ϵ-Minimax Solutions of Statistical Decision Problems via the Hedge Algorithm*

Andrés Aradillas Fernández† José Blanchet‡ José Luis Montiel Olea†
Chen Qiu† Jörg Stoye† Lezhi Tan‡

Abstract

We present an algorithm for obtaining ϵ-*minimax* solutions of statistical decision problems. We are interested in problems where i) the statistician is allowed to choose randomly among I decision rules, and ii) the statistical model may have a parameter space with infinitely many elements. The minimax solution of these problems admits a convex programming representation over the $(I-1)$-simplex, and the algorithm suggested herein to obtain an ϵ-approximation of the minimax solution is a version of *mirror subgradient descent*, initialized with uniform weights and stopped after a finite number of iterations. The resulting iterative procedure is known in the computer science literature as the *Hedge algorithm* (a particular case of the *Multiplicative Weights* update method) and it is used in algorithmic game theory as a practical tool to find approximate solutions of two-person zero-sum games. We apply the suggested algorithm to different minimax problems in the econometrics literature. An empirical application to the problem of optimally selecting *sites* to maximize the external validity of an experimental policy evaluation illustrates the usefulness of the suggested procedure.

1 Introduction

Under Wald (1950)'s *minimax* criterion different statistical decision rules are ranked based on their worst possible expected loss. Searching for a *minimax-optimal* decision rule—i.e., a rule with the

*We would like to thank Karun Adusumilli, Isaiah Andrews, Tim Armstrong, Patrik Guggenberger, Giannis Fikioris, Kei Hirano, Nicole Immorlica, Lihua Lei, Charles Manski, Guillaume Pouliot, Brenda Quesada Prallon, Vasilis Syrgkanis, David Shmoys, Sophie Sun, Yiwei Sun, Eva Tardos, Alex Torgovitsky, and Davide Viviano for helpful feedback, comments, and suggestions. We would also like to thank Rohit Kumar for excellent research assistance. We gratefully acknowledge financial support from the NSF under grant SES-2315600. First version: December 5th, 2024.

†Department of Economics, Cornell University.

‡Management Science and Engineering Department, Stanford University.

smallest worst-case expected loss—presents two well-known computational challenges. First, evaluating the worst-case performance of a given decision rule typically requires optimizing a nonlinear function (the *risk* function) over the model's parameter space. Second, minimizing the worst-case expected loss over the decision rules under consideration typically requires optimization over a high-dimensional (and possibly infinite-dimensional) space; for example, all measurable functions that map the data into actions. It is known that obtaining the minimax solution of a decision problem—and sometimes even deciding whether a minimax solution exists—is NP hard in general; see Du and Pardalos (1995); Daskalakis, Skoulakis, and Zampetakis (2021).[1]

We consider a particular class of decision problems in which the decision maker is restricted to choose from a menu of I available decision rules, all of which are assumed to have risk between zero and a known positive constant M. Our motivation is that, while it is always theoretically interesting to look for the *best* overall decision rule, there are situations in which it is equally desirable to *"evaluate the performance of relatively simple statistical decision functions that researchers use in practice"* (Dominitz and Manski, 2024) and choose optimally among them. It is known that if we allow the decision maker to choose randomly among its I options, the corresponding minimax problem can be viewed as a nonlinear convex optimization problem over the $(I-1)$-dimensional simplex (Chamberlain, 2000). The connection to convex programming is helpful, but is not a computational panacea: evaluating the objective function of the convex program associated to the minimax problem could remain computationally costly.[2] For instance, if one were to rely on a textbook *mirror descent* routine for convex problems in the simplex (Nemirovski and Yudin, 1983; Bubeck et al., 2015, Chapter 4.3), one would typically need to make infinitely many evaluations of the objective function (and its subgradient) to guarantee that a minimax solution has been found. If one is willing to measure the complexity of an iterative *first-order* algorithm for convex optimization

[1] See also Montiel Olea, Prallon, Qiu, Stoye, and Sun (2024a) for an example of a minimax problem that is NP hard but that in practice can be solved to provable (sub)optimality using solvers for linear integer programming.

[2] As we will explain later, the objective function in the convex program is the worst-case expected loss of any given random selection over the I decision rules; which, as discussed above requires optimizing a nonlinear function over the model's parameter space.

by the number of calls (or *queries*) that the procedure makes to the objective function and its subgradient—as it is done in the popular oracle model of optimization complexity of Nemirovski and Yudin (1983)—it is thus desirable to search for iterative algorithms that require as fewer queries as possible.

Our first contribution is to argue—motivated by the *oracle* complexity framework of Nemirovski and Yudin (1983)—that it is possible to make substantial progress in solving the general class of statistical decision problems considered herein if, instead of insisting in finding an *exact* minimax solution, we make our goal to find an *approximate* minimax solution. In particular, we search for an ϵ-*minimax optimal decision rule* (Ferguson, 1967, Chapter 1, Definition 4): that is, a rule that attains the smallest worst-case expected loss, but up to a given additive factor ϵ. Our second contribution, which we view as our main result, shows that we can provably obtain such an ϵ-minimax rule by using a *mirror subgradient descent routine* (with negative entropy as a mirror map) for convex optimization over the simplex (Theorem 1). The routine is initialized at uniform weights over the I original decision rules, and the step size is set to be $\eta \equiv \epsilon/M^2$. Importantly, we show that it suffices to stop the mirror descent routine after $T = \lceil 2M^2 \ln(I)/\epsilon^2 \rceil$ epochs.[3] Thus, we obtain a concrete upper bound on the computational cost of our procedure (in terms of the number of evaluations of the objective function and its subgradient). The required number of iterations scales logarithmically on the number of original decision rules, I, which means that the algorithm is still useful in problems in which I is large. Moreover, we use the results in Ben-Tal, Margalit, and Nemirovski (2001) to argue that smallest number of epochs required by *any* iterative, first-order algorithm for convex optimization over the $(I-1)$ simplex of a Lipschitz function with constant at most M (with respect to $\|\cdot\|_1$ norm) is $O(1)M^2/\epsilon^2$, provided $\epsilon \geq M/\sqrt{I}$. Thus, there is a sense in which the recommended algorithm, and the suggested number of epochs, achieve the optimal dependence on M and ϵ, up to the logarithmic factor $\ln(I)$.

[3] $\lceil \cdot \rceil$ is the ceiling function: the function that returns the smallest integer that is greater than or equal to a given number.

Importantly, the algorithm herein suggested is known in the computer science literature as the *Hedge algorithm* (a particular case of the *Multiplicative Weights* update method); see Section 2.1 in Arora, Hazan, and Kale (2012). This method is used in problems where a decision maker chooses randomly among I alternatives repeatedly (an *online decision-making* problem), but after each round he obtains a payoff for all of the I available actions.

CONTRIBUTIONS: For the sake of exposition, we provide a brief summary of the main contributions that our paper makes to the econometrics literature.

1. We show that it is possible to make substantial progress in solving a general class of statistical decision problems if, instead of insisting in finding an *exact* minimax solution, we make our goal to find an ϵ-*minimax* solution as defined in Ferguson (1967), Chapter 1, Definition 4.

2. We present theoretical results showing that the Hedge algorithm can be used to obtain an ϵ-minimax solution in the class of statistical decision problems considered herein. Interestingly, neither our theoretical results nor our algorithm need to assume that the parameter space of the statistical model has finitely many elements (or that it has been discretized). Instead, we assume that we have access to an oracle that can find the point in the parameter space associated to the worst-case performance of any random selection of the I decision rules. In other words, we assume that we have access to a computer routine that can solve the inner part of the minimax problem. We remark that the use of the Hedge algorithm in statistical decision problems is (to the best of our knowledge) novel. This is rather surprising in light of the straightforward connection between statistical decision problems and two-person zero-sum games, and the origins of Multiplicative Weights in iterative dynamics for game play—see the notion of κ-exponential fictitious play in Fudenberg and Levine (1995) and the references to the work of Blume (1993) therein. Lastly, it is important to mention that Freund and Schapire (1999) use the Hedge algorithm to approximately solve the mixed extension of two-person zero-sum games where both players have finitely many pure strategies. However, for games

in which one player has an infinitely many pure strategies, some other algorithms have been suggested in the literature; see, for example, Filar and Raghavan (1982) and our discussion of related literature below.

3. We illustrate the usefulness of the suggested algorithm by analyzing a simple and stylyzed *binary treatment choice problem with partial identification* based on the work of Stoye (2012). We use this simple, well-known example to compare the output of the Hedge algorithm with known exact solutions of two types of minimax problems: minimizing worst-case regret and solving an ex-ante Robust Bayes problem using the class of priors in Giacomini and Kitagawa (2021). Our analysis of the stylized treatment choice problem in Stoye (2012) shows that, with an appropriate selection of the I decision rules, the ϵ-minimax decision rule obtained by the Hedge algorithm is very similar to the exact solutions of the minimax problem that optimizes over *all* decision rules. We emphasize, however, that the goal of this paper is not to argue that when I is large (and ϵ is small) one can approximate the minimax value of a problem that optimizes over all decision rules. While this can be done under some assumptions (and can be verified numerically in some of our examples), the computational complexity results of Du and Pardalos (1995) and Daskalakis et al. (2021) suggest that such an approach will not scale well (and that there is no approach that does, without imposing additional restrictions). Thus, if the original minimax problem is defined over infinitely many decision rules, our suggestion is to *change the goalpost* by considering a problem that, instead, only optimizes over the (possibly randomized) choice of I candidate decision rules. As suggested in Dominitz and Manski (2024), there are some problems in which it is not difficult to think about a few, relatively simple statistical decision rules that researchers could find attractive. Once that the goalpost has been changed, one can use our algorithm to obtain an ϵ-minimax solution for the new problem.

4. Finally, we present an empirical application to the problem of optimally selecting *sites* to

maximize the external validity of an experimental policy evaluation. This *site selection problem* has been recently introduced in the work of Gechter, Hirano, Lee, Mahmud, Mondal, Morduch, Ravindran, and Shonchoy (2024). When the policy maker is restricted to select only one site for experimentation, the output of the Hedge algorithm is a selection probability for each of the sites available for experimentation. In the empirical application (we provide further details below), the ϵ-minimax solution is very different to uniform random sampling. This difference suggests that selecting uniformly at random where to experiment need not always maximize the external validity of an experimental policy evaluation.

Related Literature: Different algorithms have been suggested for approximating the solutions of minimax problems like the ones considered in this paper. Some classical references include Troutt (1978); Filar and Raghavan (1982); Kempthorne (1987); Chamberlain (2000); Elliott, Müller, and Watson (2015). One important difference between our work and this existing literature is that—once a desired approximation error ϵ has been selected, and once the bound M on the risk function has been obtained—there are no further inputs that the user needs to specify in order to run the algorithm. This means that we are explicit about the number of iterations, step size, and also the initial condition. Importantly, we are able to guarantee that, upon termination after finitely many rounds, the algorithm provably generates an ϵ-minimax rule—in the sense of Ferguson (1967)—provided our assumptions are satisfied. This is possible because the rich literature studying the Hedge algorithm (and Multiplicative Weights more broadly) and also the literature on convex optimization has explicit performance guarantees for the algorithm at any given iteration.[4]

Relatedly, there is also recent interest in approximating the solution of minimax problems in which the strategies for both the statistician and nature are parameterized via neural networks, with

[4]In this sense, one could say that our work follows closely the literature on *convergence analysis* in convex optimization; see, for example, the work of Nemirovski and Yudin (1983). This means that we try to be as explicit as we can on the computational resources that we can credibly rely on (in our case, an oracle that finds the worst-case point in the parameter space for a given decision rule) and then we try to make use of these computational resources as efficiently as possible (in our case, this means that we attempt to call the available oracle as infrequently as possible to obtain an approximation).

weights that are updated iteratively using versions of what is called *subgradient ascent-descent*; see the recent work of Luedtke, Carone, Simon, and Sofrygin (2020) and also Luedtke, Chung, and Sofrygin (2021). These algorithms where two players use subgradient descent are similar to the approaches used when optimizing Generative Adversarial Networks (GANs); see, for example, Kaji, Manresa, and Pouliot (2023). These subgradient ascent-descent algorithms are also commonly used to approximate the equilibrium of two-person zero-sum games by invoking simultaneous no-regret dynamics; see, for example, Section 3.1 in Lewis and Syrgkanis (2018) and the references therein. Convergence rates for these subgradient ascent-descent algorithms, as well some performance guarantees for a finite number of iterations, are available under some conditions. It is known, however, that the (approximate) stationary points of these gradient ascent-descent algorithms are not necessarily ϵ-minimax strategies. Instead, they are close to what the literature refers to as *local min-max* solutions; see the seminal work of Daskalakis et al. (2021). As we discuss in the conclusion, it would be interesting to further explore the differences between ϵ-minimax strategies and the notion of a local min-max point.

Empirical Application: In our main application, we study the *site selection problem* in Gechter et al. (2024) and Egami and Lee (2024). Broadly speaking, a policy maker wishes to experimentally evaluate the effects of a new policy, with the end goal of recommending its implementation on a set of different *sites*. There are two types of sites: *policy-relevant* and *experimental* sites. There are also covariates $X_s \in \mathbb{R}^d$ available for each site. The *site selection problem* asks the following question: if the policy maker can pick at most k experimental sites, what are the sites that optimize external validity?

