
When Samples Are Strategically Selected

Hanrui Zhang¹ Yu Cheng¹ Vincent Conitzer¹

Abstract

In standard classification problems, the assumption is that the entity making the decision (the *principal*) has access to *all* the samples. However, in many contexts, she either does not have direct access to the samples, or can inspect only a limited set of samples and does not know which are the most relevant ones. In such cases, she must rely on another party (the *agent*) to either provide the samples or point out the most relevant ones. If the agent has a different objective, then the principal cannot trust the submitted samples to be representative. She must set a *policy* for how she makes decisions, keeping in mind the agent's incentives. In this paper, we introduce a theoretical framework for this problem and provide key structural and computational results.

1. Introduction

In standard classification problems, the common assumption is that we have access to *all* the samples. However, in many contexts, we either do not have direct access to the samples, or we can inspect only a limited set of samples and do not know which are the most relevant ones. In such cases, we must rely on another party (the *agent*) to either provide the samples or point out the most relevant ones. This agent may have different incentives, which raises several concerns. One is that the individual samples cannot even be trusted—e.g., we ask for images but the agent manipulates the images with editing software before sending them. This is an issue that we do not consider in this paper; we assume that the agent cannot modify samples or create fake samples. (In the related research discussion below, we list some work that does not make this assumption.) But even in a context where the individual samples can be trusted (e.g., using cryptographic tools like digital signatures (Goldwasser et al.,

1988; Friedman, 1993)), there is still the concern that the agent sends only a *biased collection* of samples. If we do not know how many samples n the agent has available, then we cannot know whether the agent has submitted all of them. Moreover, we may have capacity to deal with only a fixed number m of samples.

Consider the following scenario. A faculty member (the agent) wishes to convince the chair of the department (the *principal*) to interview a particular candidate, while the chair wants to interview the strongest candidates.¹ The principal has time to read exactly three of the candidate's papers, and asks the agent, who is familiar with all of the candidate's work, which three she should read. Naturally, the agent chooses the best three papers. The principal, when reading them, then knows that these three papers may not be representative of all the candidate's work, and should make her decision with this in mind.

In fact, it is not clear that the principal should just ask for the three best papers according to a single metric. For example, she could also say: "Of the three papers, at least one must have at most two authors." Or: "At least one of the papers must introduce a new problem." Or: "In at least one of the papers, the 42nd letter must be a vowel." They also do not need to be hard constraints; she could just announce that she would appreciate it if one of the papers has at most two authors, but if that is not the case and the three papers are spectacular she would still interview the candidate. In general, the principal will set a *policy* for when she would interview a candidate, and the agent will base his choice of papers on this policy.²

Similar examples abound where an agent aims to convince a principal based on limited data. The above example can be generalized to hiring in many other professions: for example, in sports, high school coaches may want to convince a college scout to come out and take a look at a player. As before, the scout (the principal) has limited time, and she must make a decision of taking the trip or not, based on the limited video footage of that player provided by the coach

¹Department of Computer Science, Duke University, Durham, North Carolina, USA. Correspondence to: Hanrui Zhang <hrzhang@cs.duke.edu>, Yu Cheng <yucheng@cs.duke.edu>, Vincent Conitzer <conitzer@cs.duke.edu>.

¹Of course the decision may not rest with the chair (alone); if so, we may think of the body of people that needs to be convinced as the principal.

²Note that the candidate is *not* a strategic player; the candidate takes no actions in the game.

(the agent). The scout can specify a policy and the coach will choose the footage accordingly. For example, the scout may just wish to see the best moments, or have a policy that she prefers to see footage of the player in multiple roles (offense, defense, . . .). The decisions do not need to be hiring decisions, of course. For example, a city may make a bid to host an event based (in part) on a few selected photographs of the surroundings. The committee deciding on the location will probably appreciate having photographs of both the venue and of the nearby beach, rather than multiple photographs of just one of these.

1.1. Our Results

In this section, we highlight some of our structural and computational results.

We study the problem of designing an optimal classification policy when the samples are provided by a strategic agent. The agent has n samples, and while she cannot modify the samples, she can choose which m samples to submit. Our work is motivated by questions such as the following:

1. How does strategic sample selection affect the principal's classification problem? Is it easier or harder compared to when she has direct access to m samples?
2. Are there conditions under which she should just ask for the best m samples according to a single metric? Are there conditions under which this is not optimal?
3. What is an optimal policy for the principal?

To answer question (1), we give two examples (Examples 1 and 2) that show that, depending on the instance, classification based on strategically selected samples could be easier or harder for the principal compared to when she has direct access to the samples.

For question (2), we prove several structural results. We show that when a single sample is reported ($m = 1$) or when all samples are reported ($m = n$), any optimal policy should accept reports whose good vs. bad likelihood ratio (i.e., a single metric) exceeds a certain threshold (Theorems 1 and 2). However, such a characterization fails to hold for the general $1 < m < n$ case (Proposition 3).

For question (3), for the case of one good and one bad distribution, we show that for any value of m , we can design a policy that has good behavior in the limit (Corollary 2 and Theorem 3). When there are multiple good and bad distributions, we show that if $m = 1$ and the distributions are piecewise constant, we can efficiently compute policies with target accepting probabilities (Theorem 6).

