Decision Theory for Treatment Choice Problems with Partial Identification*

José Luis Montiel Olea† Chen Qiu‡ Jörg Stoye§

July 2024

Abstract

We apply classical statistical decision theory to a large class of *treatment choice problems with partial identification*, revealing important theoretical and practical challenges but also interesting research opportunities. The challenges are: In a general class of problems with Gaussian likelihood, all decision rules are admissible; it is maximin-welfare optimal to ignore all data; and, for severe enough partial identification, there are infinitely many minimax-regret optimal decision rules, all of which sometimes randomize the policy recommendation. The opportunities are: We introduce a *profiled regret* criterion that can reveal important differences between rules and render some of them inadmissible; and we uniquely characterize the minimax-regret optimal rule that least frequently randomizes. We apply our results to aggregation of experimental estimates for policy adoption, to extrapolation of Local Average Treatment Effects, and to policy making in the presence of omitted variable bias.

KEYWORDS: Statistical decision theory, treatment choice, partial identification.

*We thank Karun Adusumilli, Isaiah Andrews, Tim Armstrong, Larry Blume, Tim Christensen, Tommaso Denti, Kei Hirano, Tetsuya Kaji, Charles Manski, Douglas Miller, Francesca Molinari, Alex Tetenov, Kohei Yata, seminar audiences at Berkeley, Binghamton, and Cornell, as well as participants at the Bravo/JEA/SNSF Workshop on "Using Data to Make Decisions" (Brown University), the 2023 Cornell-Penn State Econometrics and IO Conference, 2023 Greater NY Metropolitan Area Econometrics Colloquium and CEME (Georgetown) for feedback and Krishiv Shah, Maximilian Yap, Finn Ye, Sara Yoo, Naiqi Zhang, and especially Brenda Quesada Prallon for excellent research assistance. We gratefully acknowledge financial support from the NSF under grant SES-2315600.

†Email: jlo67@cornell.edu

‡Email: cq62@cornell.edu

§Email: stoye@cornell.edu

1 Introduction

A policy maker must decide between implementing a new policy or preserving the status quo. Her data provide information about the potential benefits of these two options. Unfortunately, these data only *partially identify* payoff-relevant parameters and may therefore not reveal, even in large samples, the correct course of action. Such *treatment choice problems with partial identification* have recently received growing interest; for example, see D'Adamo (2021), Ishihara and Kitagawa (2021), Yata (2021), Christensen, Moon, and Schorfheide (2022), Kido (2022) or Manski (2022). Several interesting problems that arise in empirical research can be recast using this framework. A non-exhaustive list includes extrapolation of experimental estimates for policy adoption (Ishihara and Kitagawa, 2021; Menzel, 2023), policy-making with quasi-experimental data in the presence of omitted variable bias (Diegert, Masten, and Poirier, 2022), and extrapolation of Local Average Treatment Effects (Mogstad, Santos, and Torgovitsky, 2018; Mogstad and Torgovitsky, 2018).

In this paper, we analyze such problems in terms of Wald's (1950) Statistical Decision Theory, identifying theoretical and practical challenges as well as new research opportunities. We do so in a finite-sample framework characterized by a Gaussian likelihood and partial identification. We see our main goals as clearly articulating the challenges of applying standard decision-theoretic criteria and exploring some resulting research opportunities, not as advocating a specific criterion (e.g., minimax regret) for recommending decision rules.

Challenges: Three optimality criteria are routinely used to endorse or discard decision rules: *admissibility*, *maximin welfare*, and *minimax regret*. We highlight challenges in applying these criteria in our setting.

Admissibility. A decision rule is (welfare-)admissible if one cannot improve its expected welfare uniformly over the parameter space. This is usually considered a weak requirement for a decision rule to be considered "good". We show that, whenever problems in our setting exhibit partial identification, *every* decision rule, no matter how exotic, is admissible (Theorem 1).

Maximin Welfare. A decision rule is maximin(-welfare) optimal if it attains the highest worst-case expected welfare. Echoing earlier critiques from Savage (1951) to Manski (2004), we find that maximin decision rules will preserve the status quo regardless of the data (Theorem 2).