Two recent papers have provided an answer to this question. Gechter et al. (2024) use an elegant decision-theoretic framework to recommend a nonrandomized selection of sites with the goal of maximizing the average welfare of the policy maker. Montiel Olea et al. (2024a) use the framework of Gechter et al. (2024) to show that, under some conditions, selecting the k-sites with the most representative covariates (in a sense they make precise) is minimax (welfare) regret optimal

(restricting the policy maker to consider only nonrandomized selection of sites).

We obtain these two recommendations in the specific context of the selection of candidate *migration corridors* for conducting randomized evaluations in Bangladesh. We assume that a hypothetical policy maker is interested in selecting only one site to evaluate an encouragement design where poor rural households with family members who had migrated to a larger urban destination receive a 30–45 minute training about how to register and use a mobile banking service to send instant remittances back home; see Lee, Morduch, Ravindran, Shonchoy, and Zaman (2021) and Gechter et al. (2024).

The site selected based on average welfare and minimax (welfare) regret is not the same. We consider both of them—along with an experimental site where the encouragement design of interest has already been experimentally evaluated—to give the policy maker three concrete decision rules ($I = 3$). We then let the policy maker choose randomly over them to determine where to experiment in order to minimize worst-case regret. We note that although, in general, randomly choosing where to experiment could be viewed as contrived, the randomized selection of experimental sites is typically thought of as the first-best in applied work. For instance, Duflo, Glennerster, and Kremer (2007) note that *"the external validity of randomized evaluations for a given population (say, the population of a country) would be maximized by randomly selecting sites and, within these sites, by randomly selecting treatment and comparison groups."*[5] In this application, the parameter space consists of functions that control the treatment effect heterogeneity across sites. We impose a Lipschitz constraint on these functions (we provide details later), and consider the worst-case regret over this space.

The two main lessons from our application are the following. First, choosing uniformly at random where to experiment does not tend to be ϵ-minimax optimal. Instead, the ϵ-minimax solution seems to adjust the probability of sampling a site based on its baseline covariates. For instance, in our

[5]Although they remark that this is almost never done because *"randomized evaluations are typically performed in 'convenience' samples, with specific populations."*

application the experimental site whose covariates are closest (on average) to the covariates of the policy-relevant sites is sampled with the highest probability. Second, there seem to be some cases (for example, when one experimental site is closest to each of the policy-relevant sites) in which the ϵ-minimax solution places almost probability one such a site. This suggests that maximizing the external validity of a randomized evaluation need not be accomplished by randomly selecting sites. Instead, it is possible that using baseline covariates, experimenting on the most representative site could be useful for policy purposes. Our approach thus provides an algorithm for deciding how to *randomly* select sites to optimize external validity, taking into account information about baseline covariates.

Outline: The rest of the paper is organized as follows. Section 2 introduces notation, main assumptions, and presents the convex programming representation of the minimax problems analyzed herein. Section 3 defines an ϵ-minimax decision rule and presents the algorithm. Section 4 applies the algorithm to two illustrative examples that involve solving treatment choice problems with partial identification. Our algorithm is the used to solve for ϵ-minimax regret optimal rules; but we also argue that it can be applied to solve other minimax problems, such as (ex-ante) Robust Bayes analysis with the priors suggested by Giacomini and Kitagawa (2021). Section 5 presents the main application. Section 6 discusses some extensions. Section 7 concludes.

2 Minimax Problems

2.1 Notation

A decision maker must choose an action a that belongs to some set $\mathcal{A}$. Prior to choosing the action, he observes data: the realization of a random variable X taking values in a set $\mathcal{X}$. A data-driven choice of action is summarized by a *decision rule*: a mapping from data to actions, which is herein denoted by the function $d : \mathcal{X} \to \mathcal{A}$.

It is common to allow the decision maker to consider every (measurable) function d as a decision rule. However, we restrict our analysis to the case in which the decision maker only considers I decision rules that belong to the finite set $\mathcal{D} \equiv \{d_1, \ldots, d_I\}$. Our motivation is that, while it is always theoretically interesting to look for the *best* overall decision rule, there are situations in which it is equally desirable to *"evaluate the performance of relatively simple statistical decision functions that researchers use in practice"* (Dominitz and Manski, 2024) and choose optimally among them.

An important aspect of our analysis is that we allow the decision maker to choose randomly from the set of decision rules $\mathcal{D}$ and we represent such a random choice by an element in the $I-1$ simplex:

$$\Delta(\mathcal{D}) \equiv \left\{ (p_1, ..., p_I) \in \mathbb{R}^I \;\middle|\; \sum_{i=1}^{I} p_i = 1, p_i \geq 0 \right\}.$$

It is well known that allowing the decision maker to choose randomly is usually to his advantage.[6] Moreover, there are two additional reasons why we would like to allow for the possibility of *randomization*. The first one is that in the main application we will consider in the paper (the *site selection* problem described in Section 5), the random choice of experimental sites is viewed as the default practice in applied work. The second reason is that, as we will explain in Section 3 (Remark

[6] Consider a "matching pennies" game with two players, each with two actions: left and right. Suppose that column player gets M when matched and $-M$ when unmatched. If neither player is allowed to choose actions randomly, the worst-case payoff obtained by the column player is $-M$ regardless of the action chosen. If the column player can randomize, but the row player cannot, the worst-case payoff for the column player if he chooses each action at random with probability $1/2$ is zero.

4), allowing for random choice of actions can reduce the computational burden of selecting a good decision rule.

A *risk function* is used to summarize the performance of each decision rule $d_i \in \mathcal{D}$. This performance is contingent on the data generating process, which we parameterize by an element θ belonging to some space Θ. Thus, we write the risk function of each decision rule $d \in \mathcal{D}$ as a mapping $R : \mathcal{D} \times \Theta \to \mathbb{R}$. We refer to θ as a parameter, and to Θ as the parameter space. We are particularly interested in the case in which Θ is an infinite set; for example, when Θ equals all of $\mathbb{R}^d$. We also want to allow for the possibility that each element in the parameter space is an infinite dimensional object (for example, when θ itself is a function). We impose the following assumption on the risk function:

Assumption 1. There exists a known constant $0 < M < \infty$ such that for any $d \in \mathcal{D}$ and $\theta \in \Theta$, $0 \leq R(d, \theta) \leq M$.

In Section 4 we explain how this assumption can be verified for each of the illustrative examples we consider. We view Assumption 1 as a minimal regularity condition for the minimax problem to be well-behaved. We also note that the assumption holds if each of the I decision rules under consideration has a finite worst-case risk.

In a slight abuse of notation, we extend the original domain of the risk function—which was defined over decision rules in $\mathcal{D}$—to all possible random selections in $\Delta(\mathcal{D})$. We do this by defining, for any $p \in \Delta(\mathcal{D})$ and $\theta \in \Theta$, the function:

$$R(p, \theta) \equiv \sum_{i=1}^{I} p_i R(d_i, \theta). \tag{1}$$

We view a decision problem as a triplet $(\mathcal{D}, \Theta, R(\cdot, \cdot))$ and we define the "minimax value" of the decision problem as the scalar

$$\bar{v} \equiv \inf_{p \in \Delta(\mathcal{D})} \sup_{\theta \in \Theta} R(p, \theta). \tag{2}$$

A random selection $p^* \in \Delta(\mathcal{D})$ is said to be a *minimax* decision rule if

$$\sup_{\theta \in \Theta} R(p^\star, \theta) = \bar{v}. \tag{3}$$

The use of the minimax criterion as a solution concept in statistical decision problems is traditional, dating back to Wald (1950). Manski (2021) argues that the primary challenge to use the minimax criterion and Wald (1950)'s statistical decision theory is computational.

2.2 Exact minimax solutions via convex programming

We first show that the minimax solution of the decision problems considered in this paper can be computed via convex programming. This observation is based on an analogous result in Chamberlain (2000); see Equation 5, p. 630, and the discussion therein. The argument is as follows. For $p \in \Delta(\mathcal{D})$, define the nonlinear function

$$f(p) \equiv \sup_{\theta \in \Theta} R(p, \theta). \tag{4}$$

This function is the upper envelope—over all possible values in the parameter space—of the risk of p.

Lemma 1. *Suppose Assumption 1 holds. The function $f : \Delta(\mathcal{D}) \to \mathbb{R}$ is convex and Lipschitz continuous w.r.t. $\|\cdot\|_1$ (with constant at most M). Furthermore, fix an arbitrary $p_0 \in \Delta(\mathcal{D})$. If there exists $\theta_0 \in \Theta$ such that $R(p_0, \theta_0) = f(p_0)$, then the vector g_0 in $\mathbb{R}^I$ given by*

$$g_0 \equiv (R(d_1, \theta_0), \ldots, R(d_I, \theta_0))^\top. \tag{5}$$

is a subgradient of f at p_0.[7]

[7] If $f : \Delta(\mathcal{D}) \to \mathbb{R}$ is convex, a vector g_0 is said to be a subgradient of f at a point p_0 if $f(p) \geq f(p_0) + g_0^\top (p - p_0), \forall p \in \Delta(\mathcal{D})$. See pp. Rockafellar (1970) 214-215.

Proof. The convexity of $f(\cdot)$ follows from Chamberlain (2000). The Lipschitz continuity follows from the definition of $f(\cdot)$. We provide a detailed proof in Appendix A.1. □

Lemma 1 shows that solving the minimax problem in (2) can be viewed as a nonlinear convex program over the $(I-1)$ simplex. A popular routine for solving this type of convex optimization problems is the mirror descent method of Nemirovski and Yudin (1983). We briefly describe a textbook version of this routine (Bubeck et al., 2015, Section 4.3, p. 301). To make sure that the subgradient is well defined, we make the following assumption.

Assumption 2. For any $p \in \Delta(\mathcal{D})$, there exists $\theta_p \in \Theta$ such that

$$\sum_{i=1}^{I} p_i R(d_i, \theta_p) = \sup_{\theta \in \Theta} \sum_{i=1}^{I} p_i R(d_i, \theta).$$

The assumption says that for any $p \in \Delta(\mathcal{D})$ it is possible to find an element θ_p such that $R(p, \theta_p) = f(p)$. This means that there is an algorithm that is capable to i) evaluate the function $f(p)$ and to ii) find a maximizer that evaluates to $f(p)$.

Assumption 2 also imposes substantial regularity to search for a "worst-case" parameter in the parameter space, and requires that worst-case risk is attained. Later we discuss the extent to which this assumption can be relaxed, by requiring that we can get a δ-approximation to $f(p)$. See Remark 3.

Mirror Descent Routine: The following is a typical mirror descent routine for finding the minimum of (4) over the simplex $\Delta(D)$.[8] Initialize $w_0 \in \mathbb{R}^I$ to be the vector that contains ones in all of its entries. We will denote such vector by $\mathbf{1}$. Fix a step-size $\eta > 0$. For every $t \in \mathbb{N}$:

[8]The routine is taken from Bubeck et al. (2015) (Section 4.3, p. 301), where the mirror map is chosen to be the negative entropy $\phi(x) = \sum_{i=1}^{n} x_i \log x_i$, the routine $\nabla\phi(x_{t+1}) = \nabla\phi(x_t) - \eta\nabla f(x_t)$ becomes $x_{t+1,i} = x_{t,i} \exp(-\eta \nabla f(x_t)_i)$. We simply adjust the notation to our problem.

1. Set

$$\begin{aligned}\phi_t &\equiv \mathbf{1}^\top w_{t-1} \\ &= \sum_{i=1}^{I} w_{i,t-1}.\end{aligned}$$

2. Obtain a point $p_t \in \Delta(\mathcal{D})$ by computing for each $= 1, \ldots, I$:

$$p_{i,t} \equiv \frac{w_{i,t-1}}{\phi_t}.$$

3. Given $p_t = (p_{1,t}, ..., p_{I,t})^\top \in \Delta(\mathcal{D})$, find $\theta_t \in \Theta$ such that

$$\sum_{i=1}^{I} p_{i,t} R(d_i, \theta_t) = \sup_{\theta \in \Theta} \sum_{i=1}^{I} p_{i,t} R(d_i, \theta).$$

   Such a point in the parameter space exists by Assumption 2.

4. Define the vector

$$g_t \equiv (R(d_1, \theta_t), \ldots, R(d_I, \theta_t))^\top.$$

   By Lemma 1, this vector is a subgradient of f at p_t if Assumption 2 holds.

5. Update the weights w_{t-1} with the coordinate-by-coordinate multiplicative rule:

$$w_{i,t} \equiv w_{i,t-1} \cdot \exp\left(-\eta\, g_{i,t}\right).$$

Under this mirror descent routine, the vector p_t gets updated as $p_{i,t+1} = w_{i,t}/\phi_{t+1}$. The mirror descent routine typically uses $(1/T)\sum_{t=1}^{T} p_t$ (and not the last p obtained in the iteration) as the approximate minimizer.

We note that the connection to convex programming is helpful, but should not be viewed

as a computational panacea. Evaluating the objective function of the convex program and its subgradient could remain computationally costly. Verifying Assumption 2 requires optimizing a nonlinear function over the model's parameter space. As noted above, the solution of such nonlinear optimization problem is used to evaluate the subgradient of the objective function $f(p)$. It is known that the rate of convergence of mirror descent for convex problems in the simplex improves over regular subgradient descent (Bubeck et al., 2015, Section 4.3). However, in order to guarantee that a minimax solution has been found one would typically need to run the routine above for infinitely many epochs. This could be computational costly if one needs to fulfill Assumption 2.

