY}$. Thus, θ^* is a local minimum or stationary point of $\min_{\theta \in \Theta} \mathbb{E}_{(x,y) \sim \mathcal{D}_{XY})} \mathbb{I}\{y \neq f_{\theta}(x)\}$. However, this is not possible because $p(\min(\Theta)) < 0.5$ by the assumptions in Setup 1.

Now, we consider $\theta \neq \min(\Theta)$. In this case, as discussed above, $\mathcal{D}(\theta)$ has a point mass at θ . Roughly speaking, the only property that needs to be satisfied for θ in the interior of Θ to be a local minimum of $\mathbb{E}_{(x,y)\in\mathcal{D}(\theta)}[\mathbb{1}\{f_{\theta'}(x)=y\}]$ is that it needs to be suboptimal for the decision maker to move just above the point mass (the decision maker never benefits from moving to

 $\theta - \epsilon$ because there is a region of zero probability density underneath θ). The loss induced from the point mass at θ is $\mathbb{E}_{(x,y)\in\mathcal{D}_{XY}}[(1-p(x))\mathbb{1}\{x\in Q(\theta)\}]$, while if the decision-maker moves just above θ is $\mathbb{E}_{(x,y)\in\mathcal{D}_{XY}}[p(x)\mathbb{1}\{x\in Q(\theta)\}]$. The condition thus becomes $\mathbb{E}_{(x,y)\in\mathcal{D}_{XY}}[p(x)\mathbb{1}\{x\in Q(\theta)\}] \geq \mathbb{E}_{(x,y)\in\mathcal{D}_{XY}}[(1-p(x))\mathbb{1}\{x\in Q(\theta)\}]$, which can be written as

$$\Gamma(\theta) := \frac{\mathbb{E}_{(x,y)\in\mathcal{D}_{XY}}[p(x)\mathbb{1}\{x\in Q(\theta)\}]}{\mathbb{E}_{(x,y)\in\mathcal{D}_{XY}}[\mathbb{1}\{x\in Q(\theta)\}]} \ge 0.5.$$
(8)

It suffices to show that the set of points where (8) is satisfied is an interval of the form $[\theta_{\min}, \max(\Theta)]$.

First, we show that the set of stable points forms an interval. It suffices to show that $\Gamma(\theta)$ is continuous and strictly increasing in θ . By the assumption on c (Assumption 1), we see that the endpoints of the interval $Q(\theta)$ are strictly increasing in θ . This, coupled with the fact that ℓ is strictly increasing in x (assumed in Setup 1), implies that $\Gamma(\theta)$ is continuous and strictly increasing as desired.

Furthermore, when $p(\theta) \le 0.5$, the condition in (8) is never satisfied, and thus all stable points θ satisfy $\theta > \theta_{SL}$, and hence $\theta_{min} > \theta_{SL}$.

Lastly, we show that this interval is not nonempty, and that $c(\theta_{SL}, \theta_{min}) \le 1$. Consider θ such that $c(\theta_{SL}, \theta) = 1$ (which we know exists by Setup 1), we see that $Q_{\theta} = [\theta_{SL}, \theta]$. Using the conditions on ℓ , this means that condition (8) is satisfied and there is actually a strict equality. Using that c is valid, this means that $c(\theta_{SL}, \theta_{min}) < 1$.

We now prove that no locally stable points exist for 0 .

Proposition B.4. Consider Setup 1, where a 0 fraction of agents are non-strategic. Then, there are no locally stable points.

Proof. When $0 , we show that there are no locally stable points. Assume for sake of contradiction that <math>\theta^*$ is a locally stable point. Recall that for θ^* to be locally stable, θ^* must either be a stationary point or a local minimum of $\min_{\theta \in \Theta} \mathbb{E}_{(x,y) \sim \mathcal{D}(\theta_{PS})} \mathbb{1}\{y \neq f_{\theta}(x)\}$. We divide into three cases: $(1) \theta^* = \min(\Theta)$, $(2) \theta^* > \min(\Theta) \land p(\theta^*) \leq 0.5$, $(3) \theta^* > \min(\Theta) \land p(\theta^*) > 0.5$, and show that each results in a contradiction.

For the case (1), where $\theta^* := \min(\Theta)$, we see that $\mathcal{D}(\theta^*) = \mathcal{D}_{XY}$. Thus, θ^* is a local minimum or stationary point of $\min_{\theta \in \Theta} \mathbb{E}_{(x,y) \sim \mathcal{D}_{XY}} \mathbb{1}\{y \neq f_{\theta}(x)\}$. However, this is not possible because $p(\min(\Theta)) < 0.5$ by the assumptions in Setup 1. For the remaining two cases, we know that $\mathcal{D}(\theta^*)$ has a point mass at θ^* . This means that $\mathbb{E}_{(x,y) \in \mathcal{D}(\theta^*)} \mathbb{1}\{y \neq f_{\theta'}(x)\}$ is not differentiable at $\theta' = \theta^*$, and so θ^* must be a local minimum.

For case (2), notice that $Q(\theta^*)$ consists a nonzero density of agents for whom p(x) < 0.5, and for all agents $x \in Q(\theta^*)$, it holds that $p(x) \le 0.5$. The decision maker thus wishes to move just to the other side of the point mass. (This is possible because $\theta^* < \max(\Theta)$ based on the fact that $p(\theta^*) < 0.5$ and the assumptions in Setup 1.) In particular, $\lim_{\epsilon \to 0} \mathbb{E}_{(x,y) \in \mathcal{D}(\theta)} \mathbb{I}\{y \ne f_{\theta^*+\epsilon}(x)\} < \lim_{\epsilon \to 0} \mathbb{E}_{(x,y) \in \mathcal{D}(\theta)} \mathbb{I}\{y \ne f_{\theta^*}(x)\}$.

For case (3), notice that there exists $\epsilon > 0$ such that $p(\theta) > 0.5$ and $\theta \in Q(\theta^*)$ for all $\theta \in (\theta^* - \epsilon, \theta^*)$. The presence of non-strategic agents means that the decision-maker wishes to move to $\theta^* - \epsilon$ to achieve better performance on non-strategic agents. Since there are no strategic agents within $(\theta^* - \epsilon, \theta^*)$, this can be done without affecting the classification of strategic agents. In particular, $\mathbb{E}_{(x,y)\in\mathcal{D}(\theta)}\mathbb{1}\{y\neq f_{\theta-\epsilon}(x)\} < \mathbb{E}_{(x,y)\in\mathcal{D}(\theta)}\mathbb{1}\{y\neq f_{\theta}(x)\}$.

Now, we prove that repeated risk minimization oscillates when 0 .

Proposition B.5. Consider Setup 1, where a $0 fraction of agents are non-strategic. Repeated risk minimization will oscillate according to the following behavior. Let <math>\theta_{PS}^1$ denotes the (unique) locally performatively stable point at p = 1 and let θ_{PS}^0 denotes the minimum locally performatively stable point at p = 0. RRM will oscillate between θ_{PS}^1 and a threshold $f(p) > \theta_{PS}^1$, where f(p) is decreasing in p, approaching θ_{PS}^1 as $p \to 1$ and approaching θ_{PS}^0 as $p \to 0$.

Proof. Using Proposition B.2, we see that there is a unique performatively stable point for p = 1, given by $\theta_{PS}^1 := \theta_{SL}$. Using Proposition B.3, we see that the smallest locally stable point is given by $\theta_{PS}^0 := \theta_{min}$.

In the case of $p \in (0,1)$, the distribution map $\mathcal{D}(\theta)$ takes the form of a mixture with p weight on \mathcal{D}_{XY} and with 1-p weight on the distribution map of agents who behave according to standard microfoundations (which has a point mass at θ , zero density within Q_{θ} , and the same as the original distribution elsewhere). The main step in our proof is an analysis of the global optima of

$$B(\theta) = \operatorname{argmin}_{\theta' \in \Theta} \mathbb{E}_{(x,y) \sim \mathcal{D}(\theta)} \ \mathbb{1}\{y \neq f_{\theta'}(x)\}.$$

for each $\theta \in \Theta$. For convenience, we let

$$DPR(\theta, \theta') := \mathbb{E}_{(x,y) \sim \mathcal{D}(\theta)} \mathbb{1}\{y \neq f_{\theta'}(x)\}.$$

We split into three cases: (a) $\theta \ge \theta_{\min}$, (b) $\theta < \theta_{SL}$, and (c) $\theta_{SL} \le \theta \le \theta_{\min}$.

Case (a): $\theta \ge \theta_{\min}$. We claim that $B(\theta) = \{\theta_{SL}\}$. In this case, the proof of Proposition B.3 tells us that moving just above the point mass will incur no better risk than at θ . Moreover, since p(x) > 0.5 for all $x \ge \theta \ge \theta_{\min}$, we see that $\mathrm{DPR}(\theta, \theta') > \mathrm{DPR}(\theta, \theta)$ for all $\theta' > \theta$. Because of the presence of non-strategic agents, a p fraction of agents will be present in Q_{θ} , and these agents do not change their features. Moreover, for $\theta' < \min(Q_{\theta})$, $\mathcal{D}(\theta)$ looks like the base distribution. Since p(x) > 0.5 for $\theta_{SL} < x \le \theta_{\min}$, we see that $\mathrm{DPR}(\theta, \theta') < \mathrm{DPR}(\theta, \theta)$ for all $\theta_{SL} \le \theta' \le \theta$. Moreover, this argument actually shows that $\theta_{SL} = \mathrm{argmin}_{\theta_{SL} \le \theta' \le \theta} \mathrm{DPR}(\theta, \theta')$. Lastly, since p(x) < 0.5 for $x < \theta_{SL}$, we see that $\mathrm{DPR}(\theta, \theta') > \mathrm{DPR}(\theta, \theta_{SL})$ for all $\theta' < \theta_{SL}$.

Case (b): $\theta < \theta_{SL}$. If $\theta < \theta_{SL}$, then we claim that $B(\theta) = \{\theta_{SL}\}$. In this case, all agents x below θ_{SL} in $\mathcal{D}(\theta)$ have p(x) < 0.5. Thus, $\theta_{SL} = \operatorname{argmin}_{\theta \leq \theta' \leq \theta_{SL}} \operatorname{DPR}(\theta, \theta')$. Moreover, above θ_{SL} , $\mathcal{D}(\theta)$ looks like the base distribution. This means that $\operatorname{DPR}(\theta, \theta') > \operatorname{DPR}(\theta, \theta_{SL})$ for all $\theta' > \theta_{SL}$, as desired.