Minimax Regret. A decision rule is minimax-regret (MMR) optimal if it attains the lowest worst-case expected regret, where an action's regret is its welfare loss relative to the action that would be optimal if payoff-relevant parameters were known. In some point-identified treatment choice problems, the MMR rule is both essentially unique and nonrandomized (Canner, 1970; Stoye,

2009a; Tetenov, 2012). This may not not be true under partial identification. To make this point, we specialize our framework to the class of treatment choice problems recently studied by Yata (2021). For cases where the identified set for payoff relevant parameters is large enough, we discover *infinitely many* MMR optimal decision rules, *all of which* randomize the policy recommendation for some or all data realizations. This presents a challenge for the application of MMR, at least if one hopes for the resulting recommendation to be unique. Moreover, as we explain later by means of an example, different MMR optimal rules can lead to meaningfully different policy choices for the same data.

Opportunities: These challenges motivate new, potentially fruitful research directions. We discuss two concrete ideas that specifically refine the analysis of MMR optimal decisions.

First, we investigate whether it is useful to profile out some parameters of the risk function. We specifically analyze *profiled regret*: by considering a scalar parameter of interest—for example, a functional γ of the parameters of the model—and reporting worst-case expected regret as function of possible values of γ. We show, by means of an example, that i) not all decision rules are admissible with respect to profiled regret and ii) different MMR optimal rules can have strikingly different profiled-regret functions.

Second, we refine the set of MMR optimal rules by identifying the *least randomizing* (in a sense we make precise) MMR optimal rule. Our main motivation is the following trade-off. Recall that, for a wide range of parameter values, any MMR optimal rule must randomize for some data realizations. At the same time, despite wide adoption of randomized treatment allocations in economics and the social sciences for the purpose of experimentation, it might be difficult in many policy applications to randomize one's policy. Thus, we look for a rule that recommends such randomization as infrequently as possible. We explicitly characterize (in Theorem 4) an essentially unique least randomizing rule for the problems considered by Yata (2021); this rule also exhibits an attractive regret profile in our main example.

Applications: We illustrate the practical implications of our results for three problems that recently arose in applied work.

First, we analyze in detail a running example based on Ishihara and Kitagawa's (2021) "evidence aggregation" framework. Here, a policy maker is interested in implementing a new policy in country $i = 0$. She has access to estimates of the effect of the same policy for other countries $i = 1, ..., n$ and attempts to extrapolate results using baseline covariates. We give explicit MMR treatment rules for this example. An interesting finding is that, when the identified set is large enough, the least randomizing rule can be related to the estimated bounds on the treatment effect of interest and randomizes only (though not always) if these bounds contain both positive and negative values.

This illustrates how an estimator of the identified set can be used for optimal decision making.

Second, we study extrapolation of Local Average Treatment Effects (Mogstad et al., 2018; Mogstad and Torgovitsky, 2018) with binary instrument and no covariates. Here, the payoff-relevant parameter is a "policy-relevant treatment effect" (Heckman and Vytlacil, 2005) that corresponds to expanding the complier subpopulation. We show that Theorem 1 applies in this example, so that all decision rules are admissible. In particular, a decision rule that implements the policy for large values of the usual instrumental-variables estimator is not dominated.

Third, we consider a policy maker who uses quasi-experimental data to decide on a new policy. She is willing to assume a constant treatment effect model and unconfoundedness given a set of covariates (X, W); however, only X is observable and W is not. In this setting, Diegert et al. (2022) argue that researchers may be interested in how much selection on unobservables is required to overturn findings that are based on a feasible linear regression. The least randomizing MMR rule can inform a complementary, decision-theoretic breakdown point analysis: For a given estimated effect of the policy, what is the largest effect of unobserved confounding under which it is still optimal to adopt the seemingly better policy without any hedging? We show that this breakdown point tolerates more confounding than the one of Diegert et al. (2022).