Moreover, the use of (nonlinear) convex programming to find exact minimax solutions of general decision problems might mean that, in some cases, one incurs in a higher computational cost than needed. For example, there are some classical results in the game theory literature (Dantzig, 1951; Adler, 2013; Owen, 2013, Section III.1, p. 36) that show that when Θ has finitely many elements, it is possible to express the minimax problem in (2) as a linear program. It is possible to show that every minimax problem where the statistician chooses randomly among I alternatives can be expressed as a linear program in $I + 1$ variables with as many constraints as elements in Θ (even if Θ has infinitely many elements). This suggests that, if one is willing to make more assumptions, there might be better algorithms for solving the minimax problems of interest. There might also be better algorithms to find an ϵ-minimax decision rule; see, for example, Section 5.2, p. 317 of Bubeck et al. (2015).

3 Approximate Solutions for Minimax Problems

Computing the exact minimax rule of a decision problem could be computationally costly. This section presents a definition of an *approximate* minimax solution, and an off-the-shelf implementation of mirror descent—used routinely in different areas of computer science and machine learning—to finding it.

3.1 ϵ-Minimax Decision Rules

Definition 1. [Ferguson (1967), p. 44] A random selection $p^\star_\epsilon \in \Delta(\mathcal{D})$ is an "ϵ-minimax" decision rule for the decision problem $(\mathcal{D}, \Theta, R(\cdot,\cdot))$ if

$$\sup_{\theta\in\Theta} R(p^\star_\epsilon, \theta) \leq \inf_{p\in\Delta(\mathcal{D})} \sup_{\theta\in\Theta} R(p,\theta) + \epsilon = \bar{v} + \epsilon.$$

We note that the risk of an ϵ-minimax decision rule is smaller—up to an additive factor of size ϵ—than the worst-case risk of any other decision rule. That is:

$$R(p^\star_\epsilon, \theta) \leq \sup_{\theta\in\Theta} R(p,\theta) + \epsilon,\ \forall\theta\in\Theta,\ \forall p\in\Delta(\mathcal{D}).$$

The definition of a minimax decision rule further implies that

$$\bar{v} \leq \sup_{\theta\in\Theta} R(p^\star_\epsilon, \theta) \leq \bar{v} + \epsilon.$$

3.2 Hedge Algorithm for finding ϵ-Minimax Rules

In this section, we show that running the mirror descent routine described in Section 2.2 can be used to provably find ϵ-minimax solutions for the decision problems herein considered. To be more explicit, consider the following pseudocode for mirror descent, but stopped after T epochs.

Our concrete suggestion is to set the step size to $\eta \equiv \epsilon/M^2$ and to stop the routine after $T = \lceil 2M^2 \ln(I)/\epsilon^2 \rceil$ epochs. This routine is known as the Hedge Algorithm (a particular case of the *Multiplicative Weights* update method). The Multiplicative Weights update method is a popular algorithm in computer science that has found different applications in machine learning; see Arora et al. (2012). The specific version of the Multiplicative Weights algorithm used in this paper uses an exponential function of each of the coordinates of the gradient to update the weights and is known as the Hedge algorithm. See Section 2.1 in Arora et al. (2012).

Algorithm 1 Mirror Descent, stopped after T epochs (a.k.a. *Hedge algorithm*).

Input: Step-size $\eta > 0$; and number of epochs $T \in \mathbb{N}$.
Initialize $w_0 \in \mathbb{R}^I$ by setting $w_{i,0} = 1$ for all $i \in \{1, \ldots, I\}$.
for $t = 1, 2, \ldots$ **do**
  Compute $\phi_t := \sum_{i=1}^{I} w_{i,t-1}$
  For each $i \in \{1, \ldots, I\}$, compute

$$p_{i,t} := \frac{w_{i,t-1}}{\phi_t}$$

  Find $\theta_t \in \Theta$ such that

$$\theta_t := \arg\sup_{\theta \in \Theta} \sum_{i=1}^{I} p_{i,t} R(d_i, \theta)$$

  Define the vector

$$g_t := (R(d_1, \theta_t), \ldots, R(d_I, \theta_t))^\top$$

  Consider the multiplicative weights update:

$$w_{i,t} := w_{i,t-1} \cdot \exp(-\eta \cdot g_{i,t})$$

end for

For any nonnegative real number x, let $\lceil x \rceil$ denote the "ceiling function"; that is smallest integer larger than x. Our main result is the following.

Theorem 1. *Suppose Assumptions 1-2 hold. If $\epsilon \leq M$, $\eta \equiv \epsilon/M^2$, and $T \equiv \lceil 2M^2 \ln(I)/\epsilon^2 \rceil$, then the random choice of decision rules that assigns probability*

$$p_i^\epsilon \equiv \frac{1}{T} \sum_{t=1}^{T} p_{i,t}$$

to each decision rule d_i—where p_t corresponds to the t-th iteration of the mirror descent routine in Algorithm 1—is ϵ-minimax in the sense of Definition 1. Moreover

$$\bar{v}^\epsilon \equiv \frac{1}{T} \sum_{t=1}^{T} \left(\sum_{i=1}^{I} p_{i,t} R(d_i, \theta_t) \right),$$

where θ_t corresponds to "nature's best response" in the t-th iteration of the mirror descent routine

in Algorithm 1, is an ϵ-approximation to $\bar{v}$: that is, $\bar{v} \leq \bar{v}^{\epsilon} \leq \bar{v} + \epsilon$.

Proof. We present two different proofs of Theorem 1. First, in Appendix A.2 we adapt (and extend) the results in Arora et al. (2012) concerning the use of the Hedge algorithm to approximate the minimax solution of two-player zero-sum games where both players have finitely-many pure strategies. In adapting and extending their results, we improve the number of epochs by a factor of two. Second, in Appendix A.3 we apply results from the convex optimization literature regarding the general convergence analysis of mirror descent available; in particular, we use Theorem 4.2 from Bubeck et al. (2015). □

Theorem 1 presents a concrete computational strategy to approximately solve the statistical decision problems considered in this paper. The only "tuning" parameter that needs to be chosen is ϵ, which controls the approximation error. We note that in cases where is difficult to commit to a value of ϵ explicitly, one can *solve* for the value of ϵ if there is a specific target for the runtime of the algorithm and we know the time it takes for each iteration to run.

Theorem 1 shows that the average value of p_t (over the T rounds) provides an ϵ-minimax rule. The theorem also shows that the average worst-case payoff obtained in each round is an ϵ-approximation to the minimax value of the decision problem.

We make some remarks about Theorem 1.

Remark 1 (Optimality of the Hedge Algorithm). There exist different results in the computer science literature providing lower bounds for the *regret* of the Multiplicative Weights update method in problems where a decision maker chooses randomly among I alternatives; see Section 4 in Arora et al. (2012) and also Gravin, Peres, and Sivan (2016). We note, however, that these regret bounds do not speak directly to the question of whether there exists another iterative algorithm for convex optimization of an M-Lipschitz function over the simplex that could find an ϵ-minimizer in less epochs than the Hedge Algorithm.

Fortunately, Proposition 4.2 in Ben-Tal et al. (2001) addresses this issue. We briefly review their notation and summarize their findings.

Let $\mathcal{F}(M, I)$ denote the collection of all minimization problems of a convex, Lipschitz function f (with respect to $\|\cdot\|_1$ and with constant at most M) over the $I-1$ simplex. Since the minimization problem is indexed entirely by the function f, we denote the elements of $\mathcal{F}(M, I)$ succinctly as f. Let $\partial f(p)$ denote the subdifferential of f (the set containing all subgradients) at p. Let A be a *first-order* iterative algorithm that successively generates points $p_t(A, f) \in \Delta^{I-1}$ and approximate solutions $p^t(A, f)$. We restrict the class of algorithms by requiring both p_t and p^t to be deterministic functions of first-order information about f; namely the history of evaluations of f and its subdifferential: $\{f(p_s), \partial f(p_s)\}_{s=1}^{t-1}$. For the starting search point or *initial condition*, p_1, we require it to be chosen independently of the function f. We denote the class of deterministic, iterative, first-order algorithms as $\mathcal{A}$. Given a tolerance ϵ, define the *complexity* of the class of optimization problems $\mathcal{F}(M, I)$ with respect to algorithm A as the function

$$\text{Complexity}_A(\epsilon; \mathcal{F}(M, I)) \equiv \inf\{T \in \mathbb{N} \mid f(p^t(A, f)) - \inf_{p \in \Delta^{I-1}} f(p) \leq \epsilon, \quad \forall t \geq T, f \in \mathcal{F}(M, I)\}.$$

Define the complexity of the family of optimization problems in $\mathcal{F}(M, I)$ as

$$\text{Complexity}(\epsilon; \mathcal{F}(M, I)) \equiv \inf_{A \in \mathcal{A}} \text{Complexity}_A(\epsilon; \mathcal{F}(M, I)).$$

Proposition 4.2 in Ben-Tal et al. (2001) shows that

$$\text{Complexity}(\epsilon; \mathcal{F}(M, I)) \geq O(1) \min\{M^2/\epsilon^2, I\}.$$

Therefore, the smallest number of epochs required by *any* iterative, first-order algorithm for (convex optimization over the $(I-1)$ simplex of a Lipschitz function with constant at most M (with respect to $\|\cdot\|_1$ norm) is $O(1)M^2/\epsilon^2$, provided $\epsilon \geq M/\sqrt{I}$. Thus, there is a sense in which the recommended

algorithm, and the suggested number of epochs, are optimal up to the logarithmic factor $\ln(I)$.

Remark 2 (Finite Θ). When Θ has J elements, obtaining an exact minimax solution could be done via the linear program (Dantzig, 1951; Adler, 2013; Owen, 2013, Section III.1, p. 36). The computational cost of using the fastest solver for linear program to find an exact minimax solution can be shown to be of order $(1+J+I)^{2.055}$ time.[9] We note that Algorithm 1 makes $\lceil 2M^2 \ln(I)/\epsilon^2 \rceil$ calls to nature's oracle. Suppose that the runtime of the oracle is $r(I, J)$. In each round, the algorithm evaluates the risk of the I actions available to the decision maker. Thus, the runtime of the algorithm is of order

$$M^2 I \ln(I) r(I, J)/\epsilon^2.$$

If the calls to the oracle that computes nature's best response are not expensive, and if M/ϵ^2 is not too large, the time needed in order to compute the approximate solution to the minimax problem could be smaller than that time needed to obtain the exact solution. We also note that when Θ has J elements, there might also be better algorithms to find an ϵ-minimax decision rule; see, for example, the Saddle-Point Mirror Prox algorithm discussed in Section 5.2, p. 317 of Bubeck et al. (2015).

Remark 3 (Approximate Oracle (or approximate best response)). It is possible to extend the proof to the case that θ_t is not the exact oracle, but an approximate one. More precisely, consider θ_t^δ such that

$$\left(\sup_{\theta \in \Theta} \sum_{i=1}^{I} p_{i,t} R(d_i, \theta) \right) - \delta \leq \sum_{i=1}^{I} p_{i,t} R(d_i, \theta_t^\delta) \leq \sup_{\theta \in \Theta} \sum_{i=1}^{I} p_{i,t} R(d_i, \theta).$$

This extension can be (roughly) completed by doing the following adjustments to the proof.

Firstly, in Step 2, we say that

$$\sup_{\theta \in \Theta} \sum_{i=1}^{I} \tilde{p}_i^\star R(d_i, \theta) \leq \frac{1}{T} \sum_{t=1}^{T} \left(\sum_{i=1}^{I} p_{i,t} R(d_i, \theta_t) \right).$$

[9] Jiang, Song, Weinstein, and Zhang (2020) show that the fastest known LP solver for general (dense) linear programs can solve such a program in an order of approximate $(1+I+J)^{2.055}$ time.

We can change this to

$$\sup_{\theta\in\Theta}\sum_{i=1}^{I}\tilde{p}_i^{\star}R(d_i,\theta)\leq\frac{1}{T}\sum_{t=1}^{T}\left(\sum_{i=1}^{I}p_{i,t}R(d_i,\theta_t^{\delta})+\delta\right).$$

Then, by Step 1, this is bounded by above by

$$\frac{1}{T}\sum_{t=1}^{T}\sum_{i=1}^{I}p_{i,t}R(d_i,\theta_t^{\delta})+\frac{\epsilon}{2}+\frac{\ln(I)}{T}\left(\frac{M^2}{\epsilon}\right)+\delta.$$

Then, subbing in to Step 2, the left hand side is bounded by above by

$$\bar{v}+\frac{\epsilon}{2}+\frac{\ln(I)}{T}\left(\frac{M^2}{\epsilon}\right)+\delta.$$

Choosing T as we have done gives an $\epsilon+\delta$ approximation.

Remark 4 (Minimax Solution without randomization). Finally, note that even if one were interested in computing the minimax optimal rule among $\{d_1,\ldots,d_I\}$, one would need I calls to the oracle (one for computing the worst-case performance of each rule). Surprisingly, the ϵ-minimax solution among randomized rules calls the oracle $\lceil 2M^2\ln(I)/\epsilon^2\rceil$ times. When I is large, the difference could be substantial.