1.2. Related Research

In other research, settings in which agents strategically manipulate the samples themselves have been studied. Dekel et al. (2008) and Meir et al. (2012) consider the case where a single hypothesis must be constructed based on data held by multiple agents, and agents change the labels of their own examples in order to change the hypothesis to their benefit. Kephart & Conitzer (2015, 2016) and Hardt et al. (2016) consider settings where we design a classifier, but the entities being classified can, at some cost, change their features to obtain a better classification.

In *mechanism design with partial verification* (Green & Laffont, 1986; Yu, 2011), an agent reports a type, and the set of types that the agent can misreport depends on the agent's true type (but misreporting does not come at any cost). Our setting is similar in the sense that the set of reports that the agent can (costlessly) make is also determined by the agent's true private information (i.e., the agent's full set of samples), though in our setting the agent can never report the full private information. Our setting has much additional structure that we exploit.

In *Bayesian persuasion* (Kamenica & Gentzkow, 2011; Dughmi & Xu, 2016), a receiving agent (the receiver) must choose among a set of actions, while a sending agent (the sender) has more information about the payoffs of both players for each action. The sender seeks to maximize his own payoff by carefully revealing information to the receiver. Our setting is different from Bayesian persuasion in several ways: In Bayesian persuasion, the sender is able to commit to a signaling strategy, while in our setting, it is the recipient of the information that can commit to an acceptance policy. Moreover, we have good and bad candidates where the good candidates want to distinguish themselves from the bad ones, while in Bayesian persuasion, the sender knows the types of the candidates and (say) he wants to maximize the overall probability of acceptance (possibly by mixing together good and bad candidates).

In principle, one may try to align the agent's incentives with those of the principal by *rewarding* the agent, possibly monetarily. If so, we may even consider dispensing with the direct submission of samples and simply ask the agent, who can inspect all the samples, for a probabilistic forecast of the candidate's future performance. We can then reward the agent according to a *proper scoring rule* (Savage, 1971; Gneiting & Raftery, 2007) to ensure incentives to report truthfully. However, in many of the settings of interest, monetary payments would be impractical, inappropriate, or sufficiently limited in size by budget considerations that the agent's direct interest in the decision would outweigh the monetary reward. In such cases, cheap-talk reports will not help and we need to rely on the agent's inability to forge samples.

Di Tillio et al. (2017) consider a setting closely related to ours. They focus on a restricted case, where (1) there is one good and one bad distribution on the real line, and moreover, the two distributions are identical up to a constant shift, and (2) only one sample is observed / reported (i.e., the $m = 1$ case in this paper). While they give a remarkable characterization for this restricted case, their assumptions drastically simplify the problem. In this paper, in contrast, we consider (possibly more than two) arbitrary distributions, and the structure of the support essentially does not matter. We also consider the case where multiple samples are reported.

2. Preliminaries

The principal is interested in the underlying *type* $\theta \in \Theta$ of the entity about which she is making the decision. For example, in much of the paper we will consider a binary type space $\Theta = \{g, b\}$ (“good” and “bad”). The entity generates n samples from a sample space X ,³ where samples x are drawn i.i.d. (so with replacement) according to $P(x|\theta)$. We sometimes use the shorthand $\theta(x) = P(x|\theta)$. These n samples constitute a multiset⁴ D which is the *dataset* available to the agent. The agent must select a multiset $R \subseteq D$ ⁵ with $|R| = m$ as his *report* to the principal. The principal must decide to accept or reject, based on the report. She commits to a policy. A *randomized policy* $\Pi_m : \mathcal{R}_m \rightarrow [0, 1]$ assigns to each report a probability of acceptance. Here \mathcal{R}_m is the set of all possible reports of size m . A *deterministic policy* has $\Pi_m(R) \in \{0, 1\}$ for all R ; it is equivalently defined by the set of accepted reports, $\mathcal{S}_m = \{R : \Pi_m(R) = 1\}$.

The agent aims to maximize the probability of acceptance regardless of the true type. He will choose to report some R^* in $\operatorname{argmax}_{R \subseteq D} \Pi_m(R)$. The principal, taking into account the strategic behavior by the agent, chooses her policy Π_m to attain her own objectives. For example, she may have a *utility* $u(\theta)$ for accepting θ ; when $\Theta = \{g, b\}$, $u(b) < 0 < u(g)$. When combined with a prior over Θ , this creates a well-defined optimization problem for the principal. Alternatively, she may have *target acceptance probabilities* for each type; for example, she may want to accept b with probability at most p_b and g with probability at least p_g (corresponding to limits on the two types of errors). This does not require a prior over Θ .

For simplicity, we use the following notation in proofs interchangeably. For *one-sample* policies ($m = 1$), we some-

³For simplicity, we assume the sample space X is in some Euclidean space, i.e., $X \subseteq \mathbb{R}^d$ for some integer $d > 0$.

⁴Recall that a multiset is a set in which an element may occur more than once. One could also think of this as a vector, but the order of the elements does not matter in our context.