Case (c): $\theta_{SL} \leq \theta \leq \theta_{min}$. Using the same argument as Case (a), we see that $DPR(\theta, \theta') > DPR(\theta, \theta_{SL})$ for all $\theta' < \theta_{SL}$. Moreover, we also see that the risk obtained by the threshold right above the point mass beats any higher threshold: that is, $\lim_{\epsilon \to 0, \epsilon \geq 0} DPR(\theta, \theta + \epsilon) < DPR(\theta, \theta')$ for any $\theta' > \theta$. This is because all agents $x > \theta$ have p(x) > 0.5.

Thus, all we need to do is to compare the threshold right above the point mass with the threshold θ_{SL} . Notice that these two classifiers behave the same on strategic agents with true features $x \notin Q_{\theta}$ (this is because $\theta_{SL} \in Q_{\theta}$, because by Proposition B.3, we know that $c(\theta_{SL}, \theta) < c(\theta_{SL}, \theta_{\min}) < 1$.). Moreover, they also behave the same on non-strategic agents not in $\theta_{SL} \le x \le \theta$. Thus, we only need to focus on strategic agents with true features in Q_{θ} and non-strategic agents with $\theta_{SL} \le x \le \theta$.

Thus, we use the expression in the proof of Proposition B.3 to see that:

$$\begin{split} &\lim_{\epsilon \to 0, \epsilon > 0} \mathrm{DPR}(\theta, \theta + \epsilon) - \mathrm{DPR}(\theta, \theta_{\mathrm{SL}}) \\ &= p \cdot \mathbb{E}_{(x,y) \in \mathcal{D}_{\mathrm{XY}}} \big[\mathbb{1}\{x \in [\theta_{\mathrm{SL}}, \theta]\} p(x) \big] + (1 - p) \mathbb{E}_{(x,y) \in \mathcal{D}_{\mathrm{XY}}} \big[\mathbb{1}\{x \in Q(\theta)\} p(x) \big] \\ &- p \mathbb{E}_{(x,y) \in \mathcal{D}_{\mathrm{XY}}} \big[\mathbb{1}\{x \in [\theta_{\mathrm{SL}}, \theta]\} (1 - p(x)) \big] - (1 - p) \mathbb{E}_{(x,y) \in \mathcal{D}_{\mathrm{XY}}} \big[\mathbb{1}\{x \in Q(\theta)\} (1 - p(x)) \big] \\ &= 2p \mathbb{E}_{(x,y) \in \mathcal{D}_{\mathrm{XY}}} \big[\mathbb{1}\{x \in [\theta_{\mathrm{SL}}, \theta]\} p(x) \big] + 2(1 - p) \mathbb{E}_{(x,y) \in \mathcal{D}_{\mathrm{XY}}} \big[\mathbb{1}\{x \in Q(\theta)\} p(x) \big] \\ &- p \mathbb{E}_{(x,y) \in \mathcal{D}_{\mathrm{XY}}} \big[\mathbb{1}\{x \in [\theta_{\mathrm{SL}}, \theta]\} \big] - (1 - p) \mathbb{E}_{(x,y) \in \mathcal{D}_{\mathrm{XY}}} \big[\mathbb{1}\{x \in Q(\theta)\} \big]. \end{split}$$

The relevant quantity is:

$$Z(p) := p\left(\mathbb{E}_{(x,y) \in \mathcal{D}_{XY}}[I_{x \in [\theta_{SL},\theta]}(2p(x)-1)]\right) + (1-p)\left(\mathbb{E}_{(x,y) \in \mathcal{D}_{XY}}[I_{x \in Q(\theta)}(2p(x)-1)]\right)$$

Let's denote by θ^+ the parameter weights "right above the point mass" (that is, the parameter weights given by approaching θ from above θ , without ever reaching θ). We see that $B(\theta) = \{\theta_{SL}\}$ if and only if Z(p) > 0, $B(\theta) = \{\theta^+\}$ if and only if Z(p) < 0, and $B(\theta) = \{\theta^+, \theta_{SL}\}$ if and only if Z(p) = 0.

Now, we show that Z(p) is increasing in θ . Let p_{base} be the pdf of \mathcal{D}_{XY} . The derivative of the first term is: $p(2p(\theta)0-1)p_{\text{base}}(\theta) > 0$, and the derivative of the second term is: $(1-p)(2p(\theta)-1)p_{\text{base}}(\theta) - (1-p)(2p(\min(Q(\theta))-1)p_{\text{base}}(\min(Q(\theta))) > 0$, as desired.

Moreover, at $\theta = \theta_{SL}$, we see that Z(p) < 0; at $\theta = \theta_{min}$, on the other hand, Z(p) > 0.

Thus, repeated retraining will oscillate between θ_{SL} and f(p), where f(p) is the value such that Z(f(p)) = 0. To see that f(p) is decreasing in p, notice that Z(p) is increasing in p for all $\theta_{SL} \le \theta \le \theta_{\min}$. As $p \to 0$, it is easy to see that $f(p) \to \theta_{\min}$. As $p \to 1$, it is easy to see that $f(p) \to \theta_{SL}$.

Using the above results, we can conclude Proposition 2.

Proof of Proposition 2. When p = 1, we can apply Proposition B.2. When p = 0, we can apply Proposition B.3 to see that a locally stable point exists. When 0 , we can apply Proposition B.4 to see that no locally stable point exists. For the behavior of repeated risk minimization, we can apply Proposition B.5.

B.3 Formal Statement and Proof of Proposition 3

We give a formal statement of Proposition 3 using the technology of alternative microfoundations. Let c be a valid cost function. First, we formalize expenditure monotonicity (Property 1) in the language of alternative microfoundations.

Property 4. Let Θ be a model class of threshold functions, and let c be a cost function. A mapping $M \in \mathcal{M}$ satisfies *expenditure monotonicity* if $c(\mathcal{R}_t(x,\theta),x) \leq \gamma$ for every $\theta \in \Theta$ and every $t \in \operatorname{Image}(M)$, and if $f_{\theta}(\mathcal{R}_t(x;\theta)) = 1$, then $f_{\theta'}(\mathcal{R}_t(x;\theta')) = 1$ for all $\theta' \leq \theta$.

Let \mathcal{M}^* be the set of maps M such that every $t \in \bigcup_{(x,y) \in X} \operatorname{supp}(M(x))$ satisfies expenditure monotonicity (Property 4) and such that Assumption 3 is satisfied. Let \mathscr{D} be the set of distribution maps $\mathcal{D}(\cdot;M)$ for $M \in \mathcal{M}^*$.

Proposition B.6. Consider Setup 1. Let \mathcal{D} be the class of distribution maps defined above. Then:

$$\begin{split} \theta_{PO}(\mathcal{D}_{SM}) &\geq \theta_{PO}(\mathcal{D}) \\ \text{Burden}(\theta_{PO}(\mathcal{D}_{SM})) &\geq \text{Burden}(\theta_{PO}(\mathcal{D})). \end{split}$$

where \mathcal{D}_{SM} denotes the distribution map given by standard microfoundations, and $\theta_{PO}(\mathcal{D})$ denotes the minimal performatively optimal point associated with the distribution map \mathcal{D} .

Proof. For ease of notation, let θ_{SL} be the unique value such that $p(\theta_{SL}) = 0.5$. It is easy to see that $\theta' = \theta_{PO}(\mathcal{D}_{SM})$ is the unique point such that $c(\theta_{SL}, \theta') = 1$ and $\theta' > \theta_{SL}$.

Since $Burden(\cdot)$ is monotonic in its argument, all we need to do is to show $\theta_{PO}(\mathcal{D}_{SM}) \geq \theta_{PO}(\mathcal{D})$. It suffices to show that for $\theta > \theta_{PO}(\mathcal{D}_{SM})$ and for any $\mathcal{D} \in \mathcal{D}$, it holds that $PR(\theta) \leq PR(\theta_{PO}(\mathcal{D}_{SM}))$, where the performative risk is with respect to \mathcal{D} .

First, let's consider the set of agents $S_1 := \{(t,x,y) \mid x < \theta_{SL}\}$. For $(t,x,y) \in S_1$, notice that $c(x,\theta) > c(x,\theta_{PO}(\mathcal{D}_{SM})) > c(\theta_{SL},\theta_{PO}(\mathcal{D}_{SM})) = c(\theta_{SL},\theta_{PO}(\mathcal{D}_{RDPI})) \geq 1$. By the expenditure constraint, this means that these agents will not game on f_{θ} or $f_{\theta_{PO}(\mathcal{D}_{SM})}$: i.e, $\mathcal{R}_t(x,\theta) = x$ and $\mathcal{R}_t(x,\theta_{PO}(\mathcal{D}_{SM})) = x$ for $(t,x,y) \in S_1$. Thus, $f_{\theta_{PO}(\mathcal{D}_{SM})}(\mathcal{R}_t(x,\theta_{PO}(\mathcal{D}_{SM}))) = f_{\theta}(\mathcal{R}_t(x,\theta)) = 0$. The performative risk with respect to \mathcal{D} is thus equivalent on S_1 for f_{θ} and $f_{\theta_{PO}(\mathcal{D}_{SM})}$.

Now, let's consider the remaining set of agents $S_2 := \{(t, x, y) \mid x \ge \theta_{SL}\}$. Let $S_2 \subseteq S_2$ be the set

$$S_2' = \{(t, x, y) \in S_2 \mid f_{\theta}(\mathcal{R}_t(x, \theta)) = 1\}.$$

and

$$S_2^{\prime\prime} = \Big\{ (t,x,y) \in S_2 \mid f_{\theta_{\mathrm{PO}}(\mathcal{D}_{\mathrm{SM}})}(\mathcal{R}_t(x,\theta_{\mathrm{PO}}(\mathcal{D}_{\mathrm{SM}}))) = 1 \Big\}.$$

We claim that $S_2' \subseteq S_2''$. This is because of the second condition in expenditure monotonicity that if x was labelled positively by $f_{\theta_{\text{SL}}}$, then x is also labelled positively by f_{θ} : in particular, we can thus conclude that $f_{\theta}(R(\theta)) = 1$ implies that $f_{\theta_{\text{PO}}(\mathcal{D}_{\text{SM}})}(\mathcal{R}_t(x,\theta_{\text{PO}}(\mathcal{D}_{\text{SM}}))) = 1$.

Now, we claim that the performative risk with respect to \mathcal{D} on S_2 is no better for θ than for $\theta_{PO}(\mathcal{D}_{SM})$. This follows from the fact that $S_2' \subseteq S_2''$, and the fact that $p(x) \ge 0.5$ for $(t, x, y) \in S_2$ coupled with Assumption 3. This completes the proof.