Remark 5 (Least favorable distribution). The statistical decision problem we are interested in can be interpreted as a two-player zero-sum game: the two players are 1) the statistician who has *pure strategies* $\mathcal{D}=\{d_1,d_2,...,d_I\}$, and 2) "nature", whose set of pure strategies is given by the parameter space Θ. The payoff function is $R(d_i,\theta)$. In the mixed extension of the game, in each round, the statistician first chooses a mixed strategy $p_t\in\Delta(\mathcal{D})$, and then nature responds with a choice θ_t. Surprisingly, one implication of our proof of Theorem 1 is that the routine in Algorithm 1 not only gives an approximate minimax solution for the statistician, but also gives an approximate maximin solution for nature. In particular the empirical distribution of the sequence of nature's best responses, $\{\theta_t\}_{t=1}^{T}$, is "ϵ-maximin" strategy for the nature. See Appendix B for detailed explanation.

4 Illustrative Examples

4.1 ϵ-Minimax Regret in a treatment choice problem with partial identification

Consider the following example taken from Stoye (2012) and Yata (2021). A policy maker uses experimental data to decide whether to implement a new policy in a target population of interest. The treatment effect of action $a = 1$ is $\mu^* \in \mathbb{R}$, while the effect of action $a = 0$ is normalized to be equal to 0. Thus, the policy maker's expected payoff equals $W(a, \mu^*) \equiv a \cdot \mu^*$.

The data available to the policy maker is an estimated treatment effect, $\hat{\mu}$, for the experimental population. The policy maker assumes that

$$\hat{\mu} \sim N(\mu, \sigma^2), \tag{6}$$

where $\sigma > 0$ is known and where $\mu \in \mathbb{R}$ is the true effect of the policy in the population where the experiment was conducted. The policy maker is concerned about the external validity of the experiment at hand. This is captured by allowing the effect of the policy in the experimental population (μ) to be different from the effect in the target population (μ^*). The policy maker is willing to work under the assumption that $|\mu^* - \mu| \leq k$ for some known $k \geq 0$. In this example, $\theta = (\mu, \mu^*)^\top$ and $\Theta \equiv \{(\mu, \mu^*) \in \mathbb{R}^2 \mid |\mu - \mu^*| \leq k\} \subseteq \mathbb{R}^2$.

A decision rule for the policy maker is a mapping $d : \mathbb{R} \to [0, 1]$ from the observed experimental data (6) to an action $a \in [0, 1]$. The action is interpreted as the fraction of the target population that will be treated. Consider the regret loss associated to $W(a, \mu^*)$ given by $L(a, \theta) \equiv \mu^*[\mathbf{1}\{\mu^* \geq 0\} - a]$. Define the risk function

$$R(d, \theta) \equiv \mathbb{E}_\theta[L(d, \theta)].$$

Exact Minimax Solution Over all Decision Rules: Let $\mathcal{D}^*$ denote the set of all decision

rules. Stoye (2012) derived a solution to the minimax (regret) problem

$$\inf_{d \in \mathcal{D}^*} \sup_{\theta \in \Theta} R(d, \theta), \tag{7}$$

as a function of (σ^2, k). Stoye (2012) showed that when $k \geq \sqrt{\pi/2}\sigma$, Equation (7) equals $k/2$. Montiel Olea, Qiu, and Stoye (2024b) further showed that when $k \geq \sqrt{\pi/2}\sigma$ there are infinitely many minimax-regret optimal rules. One such solution takes the form

$$d^{\star}_{MQS}(\hat{\mu}) = \begin{cases} 0, & \hat{\mu} < -\rho^{\star}, \\ \frac{\hat{\mu}+\rho^{\star}}{2\rho^{\star}}, & -\rho^{*} \leq \hat{\mu} \leq \rho^{*}, \\ 1, & \hat{\mu} > \rho^{*}, \end{cases}$$

where $\rho^* \in (0, k)$ uniquely solves the nonlinear equation:

$$\left(\frac{1}{2k}\right)\rho^{*} - \frac{1}{2} + \Phi\left(-\frac{\rho^{*}}{\sigma}\right) = 0, \tag{8}$$

see Theorem 3 in Montiel Olea et al. (2024b).

Approximate Minimax Regret Solution over a Class of Threshold Rules: Suppose that instead of considering all decision rules, we focus on a class $\mathcal{D} \subset \mathcal{D}^*$ that contains only "threshold" rules; that is, decision rules of the form

$$d_i(\hat{\mu}) \equiv \mathbf{1}\{\hat{\mu} \geq c_i\},$$

where $c_i \in \mathbb{R}$. For concreteness, we consider 500 different values for c_i equally spaced in the interval $[-k, k]$. These threshold rules seem natural for this problem. For example, if one observes a realization of $\widehat{\mu} \geq k$, any of these rules would suggest to implement the policy at scale.

Algebra shows that, in this example, the largest worst-case risk among all threshold rules in $\mathcal{D}$

is bounded above by $M \equiv \sigma \max_{x \geq 0} x\Phi\left((2k/\sigma) - x\right)$, where $\Phi(\cdot)$ denotes the Normal c.d.f.[10] Since the expected loss is nonnegative, Assumption 1 is satisfied.

We can also show that, for a given $p \in \Delta(\mathcal{D})$, the values $(\mu, \mu^*) \in \Theta$ that verify Assumption 2 can be obtained by solving three optimization problems. Define the parameter μ_+^* to be the solution of the problem

$$\max_{\mu^* \geq 0} \mu^* \left(\sum_{i=1}^{I} p_i \Phi\left(\frac{c_i - \mu^*}{\sigma} + \frac{k}{\sigma} \right) \right),$$

and $\mu_+ \equiv \mu_+^* - k$. Define the parameter μ_-^* to be the solution of the problem

$$\max_{\mu^* \leq 0} -\mu^* \left(\sum_{i=1}^{I} p_i \Phi\left(\frac{\mu^* - c_i}{\sigma} + \frac{k}{\sigma} \right) \right),$$

and $\mu_- = \mu_-^* + k$. Set θ_p to be the maximizer of

$$\{R(p, \mu_+, \mu_+^*), R(p, \mu_-, \mu_-^*)\}.$$

Since we have verified Assumption 1 and 2, we proceed to applying Algorithm 1. We consider the case in which $\sigma = 1$ and $k = 2$. The value of the bound M is $M = 2.5294$. Since we know that the value of the problem in (7) is 1, we can set $\epsilon = .1$ (we ar willing to tolerate 10% relative error). We later discuss how to pick ϵ in more realistic problems in which there is no information about the minimax value. The number of epochs in Theorem 1 then becomes

$$T = \lceil 2M^2 \ln(I)/\epsilon^2 \rceil = 7,953.$$

The runtime of Algorithm 1 is about 30 seconds (on a personal ASUS Vivobook Pro 15 @ 2.5GHz Intel Core Ultra 9 185H). Figure (2) presents a comparison of d_{MQS}^* and the ϵ-minimax rule. The value of $\bar{v}^\epsilon$ is 1.0033.

[10]The formula corresponds to the worst-case risk of the rule that uses the threshold $c_i = k$ (or -k).

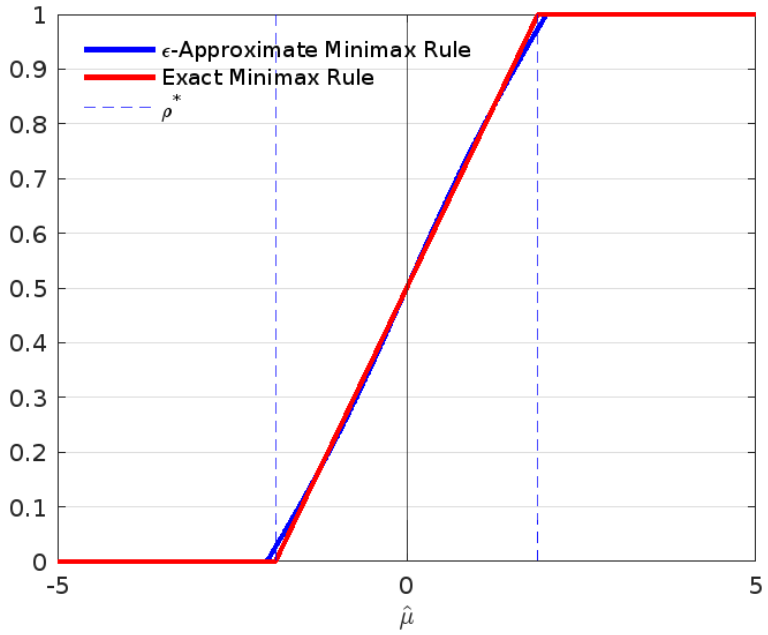

Figure 1: ϵ-Minimax Decision Rule via the Hedge algorithm. The graph is generated using $\sigma = 1$, $k = 2$. The value of ρ^* in Equation 8 is 1.8797.

4.2 ϵ-Robust Bayes rules in a treatment choice problem with partial identification

Consider the same example as in Section 4.1, but instead of focusing on minimax-regret optimality as in Stoye (2012), we are interested in computing ex-ante Robust Bayes rules as in Aradillas Fernández, Montiel Olea, Qiu, Stoye, and Tinda (2024).

Let π be a prior over (μ, μ^*). We are interested in obtaining the rule that minimizes worst-case expected risk over the class of priors suggested by Giacomini and Kitagawa (2021). We will denote this class of priors by Γ. Broadly speaking, the priors in this class fix a marginal prior over μ, but allow for arbitrary priors over $\mu^*|\mu$ (as long as the joint distribution over (μ, μ^*) is supported on Θ). For this example, we will first consider the "two-point prior" for μ analyzed in Aradillas Fernández et al. (2024). That is, we assume that the prior of μ is supported on the set $\mathcal{M} = \{-\bar{\mu}, \bar{\mu}\}$. We first assume that the policy maker has a discrete uniform prior π_μ on $\mathcal{M}$, meaning that $\pi_\mu(\mu = \bar{\mu}) = \pi_\mu(\mu = -\bar{\mu}) = 1/2$.

Just as we did in Section 4.1, we consider the regret loss $L(a,\theta) \equiv \mu^\star[\mathbf{1}\{\mu^\star \geq 0\} - a]$ and the risk function

$$R(d,\theta) \equiv \mathbb{E}_\theta[L(d,\theta)].$$

However, we are now interested in the average ("Bayesian") risk of a decision rule; which we define as

$$r(d,\pi) \equiv E_\pi[R(d,\theta)].$$

Let $\mathcal{D}^*$ be the set of all decision rules. The minimax problem of interest is thus

$$\inf_{d\in\mathcal{D}^*} \sup_{\pi\in\Gamma} r(d,\pi). \tag{9}$$

We follow the literature and refer to any decision rule that solves this problem as either ex-ante Γ-minimax or ex-ante Robust Bayes.

Aradillas Fernández et al. (2024) showed that, under some conditions, the problem in Equation (9) for the two-point priors on μ described before has infinitely many solutions. One such solution takes the form

$$d^\star(\hat{\mu}) = \begin{cases} 0, & \hat{\mu} < -\frac{\sigma^2\rho^\star}{\bar{\mu}} \\ \frac{\bar{\mu}\hat{\mu}+\sigma^2\rho^\star}{2\sigma^2\rho^\star}, & -\frac{\sigma^2\rho^\star}{\bar{\mu}} \leq \hat{\mu} \leq \frac{\sigma^2\rho^\star}{\bar{\mu}} \\ 1, & \hat{\mu} > \frac{\sigma^2\rho^\star}{\bar{\mu}} \end{cases},$$

where $\rho^\star$ uniquely solves

$$\int_0^1 \Phi\left(\frac{2\rho^\star x - \rho^\star - (\bar{\mu}/\sigma^2)}{\bar{\mu}/\sigma}\right) dx = \frac{-\bar{\mu}+k}{2k}.$$

We compare this Γ-minimax optimal rule with the ϵ-approximation obtained via the Hedge

algorithm. We again consider the class $\mathcal{D}$ of decision rules of the form

$$d_i = \mathbf{1}\{\hat{\mu} \geq c_i\},$$

where $c_i \in \mathbb{R}$. We again start with an equally spaced grid of 500 points over $[-k, k]$.

In order to apply the Hedge algorithm we extend the Bayes risk $r(d, \pi)$ to any element $p \in \Delta(\mathcal{D})$ by defining

$$r(p, \pi) \equiv \sum_{i=1}^{I} p_i r(d_i, \pi) = \sum_{i=1}^{I} p_i \mathbb{E}_\pi \left[R(d_i, \mu, \mu^\star)\right] = \mathbb{E}_\pi \left[\sum_{i=1}^{I} p_i R(d_i, \mu, \mu^\star)\right].$$

We note that Assumption 1 is satisfied with the same M as in Subsection 4.1. In order to verify Assumption 2, we note that the results in Giacomini and Kitagawa (2021) show that

$$\sup_{\pi \in \Gamma} \mathbb{E}_\pi \left[\sum_{i=1}^{I} p_i R(d_i, \mu, \mu^\star)\right]$$

equals

$$\mathbb{E}_{\pi_\mu} \left[\sup_{\mu^\star \in [\mu-k, \mu+k]} \sum_{i=1}^{I} p_i R(d_i, \mu, \mu^\star)\right] \equiv \mathbb{E}_{\pi_\mu} \left[\bar{\Lambda}(\mu, p_1, ..., p_I)\right].$$

This relation immediately gives the prior $\pi \in \Gamma$ associated to the worst-case Bayes risk of any vector $p \in \Delta(\mathcal{D})$. In particular, the subgradient used in the updates is

$$g_t^i = \pi_\mu \cdot R(d_i, \bar{\mu}, \bar{\mu}_t^\star) + \pi_\mu \cdot R(d_i, -\bar{\mu}, (-\bar{\mu})_t^\star),$$

where $\bar{\mu}_t^\star$ and $(-\bar{\mu})_t^\star$ are the corresponding values of $\mu^\star$ for $\mu = \bar{\mu}$ and $\mu = -\bar{\mu}$, that solve

$$\bar{\Lambda}(\mu, p_1, ..., p_I) \equiv \sup_{\mu^\star \in [\mu-k, \mu+k]} \sum_{i=1}^{I} p_i R(d_i, \mu, \mu^\star).$$

We can show that the solutions of μ^* (as a function of μ) are given by

$$\mu^\star = \begin{cases} \mu + k, & \frac{\mu+k}{2k} \geq \sum_{i=1}^{I} p_i \Phi\left(-\frac{c_i-\mu}{\sigma}\right) \\ \mu - k, & \frac{\mu+k}{2k} < \sum_{i=1}^{I} p_i \Phi\left(-\frac{c_i-\mu}{\sigma}\right) \end{cases}$$

We consider the case in which $\sigma = 1$, $k = 2$, and $\bar{\mu} = 0.5$. We set $\epsilon = 0.1$. The number of epochs in Theorem 1 is again

$$T = \lceil 2M^2 \ln(I)/\epsilon^2 \rceil = 7,953$$

The algorithm runs for $T = 7,953$ iterations and finishes in about 25 seconds (on a personal ASUS Vivobook Pro 15 @ 2.5GHz Intel Core Ultra 9 185H).