⁵We use “ \subseteq ” for the standard subset notion on multisets. For multisets A and B , $A \subseteq B$ iff $c_B(x) \leq c_A(x)$ for all $x \in B$, where $c_S(x)$ is the number of occurrences of x in S .

times denote the policy by its accepted set of samples S (instead of \mathcal{S} , which is a family of singleton multisets). That is, $S = \{x \mid \{x\} \in \mathcal{S}\}$. For a type, e.g., $g \in \Theta$, we abuse notation and use $g(S)$ to denote the total probability mass of g on a set $S \subseteq X$. That is, $g(S) = \Pr_{x \sim g}[x \in S]$.

2.1. Illustrative Examples

We now present a few examples that illustrate some of the key issues. The first example illustrates that the strategic selection sometimes helps the principal. In the language of our motivating example, this is the case where the principal and the agent can only classify papers as high- or low-quality, and high-quality papers are rare.

Example 1. Let $\Theta = \{g, b\}$ and $X = \{0, 1\}$. Let $g(1) = 0.05$ and $b(1) = 0.005$. Let $n = 50$ and $m = 1$. The natural policy is to accept iff the agent submits report $\{1\}$. In other words, the agent can convince the principal to accept iff at least one of the n papers has high-quality. The probability a good type is accepted is $1 - 0.95^{50} \approx 0.92$; for a bad type it is $1 - 0.995^{50} \approx 0.22$. In this example, thanks to the agent’s strategic selection, the principal can classify quite effectively while only observing a single sample; in contrast, if she had to observe samples directly, a single sample would give her very little information.

However, the next example (where high-quality papers are less rare) shows that the opposite can also happen. Strategic selection can make it much harder for the principal to distinguish between good and bad types.

Example 2. Let $\Theta = \{g, b\}$ and $X = \{0, 1\}$. Let $g(1) = 0.95$ and $b(1) = 0.05$. Let $n = 50$ and $m = 1$. Again, the natural policy is to accept iff the agent submits $\{1\}$. The probability that a good type is accepted is $1 - 0.05^{50} \approx 1$; for a bad type it is $1 - 0.95^{50} \approx 0.92$. In this example, the strategic selection of samples by the agent makes it very difficult for the principal to distinguish between g and b ; in contrast, if she could observe samples directly, a single sample would allow her to classify quite effectively.

Fortunately, there’s a workaround: the principal can effectively reduce n by specifying an irrelevant (i.e., uncorrelated with the type) requirement, such as that the 42nd letter of the paper should be a vowel.⁶ Since there will generally be nearly infinitely many irrelevant attributes of the samples, we assume that each sample is associated with a real number drawn uniformly⁷ from $(0, 1)$, representing the irrelevant

⁶One may wonder whether this creates an incentive for candidates to write their papers in a particular way; this can be avoided by choosing this requirement close to the decision. In any case, we do not consider the process that originally produces the n samples as a strategic entity in this paper.

⁷Any continuous distribution can be transformed to a uniform one, using the standard trick of applying the CDF to the drawn number first.

information.

Example 3. Let $\Theta = \{g, b\}$ and $X = \{0, 1\} \times (0, 1)$ (where the first number represents the relevant information and the second number the irrelevant information). Let $g(\{1\} \times (0, 1)) = 0.95$ and $b(\{1\} \times (0, 1)) = 0.05$. Note that this is the same example as Example 2, except for the additional irrelevant information. Let $n = 50$ and $m = 1$. Now consider the policy that accepts $\mathcal{S} = \{x : x \in \{1\} \times (0, 0.05)\}$. That is, the principal accepts only samples with good relevant information and irrelevant information that has a 1 in 20 chance of occurring. The probability that a good type is accepted is $1 - (1 - 0.95 \cdot 0.05)^{50} \approx 0.92$; for a bad type it is $1 - (1 - 0.05 \cdot 0.05)^{50} \approx 0.12$. Thus, the addition of irrelevant information allows the principal to classify much more effectively.

Example 3 illustrates that the difficulty of Example 2 is primarily due to the discreteness of the sample space.

Our last example shows that with multiple bad distributions, the optimal policy does not evaluate the quality of samples individually, but rather considers them in combination. One of the questions of interest later in this paper is under which circumstances this can happen.

Example 4. Let $\Theta = \{g, b_1, b_2\}$ and $X = \{0, 1\}$. Let $g(1) = 0.5$, $b_1(1) = 0.99$, $b_2(1) = 0.01$. Let $n = 10$ and $m = 2$. Consider the policy that accepts $\mathcal{S} = \{\{0, 1\}\}$, i.e., it accepts if both possible samples are reported. The probability that a good type is accepted is $1 - 2 \cdot 0.5^{10} \approx 0.998$; the probability that a bad type is accepted is less than $1 - 0.99^{10} \approx 0.10$. In contrast, if we accept reports that have the same sample twice, then one of the two bad types is extremely likely to succeed.

3. Basic Results

In this section, we provide some basic results that justify why we focus on deterministic policies and continuous distributions in the rest of the paper.

3.1. Deterministic vs. Randomized Policies

In this subsection, we discuss the relative power of deterministic and randomized policies, justifying our focus on deterministic policies in the rest of the paper.

We first show that in our setting, randomized policies can be decomposed into a distribution over deterministic policies.

Proposition 1. *Any randomized policy can be decomposed into a distribution over deterministic policies, such that the accepting probability of any report remains the same.*

As a simple corollary, if the principal is trying to maximize her utility, one of the deterministic policies must be optimal.