C Proofs for Section 3

C.1 Proof of Proposition 5

We prove Proposition 5. The intuition is that there is a response type for every possible agent response, and it remains to show that the appropriate choice of agent response types can "shift the mass" from \mathcal{D}_{XY} to $\mathcal{D}(\theta)$. In fact, M only needs to map the population to two different response types. Now, we formally prove this result.

Proof of Proposition 5. We prove Proposition 5 by construction and show that there is an M that can microfound any distribution map. We construct M as follows. We construct response types t_0 and t_1 , and define $M(x,0) = t_0$ for all $x \in X$ and $M(x,1) = t_1$ for all $x \in X$. In other words, we associate agents with true label 0 with the type t_0 and agents with true label 1 with the type t_1 .

In order to construct t_0 and t_1 , we define the following probability measures over the measure space $X \subseteq \mathbb{R}^D$ equipped with the Borel sigma-algebra. We consider $\mu^0(\theta)$ to be the probability measure given by the distribution over x when $(x,y) \in \mathcal{D}(\theta)$ and y=0. We define $\mu^1(\theta)$ similarly. We let μ^0_{XY} be the probability measure given by the distribution x where $(x,y) \in \mathcal{D}_{XY}$ and y=0, and we define μ^1_{XY} analogously.

First, we claim that it suffices to prove that for each $\theta \in \Theta$ there is a measurable map $f_{0,\theta}: X \to X$ that maps the probability measure μ_{XY}^0 to $\mu^0(\theta)$, and a measurable map $f_{1,\theta}: X \to X$ that maps μ_{XY}^1 to $\mu^1(\theta)$. In this case, we can define t_0 to be given by $\mathcal{R}_{t_0}(x,\theta) = f_{0,\theta}(x)$ and t_1 to be given by $\mathcal{R}_{t_1}(x,\theta) = f_{1,\theta}(x)$. Let's now consider the distribution given by $(\mathcal{R}_t(x,\theta),y)$ where $(t,x,y) \sim \mathcal{D}_{TXY}$. The condition distribution over y=0 is given by $\mu^0(\theta)$ and the conditional distribution over y=1 is given by $\mu^1(\theta)$, which means that the distribution over all is given by $\mathcal{D}(\theta)$, as desired. Moreover, the measurability requirements on $f_{0,\theta}$ and $f_{1,\theta}$ guarantee that Assumption 2 is satisfied.

Thus, it suffices to construct $f_{0,\theta}$ and $f_{1,\theta}$ for $\theta \in \Theta$ that satisfy the above conditions. To do this, we make use of Proposition 3 in [Gatzouras, 2002], which says that there exists a Borel mapping from any tight non-atomic measure to any other probability measure. Since the probability measure associated to \mathcal{D}_{XY} is non-atomic, we see that μ_{XY}^0 and μ_{XY}^1 are non-atomic as desired, and so a Borel mapping from μ_{XY}^0 to $\mu^0(\theta)$ exists and a Borel mapping from μ_{XY}^1 to $\mu^1(\theta)$ exists.

C.2 Proof of Proposition 6

Proof of Proposition 6. To prove Proposition 6 we show that $dPR_{\theta}(\theta')$ is continuous in θ and θ' . We see that

$$\mathsf{DPR}(\theta,\theta') = \int_{x' \geq \theta'} p_{\theta}((x',0)) dx' + \int_{x' < \theta'} p_{\theta}((x',1)) dx'.$$

Let's take a derivative with respect to θ' to obtain:

$$\mathsf{dPR}_{\theta}(\theta') = -p_{\theta}((\theta', 0)) + p_{\theta}((\theta', 1)).$$

The first continuity requirement tells us that this is continuous in θ' , and the second continuity requirement tells us that this continuous in θ .

C.3 Proof of Theorem 7 and Corollary 8

We first recall the definition of the decoupled performative risk [Perdomo et al., 2020]:

$$\mathsf{DPR}(\theta, \theta') := \mathbb{E}_{(x,y) \in \mathcal{D}(\theta)} [\mathbb{1} (f_{\theta'}(x) \neq y)].$$

The gradient of the decoupled performative risk plays an important role in our analysis of locally stable points. In order to take derivatives at the boundary, we consider an open set $\Theta' \supset \Theta$ that is also bounded and convex, and assume there are classifiers associated with each $\theta \in \Theta'$, although the decision maker only considers classifier weights in Θ . We use the notation:

$$\mathsf{dPR}_{\theta}(\theta') := \nabla_{\theta'} \mathsf{DPR}(\theta, \theta') = \nabla_{\theta'} \mathbb{E}_{(x,y) \sim \mathcal{D}(\theta)} [\ \mathbb{1} \{ y \neq f_{\theta'}(x) \}]$$

to denote the gradient of the decoupled performative risk with respect to the second argument. To prove Theorem 7 we show that the continuity of the derivatives of the decoupled performative risk guarantees the existence of stable points under mixtures with non-strategic agents.

Proof of Theorem 7. Our main technical ingredient is this proof is applying Brouwer's fixed point theorem on $G_{gd}(\theta) = \operatorname{Proj}_{\Theta}(\theta + \eta dPR_{\theta}(\theta))$. It thus suffices to show that the map $\theta \mapsto \operatorname{Proj}_{\Theta}(\theta + \eta dPR_{\theta}(\theta))$ is continuous.

First, we show that aggregate risk smoothness implies that $\theta \mapsto \operatorname{Proj}_{\Theta}(\theta + \eta \operatorname{dPR}_{\theta}(\theta))$ is a continuous map. By aggregate risk smoothness, we know that $\operatorname{dPR}_{\theta}(\theta)$ exists for all $\theta \in \Theta$. Moreover, for any $\theta \in \Theta$, aggregate risk smoothness tells us that:

$$\lim_{\theta' \to \theta} \left\| \mathsf{dPR}_{\theta}(\theta) - \mathsf{dPR}_{\theta'}(\theta') \right\| \leq \lim_{\theta' \to \theta} \left\| \mathsf{dPR}_{\theta}(\theta) - \mathsf{dPR}_{\theta'}(\theta) \right\| + \lim_{\theta' \to \theta} \left\| \mathsf{dPR}_{\theta'}(\theta) - \mathsf{dPR}_{\theta'}(\theta') \right\|.$$

Thus, $dPR_{\theta}(\theta)$ is continuous in θ . Moreover, since the sum of continuous functions is continuous, this means that $\theta \mapsto \theta + \eta dPR_{\theta}(\theta)$ is continuous. Now, since projection onto a convex set is a contraction map, we can conclude that $\theta \mapsto Proj_{\Theta}(\theta + \eta dPR_{\theta}(\theta))$ is continuous as desired.

We now prove Corollary 8.

Proof of Corollary 8. By Theorem 7, it suffices to establish that $dPR_{\theta}(\theta)$ exists and is continuous for the population as a whole. We use the fact that since derivatives are linear, we can break this down into a term for non-strategic agents and a term for strategic agents. Since the sum of two continuous function is continuous, it suffices to show that $dPR_{\theta}(\theta)$ exists and is continuous for non-strategic agents and for strategic agents. For strategic agents, this follows from aggregate risk smoothness. For non-strategic agents, since the (non-performative) risk $R(\theta) := \mathbb{E}_{(x,y) \in \mathcal{D}_{XY}} \mathbb{I}\{f_{\theta}(x) = y\}$ is differentiable in θ and $dPR_{\theta}(\theta) = \nabla_{\theta}R(\theta)$ is continuous in θ as desired.

C.4 Proof of Lemma 9

We consider the 1-dimensional setting where $X\subseteq\mathbb{R}$ and Θ is the class of threshold functions. First, we show a bound on the performative risk in terms of the Kolmogrov-Sminorff (KS) distance between the true distribution map and estimated distribution map. To state this bound, we introduce the following notation. We use a subscript notation $\mathcal{D}_{S(\Theta_0,c)}(\theta;M)$ to denote the aggregate response distribution $\mathcal{D}(\theta;M)$ restricted to agents with true features $x\in S(\Theta_0,c)$, where $S(\Theta_0,c)$ is defined as in (5). Let $\mathcal{D}^0_{S(\Theta_0,c)}(\theta;M)$ be the marginal distribution over x of the conditional distribution of $(x,y)\sim\mathcal{D}_{S(\Theta_0,c)}(\theta;M)$ conditional on y=0. We define $\mathcal{D}^1_{S(\Theta_0,c)}(\theta;M)$, $\mathcal{D}^0_{S(\Theta_0,c)}(\theta;\tilde{M})$, and $\mathcal{D}^1_{S(\Theta_0,c)}(\theta;\tilde{M})$ analogously.

Lemma C.1. Let Θ be a model class of posterior threshold functions, and c be an outcome-valid cost function. Suppose that M, \tilde{M} restricted to the domain $(X \setminus S(\Theta_0, c)) \times Y$ are expenditure-constrained. Then, for any $\Theta_0 \subseteq \Theta : \theta_{PO}(M) \in \Theta_0$, the predicted performative optima $\theta_{PO}(\tilde{M})$ satisfies:

$$PR_M(\theta_{PO}(\tilde{M}) \le PR_M(\theta_{PO}(M))) + 2\xi$$

where ξ is defined to be

$$\sup_{\theta} (A_{\theta} + B_{\theta})$$

where

$$A(\theta) := \mathbb{P}[x \in S(\Theta_0, c) \& y = 0] KS\left(\mathcal{D}_{S(\Theta_0, c)}^0(\theta; M), \mathcal{D}_{S(\Theta_0, c)}^0(\theta; \tilde{M})\right)$$

$$B(\theta) := \mathbb{P}[x \in S(\Theta_0, c) \& y = 1] KS\left(\mathcal{D}_{S(\Theta_0, c)}^1(\theta; M), \mathcal{D}_{S(\Theta_0, c)}^1(\theta; \tilde{M})\right).$$

Proof. Let $PR(\theta; M)$ denote the performative risk at θ on $\mathcal{D}(\theta; M)$ and let $PR(\theta; \tilde{M})$ denote the performative risk at θ on $\mathcal{D}(\theta; \tilde{M})$. It suffices to show that $|PR(\theta; M) - PR(\theta; \tilde{M})| \le \xi$ for all $\theta \in \Theta_0$ (since this would mean that $PR(\theta_{PO}(\tilde{M}); M) \le PR(\theta_{PO}(\tilde{M}); \tilde{M}) + \xi \le PR(\theta_{PO}(M); \tilde{M}) + \xi \le PR(\theta_{PO}(M); M) + 2\xi$, as desired). Notice that:

$$|\operatorname{PR}(\theta; M) - \operatorname{PR}(\theta; \tilde{M})| = \left| \mathbb{E}_{(x, y) \sim \mathcal{D}(\theta; M)} [\mathbb{1}\{y \neq f_{\theta}(x)\}] - \mathbb{E}_{(x, y) \sim \mathcal{D}(\theta; \tilde{M})} [\mathbb{1}\{y \neq f_{\theta}(x)\}] \right|.$$