Figure 2 shows the true solution versus its ϵ-approximate solution. Qualitatively, the two are very close. In fact, the minimax values are close as well, with the ϵ-approximation having a minimax value of .9377 and the true solution having a minimax value of 0.9375. Note that here, the term referred to as $\rho^\star$-adjusted is

$$\rho^\star\text{-adjusted} = \frac{\sigma^2 \rho^\star}{\bar{\mu}}. \tag{10}$$

We also consider a 500-point uniform prior (i.e., $\pi_\mu = 1/500$) supported on an equally spaced grid within $[-k, k]$. In this case, no analytical solution is available. We still keep $\sigma = 1$ and $k = 2$. Keeping $\epsilon = 0.1$, the algorithm runs for $T = 7,953$ iterations and finishes in about 440 seconds (on a personal ASUS Vivobook Pro 15 @ 2.5GHz Intel Core Ultra 9 185H). Figure 3 shows the ϵ-approximate solution.

5 Application

Lee et al. (2021) conducted a randomized controlled trial in Bangladesh to estimate the effects of encouraging rural households to receive money transfers from migrant family members. They specifically conducted an encouragement design where poor rural households with family members

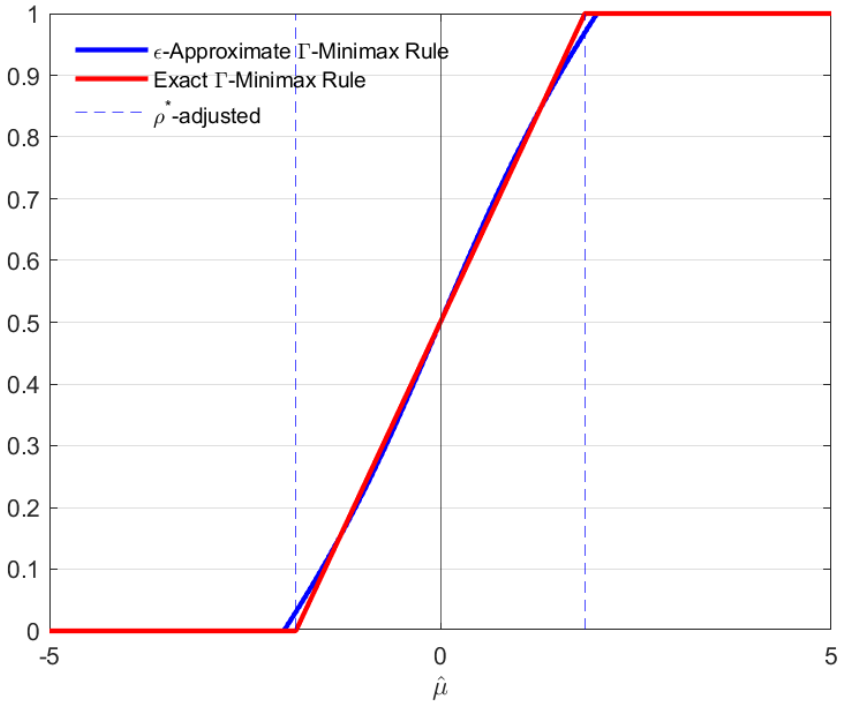

Figure 2: ϵ-Minimax Decision Rule for the 2-point Robust Bayes problem via the Hedge algorithm. The graph is generated using $\sigma = 1$, $k = 2$. The $\rho^\star$-adjusted value is about 1.8486.

who had migrated to a larger urban destination receive a 30–45 minute training about how to register and use the mobile banking service "bKash" to send instant remittances back home.

The experiment was conducted in the Gaibandha district, one of Bangladesh's poorest regions. It focused on households that had migrant workers in the Dhaka district, the administrative unit in which the capital of Bangladesh is located. Lee et al. (2021) measure several outcomes of both receiving households and sender migrants; see their Figures 3 and 4. To give a concrete example of the measured outcomes, one question of interest is whether families that adopt the mobile banking technology are more (or less) likely to declare that the *monga*—the seasonal period of hunger in September through November—was not a problem for their household. Table 9, Column 7, p. 60 in Lee et al. (2021) presents results for this specific variable showing that households that used a bKash account in the treatment group are 9.2 percentage points more likely to declare that *monga* was not a problem. The standard error of the estimator is 4.5 percentage points.

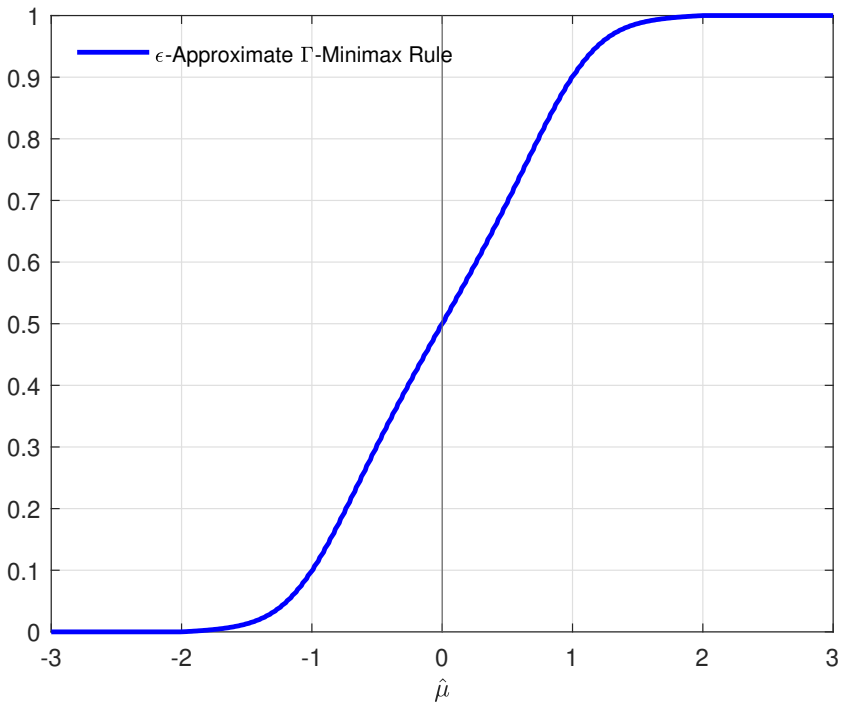

Figure 3: ϵ-Minimax Decision Rule via the Hedge algorithm for the 500-point Robust Bayes problem supported along $[-k, k]$.

Is the corridor selected by Lee et al. (2021) a good choice for a researcher who is concerned about external validity?[11] There are two recent papers that have provided an answer to this question. Gechter et al. (2024) use an elegant decision-theoretic framework to argue that the *Dhaka-Noakhali* corridor would have been a better choice from the perspective of maximizing average welfare. Montiel Olea et al. (2024a) use the framework of Gechter et al. (2024) to argue that the *Dhaka-Pabna* corridor would have been a better choice from the minimax (welfare) regret criterion perspective (restricting the policy maker to consider only nonrandomized selection of corridors). The *Dhaka-Pabna* corridor is also recommended by the *synthetic purposive sampling* approach in Egami and Lee (2024). One important comment is that the *Dhaka-Pabna* corridor is the most representative in terms of covariates, in the sense that it minimizes the average distance (measured using the euclidean distance between covariates) to the 41 migration corridors analyzed in Gechter

[11]Following Gechter et al. (2024), we name the corridors using a destination-origin format; for example, the migration corridor studied in Lee et al. (2021) is "Dhaka-Gaibandha".

et al. (2024).

In our application, we consider a situation where a policy maker is considering the three sites mentioned above to run an experiment: Dhaka-Gaibandha (the original site in Lee et al. (2021)), Dhaka-Noakhali (the site suggested by Gechter et al. (2024)) and Dhaka-Pabna (the site suggested in Montiel Olea et al. (2024a)). Each of these sites (migration corridor) have site characteristics $X_s \in \mathbb{R}^d$, with $d = 13$.[12] We index these three sites by $1, 2, 3$ respectively and refer to the set $\mathcal{S}_E \equiv \{1, 2, 3\}$ as the set of experimental sites. Once we exclude these three sites, we have 38 migration corridors. We use the distance between the covariates of each of these sites and Dhaka-Pabna to order them in increasing order and index them with integers 4 to 41. Figure 4 presents the distances. The figure shows that for most of the sites the corridor Dhaka-Pabna is the "closest" in terms of the Euclidean distance between covariates.

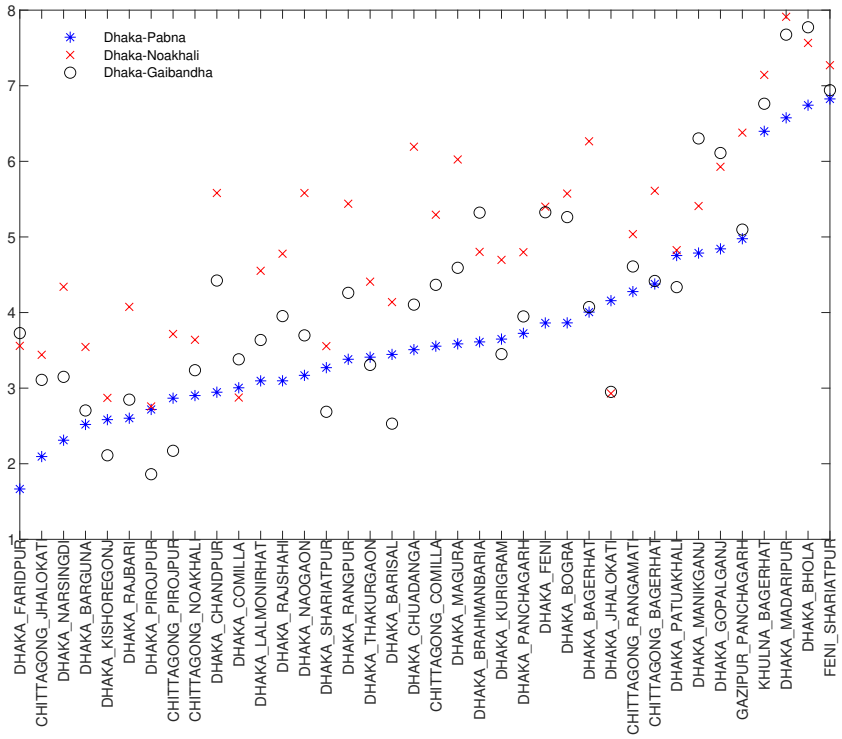

Figure 4: Distances from each of the experimental sites to each of the policy-relevant sites.

[12]The covariates include mean household income, mean household size, migrant density, mean remittances. See Figure 2 in Montiel Olea et al. (2024a).

We assume that the sites $\mathcal{S}_p \equiv \{4, \ldots, 41\}$ in the x-axis of Figure 4 are the policy-relevant sites. This means that policy maker is interested in deciding whether the training program discussed in Lee et al. (2021) should be rolled out in these sites. We assume that the outcome variable of interest for the policy maker is the likelihood that the households declare that the *monga* was not a problem.

Treatment Effect Heterogeneity: Treatment effect heterogeneity is allowed, but only via the observable site characteristics. The effects of the policy in each site, denoted by τ_s, are restricted to be a Lipschitz function (with respect to a Euclidean norm $||\cdot||$) with known constant C; that is, $\tau_s = \tau(X_s)$ where

$$|\tau(x) - \tau(x')| \leq C||x - x'||, \quad \forall x, x' \in \mathbb{R}^{13}.$$

One first issue that we need to address in order to conduct our exercise is the value of C that will be used in our application. We do this by using the available point estimates of the treatment effect of the program in Lee et al. (2021). Let x_{DG} denote the covariates of the corridor Dhaka-Gaibandha. Assume that the we entertained the possibility that the true effect, $\tau(x_{DG})$, coincides with the estimated effect 9.2. We consider a "low C" regime.