Corollary 1. *Assume there is a prior over Θ and the prin-*

cipal has a utility $u(\theta)$ for accepting type $\theta \in \Theta$. If the principal wishes to maximize her expected utility, there exists a deterministic policy that is optimal.

We defer the proofs of Proposition 1 and Corollary 1 to the appendix.

It should be noted that Corollary 1 is generally not true in mechanism design. For example, for designing revenue-maximizing auctions, it is well-known that the optimal mechanism may require randomization (Hart & Reny, 2015). In fact, even in our problem, if the principal has target acceptance probabilities for each type, for example “I want to accept at least 90% of good candidates and at most 10% of bad candidates,” then it is possible that only randomized policies obtain these goals simultaneously, as the following example demonstrates.

Example 5. Let $\Theta = \{g, b\}$ and $X = \{0, 1\}$. Let $g(0) = 1/2$, $g(1) = 1/2$, $b(0) = 1$, $b(1) = 0$, and $n = m = 1$. If we wish to accept the good distribution at least $3/4$ of the time, and accept the bad distribution at most $1/2$ of the time, we have to use a randomized policy: always accept $\{1\}$ and accept $\{0\}$ with probability $1/2$.

On the other hand, this example heavily relies on the discrete nature of the sample space; if we make the sample space continuous—say, $\{0, 1\} \times (0, 1)$ where the second number is drawn uniformly at random, as in Example 3—then we can achieve the same result with a deterministic mechanism, by effectively using the second number to generate the randomness.

In practice, deterministic policies have several other advantages as well. They are straightforward to implement, transparent, fair, and not subject to willful manipulation of the random numbers. For all these reasons, we focus on deterministic policies in the rest of this paper.

3.2. Continuous vs. Discrete Distributions

In this subsection, we compare continuous and discrete distributions, justifying our focus on continuous distributions in the rest of the paper. As we will prove in the next proposition, given our focus on deterministic policies, for discrete distributions we face NP-hardness, simply due to having to solve a knapsack problem. We defer the proof of Proposition 2 to the appendix.

Proposition 2. *Given two discrete distributions g and b , and two target acceptance probabilities p_g and p_b , it is NP-hard to decide if there exists a deterministic policy that accepts g with probability at least p_g and accepts b with probability at most p_b . This is true even when $n = m = 1$.*

In contrast, Theorem 6 states that for continuous distributions, we can solve the decision problem of Proposition 2 efficiently in much more general settings.

As we discussed in Example 3, in practice, we can often make the distribution over the sample space effectively continuous, by considering irrelevant information. For example, when considering video footage of an athlete, we may be able to summarize all the relevant information in discrete terms (did the athlete make the shot or not, etc.). Meanwhile, the background of the video provides almost unlimited irrelevant information (is someone eating popcorn in the background, etc.). Imagine that we can only watch a single video clip of a basketball player. If our policy is to only accept if the player (say) makes a 3-pointer, then even mediocre players will be able to produce such a clip. But if our policy is to only accept if the player makes a 3-pointer while a girl is eating popcorn in the background, then only players who frequently score 3-pointers are likely to be able to produce such a clip. From a technical viewpoint, we can assume that each sample is associated with a random real number drawn uniformly from $(0, 1)$, as in Example 3.

For all these reasons, we focus on continuous distributions in the rest of this paper. We now move on to study strategic sample selection in these more restricted settings.

4. One Good and One Bad Distribution

In this section, we consider the setting in which the type space is binary: $\Theta = \{g, b\}$.

4.1. One Sample

We first consider the special case where the agent submits only one sample ($m = 1$). Our main structural result (Theorem 1) states that any Pareto optimal policy takes the following form: accept all reports $\{x\}$ such that the *likelihood ratio* $g(x)/b(x)$ is greater than some threshold.

As a corollary of Theorem 1, when the agent has sufficiently many samples ($m = 1$ and $n \rightarrow \infty$), we can characterize the optimal tradeoff between the accepting probabilities p_g and p_b (Corollary 2). This tradeoff is quantitatively governed by the maximum likelihood ratio $\max_{x \in X} \frac{g(x)}{b(x)}$ over the entire sample space $x \in X$. This is in contrast to the distribution learning literature (Pearson, 1895; Batu et al., 2000; Chan et al., 2014), where this tradeoff is often determined by some global measure (e.g., total variation distance) between the two distributions.

Let Π, Π' be two policies that accept the good distribution with probability p_g and p'_g , and accept the bad distribution with probability p_b and p'_b respectively. We say Π' is *strictly better* than Π if $p'_g \geq p_g$ and $p'_b \leq p_b$ and at least one of the two inequalities holds strictly. We say Π is *Pareto optimal* if there is no other policy Π' strictly better than Π .

Theorem 1. *Suppose $m = 1$ and we have continuous distributions g and b . Consider any optimal deterministic policy*

II. Let $S \subseteq X$ be the accepting region of Π . Then, for any point x_1 strictly inside S and any x_2 strictly outside of S ,

$$g(x_1)/b(x_1) \geq g(x_2)/b(x_2).$$

Before proving Theorem 1, we first discuss its conditions and implications. We can always change the policy on a set of measure zero, and the resulting policy is equivalent to the original one. Therefore, the condition in Theorem 1 only applies to the interior of S and $X \setminus S$. An immediate consequence of Theorem 1 is that, if the principal's utility only depends on the accepting probabilities, and is strictly monotonically increasing in p_g and strictly monotonically decreasing in p_b , then the optimal policy must be Pareto optimal and hence it must satisfy the condition in Theorem 1.