Related literature. The econometric literature on treatment choice has grown rapidly since Manski (2004) and Dehejia (2005). When welfare is partially identified, Manski (2000, 2005, 2007a) and Stoye (2007) provide optimal treatment rules assuming the true distribution of the data is known. Stoye (2012a,b), Yata (2021), and Ishihara and Kitagawa (2021) focus on finite sample MMR optimal rules, and Aradillas Fernández, Montiel Olea, Qiu, Stoye, and Tinda (2024) on multiple prior MMR rules, in such settings. For different settings with point-identified welfare, finite- and large-sample results on optimal treatment choice rules were derived by Canner (1970), Chen and Guggenberger (2024), Hirano and Porter (2009, 2020), Kitagawa, Lee, and Qiu (2022), Schlag (2006), Stoye (2009a), and Tetenov (2012). Christensen et al. (2022) extend Hirano and Porter's (2009) limit experiment framework to a class of partially identified settings; see on this also Kido (2023). Treatment choice is furthermore related to a large literature on optimal policy learning that contains many results for point identified (Bhattacharya and Dupas, 2012; Kitagawa and Tetenov, 2018, 2021; Mbakop and Tabord-Meehan, 2021; Kitagawa and Wang, 2020; Athey and Wager, 2021; Kitagawa et al., 2021; Ida et al., 2022) as well as partially identified (Kallus and Zhou, 2018; Ben-Michael et al., 2021, 2022; D'Adamo, 2021; Christensen et al., 2022; Adjaho and Christensen, 2022; Kido, 2022; Lei et al., 2023) treatment choice that may condition on covariates. Bayesian aspects of treatment choice are discussed in Chamberlain (2011).

The remainder of this paper is organized as follows: Section 2 introduces the formal framework and the running example. Section 3 is devoted to the aforementioned challenges, whereas Section 4 focuses on the opportunities. The applications are presented in Section 5. Section 6 concludes. Proofs and additional technical results can be found in Appendix A and Online Appendix B.

2 Notation and Framework

Statistical decision theory calls for three ingredients: the menu of actions available, their consequences as a function of an unknown state of the world, and a statistical model of how the data distribution depends on that state. We now present these elements and lay out an example that will be used to illustrate objects, terms, and results throughout.

The policy maker can choose an *action* $a \in [0,1]$, which we interpret as the proportion of a population that will be randomly assigned to the new policy. Thus, $a = 1$ means that everyone is exposed to the new policy and $a = 0$ means that the status quo is preserved. Under this interpretation, $a = .5$ means that 50% of the population will be exposed at random to the new policy; however, the formal development equally applies to either individual or population-level randomization. Our interpretation abstracts from integer issues arising with small populations.

The payoff for the policy maker when taking action $a \in [0,1]$ is captured by a welfare function

$$W(a,\theta) := aW(1,\theta) + (1-a)W(0,\theta), \tag{2.1}$$

where $\theta \in \Theta$ is an unknown state of the world or *parameter* (possibly of infinite dimension) and $W(1,\cdot) : \Theta \to \mathbb{R}$ and $W(0,\cdot) : \Theta \to \mathbb{R}$ are known functions. Thus, welfare is linear in the action, a common assumption in the literature.[1] The form of (2.1) also implies that if θ were known to the policy maker, her optimal choice of action would simply be

$$\mathbf{1}\{U(\theta) \geq 0\}, \quad \text{where} \quad U(\theta) := W(1,\theta) - W(0,\theta). \tag{2.2}$$

Following Hirano and Porter (2009), we refer to $U(\theta)$ as the *welfare contrast* at θ. Thus, the policy maker's optimal action in (2.2) is to expose the whole population to the new policy if the welfare

[1] In particular, this applies if $W(\cdot,\theta)$ is an expectation and, for the case where randomization is interpreted as fractional assignment, there are no spillover effects or externalities. These assumptions are the default in the literature. An exception is Manski and Tetenov (2007), who consider the welfare of an action to be a concave transformation of $W(\cdot,\theta)$.

contrast is nonnegative and to preserve the status quo otherwise.

The policy maker observes a realization of $Y \in \mathbb{R}^n$ distributed as

$$Y \sim N(m(\theta), \Sigma). \tag{2.3}$$

Here, the function $m(\cdot) : \Theta \to \mathbb{R}^n$ and the positive definite covariance matrix Σ are known. However, $m(\cdot)$ need not be injective: $m(\theta) = m(\theta')$ does *not* imply $\theta = \theta'$. As a result, even perfectly identifying $m(\theta)$ (from infinite amount of data) need not imply identifying the optimal action.