Suppose that we want to consider a value of C that imposes that if 9.2 were the true effect, then even the corridor that is the most different (in terms of covariates) to Dhaka-Gaibandha the effect of the program must be nonnegative. Intuitively, this means that even for the most different corridor, the effect must remain nonnegative. Dhaka-Bhola is the most different site and $\|x_{\text{DG}} - x_{\text{DB}}\| = 7.7736$. Since the Lipschitz restriction imposes that

$$\tau(x_{\text{DG}}) - C\|x_{\text{DG}} - x_{\text{DB}}\| \leq \tau(x_{\text{DB}}),$$

we could pick C as

$$C = 9.2/7.7736 \approx 1.1834.$$

Treatment Rules: The policy maker makes two choices. First, the policy maker must pick one site on which to experiment. Second, the policy maker must decide how to make treatment choices in all the sites of interest given the available data. We assume that if the policy maker decides to experiment on site s, the available data becomes $\widehat{\tau}_s$, with

$$\widehat{\tau}_s \sim \mathcal{N}(\tau_s, \sigma_s^2) \tag{11}$$

and, as in Gechter et al. (2024), we assume σ_s^2 is known. In order to conduct our exercise, we assume that σ_s is the same for all experimental sites, and that it matches the standard error of the estiamted effect of the program in the Dhaka-Gaibhanda site. That is $\sigma_s = 4.5$ for all $s \in \mathcal{S}_E$.

The treatment rule is a mapping $T : \mathbb{R} \to [0,1]^{\#\mathcal{S}_P}$. For $s \in \mathcal{S}_E$ we further denote by T_s the specific policy choice for site s. We refer to a tuple (s, T) as a policy, and we use d to denote it. We consider three nonrandomized policies

$$\mathcal{D} \equiv \{d_1, d_2, d_3\}.$$

Under policy d_s, the policy maker experiments on site $s \in \mathcal{S}_E$ and its recommendation for any policy relevant site is $\mathbf{1}\{\widehat{\tau}_s \geq 0\}$. That is, the policy maker makes the same policy recommendation for every policy-relevant site depending on the sign of $\widehat{\tau}_s$. [13] We focus on this special form of policy rule because we think it captures the policy recommendations that are given based on randomized controlled trials.

We consider the following regret function for the policy d_s,

$$R(d_s, \tau) \equiv \frac{1}{\#\mathcal{S}_P} \sum_{s' \in \mathcal{S}_P} \left(\tau(X_{s'}) (\mathbf{1}\{\tau(X_{s'}) \geq 0\} - \mathbb{E}_{\tau(X_s)} [\mathbf{1}\{\widehat{\tau}_s \geq 0\}]) \right). \tag{12}$$

[13] The results in Montiel Olea et al. (2024a) suggest that this type of policy is likely to be suboptimal. The policy maker could improve its welfare by allowing the treatmetment choice to be randomnly selected, depending on the distance between the policy-relevant site and the experimental site.

This expression can be simplified to:

$$R(d_s, \tau) \equiv \frac{1}{\#\mathcal{S}_P} \sum_{s' \in \mathcal{S}_P} \left(\tau(X_{s'}) \left(\mathbf{1}\{\tau(X_{s'}) \geq 0\} - \Phi\left(\frac{\tau(X_s)}{\sigma_s} \right) \right) \right). \tag{13}$$

The minimax (regret) problem that we are interested in solving is

$$\inf_{p \in \Delta^2} \sup_{\tau \in \text{Lip}_C(\mathbb{R}^{13})} \sum_{s=1}^{3} p_s R(d_s, \tau), \tag{14}$$

where $\text{Lip}_C(\mathbb{R}^{13})$ refers to the space of all Lipschitz functions $f : \mathbb{R}^{13} \to \mathbb{R}$ with constant C.

5.1 Results

We report results for the case in which $C = 1.1834$. We consider four different scenarios that vary in terms of the number of sites that are policy relevant. The scenarios we consider have either 1, 5, 15, or 38 policy-relevant sites. In each of these cases, we choose to include the sites that are closest to Dhaka-Pabna. For example, when we include only one policy-relevant site we include Dhaka-Faridpur. We do this because, in light of the results in Montiel Olea et al. (2024a), the best nonrandomized choice of experimental site is Dhaka-Pabna. And we would like to use this application to understand how the probability of selecting this site changes as we include sites that perhaps are closer to some of the other experimental sites under consideration.

Figure 5 presents the ϵ-minimax selection of sites obtained via the Hedge algorithm. Note first that when there is only one policy-relevant site (and this site is closest to Dhaka-Pabna) the probability of choosing Dhaka-Pabna is close to 1. This is measured by the height of the first yellow bar in Figure 5. We think this is an interesting result as it shows that even if randomization is allowed, it is possible that choosing the site that is most representative for the policy-relevant site is still approximately minimax regret optimal.

The results with five policy relevant sites are also worth discussing. By construction, the five

policy-relevant sites that we considered are those that are closest to Dhaka-Pabna. According to Figure 4, Dhaka-Pabna is the nearest neighbor for all of them, with the exception of Dhaka-Kishoregonj. For the latter site, the nearest neighbor is Dhaka-Gaibandha. Figure 5 shows that with 5 sites the ϵ-minimax selection of experimental sites places probability slightly above .2 to Dhaka-Gaibandha (corresponding to the height of the second blue bar) and probability close to .7 to Dhaka-Pabna (corresponding to the height of the second yellow bar).

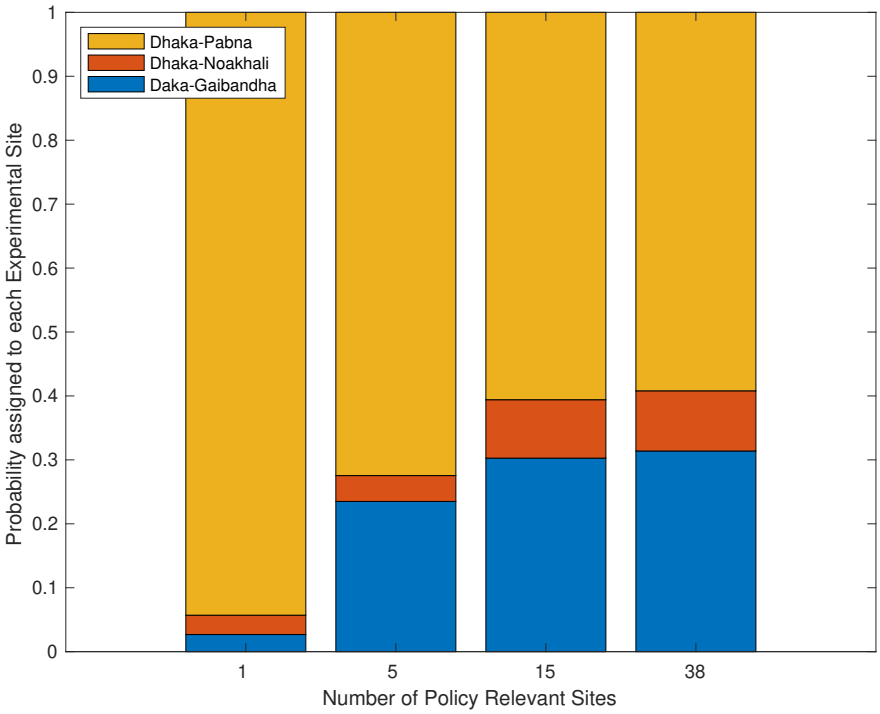

Figure 5: ϵ-Minimax decision rule for the Site Selection Problem via the Hedge algorithm. The graph is generated using $C = 1.1834$, $\sigma = 4.5$, and $\epsilon = .1$.

Finally, we discuss the cases in which the experimental sites are 15 and 38. We note that in both cases the ϵ-minimax solutions are very similar, though the time required by the algorithm and the number of iterations are not; see Tables 1. The recommended probability of experimenting on Dhaka-Pabna is close to .6 (height of the last yellow bar). The recommended probability of experimenting on Dhaka-Noakhali is close to .1. Interestingly, we note that the ordering of the probabilities is also consistent with the ordering of the three experimental sites (in terms of how

frequently they are the nearest neighbor of each of the policy-relevant sites).

Number of Sites	Runtime (seconds)	Iterations
1	1,765	14,063
5	1,554	11,449
15	7,427	12,974
38	38,741	21,248

Table 1: Runtime (seconds) and Number of Iterations ($C = 1.1834$).

6 Conclusion

This paper presented an algorithm for obtaining ϵ-*minimax* solutions of statistical decision problems where the statistician is allowed to choose randomly among I decision rules. The notion of an ϵ-minimax decision rule was taken from Ferguson (1967) (Chapter 1, Definition 4) and it refers to a decision rule whose worst-case expected loss exceeds the minimax value of the decision problem by at most an additive factor of ϵ.[14]

Once we allow for randomized selection over the I decision rules, the minimax problem admits a convex programming representation over the $(I-1)$-simplex, an observation which has been previously documented in the literature by Chamberlain (2000). Both the objective function and the subgradient of this problem are in general difficult to evaluate. The reason being that the objective function of the convex problem involves solving a nonconvex maximization problem to find the worst-case performance (over the model's parameter space) of a specific randomized selection over the I rules. This type of problem arises commonly in the convex optimization literature; see Bubeck et al. (2015) and the seminal work of Nemirovski and Yudin (1983). The algorithm herein suggested is a version of *mirror subgradient descent*, initialized with uniform weights and stopped after a finite number of iterations. The early stopping of the algorithm tries to minimize the number

[14]We note that the definition given in Ferguson (1967) differs of the usage of ϵ-minimax decision rules in other contexts. Most notably, from the work of Manski and Tetenov (2016), who use the term ϵ-minimax to refer to a decision rule whose worst-case regret is at most ϵ.

of calls to the objective function and its subgradient.

The iterative procedure arising from this mirror descent routine described in this paper is known in the computer science literature as the *Multiplicative Weights* update method, and it is used in algorithmic game theory as a practical tool to find approximate solutions of two-person zero-sum games. The paper applies the suggested algorithm to different minimax problems in the econometrics literature. In some of these problems, the minimax solution is known, and we show numerically that in these examples the ϵ-minimax solution is practically the same as the true minimax solution.

Finally, we apply the algorithm to the *site selection problem* of Gechter et al. (2024); namely, how to optimally selecting *sites* to maximize the external validity of an experimental policy evaluation. Our algorithm allows the researcher to choose randomly where to experiment, but adjusting optimally for the available baseline covariate information.

We think there are several interesting areas for future work, both from an applied and from a more theoretical perspective. From a purely applied angle, applying this algorithm could be helpful in settings in which the space of decision rules is already discrete, but large. For example, as in the recent interesting work of Christy and Kowalski (2024). Relatedly, the algorithm suggested here might also be helpful in applications in which the parameter space is "large". For example, in the site selection application we considered the parameter space is a subset of $\mathbb{R}^S$ where S is the total number of experimental and policy-relevant sites. This means that that our algorithm could be useful in extending the scope of certain minimax problems, such as the one described in the recent work of Armstrong, Kline, and Sun (2024). More generally, the algorithm we presented is part of large literature in computer science and algorithmic game theory. We think that there is an opportunity to apply some of the procedures suggested in this literature (including some of the recent procedures that parameterize the strategy space using neural networks).

From a more theoretical perspective, it would be interesting to further explore the differences between ϵ-minimax strategies and the notion of a local min-max point in Daskalakis et al. (2021).

There are very interesting results about the relation between this notion and the stationary points of sugbradient ascent-descent dynamics. But it would be great to understand, theoretically and empirically, what are the potential benefits of searching for these type of points as opposed to ϵ-minimax strategies.

References

Adler, I. (2013): "The equivalence of linear programs and zero-sum games," *International Journal of Game Theory*, 42, 165–177.

Aradillas Fernández, A., J. L. Montiel Olea, C. Qiu, J. Stoye, and S. Tinda (2024): "Robust Bayes Treatment Choice with Partial Identification," *arXiv e-prints*, arXiv–2408.

Armstrong, T., P. M. Kline, and L. Sun (2024): "Adapting to misspecification," Tech. rep., National Bureau of Economic Research.

Arora, S., E. Hazan, and S. Kale (2012): "The multiplicative weights update method: a meta-algorithm and applications," *Theory of computing*, 8, 121–164.

Ben-Tal, A., T. Margalit, and A. Nemirovski (2001): "The ordered subsets mirror descent optimization method with applications to tomography," *SIAM Journal on Optimization*, 12, 79–108.

Blackwell, D. A. and M. A. Girshick (1954): *Theory of games and statistical decisions*, John Wiley, New York.

Blume, L. E. (1993): "The statistical mechanics of strategic interaction," *Games and economic behavior*, 5, 387–424.

Bubeck, S. et al. (2015): "Convex optimization: Algorithms and complexity," *Foundations and Trends® in Machine Learning*, 8, 231–357.

Chamberlain, G. (2000): "Econometric applications of maxmin expected utility," *Journal of Applied Econometrics*, 15, 625–644.

Christy, N. and A. E. Kowalski (2024): "Starting Small: Prioritizing Safety over Efficacy in Randomized Experiments Using the Exact Finite Sample Likelihood," *arXiv preprint arXiv:2407.18206*.

Dantzig, G. B. (1951): "A proof of the equivalence of the programming problem and the game problem," *Activity analysis of production and allocation*, 13.

Daskalakis, C., S. Skoulakis, and M. Zampetakis (2021): "The complexity of constrained min-max optimization," in *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, 1466–1478.