Proof. Let p_g, p_b be the probabilities that Π accepts the good and bad distributions respectively.

$$p_g = \Pr_{D \sim g^n} [D \cap S \neq \emptyset] = 1 - (1 - g(S))^n,$$

$$p_b = \Pr_{D \sim b^n} [D \cap S \neq \emptyset] = 1 - (1 - b(S))^n.$$

Suppose there is x_1 strictly in S and x_2 strictly outside S , where

$$g(x_1)/b(x_1) < g(x_2)/b(x_2).$$

Pick neighborhoods N_1 and N_2 of x_1 and x_2 , in S and $X \setminus S$ respectively, such that $g(N_1) = g(N_2) > 0$ and $g(N_1)/b(N_1) < g(N_2)/b(N_2)$. Such neighborhoods exist because the likelihood ratio $g(x)/b(x)$ is a continuous function, and both x_1 and x_2 are not on the boundary of S .

We will show that a different policy Π' with accepting region $S' = (S \setminus N_1) \cup N_2$ is a better policy. Let p'_g and p'_b be the accepting probabilities of Π' . Since $g(S') = g(S)$ and $b(S') = b(S) - b(N_1) + b(N_2) < b(S)$, we have

$$p'_g = 1 - (1 - g(S'))^n = 1 - (1 - g(S))^n = p_g, \text{ and}$$

$$p'_b = 1 - (1 - b(S'))^n < 1 - (1 - b(S))^n = p_b. \quad \square$$

Corollary 2. *Fix continuous distributions g and b . Let $r = \sup_{x \in X} (g(x)/b(x))$ be the maximum likelihood ratio over the entire sample space X .⁸ When $m = 1$ and $n \rightarrow \infty$, any Pareto optimal deterministic policy Π satisfies*

$$p_g + (1 - p_b)^r = 1,$$

where p_g and p_b are the probabilities that Π accepts g and b respectively.

Proof. Fix any $0 < \varepsilon < 1$. Because the likelihood ratio function $g(x)/b(x)$ is continuous and its supremum

⁸For simplicity, we assume the maximum likelihood ratio $r = \sup_x (g(x)/b(x))$ exists and is finite. A similar argument shows that when $r = \infty$, we can get policies with $p_g = 1$ and $p_b = 0$.

is r , there exists a small neighborhood $S \subseteq X$ such that $g(x)/b(x) \geq (1 - \varepsilon)r$ for all $x \in S$. Recall that $g(S) = \Pr_{x \sim g}[x \in S] > 0$ is the probability that a random sample from the good distribution $x \sim g$ is in S . By the definition of S , we know that $b(S) \leq g(S)/(r(1 - \varepsilon))$.

Consider the policy that accepts iff the agent reports a sample in S . We will show that S is essentially optimal as $n \rightarrow \infty$. Let p_g and p_b denote the accepting probabilities of S . For notational convenience let $\delta = g(S)$. We have

$$p_g = (1 - (1 - g(S))^n) = 1 - (1 - \delta)^n, \text{ and}$$

$$p_b = (1 - (1 - b(S))^n) \leq 1 - \left(1 - \frac{\delta}{r(1 - \varepsilon)}\right)^n.$$

We can rewrite the inequality on p_b by substituting δ . Because the inequality holds for any ε and n , we can let $\varepsilon \rightarrow 0$ and $n \rightarrow \infty$ and get

$$p_b \leq 1 - \left(1 - \frac{1 - (1 - p_g)^{1/n}}{r(1 - \varepsilon)}\right)^n \rightarrow 1 - (1 - p_g)^{1/r}.$$

On the other hand, the upper bound on p_b is tight when $\varepsilon = 0$. This is because the acceptance region S' of any deterministic policy can have likelihood ratio at most r , and thus $b(S') \geq g(S')/r$ and a similar calculation gives the same lower bound on p_b . Therefore, we can conclude that for any Pareto optimal policy, $p_g + (1 - p_b)^r = 1$ as $n \rightarrow \infty$. \square

4.2. Multiple Samples

We now move on to the case where the agent submits m samples ($m > 1$). We first generalize the notion of the likelihood ratio to reports of multiple samples.

Definition 1. We define the likelihood ratio of a report R to be the product of the samples' likelihood ratios $\prod_{x \in R} \frac{g(x)}{b(x)}$, as if the samples in R are drawn i.i.d. from the distribution.

The following theorem states that when $m = n$, any Pareto optimal policy essentially accepts reports whose likelihood ratio exceeds some threshold.

Theorem 2. *Suppose $m = n$ and we have continuous distributions g and b . Consider any Pareto optimal deterministic policy Π which accepts all and only reports in \mathcal{S} . Then, for any report R_1 strictly inside \mathcal{S} and any R_2 strictly outside of \mathcal{S} , we have*

$$\prod_{x \in R_1} g(x)/b(x) \geq \prod_{x \in R_2} g(x)/b(x).$$

We defer the proof of Theorem 2 to the appendix.