In economics applications, the normality assumption in (2.3) is unlikely to hold exactly; however, the data can often be summarized by statistics that are asymptotically normal and whose asymptotic variances can be estimated. Treating the limiting model as a finite-sample statistical model then eases exposition and allows us to focus on the core features of the policy problem. Indeed, working directly with such a limiting model is common in applications of statistical decision theory to econometrics; see Müller (2011) and the references therein for theoretical support and applications in the context of testing problems and Ishihara and Kitagawa (2021), Stoye (2012a), or Tetenov (2012) for precedents in closely related work.

We finally define a *decision rule*, $d : \mathbb{R}^n \to [0, 1]$, as (measurable) mapping from the data Y to the unit interval. We let $\mathcal{D}_n$ denote the set of all decision rules. We call $d \in \mathcal{D}_n$ *non-randomized* if $d(y) \in \{0, 1\}$ for (Lebesgue) almost every $y \in \mathbb{R}^n$ and *randomized* otherwise.

Running Example: Our running example is a special case of Ishihara and Kitagawa's (2021; see also Manski (2020)) "evidence aggregation" framework. A policy maker is interested in implementing a new policy in country $i = 0$ and observes estimates of the policy's effect for countries $i = 1, ..., n$. Let $Y = (Y_1, ..., Y_n)^\top \in \mathbb{R}^n$ denote such experimental estimates and let $(x_0, \dots, x_n)$ be nonrandom, d-dimensional baseline covariates. The policy maker is willing to extrapolate from her data by assuming that the welfare contrast of interest equals $U(\theta) = \theta(x_0)$ and that

$$Y = \begin{pmatrix} Y_1 \\ \vdots \\ Y_n \end{pmatrix} \sim N(m(\theta), \Sigma), \quad m(\theta) = \begin{pmatrix} \theta(x_1) \\ \vdots \\ \theta(x_n) \end{pmatrix}, \quad \Sigma = \operatorname{diag}(\sigma_1^2, \dots, \sigma_n^2),$$

where $\theta : \mathbb{R}^d \to \mathbb{R}$ is an unknown Lipschitz function with known constant C. For simplicity, in the example we write μ_i for $\theta(x_i)$ henceforth, thus $Y_i \sim N(\mu_i, \sigma_i^2)$. We also assume w.l.o.g. that countries are arranged in nondecreasing order of $\|x_i - x_0\|$. While our analysis extends to $x_1 = x_0$,

we focus on the case of $x_1 \neq x_0$, so that the sign of μ_0 is not necessarily identified. Finally, we assume that $(x_1, \ldots, x_n)$ are distinct. Even if this were not the case in raw data, one would presumably want to induce it (by adding fixed effects, whose size can be bounded) because the Lipschitz constraint would otherwise imply that countries with same x exhibit no heterogeneity whatsoever. □

3 Challenges

In statistical decision theory, three criteria are commonly used to recommend decision rules: *admissibility*, *maximin welfare*, and *minimax regret*.[2] In this section, we show that the application of these criteria to our setting presents nontrivial challenges.

3.1 Everything is Admissible

Let $\mathbb{E}_{m(\theta)}[\cdot]$ denote expectation with respect to $Y \sim N(m(\theta), \Sigma)$. Recall the following definition:

Definition 1. A rule $d \in \mathcal{D}_n$ is *(welfare-)admissible* if there does not exist $d' \in \mathcal{D}_n$ such that

$$\mathbb{E}_{m(\theta)}[W(d'(Y), \theta)] \geq \mathbb{E}_{m(\theta)}[W(d(Y), \theta)], \quad \forall \theta \in \Theta,$$

with strict inequality for some $\theta \in \Theta$.

Thus, a rule is admissible if it is not dominated (in the usual sense of weak dominance everywhere and strict dominance somewhere). This is generally considered a minimal but compelling requirement for a decision rule to be "good" and goes back at least to Wald (1950). Admissibility can be used to recommend classes of rules whose payoff cannot be uniformly improved and/or whose members improve uniformly on non-members, as in complete class theorems (Karlin and Rubin, 1956; Manski and Tetenov, 2007); conversely, one may use it to caution against particular (classes of) decision rules, as was recently done by Andrews and Mikusheva (2022).