Let's let \mathcal{D}_{TXY} be the distribution of (t, x, y) where $(x, y) \sim \mathcal{D}_{XY}$ and $t \sim M(x, y)$. Similarly, let $\tilde{\mathcal{D}}_{TXY}$ be the distribution of (t, x, y) where $(x, y) \sim \mathcal{D}_{XY}$ and $t \sim \tilde{M}(x, y)$. Notice that:

$$\left| \operatorname{PR}(\theta; M) - \operatorname{PR}(\theta; \tilde{M}) \right| = \left| \mathbb{E}_{(t, x, y) \sim \mathcal{D}_{TXY}} \left[\ \mathbb{1}\{y \neq f_{\theta}(R_t(x, \theta))\} \right] - \mathbb{E}_{(t, x, y) \sim \tilde{\mathcal{D}}_{TXY}} \left[\ \mathbb{1}\{y \neq f_{\theta}(R_t(x, \theta))\} \right] \right|.$$

Now, we claim that for any agent (t,x) where $x \notin S(\Theta_0,c)$ and for $t \in \text{supp}(\mathcal{D}_{TXY}) \cup \text{supp}(\tilde{\mathcal{D}}_{TXY})$, it holds that $f_{\theta}(R_t(x,\theta)) = f_{\theta}(x)$ for every $\theta \in \Theta_0$. Note that since M satisfies the expenditure

constraint with respect to c, then we know that if $x \notin S_{\theta}$, it holds that $f_{\theta}(R_t(x,\theta)) = f_{\theta}(x)$. Moreover, note that since $S_{\theta} \subseteq S(\Theta_0,c)$ by definition, this yields the desired statement. Thus we have that:

$$\begin{split} & \left| \mathbb{E}_{(t,x,y) \sim \mathcal{D}_{TXY}} \left[\ \mathbb{1} \{ y \neq f_{\theta}(R_{t}(x,\theta)) \} \right] - \mathbb{E}_{(t,x,y) \sim \tilde{\mathcal{D}}_{TXY}} \left[\ \mathbb{1} \{ y \neq f_{\theta}(R_{t}(x,\theta)) \} \right] \right| \\ & \leq \left| \mathbb{E}_{(t,x,y) \sim \mathcal{D}_{TXY}} \left[\ \mathbb{1} \{ y \neq f_{\theta}(R_{t}(x,\theta)) \} \ \mathbb{1} \{ x \notin S(\Theta_{0},c) \} \right] - \mathbb{E}_{(t,x,y) \sim \tilde{\mathcal{D}}_{TXY}} \left[\ \mathbb{1} \{ y \neq f_{\theta}(R_{t}(x,\theta)) \} \ \mathbb{1} \{ x \notin S(\Theta_{0},c) \} \right] \right| \\ & + \left| \mathbb{E}_{(t,x,y) \sim \mathcal{D}_{TXY}} \left[\ \mathbb{1} \{ y \neq f_{\theta}(R_{t}(x,\theta)) \} \ \mathbb{1} \{ x \in S(\Theta_{0},c) \} \right] - \mathbb{E}_{(t,x,y) \sim \tilde{\mathcal{D}}_{TXY}} \left[\ \mathbb{1} \{ y \neq f_{\theta}(R_{t}(x,\theta)) \} \ \mathbb{1} \{ x \in S(\Theta_{0},c) \} \right] \right| \\ & = \left| \mathbb{E}_{(t,x,y) \sim \mathcal{D}_{TXY}} \left[\ \mathbb{1} \{ y \neq f_{\theta}(x) \} \ \mathbb{1} \{ x \in S(\Theta_{0},c) \} \right] - \mathbb{E}_{(x,y) \sim \tilde{\mathcal{D}}_{TXY}} \left[\ \mathbb{1} \{ y \neq f_{\theta}(x) \} \ \mathbb{1} \{ x \in S(\Theta_{0},c) \} \right] \right| \\ & = \left| \mathbb{E}_{(x,y) \sim \mathcal{D}(\theta)} \left[\ \mathbb{1} \{ y \neq f_{\theta}(x) \} \ \mathbb{1} \{ x \in S(\Theta_{0},c) \} \right] - \mathbb{E}_{(x,y) \sim \tilde{\mathcal{D}}(\theta)} \left[\ \mathbb{1} \{ y \neq f_{\theta}(x) \} \ \mathbb{1} \{ x \in S(\Theta_{0},c) \} \right] \right| . \end{split}$$

We can break this into terms where y = 0 and terms where y = 1. Thus, it suffices to bound:

$$\left| \mathbb{E}_{(x,y) \sim \mathcal{D}(\theta;M)} [\ \mathbb{1}\{y \neq f_{\theta}(x)\} \ \mathbb{1}\{x \in S(\Theta_{0},c)\} \ \mathbb{1}\{y = 0\}] - \mathbb{E}_{(x,y) \sim \mathcal{D}(\theta;\tilde{M})} [\ \mathbb{1}\{y \neq f_{\theta}(x)\} \ \mathbb{1}\{x \in S(\Theta_{0},c)\} \ \mathbb{1}\{y = 0\}] \right|$$
 and

$$\left| \mathbb{E}_{(x,y) \sim \mathcal{D}(\theta; M)} [\ \mathbb{1}\{y \neq f_{\theta}(x)\} \ \mathbb{1}\{x \in S(\Theta_0, c)\} \ \mathbb{1}\{y = 1\}] \mathbb{E}_{(x,y) \sim \mathcal{D}(\theta; \tilde{M})} [\ \mathbb{1}\{y \neq f_{\theta}(x)\} \ \mathbb{1}\{x \in S(\Theta_0, c)\} \ \mathbb{1}\{y = 1\}] \right|.$$

It suffices to show that the first term is upper bounded by $A(\theta)$ and the second term is upper bounded by $B(\theta)$. Since these two bounds follow from analogous arguments, we only present the proof of the first bound.

$$\begin{split} & \left| \mathbb{E}_{(x,y) \sim \mathcal{D}(\theta;M)} [\ \mathbb{1}\{y \neq f_{\theta}(x)\} \ \mathbb{1}\{x \in S(\Theta_{0},c)\} \ \mathbb{1}\{y = 0\}] - \mathbb{E}_{(x,y) \sim \mathcal{D}(\theta;\tilde{M})} [\ \mathbb{1}\{y \neq f_{\theta}(x)\} \ \mathbb{1}\{x \in S(\Theta_{0},c)\} \ \mathbb{1}\{y = 0\}] \right| \\ & = \mathbb{P}[x \in S(\Theta_{0},c) \ \& \ y = 0] \left| \mathbb{E}_{(x,y) \sim \mathcal{D}_{S(\Theta_{0},c)}^{0}(\theta;M)} [\ \mathbb{1}\{p(x) \geq \theta\}] - \mathbb{E}_{(x,y) \sim \mathcal{D}_{S(\Theta_{0},c)}^{0}(\theta;\tilde{M})} [\ \mathbb{1}\{p(x) \geq \theta\}] \right| \\ & = \mathbb{P}[x \in S(\Theta_{0},c) \ \& \ y = 0] \left| \mathbb{E}_{l \sim \mathcal{D}_{S(\Theta_{0},c)}^{0}(\theta;M)} [\ \mathbb{1}\{l \geq \theta\}] - \mathbb{E}_{l \sim \mathcal{D}_{S(\Theta_{0},c)}^{p,0}(\theta;\tilde{M})} [\ \mathbb{1}\{l \geq \theta\}] \right| \\ & \leq \mathbb{P}[x \in S(\Theta_{0},c) \ \& \ y = 0] \mathrm{KS} \left(\mathcal{D}_{S(\Theta_{0},c)}^{p,0}(\theta;M), \mathcal{D}_{S(\Theta_{0},c)}^{0}(\theta;\tilde{M}) \right). \end{split}$$

Now, we are ready to prove Lemma 9.

Proof of Lemma 9. Let Θ_{net} be an ϵ net of Θ_0 . The decision-maker uses the agent response oracle as follows. For each $\theta \in \Theta_{\text{net}}$, they can generate n_0 samples as follows: draw a sample $(x,y) \sim \mathcal{D}_{XY}$ conditioned on y=0. If $x \in S(\Theta_0,c)$, then query the agent response oracle on x at θ . It is easy to see that these samples are distributed as n_0 independent samples from $\mathcal{D}^{p,0}_{S(\Theta_0,c)}(\theta)$. Similarly, the decision-maker uses the agent response oracle to draw n_1 samples that are distributed as n_1 independent samples from $\mathcal{D}^{p,1}_{S(\Theta_0,c)}(\theta)$. (We will specify the values of n_0 and n_1 later.)

First, we define a distribution map $\tilde{\mathcal{D}}$ using these samples and the base distribution. Let's define $D_0(\theta)$ to be the empirical distribution of the n_0 samples, and let $D_1(\theta)$ be the empirical distribution of the n_1 samples. Let $D'(\theta)$ be the distribution given by a mixture of (x,0) where $x \sim D_0(\theta)$ with probability $\mathbb{P}_{\mathcal{D}_{XY}}[y=0 \mid x \in S(\Theta_0,c)]$ and $x \sim D_1(\theta)$ with probability $\mathbb{P}_{\mathcal{D}_{XY}}[y=1 \mid x \in S(\Theta_0,c)]$. Let $D''(\theta)$ be the distribution given by (x,y) drawn from the conditional distribution of \mathcal{D}_{XY} given $x \notin S(\Theta_0,c)$. We let $\tilde{\mathcal{D}}$ be the distribution given by a mixture of $D'(\theta)$ with probability $\mathbb{P}_{\mathcal{D}_{XY}}[x \in S(\Theta_0,c)]$ and $D''(\theta)$ with probability $1 - \mathbb{P}_{\mathcal{D}_{XY}}[x \in S(\Theta_0,c)]$.

We can microfound $\tilde{\mathcal{D}}$ with a map \tilde{M} as follows. Let $\tilde{M}(x,y) = x$ when $x \notin S(\Theta_0,c)$. Let \tilde{M} on $S(\Theta_0,c) \times Y$ be defined in such any way it microfounds $D'(\theta)$ (this is possible because of Proposition

5). It is easy to see that \tilde{M} microfounds $\tilde{\mathcal{D}}$ and that \tilde{M} restricted to the domain $(X \setminus S(\Theta_0, c)) \times Y$ is expenditure-constrained. This means that we can apply Lemma C.1.