Note that for the case $m = 1$, Theorem 1 states that in that case, too, the agent should report the sample that maximizes

the likelihood ratio. Thus, it is natural to conjecture that this continues to hold when $m < n$. This, however, is false. In fact, we can show that the optimal policy does not admit even the following weaker structural property.

Definition 2. A policy *orders the sample space* if there exists an ordering on the elements of the sample space, such that an optimal response for the agent is to always report his highest samples in this ordering.

In the $m = 1$ and $m = n$ cases, the optimal policy based on the likelihood ratio clearly satisfies this property, ordering the sample space by likelihood ratio. (In the $m = n$ case, this is because the likelihood ratio is the product of the likelihood ratios of the individual samples, and this product is maximized by choosing the samples that maximize that ratio.) The following proposition shows that this does not hold in general for $1 < m < n$.

Proposition 3. *When $1 < m < n$, sometimes a Pareto optimal policy does not order the sample space.*

Proof. For simplicity, we present a counter example that is a discrete distribution. It can be easily changed to a continuous distribution without affecting any part of the argument.

Let $\Theta = \{g, b\}$, $X = \{0, 1, 2\}$, and

$$g(0) = 0, \quad g(1) = 0.1, \quad g(2) = 0.9, \quad \text{and}$$

$$b(0) = 0.8, \quad b(1) = 0.1, \quad b(2) = 0.1.$$

Let $m = 2$ and $n = 3$, i.e., the agents has 3 i.i.d. samples and chooses 2 of them to submit.

We claim a policy Π that accepts reports $\{1, 1\}$ and $\{2, 2\}$ is Pareto optimal. First notice that Π accepts g with probability 1. Since an agent does not draw $\{0\}$ from g , so by the pigeonhole principle, among the $n = 3$ samples there must be either two copies of $\{1\}$ or two copies of $\{2\}$. On the other hand, the principal must accept these two reports with probability 1 if she wants to always accept g . This is because when $\theta = g$, the agent's data D could be $\{1, 1, 1\}$ (or $\{2, 2, 2\}$), in which case he is forced to report $R = \{1, 1\}$ (or resp. $\{2, 2\}$). Therefore, among all policies, Π has the smallest probability of accepting b .

The above example rules out structural results in the form of Theorems 1 and 2, because the report $\{1, 2\}$ has higher likelihood ratio than $\{1, 1\}$. Furthermore, note that any policy that accepts $\{1, 1\}$ and $\{2, 2\}$ but not $\{1, 2\}$ cannot order the sample space; for example, if 1 were ordered at least as high as 2, then an agent with data $\{1, 2, 2\}$ can report $\{1, 2\}$ instead of $\{2, 2\}$ according to the ordering, but this is suboptimal because $\{1, 2\}$ is rejected while $\{2, 2\}$ is accepted. \square

It of course remains possible that there is an elegant way to describe the optimal policy in this context, but Proposition 3 rules out many natural possibilities. However, if we are willing to give up on exact optimality, then we can still define a policy that performs reasonably well in the limit. Theorem 3 gives a policy whose error probability (probability of rejecting g or accepting b) decreases exponentially in m . This error guarantee is similar to the setting where the principal has direct access to m samples, in which case the failing probability also decreases exponentially in m . The difference is that, as in Corollary 2, the coefficient in the exponent depends on the maximum likelihood ratio r rather than (say) the total variation distance between g and b .

Theorem 3. Fix $m \geq 1$ and two continuous distributions g and b . Let $r = \sup_x \frac{g(x)}{b(x)} > 1$ denote the maximum likelihood ratio. As $n \rightarrow \infty$, there is a deterministic policy whose error probability is at most $\exp\left(-\frac{1}{2}(1-r^{-1/2})^2 m\right)$.

The policy that achieves Theorem 3 focuses on a small region S where $g(S) \gg m/n$ and $b(S) \ll m/n$. This way, for n samples drawn from g , in expectation, $ng(S) \gg m$ samples are from S ; and for n samples drawn from b , $nb(S) \ll m$ samples are from S . Therefore, if we accept all reports with m samples from S , we can distinguish g from b . We defer the proof of Theorem 3 to the appendix.

5. One Good and Multiple Bad Distributions

For men are good in but one way, but bad in many.

— Aristotle, *Nicomachean Ethics*

We now consider the case where there are multiple bad distributions, but still only a single good one.

5.1. One Sample

Again, we first investigate the single-sample case ($m = 1$). We first show there are cases in which no policy can perform well across all possible priors over Θ .

Example 6. Let $\Theta = \{g, b_1, b_2\}$ and $X = \{0, 1\} \times (0, 1)$. Let $g(\{0\} \times (0, 1)) = g(\{1\} \times (0, 1)) = 0.5$, $b_1(\{0\} \times (0, 1)) = 1$, and $b_2(\{1\} \times (0, 1)) = 1$. There is no policy which makes the right decision with probability larger than 0.5 against any prior. This is because any deterministic policy has the following form: for some $p, q \in [0, 1]$, the policy accepts all reports in $\mathcal{S} = \{\{x\} \mid x \in (\{0\} \times (0, p)) \cup (\{1\} \times (0, q))\}$. To accept g w.p. larger than $1/2$, we need $p + q > 1$. W.l.o.g. assume $p > 1/2$, but then the policy accepts b_1 w.p. $p > 1/2$.