Our first result shows that, under mild assumptions, admissibility cannot serve either purpose in our problem. This is because any decision rule is admissible. To formalize this, let

$$M := \{\mu \in \mathbb{R}^n : m(\theta) = \mu, \theta \in \Theta\} \tag{3.1}$$

[2]See Stoye (2012b) and references therein for theoretical trade-offs between these criteria. For example, maximin violates von Neumann-Morgenstern Independence, whereas minimax regret violates a menu independence axiom.

collect all means that can be generated as θ ranges over Θ. We refer to elements $\mu \in M$ as *reduced-form* parameters because they are identified in the statistical model (2.3).[3] Define the *identified set* for the welfare contrast as function of μ as

$$I(\mu) := \{u \in \mathbb{R} : U(\theta) = u, m(\theta) = \mu, \theta \in \Theta\} \tag{3.2}$$

and the corresponding upper and lower bounds as

$$\overline{I}(\mu) := \sup I(\mu), \quad \underline{I}(\mu) := \inf I(\mu). \tag{3.3}$$

When we refer to models as "partially identified," we henceforth mean that partial identification obtains on an open set in parameter space (i.e., not almost nowhere).[4]

Definition 2 (Nontrivial partial identification)**.** The treatment choice problem with payoff function (2.1) and statistical model (2.3) exhibits *nontrivial partial identification* if there exists an open set $\mathcal{S} \subseteq M \subseteq \mathbb{R}^n$ such that

$$\underline{I}(\mu) < 0 < \overline{I}(\mu), \text{ for all } \mu \in \mathcal{S}.$$

Running Example—Continued: The identified set for the welfare contrast μ_0 is

$$I(\mu) = \{u \in \mathbb{R} : |\mu_i - u| \leq C \|x_i - x_0\|, \quad i = 1, \ldots, n\}.$$

Its extrema can be written as intersection bounds:

$$\underline{I}(\mu) = \max_{i=1,\ldots,n} \{\mu_i - C \|x_i - x_0\|\}, \quad \overline{I}(\theta) = \min_{i=1,\ldots,n} \{\mu_i + C \|x_i - x_0\|\}.$$

For μ sufficiently close to the zero vector, we therefore have nontrivial partial identification.

□

We are now ready to state the first main result.

Theorem 1. *If a treatment choice problem with payoff function (2.1) and statistical model (2.3) exhibits nontrivial partial identification in the sense of Definition 2, then every decision rule $d \in \mathcal{D}_n$*

[3] The notation is consistent with our running example, in which the observable moments are $(\mu_1, \ldots, \mu_n)$.

[4] For simplicity of exposition, Definition 2 is stated in a way that forces M to have full measure in $\mathbb{R}^n$. This would, for example, exclude equality constraints. For the purpose of Theorem 3 below, we can weaken the assumption to allow such cases as long as $\mathcal{S}$ is rich within M.

is (welfare-)admissible.

Proof. See Appendix A.1. □

For a proof sketch, suppose by contradiction that some rule d is inadmissible. Then some other rule d' dominates it. This d' must perform weakly better at every $\theta \in m^{-1}(\mathcal{S})$, where $\mathcal{S}$ is the set that appears in Definition 2. Because of nontrivial partial identification, all $\theta \in m^{-1}(\mathcal{S})$ are compatible with positive and negative welfare contrast $U(\theta)$. This implies that

$$\mathbb{E}_{m(\theta)}[d(Y)] = \mathbb{E}_{m(\theta)}[d'(Y)] \quad \text{for each } \theta \in m^{-1}(\mathcal{S}).$$

By i) completeness of the Gaussian statistical model in (2.3) and ii) mutual absolute continuity of the Gaussian and Lebesgue measures in $\mathbb{R}^n$, we then have $d(\cdot) = d'(\cdot)$ (Lebesgue) almost everywhere in $\mathbb{R}^n$, a contradiction.[5]

Remark 1. To see that partial identification is essential in this result, let $n = 1$, $\Theta = \mathbb{R}$, $m(\theta) = W(1, \theta) = \theta$, and $W(0, \theta) = 0$. In this problem, the policy maker observes a noisy signal, $Y \sim N(\theta, \sigma^2)$, of the payoff relevant parameter $\theta \in \mathbb{R}$. There is no partial identification, and θ determines the optimal choice of action. By Karlin and Rubin's (1956, Theorem 2) classic result, any decision rule that is not a threshold rule (i.e., is not of form $\mathbf{1}\{Y > c\}$ for some fixed $c \in \mathbb{R} \cup \{-\infty, \infty\}$) is dominated. In fact, the class just defined is complete.