Dominitz, J. and C. F. Manski (2024): "Comprehensive OOS Evaluation of Predictive Algorithms with Statistical Decision Theory," Tech. rep., National Bureau of Economic Research.

Du, D.-Z. and P. M. Pardalos (1995): *Minimax and applications*, vol. 4, Springer Science & Business Media.

Duflo, E., R. Glennerster, and M. Kremer (2007): "Using randomization in development economics research: A toolkit," *Handbook of development economics*, 4, 3895–3962.

Egami, N. and D. D. I. Lee (2024): "Designing Multi-Context Studies for External Validity: Site Selection via Synthetic Purposive Sampling," *Working Paper*.

Elliott, G., U. K. Müller, and M. W. Watson (2015): "Nearly optimal tests when a nuisance parameter is present under the null hypothesis," *Econometrica*, 83, 771–811.

Ferguson, T. (1967): *Mathematical Statistics: A Decision Theoretic Approach*, vol. 7, Academic Press New York.

Filar, J. A. and T. Raghavan (1982): "An algorithm for solving S-games and differential S-games," *SIAM Journal on Control and Optimization*, 20, 763–769.

Freund, Y. and R. E. Schapire (1999): "Adaptive game playing using multiplicative weights," *Games and Economic Behavior*, 29, 79–103.

Fudenberg, D. and D. K. Levine (1995): "Consistency and cautious fictitious play," *Journal of Economic Dynamics and Control*, 19, 1065–1089.

Gechter, M., K. Hirano, J. Lee, M. Mahmud, O. Mondal, J. Morduch, S. Ravindran, and A. S. Shonchoy (2024): "Selecting Experimental Sites for External Validity," *Working Paper*.

Giacomini, R. and T. Kitagawa (2021): "Robust Bayesian inference for set-identified models," *Econometrica*, 89, 1519–1556.

Gravin, N., Y. Peres, and B. Sivan (2016): "Tight lower bounds for multiplicative weights algorithmic families," *arXiv preprint arXiv:1607.02834*.

Jiang, S., Z. Song, O. Weinstein, and H. Zhang (2020): "Faster dynamic matrix inverse for faster lps," *arXiv preprint arXiv:2004.07470*.

Kaji, T., E. Manresa, and G. Pouliot (2023): "An adversarial approach to structural estimation," *Econometrica*, 91, 2041–2063.

Kempthorne, P. J. (1987): "Numerical specification of discrete least favorable prior distributions," *SIAM Journal on Scientific and Statistical Computing*, 8, 171–184.

Lee, J. N., J. Morduch, S. Ravindran, A. Shonchoy, and H. Zaman (2021): "Poverty and migration in the digital age: Experimental evidence on mobile banking in Bangladesh," *American Economic Journal: Applied Economics*, 13, 38–71.

Lewis, G. and V. Syrgkanis (2018): "Adversarial generalized method of moments," *arXiv preprint arXiv:1803.07164*.

Luedtke, A., M. Carone, N. Simon, and O. Sofrygin (2020): "Learning to learn from data: Using deep adversarial learning to construct optimal statistical procedures," *Science Advances*, 6, eaaw2140.

Luedtke, A., I. Chung, and O. Sofrygin (2021): "Adversarial Monte Carlo meta-learning of optimal prediction procedures," *Journal of Machine Learning Research*, 22, 1–67.

Manski, C. F. (2021): "Econometrics for decision making: Building foundations sketched by Haavelmo and Wald," *Econometrica*, 89, 2827–2853.

Manski, C. F. and A. Tetenov (2016): "Sufficient trial size to inform clinical practice," *Proceedings of the National Academy of Sciences*, 113, 10518–10523.

Montiel Olea, J. L., B. Prallon, C. Qiu, J. Stoye, and Y. Sun (2024a): "Externally Valid Selection of Experimental Sites via the k-Median Problem," *https://arxiv.org/abs/2408.09187*.

Montiel Olea, J. L., C. Qiu, and J. Stoye (2024b): "Decision Theory for Treatment Choice Problems with Partial Identification," *arXiv preprint arXiv:2312.17623*.

Nemirovski, A. and D. Yudin (1983): *Problem Complexity and Method Efficiency in Optimization*, A Wiley-Interscience publication, Wiley.

Owen, G. (2013): *Game theory*, Emerald Group Publishing.

Rockafellar, R. T. (1970): "Convex analysis," .

Stoye, J. (2012): "Minimax regret treatment choice with covariates or with limited validity of experiments," *Journal of Econometrics*, 166, 138–156.

Troutt, M. D. (1978): "Algorithms for non-convex S n-games," *Mathematical Programming*, 14, 332–348.

Wald, A. (1950): *Statistical Decision Functions*, New York: Wiley.

Yata, K. (2021): "Optimal Decision Rules Under Partial Identification," ArXiv:2111.04926 [econ.EM], https://doi.org/10.48550/arXiv.2111.04926.

A Proofs of Main Results

A.1 Proof of Lemma 1

Take $p, p' \in \Delta(\mathcal{D})$. Note that

$$\begin{aligned} f(\alpha p + (1-\alpha)p') &= \sup_{\theta\in\Theta} R(\alpha p + (1-\alpha)p', \theta) \\ &= \sup_{\theta\in\Theta} \sum_{i=1}^{I} (\alpha p_i + (1-\alpha)p'_i) R(d_i, \theta) \\ &= \sup_{\theta\in\Theta} \left(\alpha \sum_{i=1}^{I} p_i R(d_i, \theta) + (1-\alpha) p'_i R(d_i, \theta) \right) \\ &\leq \alpha \sup_{\theta\in\Theta} \sum_{i=1}^{I} p_i R(d_i, \theta) + (1-\alpha) \sup_{\theta\in\Theta} p'_i R(d_i, \theta) \\ &= \alpha f(p) + (1-\alpha) f(p'). \end{aligned}$$

Thus, the function $f(\cdot)$ is convex.

Now we establish the Lipschitz continuity of $f(\cdot)$. Take any $p, p' \in \Delta(\mathcal{D})$. Note that

$$\begin{aligned} |f(p) - f(p')| &\leq |\sup_{\theta\in\Theta} \sum_{i=1}^{I} (p_i - p'_i) R(d_i, \theta)| \\ &\leq M \sum_{i=1}^{I} |p_i - p'_i| \\ &= M \|p - p'\|_1, \end{aligned}$$

where the second inequality applies Assumption 1. This shows that $f(\cdot)$ is Lipschitz continuous with constant at most M.

We now show that g_0 is a subgradient of f at p_0. That is, that for any $p \in \Delta(\mathcal{D})$:

$$f(p) \geq f(p_0) + g_0^\top (p - p_0).$$

Let p be an arbitrary point in $\Delta(\mathcal{D})$. By definition

$$\begin{aligned} f(p_0) &\equiv \sup_{\theta\in\Theta} \sum_{i=1}^{I} p_{0,i} R(d_i, \theta) \\ &= \sum_{i=1}^{I} p_{0,i} R(d_i, \theta_0) \\ &= \sum_{i=1}^{I} (p_{0,i} - p_i) R(d_i, \theta_0) + \sum_{i=1}^{I} p_i R(d_i, \theta) \\ &= g_0^\top (p_0 - p) + \sum_{i=1}^{I} p_i R(d_i, \theta) \\ &\leq g_0^\top (p_0 - p) + f(p). \end{aligned}$$

Thus, g_0 is a subgradient of $f(\cdot)$ at p_0.

A.2 Proof of Theorem 1 via Arora et al. (2012)

We extend Theorem 2.1 and Theorem 2.3 in Arora et al. (2012). For the sake of exposition, we divide our proof in three steps.

STEP 1: Fix the step-size η. First, we show that after T rounds the algorithm guarantees that, for all decision rules $d_i \in \{d_1, ..., d_I\}$, we have obtained average payoffs bounded above by the average payoff of any decision rule d_i plus an error term. More precisely:

$$\frac{1}{T}\sum_{t=1}^{T}\left(\sum_{i=1}^{I} p_{i,t} R(d_i, \theta_t)\right) \leq \frac{1}{T}\sum_{t=1}^{T} R(d_i, \theta_t) + \frac{M^2\eta}{2} + \frac{\ln(I)}{T\eta} \tag{15}$$

To show this, we use a similar argument to Theorem 2.1 in Arora et al. (2012). Let $M > 0$ be the bound on the risk function in Assumption 1. We define the normalized sub-gradient, $m_t \equiv g_t/M$.

Then, recall the definition of ϕ_t, we have that

$$\begin{aligned}
\phi_{t+1} &= \sum_{i=1}^{I} w_{i,t} \\
&= \sum_{i=1}^{I} w_{i,t-1} \exp\left(-\eta g_{i,t}\right) \\
&= \sum_{i=1}^{I} w_{i,t-1} \exp\left(-\eta M m_{i,t}\right) \\
&\leq \sum_{i=1}^{I} w_{i,t-1} \left(1 - \eta M m_{i,t} + \frac{\eta^2 M^2 m_{i,t}^2}{2}\right) \\
&= \phi_t - \phi_t \eta M \sum_{i=1}^{I} p_{i,t} m_{i,t} + \phi_t \frac{\eta^2 M^2}{2} \sum_{i=1}^{I} p_{i,t} m_{i,t}^2 \\
&= \phi_t \left(1 - \eta M \sum_{i=1}^{I} p_{i,t} m_{i,t} + \frac{\eta^2 M^2}{2} \sum_{i=1}^{I} p_{i,t} m_{i,t}^2\right) \\
&\leq \phi_t \exp\left(-\eta M \sum_{i=1}^{I} p_{i,t} m_{i,t} + \frac{\eta^2 M^2}{2} \sum_{i=1}^{I} p_{i,t} m_{i,t}^2\right).
\end{aligned}$$

The first inequality follows from the fact that $\exp(-x) \leq 1 - x + x^2/2$, for $|x| \leq 1$.[15] The last inequality follows from $1 - x \leq e^{-x}$ for all $x \in \mathbb{R}$. By induction after T rounds, and using the fact

[15]Note that:

$$\begin{aligned}
\exp(-x) &= 1 - x + \frac{x^2}{2!} - \frac{x^3}{3!} + ... \\
&\leq 1 - x + \frac{x^2}{2},
\end{aligned}$$

if and only if:

$$0 \leq \frac{x^3}{3!} - \frac{x^4}{4!} + ...$$

Note that $x^n \geq x^{n+1}$ for all $|x| \leq 1$ and $n \in \mathbb{N}$, and so the statement above holds.

that w_0 was initialized to be a vector of ones (i.e., $w_0 := \mathbf{1}$, we have

$$\begin{aligned}\phi_{T+1} &\leq \phi_1 \exp\left(-\eta M \sum_{t=1}^{T}\sum_{i=1}^{I} p_{i,t} m_{i,t} + \frac{\eta^2 M^2}{2} \sum_{t=1}^{T}\sum_{i=1}^{I} p_{i,t} m_{i,t}^2\right) \\ &= I \exp\left(-\eta M \sum_{t=1}^{T}\sum_{i=1}^{I} p_{i,t} m_{i,t} + \frac{\eta^2 M^2}{2} \sum_{t=1}^{T}\sum_{i=1}^{I} p_{i,t} m_{i,t}^2\right). \end{aligned} \tag{16}$$

Also notice that

$$\phi_{T+1} \geq w_{i,t+1} = \prod_{t=1}^{T} \exp\left(-\eta g_{i,t}\right), \tag{17}$$

After taking logs of both sides in (16) and (17), we have

$$-\sum_{t=1}^{T} g_{i,t} \leq \frac{\ln(I)}{\eta} - M \sum_{t=1}^{T}\sum_{i=1}^{I} p_{i,t} m_{i,t} + \frac{\eta M^2}{2} \sum_{t=1}^{T}\sum_{i=1}^{I} p_{i,t} m_{i,t}^2.$$

Since $m_{i,t} = g_{i,t}/M = R(d_i, \theta_t)/M$, we have

$$\begin{aligned}\frac{1}{T}\sum_{t=1}^{T}\sum_{i=1}^{I} p_{i,t} R(d_i, \theta_t) \leq& \frac{1}{T}\sum_{t=1}^{T} R(d_i, \theta_t) + \frac{M^2\eta}{2T} \sum_{t=1}^{T}\sum_{i=1}^{I} p_{i,t} m_{i,t}^2 + \frac{\ln(I)}{T\eta} \\ \leq& \frac{1}{T}\sum_{t=1}^{T} R(d_i, \theta_t) + \frac{\eta M^2}{2} + \frac{\ln(I)}{T\eta} \\ & \text{(since } m_{i,t}^2 \leq 1\text{)} \\ \leq& \frac{1}{T}\sum_{t=1}^{T} R(d_i, \theta_t) + \frac{\epsilon}{2} + \frac{\ln(I) M^2}{T\epsilon} \\ & \text{(since } \eta = \epsilon/M^2\text{)}. \end{aligned}$$

STEP 2: Let $p_i^\epsilon \equiv \frac{1}{T}\sum_{t=1}^{T} p_{i,t}$. We show that after T rounds, we have that:

$$\bar{v} \leq \sup_{\theta\in\Theta} \sum_{i=1}^{I} p_i^\epsilon R(d_i, \theta) \leq \frac{1}{T}\sum_{t=1}^{T}\left(\sum_{i=1}^{I} p_{i,t} R(d_i, \theta_t)\right) \leq \bar{v} + \frac{\epsilon}{2} + \frac{\ln(I)}{T}\left(\frac{M^2}{\epsilon}\right),$$

where $\bar{v}$ is the minimax value of the decision problem.