When there is a prior over Θ and the principal has utilities for accepting each type $\theta \in \Theta$, the next theorem characterizes the behavior of optimal policies in the limit. Note

that Theorem 4 does not hold if we have specific target acceptance probabilities for individual bad distributions.

Theorem 4. Fix $m = 1$ and a partition of the sample space X into t pieces. Let $\Theta = \{g, b_1, \dots, b_k\}$. Assume every distribution is constant on every piece, the principal has utility $u(\theta)$ for accepting type θ , and there is a prior q over Θ . Then, for sufficiently large n , there is a utility-maximizing policy that accepts only reports in a subset of one single piece (modulo accepting any report that has measure zero).

Proof. Suppose the optimal policy accepts all reports in S and S overlaps with multiple pieces P_1, \dots, P_t . Let $(g(S_j), b_1(S_j), \dots, b_k(S_j))$ denote the ‘‘cumulative’’ probabilities of the types in $S_j = S \cap P_j$. Let α_i^g be the density of distribution θ on P_i , and let $\beta_{ij} = b_i(S_j)/g(S_j) = \alpha_j^{b_i}/\alpha_j^g$ be the likelihood ratio $b_i(x)/g(x)$ on piece j . We argue that moving all the mass, measured by g , to one of the k pieces achieves at least the same probability of success. Since n is large enough, $\sum_j g(S_j)$ can be contained in any of the t pieces. The expected utility of S is

$$q_g u(g) \left(1 - \left(1 - \sum_j g(S_j) \right)^n \right) + \sum_i q_{b_i} u(b_i) \left(1 - \left(1 - \sum_j \beta_{ij} g(S_j) \right)^n \right).$$

Let $\beta_i = (\beta_{i1}, \dots, \beta_{it})$, $\gamma = (g(S_1), \dots, g(S_t))$. The principal’s expected utility can be written as

$$u(g)q_g(1 - (1 - \|\gamma\|_1)^n) + \sum_i u(b_i)q_{b_i}(1 - (1 - \beta_i^\top \gamma)^n).$$

Note that since $0 \leq \beta_i^\top \gamma \leq 1$, $(1 - \beta_i^\top \gamma)^n$ is convex in γ for any i . Fixing $\|\gamma\|_1$, since $q_{b_i} \geq 0$ and $u(b_i) < 0$, the overall utility is also convex in γ . Therefore, the maximum utility is achieved when γ has only one non-zero entry. Equivalently, the optimal policy should focus on a single piece. \square

Theorem 4 crucially relies on there being only a single good distribution, as the following example demonstrates.

Example 7 (non-locality with multiple good distributions). Let $\Theta = \{g_1, g_2, b\}$ and $X = \{0, 1, 2\} \times (0, 1)$. Let $b(\{0\} \times (0, 1)) = g_1(\{1\} \times (0, 1)) = g_2(\{2\} \times (0, 1)) = 1$. Let $m = 1$. Even as $n \rightarrow \infty$, we will need to accept points from both the pieces $\{1\} \times (0, 1)$ and $\{2\} \times (0, 1)$ in order to accept both good distributions.

5.2. Multiple Samples

When $m > 1$ and there are multiple bad distributions, it turns out that sometimes the *agent* faces an NP-hard problem. This is because an individual sample may rule out

several bad distributions, and to convince the principal to accept, the agent may have to judiciously choose his m reported samples to cover all the bad distributions, in terms of ruling them out. The following theorem makes this precise.

Theorem 5. *With one good distribution and k bad distributions, it is NP-hard for the agent to determine, given dataset D , whether it is possible to report $R \subseteq D$, such that the optimal policy accepts R .*

Proof. We reduce from the decision version of Set Cover. Given a set cover instance with elements U , n' sets $\{S_j\}_{j \in [n']}$, and a target number m' , we know it is NP-hard to decide whether all elements can be covered with m' sets.

We construct a strategic sample selection instance as follows. We partition the sample space into n' pieces. Each piece P_j corresponds to a set S_j in the set cover instance. The good distribution g is the uniform distribution. For each element $i \in U$, we create a bad distribution b_i . We set the probability density of b_i to 0 on P_j if $S_j \ni i$, and set it equally on all other pieces. (W.l.o.g., we can assume there is no element that is contained in all sets.)

Consider an agent with $n = n'$ samples, one for each piece. Let $m = m'$ be the number of samples he can report. Suppose that the principal will accept if and only if she is sure the underlying distribution is g (say she has extremely negative utility for accepting a bad type).

Suppose a set cover of size m exists. Then, the agent can report the corresponding samples. For each b_i , there is a sample in the report that has probability 0 under b_i . Hence the principal can rule out every bad distribution.

Conversely, suppose the agent has a report that will get accepted. For each b_i , there must be a sample in the report that has probability density 0 under b_i ; this sample corresponds to some $S_j \ni i$. Hence, the agent's report produces a set cover of size m . \square

6. Multiple Good/Bad Distributions

We now allow multiple good and multiple bad distributions. The hardness result from Theorem 5 still applies here when $m > 1$, so we focus on $m = 1$. Even so, Example 7 shows that we will get nonlocality in the optimal policy, so we do not prove a structural result. This leaves the question of whether we can efficiently compute policies with target accepting probabilities when $m = 1$.