Remark 2. While completeness can be weakened, it cannot be dropped. To see this, suppose that all elements of Σ are zero, so that $Y = m(\theta)$ and a decision rule is really a mapping $d : M \to [0, 1]$. Then any rule that does not set $d(\mu) = 1$ ($d(\mu) = 0$) when $\underline{I}(\mu) > 0$ ($\overline{I}(\mu) < 0$) is dominated.

Theorem 1 admits an optimistic though also a pessimistic reading. The positive interpretation is that the procedures suggested in the related literature will all perform well relative to one another in some parts of parameter space. This is an important observation because some of these suggestions contain novel or nonstandard components. For example, Ishihara and Kitagawa (2021) place ex-ante restrictions on the class of decision rules, while Christensen et al. (2022) transform the original loss function by profiling out partially identified parameters. By Theorem 1, all these approaches are at least admissible. This was arguably obvious in the former case (since for any linear threshold rule, it is easy to find a prior that it uniquely best responds to) but certainly not the latter one.

[5] A family $\mathcal{P}$ of distributions P is complete if $\mathbb{E}_P[f(X)] = 0$ for all $P \in \mathcal{P}$ implies $f(x) = 0$ P-almost everywhere, for every $P \in \mathcal{P}$. See, for example, Lehmann and Romano (2005) p.115. The proof extends to any sampling distribution satisfying i) and ii). Moreover, because decision rules are bounded, it suffices to verify bounded completeness.

The negative interpretation is that no rule, regardless how eccentric it may appear, is dominated. This makes it difficult to recommend a rule or a class of rules, at least without commitment to a more specific decision-theoretic criterion or focus on a particular part of parameter space.

3.2 Maximin Welfare is Ultra-Pessimistic

We next analyze the *maximin welfare* criterion. Our main result echoes earlier findings by Savage (1951) and Manski (2004): Maximin typically leads to "no-data rules" that preserve the status quo.

Definition 3. A rule $d_{\text{maximin}} \in \mathcal{D}_n$ is maximin optimal if

$$\inf_{\theta\in\Theta} \mathbb{E}_{m(\theta)}[W(d_{\text{maximin}}(Y),\theta)] = \sup_{d\in\mathcal{D}_n} \inf_{\theta\in\Theta} \mathbb{E}_{m(\theta)}[W(d(Y),\theta)].$$

Theorem 2. *Suppose that there exists $\theta \in \Theta$ such that $U(\theta) \leq 0$. If*

$$\inf_{\theta\in\Theta} W(0,\theta) = \inf_{\theta\in\Theta: U(\theta)\leq 0} W(0,\theta), \tag{3.4}$$

then the no-data rule $d_{\text{no-data}}(y) := 0$ is maximin optimal. The maximin value is

$$\inf_{\theta\in\Theta} W(0,\theta).$$

Proof. See Appendix A.2. □

This result can be seen as follows. When $U(\theta) \leq 0$, it is optimal to preserve the status quo; substituting in for this response, we find that $\inf_{\theta\in\Theta} \mathbb{E}_{m(\theta)}[W(d(Y),\theta)] \leq \inf_{\theta\in\Theta, U(\theta)\leq 0} W(0,\theta)$ for any rule $d \in \mathcal{D}_n$. Under condition (3.4), this upper bound is attained by $d_{\text{no-data}}$.