Now, we bound the performative risk $PR(\theta_{PO}(\tilde{M}))$, where:

$$\theta_{\text{PO}}(\tilde{M}) = \operatorname{argmin}_{\theta \in \Theta_0} \mathbb{E}_{(x,y) \sim \mathcal{D}(\theta;\tilde{M})} [\mathbb{1}\{y \neq f_{\theta}(x)\}] = \operatorname{argmin}_{\theta \in \Theta_{\text{not}}} \mathbb{E}_{(x,y) \sim \mathcal{D}(\theta;\tilde{M})} [\mathbb{1}\{y \neq f_{\theta}(x)\}].$$

In order to apply Lemma C.1, we need to bound:

$$\sup_{\theta \in \Theta_{\text{net}}} \{ A(\theta) + B(\theta) \} \tag{9}$$

where:

$$\begin{split} A(\theta) &:= \mathbb{P}[x \in S(\Theta_0, c) \& y = 0] \cdot \mathsf{KS}\left(\mathcal{D}^{p,0}_{S(\Theta_0, c)}(\theta; M), \mathcal{D}^{p,0}_{S(\Theta_0, c)}(\theta; \tilde{M})\right) \\ B(\theta) &:= \mathbb{P}[x \in S(\Theta_0, c) \& y = 1] \cdot \mathsf{KS}\left(\mathcal{D}^{p,1}_{S(\Theta_0, c)}(\theta; M), \mathcal{D}^{p,1}_{S(\Theta_0, c)}(\theta; \tilde{M})\right). \end{split}$$

To bound (9), we union bound over Θ_{net} . This set has cardinality $O(1/\epsilon)$. Notice that with probability $\geq 1 - \alpha$, we know that:

$$KS\left(\mathcal{D}_{S(\Theta_{0},c)}^{p,0}(\theta;M),\mathcal{D}_{S(\Theta_{0},c)}^{p,0}(\theta;\tilde{M})\right) \leq \sqrt{\frac{\ln(2/\alpha)}{2n_{0}}}.$$

$$KS\left(\mathcal{D}_{S(\Theta_{0},c)}^{p,1}(\theta;M),\mathcal{D}_{S(\Theta_{0},c)}^{p,1}(\theta;\tilde{M})\right) \leq \sqrt{\frac{\ln(2/\alpha)}{2n_{0}}}.$$

We can now set $\alpha = \Theta(\epsilon/100)$ in the previous result to obtain that with probability $\geq 99/100$, the expression in (9) is bounded by:

$$E := O\left(\mathbb{P}_{\mathcal{D}_{XY}}[x \in S(\Theta_0, c) \land y = 0]\sqrt{\frac{\ln(2/\epsilon)}{2n_0}}\right) + O\left(\mathbb{P}_{\mathcal{D}_{XY}}[x \in S(\Theta_0, c) \land y = 1]\sqrt{\frac{\ln(2/\epsilon)}{2n_1}}\right).$$

We can now apply Lemma C.1 to Θ_{net} to see that:

$$PR(\theta_{PO}(\tilde{M})) \le E + \min_{\theta \in \Theta_{rot}} \mathbb{E}_{(x,y) \sim \mathcal{D}(\theta;M)} [\mathbb{1}\{y \ne f_{\theta}(x)\}],$$

Now, let's use the Lipschitz requirement on the distribution map to move to the set Θ_0 . Let's consider a distribution map \mathcal{D}' that is defined as follows: for $\theta \in \Theta_{\text{net}}$, we take $\mathcal{D}'(\theta) := \mathcal{D}(\theta)$, and for $\theta \notin \Theta_{\text{net}}$, we take $\mathcal{D}'(\theta) := \mathcal{D}(\theta')$ where θ' is the closest element in Θ_{net} to θ . Now, let's apply Lemma D.2 to \mathcal{D} and \mathcal{D}' on Θ_0 to obtain that:

$$\begin{split} \min_{\theta \in \Theta_{\text{net}}} \mathbb{E}_{(x,y) \sim \mathcal{D}(\theta;M)} \left[\mathbb{1} \{ y \neq f_{\theta}(x) \} \right]) &\leq \epsilon + \min_{\theta \in \Theta_{0}} \mathbb{E}_{(x,y) \sim \mathcal{D}(\theta;M)} \left[\mathbb{1} \{ y \neq f_{\theta}(x) \} \right] \\ &= \epsilon + \min_{\theta \in \Theta} \mathbb{E}_{(x,y) \sim \mathcal{D}(\theta;M)} \left[\mathbb{1} \{ y \neq f_{\theta}(x) \} \right]. \end{split}$$

This means that

$$PR(\theta_{PO}(\tilde{M})) \le E + \epsilon + PR(\theta_{PO}(M)).$$

Thus, it suffices to bound E and set n_0 and n_1 appropriately. Suppose that

$$n_0 = \Theta\left(\mathbb{P}_{\mathcal{D}_{XY}}[x \in S(\Theta_0, c) \& y = 0]^2 \frac{\ln(1/\epsilon)}{2\epsilon^2}\right)$$

and

$$n_1 = \Theta\left(\mathbb{P}_{\mathcal{D}_{XY}}[x \in S(\Theta_0, c) \& y = 1]^2 \frac{\ln(1/\epsilon)}{2\epsilon^2}\right).$$

Plugging in these expressions into the expression for E, we obtain the desired bounds. Moreover, notice that the total number of queries to the oracle is $\Theta(1/\epsilon) \cdot (n_0 + n_1) \leq \Theta\left(\mathbb{P}_{\mathcal{D}_{XY}}[x \in S(\Theta_0, c)]^2 \frac{\ln(1/\epsilon)}{2\epsilon^3}\right)$.

D Proofs for Section 4

D.1 Proof of Theorem 10

Proof of Theorem 10. We use the following notation for this proof. Let's extend the cost function to be defined and valid on all of \mathbb{R} rather than just X. For $x \in X$, let's use the notation $l_x \in \mathbb{R}$ to denote the unique value such that $l_x < x$ and $c(l_x, x) = 1$. Similarly, let $u_x \in \mathbb{R}$ denote the unique value such that $u_x > x$ and $c(x, u_x) = 1$. These values are unique by the definition of a valid cost function.

Fix $\sigma \in (0, \infty)$, and $x' \in X$. Let's characterize the agents who will change their features to x' when the threshold is θ . Either the agents' true features are equal to x' and their perception function $P(\theta) \notin (x', u_{x'}]$, or the agents' perception function $P(\theta) = x'$ and their true features x are in $[l_{x'}, x']$. Since the base distribution and the noise distribution are continuous, this means that there are no point masses in the distribution. To see that a probability density function exists everywhere and is continuous, let's compute the density. Let p_{base} denote the pdf of the base distribution (which is assumed to exist and be continuous since \mathcal{D}_{XY} is a continuous distribution), and let p_{noise} denote the pdf of D (which is continuous since it is the pdf of a gaussian). Notice that the probability density of $\mathcal{D}(\theta)$ at (x', y') is

$$p_{\text{base}}((x',y')) \cdot \mathbb{P}_D[\eta \notin (x'-\theta,u_{x'}-\theta)] + p_{\text{noise}}(x'-\theta) \cdot \mathbb{P}_{\mathcal{D}_{xy}}[x \in [l_{x'},x'],y=y'].$$

This is continuous in x' because $u_{x'}$ and $l_{x'}$ are continuous in x'. Moreover, this is nonzero on all x' because for all $x' \in X$, we see that $p_{\text{base}}((x',y')) > 0$ and $\mathbb{P}_D[\eta \notin (x'-\theta,u_{x'}-\theta)] > 0$ as well.

Now, we show aggregate smoothness. We see that the probability density $p_{\theta}((x', y'))$ at (x', y') is continuous in x' because each term is continuous in x'. Similarly, we see that this is continuous in θ because each term is continuous in θ . By Proposition 6, this implies aggregate smoothness. \square

D.2 Proof of Lemma 11

Proof of Lemma 11. For a given classifier f_{θ} , consider the product distribution $\mathcal{D}_{\text{prod}}(\sigma,\theta)$ of the base distribution \mathcal{D}_{XY} and the multivariate gaussian distribution $\mathcal{N}(\theta,\sigma\cdot\mathbf{I})$. This is a distribution over $X\times\mathbb{R}^d$ that will describe the distribution over noise vectors and features vectors. That is, an agent A with features x_A and noise η_A corresponds to $(x_A,\theta+\eta_A)$. If we apply the function in (6) to $\mathcal{D}_{\text{prod}}$ so that $(x_A,\theta_A')\mapsto \arg\max_{x'\in\mathbb{R}^d}\left[\gamma\cdot f_{\theta_A'}(x')-c(x_A,x')\right]$, then it is easy to see that we obtain the distribution $\mathcal{D}(\theta)$.

Since the total variation distance can only decrease when we apply a function to the distributions, we know that $TV(\mathcal{D}(\theta_1), \mathcal{D}(\theta_2)) \leq TV(\mathcal{D}_{prod}(\sigma, \theta_1), \mathcal{D}_{prod}(\sigma, \theta_2))$. Thus, it suffices to bound $TV(\mathcal{D}_{prod}(\sigma, \theta_1), \mathcal{D}_{prod}(\sigma, \theta_2))$. Using the properties of product distributions, we see that:

$$\begin{split} \text{TV}(\mathcal{D}_{\text{prod}}(\sigma, \theta_1), \mathcal{D}_{\text{prod}}(\sigma, \theta_2)) &\leq \text{TV}(\mathcal{N}(\theta_1, \sigma \cdot \mathbf{I}), \mathcal{N}(\theta_2, \sigma \cdot \mathbf{I})) \\ &= \text{TV}(\mathcal{N}(\theta_1/\sigma, \mathbf{I}), \mathcal{N}(\theta_2/\sigma, \mathbf{I})) \\ &\leq \frac{1}{\sqrt{2}\sigma} \|\theta_1 - \theta_2\|_2 \,. \end{split}$$

Now, the result follows from a bound on the total variation distance between two multivariate gaussians (i.e. see Corollary 2.13 in [Diakonikolas et al., 2019]).

D.3 Social burden of noisy responses in general

We show that for any valid cost function, noisy responses results in an optimal point with no higher social burden than the optimal point deduced from standard microfoundations.

Proposition D.1. Let $\sigma \in (0, \infty)$, and let c be a valid cost function. Consider a 1-dimensional setting where $X \subseteq \mathbb{R}$ and Θ is a model class of threshold functions. Then, the following holds:

$$\theta_{PO}(M_{SM}) \ge \theta_{PO}(M_{\sigma})$$

Burden $(\theta_{PO}(M_{SM})) \ge$ Burden $(\theta_{PO}(M_{\sigma}))$,

where M_{SM} is the mapping induced by standard microfoundations.