Theorem 6. *Fix $m = 1$, $n \geq 1$, and a partition of the sample space X into t pieces. Assume we are given distributions g_1, \dots, g_k , and b_1, \dots, b_ℓ such that every distribution is constant on every piece. Then, given a vector of target accepting probabilities $(p_{g_1}, \dots, p_{g_k}, p_{b_1}, \dots, p_{b_\ell})$, we can decide in $\text{poly}(k, \ell, t, n)$ time whether there is a policy that*

can achieve these requirements.

We defer the proof of Theorem 6 to the appendix.

7. Conclusion

We have introduced the problem of designing an optimal classification policy when the samples are selected by a strategic agent who favors a specific outcome.

We proved several basic structural results. If the principal aims to maximize expected utility, where she associates utilities with individual outcomes, then there is no benefit to randomization (Corollary 1). When distinguishing a single good from a single bad distribution, if only a single sample is reported, then the optimal policy is to accept samples whose good/bad likelihood ratio exceeds some threshold (Theorem 1). Moreover, in the limit as $n \rightarrow \infty$, our success probability is solely determined by the highest likelihood ratio in the sample space (Corollary 2). While a result similar to Theorem 1 holds when $m = n$ (Theorem 2), unfortunately nothing like it holds for the case $1 < m < n$ (Proposition 3). Still, we can design a policy that has good behavior in the limit for this case (Theorem 3). Moving on to the case of multiple bad distributions, we show that for $m = 1$, in the limit our optimal policy focuses on a single piece (Theorem 4)—but this is not true with multiple good distributions (Example 7).

We also proved basic computational results. In the discrete, deterministic case, determining whether a combination of a given false positive and a given false negative rate can be obtained is NP-hard even with $m = n = 1$ (Proposition 2). However, if we restrict ourselves to piecewise-constant distributions, then we can obtain an efficient algorithm even with multiple good and bad distributions (but still $m = 1$; Theorem 6). When $1 < m < n$ and there are multiple bad distributions, the agent's problem of best-responding to the optimal policy becomes NP-hard (Theorem 5).

There are several open questions. Perhaps the most significant open questions are in the setting where there is a single good and a single bad distribution, and $1 < m < n$. We have shown that for optimal policies in this case, it is not always true that the agent should just report the “best” samples according to a single criterion. Still, do optimal policies in this case have some natural structure? Can they be computed efficiently?

Acknowledgements

We are thankful for support from NSF under awards IIS-1814056 and IIS-1527434. We also thank anonymous reviewers for helpful comments.

References

- Batu, T., Fortnow, L., Rubinfeld, R., Smith, W. D., and White, P. Testing that distributions are close. In *Proceedings of the 41st IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 259–269, 2000.
- Chan, S.-O., Diakonikolas, I., Valiant, P., and Valiant, G. Optimal algorithms for testing closeness of discrete distributions. In *Proceedings of the 25th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 1193–1203, 2014.
- Dekel, O., Fischer, F., and Procaccia, A. D. Incentive compatible regression learning. In *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 884–893, Philadelphia, PA, USA, 2008. Society for Industrial and Applied Mathematics.
- Di Tillio, A., Ottaviani, M., and Sorensen, P. N. Strategic sample selection. 2017.
- Dughmi, S. and Xu, H. Algorithmic bayesian persuasion. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC)*, pp. 412–425, 2016.
- Friedman, G. L. The trustworthy digital camera: Restoring credibility to the photographic image. *IEEE Transactions on consumer electronics*, 39(4):905–910, 1993.
- Gneiting, T. and Raftery, A. E. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.
- Goldwasser, S., Micali, S., and Rivest, R. L. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM Journal on Computing*, 17(2):281–308, 1988.
- Green, J. and Laffont, J.-J. Partially verifiable information and mechanism design. *Review of Economic Studies*, 53: 447–456, 1986.
- Hardt, M., Megiddo, N., Papadimitriou, C., and Wootters, M. Strategic classification. In *Innovations in Theoretical Computer Science (ITCS)*, Cambridge, MA, USA, 2016.
- Hart, S. and Reny, P. J. Maximal revenue with multiple goods: Nonmonotonicity and other observations. *Theoretical Economics*, 10(3):893–922, 2015.
- Kamenica, E. and Gentzkow, M. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.
- Kephart, A. and Conitzer, V. Complexity of mechanism design with signaling costs. In *Proceedings of the Fourteenth International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pp. 357–365, Istanbul, Turkey, 2015.
- Kephart, A. and Conitzer, V. The revelation principle for mechanism design with reporting costs. In *Proceedings of the Seventeenth ACM Conference on Economics and Computation (EC)*, pp. 85–102, Maastricht, the Netherlands, 2016.
- Meir, R., Procaccia, A. D., and Rosenschein, J. S. Algorithms for strategyproof classification. *Artificial Intelligence*, 186:123–156, 2012.
- Pearson, K. Contributions to the mathematical theory of evolution. ii. skew variation in homogeneous material. *Philosophical Transactions of the Royal Society of London*, 186(Part I):343–424, 1895.
- Savage, L. J. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, 66:783–801, 1971.
- Yu, L. Mechanism design with partial verification and revelation principle. *Autonomous Agents and Multi-Agent Systems*, 22(1):217–223, 2011.