Running Example—Continued: Theorem 2 applies to the running example. In particular, the example's maximin welfare equals 0 and is achieved by never assigning the new policy. □

A similar result was shown by Manski (2004) for testing an innovation with point-identified welfare contrast (a result that we generalize[6]), and the concern can be traced back at least to Savage (1951). There is a discussion of whether such "ultrapessimism" occurs, in a technical sense, more generically with maximin utility versus minimax regret (Parmigiani, 1992; Sadler, 2015). However, a string of more optimistic results regarding MMR (Canner, 1970; Stoye, 2009a, 2012b; Tetenov, 2012; Yata, 2021) suggests that, with state spaces that describe real-world decision problems, the

[6] In Manski's (2004) example, $W(0,\theta)$ does not depend on θ, so that condition (3.4) is trivially satisfied.

concern is more salient for maximin. In particular, a result like this is expected whenever there exists a "uniformly awful" parameter value.[7]

3.3 Minimax Regret Admits Many Solutions

In view of our results so far, it seems natural to consider the *minimax regret* optimality criterion. Indeed, some point-identified treatment choice problems admit nonrandomized, essentially unique MMR optimal rules. Unfortunately, this extends to treatment choice with partial identification only in special cases. To make this point, we consider a relatively simple (though rich enough to cover our applications) class of problems where finite-sample minimax optimal results are available. Within this class, we show that, if the identified set for the welfare contrast is sufficiently large relative to sampling error, then (i) we can find infinitely many MMR optimal decision rules, all of which randomize and (ii) under weak additional conditions, all MMR rules that depend on data through a linear index must be randomized.

The expected regret of a decision rule $d \in \mathcal{D}_n$ in state θ is its expected welfare loss compared to the oracle rule:

$$\begin{aligned} R(d,\theta) &:= \sup_{a\in[0,1]} W(a,\theta) - \mathbb{E}_{m(\theta)}[W(d(Y),\theta)] \\ &= \max\{W(0,\theta), W(1,\theta)\} - \mathbb{E}_{m(\theta)}[W(d(Y),\theta)] \\ &= U(\theta)\left\{\mathbf{1}\{U(\theta)\geq 0\} - \mathbb{E}_{m(\theta)}[d(Y)]\right\}. \end{aligned} \tag{3.5}$$

Definition 4. A rule $d^* \in \mathcal{D}_n$ is minimax regret (MMR) optimal if

$$\sup_{\theta\in\Theta} R(d^*,\theta) = \inf_{d\in\mathcal{D}_n} \sup_{\theta\in\Theta} R(d,\theta). \tag{3.6}$$

Solving minimax regret problems is often hard. Algorithms exist for certain cases (Yu and Kouvelis, 1995; Chamberlain, 2000), but we are not aware of a general purpose algorithm. In important, recent work, Yata (2021) characterizes MMR optimal rules for a large class of binary action problems. He imposes the following restrictions on the parameter space and welfare function.

Assumption 1. (i) Θ is convex, centrosymmetric (i.e., $\theta \in \Theta$ implies $-\theta \in \Theta$) and nonempty.

[7] Given that we next elaborate on multiplicity of MMR rules, a quick clarification: While the maximin rule may be unique here, this is due to treating the status quo as known. In settings where "treatment 0" and "treatment 1" are equally unknown ex ante, the typical result is that all decision rules are maximin.

(ii) $m(\cdot)$ and $U(\cdot)$ are linear.

RUNNING EXAMPLE—CONTINUED: Our example satisfies Assumption 1. In particular, the space of C-Lipschitz functions, $\theta : \mathbb{R}^d \to \mathbb{R}$ is convex as well as centrosymmetric, and the functions $U(\cdot)$ and $m(\cdot)$ are linear in θ as they simply report the values of θ at points $(x_0, \ldots, x_n)$. □

Assumption 1 is based on the literature on minimax estimation and inference on a linear functional in nonparametric problems (Donoho, 1994; Low, 1995; Armstrong and Kolesár, 2018). Under Assumption 1, Yata (2021) shows existence of an MMR rule that depends on the data only through $(w^*)^\top Y$, where the unit vector w^* can be approximated by solving a sequence of tractable optimization problems. When the identified set for the welfare contrast at $\mu = \mathbf{0} := 0_{n\times 1}$ is large enough, Yata's (2021) MMR rule can be expressed as

$$d^*_{\mathrm{RT}}((w^*)^\top Y) := \Phi\big((w^*)^\top Y/\tilde{\sigma}\big) \tag{3.7}$$

for some uniquely characterized $\tilde{\sigma} > 0$. Moreover, it then has two important algebraic properties:

$$\mathbb{E}_{m(\theta)}[d^*_{\mathrm{RT}}(Y)] = 1/2 \quad \text{when } m(\theta) = \mathbf{0} \tag{3.8}$$

and

$$\sup_{\theta\in\Theta} R(d^*_{\mathrm{RT}}, \theta) = \sup_{\theta\in\Theta, m(\theta)=\mathbf{0}} R(d^*_{\mathrm{RT}}, \theta). \tag{3.9}$$