Proof. By Proposition B.6, it suffices to show that M_{σ} satisfies expenditure monotonicity and Assumption 3. The fact that M_{σ} satisfies Assumption 3 follows from its definition. For expenditure monotonicity, note that the first condition follows from the fact that the optimization problem in (6) tells us that fuzzy perception agents never exceed their utility of a positive outcome from manipulation expenditure. We now show that the second condition is satisfied. Note that each agents' perception function takes the form $P(\theta) = \theta + \eta$ for some fixed η . Thus, any given agent either consistently overshoots or consistently undershoots the threshold. If $\eta < 0$, then the agent will only be positively classified if and only if $\theta \le x$ where x are the agent's true features. If $\eta > 0$, then the agent will be positively classified if and only if $c(x, \theta + \eta) \le 1$ or $\theta \le x$. This proves the desired statement.

D.4 Proof of Proposition 12

Proof of Proposition 12. By Proposition D.1, we see that $\theta_{PO}(\mathcal{D}_{SM}) \geq \theta_{PO}(\mathcal{D})$. It thus suffices to show that $\theta_{PO}(\mathcal{D}_{SM}) > \theta_{PO}(\mathcal{D})$. To show this, it suffices to show that the derivative of the performative risk exists and is nonzero at $\theta_{PO}(\mathcal{D}_{SM})$. Like in the proof of Theorem 10, we use the notation l_x , u_x , p_{base} , and \mathbb{P}_D . We write the performative risk in a form that is more convenient to differentiate.

$$\begin{split} \operatorname{PR}(\theta) &= \int_{\theta}^{\infty} p_{\operatorname{base}}((x,0)) dx + \int_{-\infty}^{l_{\theta}} p_{\operatorname{base}}((x,0)) dx + \int_{\theta-1}^{\theta} p_{\operatorname{base}}((x',0)) \mathbb{P}_{D}[\eta \in [0,u_{x'}-\theta]] dx' \\ &+ \int_{\theta-1}^{\theta} p_{\operatorname{base}}((x',1)) \mathbb{P}_{D}[\eta \in [0,u_{x'}-\theta]] dx' \\ &= \int_{\theta}^{\infty} p_{\operatorname{base}}((x,0)) dx + \int_{-\infty}^{\theta-1} p_{\operatorname{base}}((x,0)) dx + \int_{\theta-1}^{\theta} p_{\operatorname{base}}((x',0)) \mathbb{P}_{D}[\eta \in [0,x'+1-\theta]] dx' \\ &+ \int_{\theta-1}^{\theta} p_{\operatorname{base}}((x',1)) \mathbb{P}_{D}[\eta \in [0,x'+1-\theta]] dx' \\ &= \int_{\theta}^{\infty} p_{\operatorname{base}}((x,0)) dx + \int_{-\infty}^{\theta-1} p_{\operatorname{base}}((x,0)) dx + \int_{0}^{1} p_{\operatorname{base}}((\theta-1+x,0)) \mathbb{P}_{D}[\eta \in [0,x] dx \\ &+ \int_{0}^{1} p_{\operatorname{base}}((\theta-1+x,1)) \mathbb{P}_{D}[\eta \notin [0,x] dx \\ &= \int_{\theta}^{\infty} p_{\operatorname{base}}((x,0)) dx + \int_{-\infty}^{\theta-1} p_{\operatorname{base}}((x,0)) dx + \mathbb{P}_{\mathcal{D}_{XY}}[x \in (\theta-1,\theta), y=0] \\ &- \int_{0}^{1} (p_{\operatorname{base}}((\theta-1+x,1)) - p_{\operatorname{base}}((\theta-1+x,0))) \mathbb{P}_{D}[\eta \notin [0,x]] dx. \end{split}$$

Let's write $\int_0^1 (p_{\text{base}}((\theta-1+x,1)) - p_{\text{base}}((\theta-1+x,0))) \mathbb{P}_D[\eta \notin [0,x]] dx$ in a slightly different form.

$$\begin{split} &\int_{0}^{1}(p_{\text{base}}((\theta-1+x,1))-p_{\text{base}}((\theta-1+x,0)))\mathbb{P}_{D}[\eta\notin[0,x]]dx\\ &=(\mathbb{P}_{D}[\eta\in[-\infty,0]]+\mathbb{P}_{D}[\eta\in[1,\infty]])\int_{0}^{1}(p_{\text{base}}((\theta-1+x,1))-p_{\text{base}}((\theta-1+x,0)))dx\\ &+\int_{0}^{1}(p_{\text{base}}((\theta-1+x,1))-p_{\text{base}}((\theta-1+x,0)))\mathbb{P}_{D}[\eta\in[x,1]]dx\\ &=(\mathbb{P}_{D}[\eta\in[-\infty,0]]+\mathbb{P}_{D}[\eta\in[1,\infty]])(\mathbb{P}_{\mathcal{D}_{XY}}[x\in(\theta-1,\theta),y=1]-\mathbb{P}_{\mathcal{D}_{XY}}[x\in(\theta-1,\theta),y=0])\\ &+\int_{0}^{1}(p_{\text{base}}((\theta-1+x,1))-p_{\text{base}}((\theta-1+x,0)))\mathbb{P}_{D}[\eta\in[x,1]]dx. \end{split}$$

We can rewrite:

$$\begin{split} &\int_{0}^{1} (p_{\text{base}}((\theta - 1 + x, 1)) - p_{\text{base}}((\theta - 1 + x, 0))) \mathbb{P}_{D}[\eta \in [x, 1]] dx \\ &= \int_{0}^{1} \int_{x}^{1} (p_{\text{base}}((\theta - 1 + x, 1)) - p_{\text{base}}((\theta - 1 + x, 0))) p_{\text{noise}}(z) dz dx \\ &= \int_{0}^{1} p_{\text{noise}}(z) \int_{0}^{z} (p_{\text{base}}((\theta - 1 + x, 1)) - p_{\text{base}}((\theta - 1 + x, 0))) dx dz \\ &= \int_{0}^{1} p_{\text{noise}}(z) (\mathbb{P}_{\mathcal{D}_{XY}}[x \in ((\theta - 1, \theta - 1 + z)), y = 1] - \mathbb{P}_{\mathcal{D}_{XY}}[x \in ((\theta - 1, \theta - 1 + z))), y = 0]) dz \end{split}$$

When we take a derivative with respect to θ , we obtain:

$$\begin{split} \frac{\partial \text{PR}(\theta)}{\partial \theta} &= -p_{\text{base}}((\theta, 0)) + p_{\text{base}}((\theta - 1, 1)) - p_{\text{base}}((\theta - 1, 0)) + p_{\text{base}}((\theta, 0)) \\ &+ (\mathbb{P}_D[\eta \in [-\infty, 0]] + \mathbb{P}_D[\eta \in [1, \infty]])(p_{\text{base}}((\theta, 1)) - p_{\text{base}}((\theta - 1, 1)) - p_{\text{base}}((\theta, 0)) + p_{\text{base}}((\theta - 1, 0))) \\ &+ \int_0^1 p_{\text{noise}}(z)(p_{\text{base}}((\theta - 1 + z, 1)) - p_{\text{base}}((\theta - 1, 1)) - p_{\text{base}}((\theta - 1 + z, 0)) + p_{\text{base}}((\theta - 1, 0))) dz. \end{split}$$

Let's analyze this expression at $\theta = \theta_{PO}(\mathcal{D}_{SM})$. By the assumptions on the cost function, and using that $\theta_{SL} + 1 \in \Theta \cap X$, we see that $\theta_{PO}(\mathcal{D}_{SM}) = \theta_{SL} + 1$, so $\theta - 1 = \theta_{SL}$. This means that $p_{base}((\theta - 1, 1)) - p_{base}((\theta - 1, 0)) = p_{base}((\theta_{SL}, 1)) - p_{base}((\theta_{SL}, 0)) = 0$. Thus, the expression simplifies to:

$$\frac{\partial \text{PR}(\theta)}{\partial \theta} = (\mathbb{P}_D[\eta \in [-\infty, 0]] + \mathbb{P}_D[\eta \in [1, \infty]])(p_{\text{base}}((\theta, 1)) - p_{\text{base}}((\theta, 0)))$$
$$+ \int_0^1 p_{\text{noise}}(z)(p_{\text{base}}((\theta - 1 + z, 1)) - p_{\text{base}}((\theta - 1 + z, 0)))dz.$$

We see that $p_{\text{base}}((\theta',1)) > p_{\text{base}}((\theta',0))$ for all $\theta' \geq \theta_{\text{SL}}$ by the assumption on μ in Setup 1. This implies that the first term is positive and the second term is nonnegative, so $\frac{\partial PR(\theta)}{\partial \theta}$ is positive as desired.

D.5 Proof of Lemma 13

Proof of Lemma 13. Consider the product distribution $\mathcal{D}_{prod}(\sigma)$ of the base distribution \mathcal{D}_{XY} and the multivariate Gaussian distribution $\mathcal{N}(\mathbf{0}, \sigma \cdot \mathbf{I})$. This is a distribution over $X \times \mathbb{R}^d$ that will describe the distribution over noise vectors and features vectors. That is, an agent a with features x_a and noise η_A corresponds to (x_A, η_A) . If we apply the function in (6) to \mathcal{D}_{prod} so that $(x_A, \eta_A) \mapsto \arg\max_{x' \in \mathbb{R}^d} \left[\gamma \cdot f_{\theta + \eta_A}(x') - c(x_A, x') \right]$, then it is easy to see that we obtain the distribution $\mathcal{D}(\theta)$. Since the total variation distance can only decrease when we apply a function to the dis-

Since the total variation distance can only decrease when we apply a function to the distributions, we know that $TV(\mathcal{D}(\theta), \mathcal{D}(\tilde{\theta}) \leq TV(\mathcal{D}_{prod}(\sigma), \mathcal{D}_{prod}(\tilde{\sigma}))$. Thus, it suffices to bound $TV(\mathcal{D}_{prod}(\sigma), \mathcal{D}_{prod}(\tilde{\sigma}))$. Using the properties of product distributions, we see that

$$\mathrm{TV}(\mathcal{D}_{\mathrm{prod}}(\sigma),\mathcal{D}_{\mathrm{prod}}(\tilde{\sigma})) \leq \mathrm{TV}(\mathcal{N}(\mathbf{0},\sigma\cdot\mathbf{I}),\mathcal{N}(\mathbf{0},\tilde{\sigma}\cdot\mathbf{I})).$$

Now, the result follows from the standard bound on the total variation distance between two multivariate gaussians.