To show this, note that the lower bound holds by definition. For the upper bound, note:

$$\begin{aligned}\sup_{\theta\in\Theta}\sum_{i=1}^{I}p_i^{\epsilon}R(d_i,\theta) &= \sup_{\theta\in\Theta}\sum_{i=1}^{I}\left(\frac{1}{T}\sum_{t=1}^{T}p_{i,t}\right)R(d_i,\theta)\\ &= \sup_{\theta\in\Theta}\frac{1}{T}\sum_{t=1}^{T}\left(\sum_{i=1}^{I}p_{i,t}R(d_i,\theta)\right)\\ &\leq \frac{1}{T}\sum_{t=1}^{T}\left(\sum_{i=1}^{I}p_{i,t}R(d_i,\theta_t)\right), \end{aligned} \tag{18}$$

where the inequality uses the fact that θ_t is nature's best response to p_t.

Lemma 1 showed that $f(\cdot)$ is a continuous function on the closed set $\Delta(\mathcal{D})$. Therefore, the minimax strategy of the decision problem exists, and we denote it as

$$p^* \in \arg\min_{p\in\Delta(\mathcal{D})} f(p).$$

By Step 1, the right hand side of the Equation (18) is bounded by above by Equation (15) for any d_i. It then follows that for any $p_i \in \Delta^{I-1}$, 18 is bounded above by

$$\frac{1}{T}\sum_{t=1}^{T}p_iR(d_i,\theta_t) + \frac{\epsilon}{2} + \frac{\ln(I)}{T}\left(\frac{M^2}{\epsilon}\right).$$

In particular, we can use p^* and further use the bound

$$\sum_{i=1}^{I}p_i^{\star}R(d_i,\theta_t) \leq \sup_{\theta\in\Theta}\sum_{i=1}^{I}p_i^{\star}R(d_i,\theta) = \bar{v}.$$

Consequently,

$$\sup_{\theta\in\Theta}\sum_{i=1}^{I}p_i^{\epsilon}R(d_i,\theta) \leq \frac{1}{T}\sum_{t=1}^{T}\left(\sum_{i=1}^{I}p_{i,t}R(d_i,\theta_t)\right) \leq \bar{v} + \frac{\epsilon}{2} + \frac{\ln(I)}{T}\left(\frac{M^2}{\epsilon}\right).$$

This gives the desired result.

STEP 3: By taking the smallest integer T such that

$$\frac{\ln(I)}{T}\left(\frac{M^2}{\epsilon}\right) \leq \frac{\epsilon}{2},$$

or, equivalently,

$$T = \lceil 2M^2 \ln(I)/\epsilon^2 \rceil.$$

We then conclude that

$$\bar{v} \leq \sup_{\theta\in\Theta}\left(\sum_{i=1}^{I} p_i^{\epsilon} R(d_i,\theta)\right) \leq \frac{1}{T}\sum_{t=1}^{T}\left(\sum_{i=1}^{I} p_{i,t} R(d_i,\theta_t)\right) \leq \bar{v} + \frac{\epsilon}{2} + \frac{\epsilon}{2} = \bar{v} + \epsilon,$$

where $p_i^{\epsilon} \equiv \frac{1}{T}\sum_{t=1}^{T} p_{i,t}$ (as defined before). Since

$$p^{\epsilon} \equiv (p_1^{\epsilon}, ..., p_I^{\epsilon}) \in \Delta^{I-1},$$

we conclude that $\tilde{p}^{\star}$ is an ϵ-minimax decision rule and that $\bar{v}^{\epsilon} \equiv \frac{1}{T}\sum_{t=1}^{T}\left(\sum_{i=1}^{I} p_{i,t} R(d_i,\theta_t)\right)$ is an ϵ-approximation to $\bar{v}$.

A.3 Proof of Theorem 1 via Bubeck et al. (2015)

We apply Theorem 4.2 from Bubeck et al. (2015). In order to do so, we note that:

1. The mirror map used in our mirror-descent routine is the negative entropy $\Phi(p) = \sum_{i=1}^{I} p_i \log p_i$, defined for any $p \in \text{int}(\Delta^{I-1})$. By Pinsker's inequality, we have that for any $p, q \in \text{int}(\Delta^{I-1})$:

$$\sum_{i=1}^{I} p_i \log p_i - \sum_{i=1}^{I} p_i \log q_i \geq \frac{1}{2}\left(\sum_{i=1}^{I} |p_i - q_i|\right)^2.$$

This can be written as

$$\begin{aligned}\Phi(q) - \Phi(p) &\leq \sum_{i=1}^{I}(q_i - p_i)\log q_i - \frac{1}{2}\|p-q\|_1^2 \\ &= \sum_{i=1}^{I}(q_i - p_i)\left(1 + \log q_i\right) - \frac{1}{2}\|p-q\|_1^2.\end{aligned}$$

Since the gradient of $\Phi(\cdot)$ with respect to q is $(1+\ln q_1, \ldots, 1+\ln q_I)^\top$, this means that the negative entropy is ρ-strongly convex over $\text{int}(\Delta^{I-1})$ (with respect to $\|\cdot\|_1$) with parameter $\rho = 1$.

2. We calculate the "radius" of $\text{int}(\Delta^{I-1})$ defined as $R^2 \equiv \sup_{p\in\text{int}(\Delta^{I-1})}\Phi(p) - \Phi(p_0)$. We show that $R^2 = \ln(I)$. To do this, note that

$$\sup_{p\in\text{int}(\Delta^{I-1})}\Phi(p) = 0.$$

The supremum is attained by any sequence of distributions that converges to a discrete distribution that places all of its mass in one of point. Then, since $p_0 = (1/I, ..., 1/I)$ in our setting, we have that $\Phi(p_0) = -\ln(I)$. This implies that $R^2 = \ln(I)$.

3. Lemma 1 has implied that the objective function f is convex and L-Lipschitz continuous w.r.t. $\|\cdot\|_1$, with parameter $L = M$. Let p^* be a minimizer of f over Δ^{I-1}.

In light of 1-2-3, the conditions of Theorem 4.2 in Bubeck et al. (2015) are verified. The theorem then implies that mirror descent with step size η satisfies

$$f\left(\frac{1}{T}\sum_{t=1}^{T}p_t\right) - f(p^\star) \leq \frac{1}{T}\sum_{t=1}^{T}f(p_t) - f(p^\star) \leq \frac{\ln(I)}{T\eta} + \frac{\eta M^2}{2}$$

Notice that $f(p^\star) = \bar{v}$. For a given error ϵ, we choose $\eta = \epsilon/M^2$ and $T \geq 2M^2 \ln(I)/\epsilon^2$ such that the right hand side is smaller than ϵ:

$$\frac{\ln(I)}{T\eta} + \frac{\eta M^2}{2} = \frac{\ln(I)M^2}{T\epsilon} + \frac{\epsilon}{2} \leq \epsilon.$$

Therefore, we conclude with

$$f(p_i^\epsilon) - \bar{v} \leq \bar{v}^\epsilon - \bar{v} \leq \epsilon.$$

B Connection to S Games

The minimax problem we study is closely related to what Blackwell and Girshick (1954) call an S-Game. Player I, the statistician in our case, has a finite number of pure strategies $d \in \mathcal{D} \equiv \{d_1, \ldots, d_I\}$. Player II, nature, may have infinitely many pure strategies. For each $\theta \in \Theta$, the strategy of nature can be represented as a vector in $\mathbb{R}^I$, $s(\theta) = (R(d_1, \theta), R(d_2, \theta), ..., R(d_I, \theta))$. Define

$$S := \{(R(d_1, \theta), R(d_2, \theta), ..., R(d_I, \theta)) \in \mathbb{R}^I | \theta \in \Theta\} \tag{19}$$

and $M(i, s) = s_i$, then $\Gamma_P = (\mathcal{D}, S, M)$ is a S-game with payoff matrix $M(i, s)$. The index P stands for *pure* as we are only considering pure strategies. The mixed extension of the S-game is equivalent to $\Gamma_m = (\Delta, R, M)$, where Δ is the set of discrete probability distribution over $\mathcal{D}$, and R is the set of all countable convex linear combination of points in S. When S is closed and convex, our minimax problem is exactly to solve the best mixed strategy for player I.

(Blackwell and Girshick, 1954, Theorem 2.4.2) states that i) Every S game has a value, and the first player has a good (a minimax) strategy; and ii) If S is closed and convex, player II has a pure good strategy. Further, (Blackwell and Girshick, 1954, Theorem 1.8.1) indicates that if game Γ_P has a pure value, then its mixed extension Γ_m also has the same value.

Combining these results, we have that i) there exists $\bar{v}$,

$$\inf_{d\in\mathcal{D}}\sup_{\theta\in\Theta} R(d,\theta) = \sup_{\theta\in\Theta}\inf_{d\in\mathcal{D}} R(d,\theta) = \bar{v}, \tag{20}$$

$$\text{and} \quad \inf_{p\in\Delta(\mathcal{D})}\sup_{\theta\in\Theta} R(p,\theta) = \sup_{\theta\in\Theta}\inf_{p\in\Delta(\mathcal{D})} R(p,\theta) = \bar{v}, \tag{21}$$

and there exists *minimax* decision rule $p^* \in \Delta(D)$ such that

$$\sup_{\theta\in\Theta} R(p^*,\theta) = \bar{v}.$$

Further, consider mixed strategies for the nature,

$$\inf_{p\in\Delta(D)}\sup_{q\in\mathcal{P}(\Theta)}\int_\Theta R(p,\theta)dq(\theta) = \sup_{q\in\mathcal{P}(\Theta)}\inf_{p\in\Delta(D)}\int_\Theta R(p,\theta)dq(\theta) = \bar{v}, \tag{22}$$

where $\mathcal{P}(\Theta)$ is the set of all mixed strategies of the nature. This means the assumption in Theorem 1 can be verified from our Assumptions. ii) If the set S defined in (19) is convex and closed, there exists exactly one θ^* that is *maxmin* strategy for nature, i.e.,

$$\inf_{p\in\Delta(D)} R(p,\theta^*) = \bar{v}.$$

Otherwise, the results in Blackwell and Girshick (1954) show that that the *maxmin* strategy for nature is supported on at most I points.

B.1 Maximin Strategy (least-favorable distribution)

Our proof for Theorem 1 also gives a surprising side result: When the game has a value, i.e. (22) holds, we can derive approximate max-min strategy for the nature.

Definition 2. For simplicity, we denote

$$R(p,q) := \int_{\Theta} R(p,\theta)dq(\theta).$$

A distribution $q_\epsilon^* \in \mathcal{P}(\Delta)$ is called an "ϵ-maximin" strategy for the game $(\Delta(\mathcal{D}), \mathcal{P}(\Theta), R(\cdot,\cdot))$ with value $\bar{v}$ if

$$\inf_{p\in\Delta(D)} R(p,q_\epsilon^*) \geq \sup_{q\in\mathcal{P}(\Theta)} \inf_{p\in\Delta(D)} R(p,q) - \epsilon = \bar{v} - \epsilon.$$

In our proof, we derived an intermediate result that (15),

$$\frac{1}{T}\sum_{t=1}^{T}\left(\sum_{i=1}^{I} p_{i,t}R(d_i,\theta_t)\right) \leq \frac{1}{T}\sum_{t=1}^{T}\sum_{i=1}^{I} p_i R(d_i,\theta_t) + M^2\eta + \frac{\ln(I)}{T\eta}, \forall p \in \Delta(\mathcal{D})$$

Recall assumption 2 states that for all t, $\sum_{i=1}^{I} p_{i,t}R(d_i,\theta_t) = \sup_{\theta\in\Theta} R(p_t,\theta)$, so

$$\frac{1}{T}\sum_{t=1}^{T}\left(\sum_{i=1}^{I} p_{i,t}R(d_i,\theta_t)\right) = \frac{1}{T}\sum_{t=1}^{T}\sup_{\theta\in\Theta} R(p_t,\theta) \geq \inf_{p\in\Delta(\mathcal{D})}\sup_{\theta\in\Theta} R(p,\theta) = \bar{v},$$

we get

$$\bar{v} \leq \frac{1}{T}\sum_{t=1}^{T}\sum_{i=1}^{I} p_i R(d_i,\theta_t) + M^2\eta + \frac{\ln(I)}{T\eta}, \forall p \in \Delta(\mathcal{D})$$

By taking $\eta = \epsilon/2M^2, T = \lceil 4M^2\ln(I)/\epsilon^2 \rceil$, we get

$$\bar{v} - \epsilon \leq \frac{1}{T}\sum_{t=1}^{T}\sum_{i=1}^{I} p_i R(d_i,\theta_t), \forall p \in \Delta(\mathcal{D}).$$

Now, if we choose q_ϵ^* to be a discrete distribution that

$$q_\epsilon^*(\theta) = \frac{|\{t \in [T] : \theta_t = \theta\}|}{T},$$

Then,

$$\inf_{p\in\Delta(\mathcal{D})} R(p, q_\epsilon^*) = \frac{1}{T}\sum_{t=1}^{T}\sum_{i=1}^{I} p_i R(d_i, \theta_t) \geq \bar{v} - \epsilon = \sup_{q\in\mathcal{P}(\Theta)} \inf_{p\in\Delta(\mathcal{D})} R(p, q) - \epsilon,$$

which means q_ϵ^* is an "ϵ-maximin" strategy for the nature.