D.6 Proof of Corollary 14

In order to prove Corollary 14, we show that if the estimated distribution map is sufficiently close to the true distribution map, then the optimal point computed using this model will achieve near-optimal performative risk.

Lemma D.2. Let \tilde{M} be an estimate of the true distribution map M. Then the suboptimality of the performative risk of $\theta_{PO}(M)$ as per (4) is bounded by: $PR(\theta_{PO}(\tilde{M})) - PR(\theta_{PO}(M)) \le 2 \sup_{\theta} \{TV(\mathcal{D}(\theta;M), \mathcal{D}(\theta;\tilde{M}))\}$, where $PR(\theta) := \mathbb{E}_{(x,v) \sim \mathcal{D}(\theta)}[\mathbb{1}\{y \neq f_{\theta}(x)\}]$ denotes the performative risk with respect to M.

Proof of Lemma D.2. Let $\xi = \{ TV(\mathcal{D}(\theta; M), \mathcal{D}(\theta; \tilde{M})) \}$. Let $PR(\theta; M)$ denote the performative risk at θ on $\mathcal{D}(\theta; M)$ and let $PR(\theta; \tilde{M})$ denote the performative risk at θ on $\mathcal{D}(\theta; \tilde{M})$. It suffices to show that $|PR(\theta; M) - PR(\theta; \tilde{M})| \le \xi$ (since this would mean that $PR(\theta_{PO}(\tilde{M}); M) \le PR(\theta_{PO}(\tilde{M}); \tilde{M}) + \xi \le PR(\theta_{PO}(M); \tilde{M}) + \xi \le PR(\theta_{PO}(M); M) + 2\xi$, as desired). Notice that:

$$|\operatorname{PR}(\theta; M) - \operatorname{PR}(\theta; \tilde{M})| = \left| \mathbb{E}_{(x, y) \sim \mathcal{D}(\theta; M)} [\mathbb{1}\{y \neq f_{\theta}(x)\}] - \mathbb{E}_{(t, x, y) \sim \mathcal{D}(\theta; \tilde{M})} [\mathbb{1}\{y \neq f_{\theta}(x)\}] \right|.$$

Since the indicator variables are always constrained between 0 and 1, we can immediately obtain an upper bound of $TV(\mathcal{D}(\theta;M),\mathcal{D}(\theta;\tilde{M}))$.

We can easily deduce Corollary 14 from these above facts.

Proof of Corollary 14. We use Lemma 13 to see that

$$\mathrm{TV}(\mathcal{D}(\theta; M_\sigma), \mathcal{D}(\theta, M_{\tilde{\sigma}})) \leq \frac{3}{2} \sqrt{\frac{\epsilon m}{\min(\sigma^2, \tilde{\sigma}^2}}.$$

Then we Lemma D.2 to see that:

$$\Pr(\theta_{\text{PO}}(M_{\tilde{\sigma}})) - \Pr(\theta_{\text{PO}}(M_{\sigma})) \leq 2\sup_{\theta} \left\{ \text{TV} \Big(\mathcal{D}(\theta; M_{\sigma}), \mathcal{D}(\theta; M_{\tilde{\sigma}}) \Big) \right\} \leq 3\sqrt{\frac{\epsilon m}{\min(\sigma^2, \tilde{\sigma}^2}}.$$

E Lipschitzness in Wasserstein distance

In existing performative prediction approaches [Perdomo et al., 2020; Mendler-Dünner et al., 2020; Brown et al., 2020], it is assumed that the distribution map is Lipschitz with respect to changes in θ . In particular,

$$\mathcal{W}(\mathcal{D}(\theta), \mathcal{D}(\theta')) \le C \|\theta - \theta'\|_{2}$$

for some constant C > 0, where W is the Wasserstein-1 distance.

E.1 Lipschitzness is not sufficient for stability

We show that for binary classification, Lipschitzness in Wasserstein distance is not sufficient to guarantee the existence of stable points. We construct a simple example of a Lipschitz distribution map for which stable points do not exist. Consider our counterexample in Setup 1 with $p \in (0,1)$: let the cost function of the best responding agents be linear, and consider a uniform base distribution; this results in a Lipschitz distribution map. By Proposition 2, stable points do not exist.

E.2 Lipschitzness can be restrictive

In the context of binary classification, we show that this requirement is quite restrictive on the cost function and the base distribution, even in the context of a 1-d setting and standard microfoundations. Example 3 provides a simple 1-d setting where the aggregate response distribution $\mathcal{D}(\theta)$ induced by best-responding agents does not satisfy Lipschitzness:

Example 3. Suppose that $\Theta = [0,1]$ and X = [-10,10]. Let the marginal distribution of the features for \mathcal{D}_{XY} be uniform on [0,1] and consider the cost function $c(x,y) = |x^2 - y^2|$. The distribution map $\mathcal{D}_{NS}(\theta)$ given by agents who follow standard microfoundations is not Lipschitz in Wasserstein distance.

To prove Example 3, we show the following lemma.

Lemma E.1 (Violation of Lipschitzness). Assume that Θ is 1-dimensional. Suppose that $c(x,y) \le k|x-y|$ for some constant k > 0. For each $\theta \in \Theta$, let S_{θ} be the salient part. If the following condition holds:

$$\sup_{\theta'>\theta>0} \frac{\mathbb{P}_{\mathcal{D}_{XY}}[S_{\theta} \setminus S_{\theta'}]}{|\theta'-\theta|} = \infty,$$

then $\mathcal{D}(\theta)$ is not Lipschitz in Wasserstein distance.

Proof of Lemma E.1. First, we show that when $\theta' \leq \theta - 1/k$ and $\theta' > \theta > 0$, it holds that $\mathcal{W}(\mathcal{D}(\theta), \mathcal{D}(\theta')) \geq \frac{1}{2k} \mathbb{P}_{\mathcal{D}_{XY}}[Q_{\theta'} \setminus Q_{\theta}]$. The distribution map $\mathcal{D}(\theta)$ corresponds to \mathcal{D}_{XY} with all of the agents with type $x \in Q_{\theta}$ moving up to θ . This means that $W(\mathcal{D}(\theta'), \mathcal{D}(\theta))$ must move all of the agents with features $x \in Q_{\theta'} \setminus Q_{\theta}$ up to at least θ . Notice that $\min_{x \in Q_{\theta'} \setminus Q_{\theta}} |x - \theta| = |\sup(Q_{\theta'}) - \theta| \geq |\sup(Q_{\theta'}) - \theta'| - |\theta' - \theta| \geq \frac{1}{k} c(\sup(Q_{\theta'}), \theta') - |\theta' - \theta| \geq \frac{1}{k} - |\theta' - \theta| \geq \frac{1}{2k}$. This means that $W(\mathcal{D}(\theta), \mathcal{D}(\theta')) \geq \frac{1}{2k} \mathbb{P}_{\mathcal{D}_{XY}}[Q_{\theta'} \setminus Q_{\theta}]$, as desired.

Now, suppose that the condition $\sup_{\theta'>\theta>0}\frac{\mathbb{P}_{\mathcal{D}_{XY}}[Q_{\theta'}\setminus Q_{\theta}]}{|\theta'-\theta|}=\infty$ holds. Notice that this implies that $\sup_{\theta'>\theta>0,\theta'\leq\theta-1/k}\frac{\mathbb{P}_{\mathcal{D}_{XY}}[Q_{\theta'}\setminus Q_{\theta}]}{|\theta'-\theta|}=\infty$. This implies that $\sup_{\theta'>\theta>0,\theta'\leq\theta-1/k}\frac{\mathcal{W}(\mathcal{D}(\theta),\mathcal{D}(\theta'))}{|\theta'-\theta|}=\infty$, and so the Lipschitzness constraint is violated.

Now, we prove Example 3 from Lemma E.1.

Proof of Example 3. We apply Lemma E.1. First, observe that $c(x,y) = |x^2 - y^2| = |x - y||x + y| \le 20|x - y|$ as desired. Thus, it suffices to show that

$$\sup_{\theta'>\theta>0} \frac{\mathbb{P}_{\mathcal{D}_{XY}}[S_{\theta} \setminus S_{\theta'}]}{|\theta'-\theta|} = \infty.$$

Let's take $\theta=1$ and $\theta'=1+\epsilon$. Because of the uniform density, it suffices to show that $\lim_{\epsilon \to 0} \frac{|S_1 \setminus S_{1+\epsilon}|}{\epsilon} = \infty$. Notice that for sufficiently small ϵ , we see that $S_{\theta} \setminus S_{\theta'} = [0, \sqrt{\epsilon^2 + 2\epsilon}]$, and so $|S_{\theta} \setminus S_{\theta'}| = \sqrt{\epsilon^2 + 2\epsilon}$. We see that $\frac{|S_{\theta} \setminus S_{\theta'}|}{\theta' - \theta} = \frac{\sqrt{\epsilon^2 + 2\epsilon}}{\epsilon} = 1 + \frac{2}{\sqrt{\epsilon}}$. This approaches ∞ as $\epsilon \to 0$, as desired.

F Additional results on the implications of the expenditure constraint

In Appendix F.1, we explicitly construct Θ_0 and $S(\Theta_0, c)$ in an example setting. In Appendix F.2, we demonstrate that our insights about how the expenditure constraint results can reduce empirical burden extend to more general model classes.

F.1 Example construction of Θ_0 and $S(\Theta_0, c)$

In this section, we discuss an example setting where the decision-maker can prune the search space Θ_0 using the expenditure constraint, and the resulting set $S(\Theta_0, c)$ is much smaller than X. We first present an informal version of this example:

Example 4 (Informal). Consider Setup 1. Define the cost function to be linear: $c(x, x') = \alpha |x - x'|$ for some $\alpha > 0$. Then, if M satisfies the expenditure constraint and Assumption 3, then:

$$\theta_{\mathrm{PO}} \in \Theta_0 := \left[\theta_{\mathrm{SL}} - 3/\alpha, \theta_{\mathrm{SL}} + 3/\alpha\right] \quad S(\Theta_0, c) = \left\{x \mid \mu(x) \in \left[\theta_{\mathrm{SL}} - 4/\alpha, \theta_{\mathrm{SL}} + 4/\alpha\right]\right\},$$

where θ_{SL} is such that $\mu(\theta) = 0.5$ (where μ is defined in Setup 1).

Example 4 demonstrates the salient part consists of agents who are sufficiently close to the supervised learning threshold, where closeness is measured by the cost function. As a result, the salient part shrinks as costs increase.

Let's now formally construct $\Theta_0 := [l, u]$. We first define l: let $s' < \theta_{SL}$ be the value such that $c(\theta_{SL}, s') = 1$, let s'' < s' be the value such that c(s'', s') = 1, and let l < s be the value such that c(s'', l) = 1. We define u similarly: let $t' > \theta_{SL}$ be the value such that $c(\theta_{SL}, t') = 1$, let t'' > t be the value such that c(t'', t') = 1, and let u > t'' be the value such that c(t'', u) = 1. (To be precise, if we ever reach a stage when defining l where no such point exists, then we take l = 0; similarly, if we ever reach a stage when defining u where no such point exists, then we take u = 1.)

Proposition F.1. Suppose that the agent response types $T \subseteq \mathcal{T}$ are expenditure-constrained with respect to cost function c that is an outcome-valid cost function. The set Θ_0 (defined above) contains a performatively optimal point θ_{PO} of \mathcal{D}_{TXY} .

Proof. Since $s', t' \in \Theta_0$, it suffices to show that $PR(t') \leq PR(u)$ and $PR(s') \leq PR(l)$.

First, we show that $PR(t') \leq PR(u)$. For both classifiers, by the expenditure constraint, all agents with true features x such that $x < \theta_{SL}$ will necessarily be classified as 0 by both $f_{t'}$ and f_u . Thus, we only need to consider x such that $x \geq \theta_{SL}$. By the expenditure constraint, all agents with features x such that x > t'' will necessarily be classified as 1 by $f_{t'}$. By Assumption 3, coupled with the fact that $x \geq \theta_{SL}$ for these agents, the classifier f_u cannot achieve a better loss for these agents. We can ignore agents with x = t'' since these agents form a measure 0 set and \mathcal{D}_{XY} is continuous. For agents

with true features x such that $\theta_{SL} \le x < t''$, notice that these agents will necessarily be classified as 0 by f_u due to the expenditure constraint. Thus, by Assumption 3, coupled with the fact that $x \ge \theta_{SL}$ for these agents, the classifier $f_{t'}$ will not achieve a worse loss for these agents than f_u .

We use a similar argument to show that $PR(s') \leq PR(l)$. For both classifiers, by the expenditure constraint, all agents with true features x such that $x \geq \theta_{SL}$ will necessarily be classified as 0 by both $f_{t'}$ and f_u . Thus, we only need to consider x such that $x < \theta_{SL}$. By the expenditure constraint, all agents with features x such that x < s'' will necessarily be classified as 0 by $f_{s'}$. By Assumption 3, coupled with the fact that $x \leq \theta_{SL}$ for these agents, the classifier f_l cannot achieve a better loss for these agents. We can ignore agents with x = s'' since these agents form a measure 0 set and \mathcal{D}_{XY} is continuous. For agents with true features x such that $x'' < x \leq \theta_{SL}$, notice that these agents will necessarily classified as 1 by f_l due to the expenditure constraint, and so by Assumption 3, coupled with the fact that $x \leq \theta_{SL}$ for these agents, the classifier $f_{s'}$ will not achieve a worse loss for these agents.

We now show how to formally construct $S(\Theta_0, c)$. Let's construct a set $S'(\Theta_0, c) := [l', u']$ as follows. We can l' and u' in terms of l and u (the upper and lower endpoints of Θ_0). Let l' < l be the value such that c(l', l) = 1 and let u' > u be the value such that c(u, u') = 1. Then $S'(\Theta_0, c) = [l', u']$. (To be precise, if no such l' exists, then we take l = 0; similarly, if no such u' exists, then we take u' = 1.) We now show that $S'(\Theta_0, c) = S(\Theta_0, c)$.

Proposition F.2. Let $S'(\Theta_0, c)$ be defined as above. Then, $S(\Theta_0, c) = S'(\Theta_0, c)$.

Proof. It suffices to show that $S'(\Theta_0, c) = \bigcup_{\theta} S_{\theta}$, where $S_{\theta} = ...$

Suppose that $x \in S_{\theta}$. Then there exists x' such that $c(x,x') \le 1$ and $f_{\theta}(x') \ne f_{\theta}(x)$ (which is equivalent to θ is between x and x'). Using Assumption 1, this implies that $c(x,\theta) \le 1$. Using Assumption 1 again, we see that since $\theta \in [l,u]$, this means that either $x \in [l,u]$, or $c(x,l) \le 1$, or $c(x,u) \le 1$ must be true. This implies that $x \in S'(\Theta_0,c)$.

Suppose that $x \in S'(\Theta_0, c)$. If $x \in \Theta_0$, then we know that $x \in S_x$. If $p(x) \notin \Theta_0$ and x < l, then we know $x \in S_l$. If $x \notin \Theta_0$ and x > u, then we know that $x \in S_u$.

F.2 General model classes

While we focused on the 1-dimensional setting in Section 3.3, we now demonstrate that the expenditure constraint can reduce empirical burden on the decision-maker in general settings. The following lemma formalizes the intuition that an appropriate estimate of agent's response types on $S(\Theta_0,c)$ is sufficient to achieve near-optimal performative risk. We use a subscript notation $\mathcal{D}_{S(\Theta_0,c)}(\theta)$ to denote the aggregate response distribution $\mathcal{D}(\theta)$ restricted to agents with true features $x \in S(\Theta_0,c) \subseteq X$.

Lemma F.3. Let c be a valid cost function, let M, \tilde{M} be mappings that satisfy the expenditure constraint. Then, for any $\Theta_0 \subseteq \Theta : \theta_{PO}(M) \in \Theta_0$, it holds that:

$$\mathrm{PR}(\theta_{\mathrm{PO}}(M)) \leq \mathrm{PR}(\theta_{\mathrm{PO}}(\tilde{M})) + 2\xi$$

with
$$\xi := \sup_{\theta} \left\{ \mathbb{P}_{\mathcal{D}_{XY}} \left[x \in S(\Theta_0, c) \right] \cdot \mathrm{TV} \left(\mathcal{D}_{S(\Theta_0, c)}(\theta; M), \mathcal{D}_{S(\Theta_0, c)}(\theta; \tilde{M}) \right) \right\}$$

In words, to achieve a small error ξ , we need to accurately estimate the distribution map for all $x \in S(\Theta_0, c)$, and the larger $S(\Theta_0, c)$, the more accurate the estimate must be.

We prove Lemma F.3.

Proof of Lemma F.3. Our starting point is similar to the proof of Lemma D.2. Like in that proof, let $PR_{\mathcal{D}_{TXY}}(\theta)$ denote the performative risk at θ on \mathcal{D}_{TXY} and let $PR_{\tilde{\mathcal{D}}_{TXY}}(\theta)$ denote the performative risk at θ on $\tilde{\mathcal{D}}_{TXY}$. As in that proof, it suffices to show that $|PR_{\mathcal{D}_{TXY}}(\theta) - PR_{\tilde{\mathcal{D}}_{TXY}}(\theta)| \leq \xi$. At θ , notice that this difference is:

$$\left| \mathbb{E}_{(t,x,y) \sim \mathcal{D}_{\mathcal{T}XY}} \left[\mathbb{1} \left\{ 1 = f_{\theta}(R_t(x,\theta)) \right\} \right] - \mathbb{E}_{(t,x,y) \sim \tilde{\mathcal{D}}_{\mathcal{T}XY}} \left[\mathbb{1} \left\{ 1 = f_{\theta}(R_t(x,\theta)) \right\} \right] \right|.$$

First, we note that we only need to show that $|PR_{\mathcal{D}_{TXY}}(\theta) - PR_{\tilde{\mathcal{D}}_{TXY}}(\theta)| \le \xi$ for $\theta \in \Theta_0$. This is because we know that $\theta_{PO} \in \Theta_0$, and also the decision maker is only searching within Θ_0 .

Now, we claim that for any agent (t,x) where $x \notin S(\Theta_0,c)$ and for $t \in \text{supp}(\mathcal{D}_{TXY}) \cup \text{supp}(\tilde{\mathcal{D}}_{TXY})$, it holds that $f_{\theta}(R_t(x,\theta)) = f_{\theta}(x)$ for every $\theta \in \Theta_0$. Note that since t is expenditure-constrained with respect to c, then we know that if $x \notin S_{\theta}$, it holds that $f_{\theta}(R_t(x,\theta)) = f_{\theta}(x)$. Moreover, note that since $S_{\theta} \subseteq S(\Theta_0,c)$ by definition, this yields the desired statement.

Thus, we have that:

$$\begin{split} & \left| \mathbb{E}_{(t,x,y) \sim \mathcal{D}_{TXY}} [\ \mathbb{I}\{y = f_{\theta}(R_{t}(x,\theta))\}] - \mathbb{E}_{(t,x,y) \sim \tilde{\mathcal{D}}_{TXY}} [\ I\{y = f_{\theta}(R_{t}(x,\theta))\}] \right| \\ & = \left| \mathbb{E}_{(t,x,y) \sim \mathcal{D}_{TXY}} [\ \mathbb{I}\{y = f_{\theta}(R_{t}(x,\theta))\} \ \mathbb{I}\{x \in S(\Theta_{0},c)\}] - \mathbb{E}_{(t,x,y) \sim \tilde{\mathcal{D}}_{TXY}} [\ \mathbb{I}\{y = f_{\theta}(R_{t}(x,\theta))\} \ \mathbb{I}\{x \in S(\Theta_{0},c)\}] \right| \\ & = \left| \mathbb{E}_{(x,y) \sim \mathcal{D}(\theta)} [\ \mathbb{I}\{y = f_{\theta}(x)\} \ \mathbb{I}\{x \in S(\Theta_{0},c)\}] - \mathbb{E}_{(x,y) \sim \tilde{\mathcal{D}}(\theta)} [\ \mathbb{I}\{y = f_{\theta}(x)\} \ | \ x \in S(\Theta_{0},c)] \right| \\ & = \mathbb{P}[x \in S(\Theta_{0},c)] \left(\left| \mathbb{E}_{(x,y) \sim \mathcal{D}_{S(\Theta_{0},c)}(\tilde{c})} [\ \mathbb{I}\{y = f_{\theta}(x)\}] - \mathbb{E}_{(x,y) \sim \tilde{\mathcal{D}}_{S(\Theta_{0},c)}(\theta)} [\ \mathbb{I}\{y = f_{\theta}(x)\}] \right| \right) \end{split}$$

The bound of $\mathbb{P}[x \in S(\Theta_0, c)] \text{TV} \left(\mathcal{D}_{S(\Theta_0, c)}(\theta), \tilde{\mathcal{D}}_{S(\Theta_0, c)}(\theta) \right)$ now directly follows.