In words, these features are as follows: First, if the mean function $m(\cdot)$ equals zero, expected exposure to the new policy is 1/2. Second, worst-case regret occurs precisely at this point. The latter is due to careful calibration of the MMR decision rule, and one might conjecture that it renders this rule unique. However, the following result establishes the contrary.

Theorem 3. *Consider a treatment choice problem with payoff function (2.1) and statistical model (2.3) that exhibits nontrivial partial identification in the sense of Definition 2. Suppose that Assumption 1 holds and that there is a MMR optimal rule d^* that depends on the data only through $(w^*)^\top Y$ and that satisfies (3.8) and (3.9). If there exists $\mu \in M$ such that $\overline{I}(\mu) > \overline{I}(\mathbf{0})$ and $\overline{I}(\mathbf{0})$ is large enough, then*

(i) There are infinitely many MMR optimal rules.

(ii) Any MMR rule that depends on the data only through $(w^*)^\top Y$ (and is weakly increasing in this argument) must randomize for some data realizations.

(iii) If $\overline{I}(\mu)$ is differentiable at $\mu = \mathbf{0}$, then no linear threshold rule, i.e., no rule of form $\mathbf{1}\{w^\top Y \geq c\}$ for some $w \in \mathbb{R}^n$ and $c \in \mathbb{R} \cup \{-\infty, \infty\}$, is MMR optimal.

Proof. See Appendix A.3. □

To be clear, this finding applies if the problem is sufficiently far from point identification, with an exact condition given in the proof.[8] Close to point identification and under mild additional assumptions, Yata (2021) shows MMR optimality of a linear threshold rule.

Part (i) of Theorem 3 is established constructively: We show that, whenever d^*_{RT} is MMR optimal, then so is the *piecewise linear* rule

$$d^*_{\mathrm{linear}}((w^*)^\top Y) := \begin{cases} 0, & (w^*)^\top Y < -\rho^*, \\ \frac{(w^*)^\top Y + \rho^*}{2\rho^*}, & -\rho^* \leq (w^*)^\top Y \leq \rho^*, \\ 1, & (w^*)^\top Y > \rho^*, \end{cases} \tag{3.10}$$

where $\rho^* > 0$ is characterized in Appendix A.3, Equation (A.19). This implies existence of infinitely many MMR rules because the set of such rules is closed under convex combination.

Next, if the identified set is large enough for given Σ (or as Σ vanishes for given identified set), all of the above rules randomize for some data realizations. A natural question to ask is whether this feature is shared by all MMR rules. Parts (ii) and (iii) give qualified affirmative answers: If we focus on decision rules that increase in $(w^*)^\top Y$ and if $\overline{I}(\mathbf{0})$ is large enough, then randomization is necessary for MMR optimality; if bounds are furthermore differentiable in reduced-form parameters at $\mathbf{0}$, randomization is necessary for any MMR rule that depends on a linear index of the data.[9] In particular, the threshold rule

$$d^*_0((w^*)^\top Y) = \mathbf{1}\{(w^*)^\top Y \geq 0\} \tag{3.11}$$

is not MMR optimal.

Running Example—Continued: Theorem 3 applies to our running example. We next improve on this observation by providing explicit MMR optimal rules for the example. Broadly speaking,

[8] The relevant threshold will always be crossed as (other things equal) Σ vanishes. To the extent that small Σ models large sample size, multiplicity therefore becomes relevant for any partially identified DGP as the sample becomes large. Note, however, that the decision problems considered here give rise to stable asymptotic experiments only if the extent of partial identification also vanishes, and so these limit experiments may fall into either regime.

[9] The class of nonrandomized but otherwise unrestricted rules is not interestingly different from the class of all rules because high decimal places of Y_1 can be used to approximate randomization. Similarly, the differentiability condition is needed to preclude that some component of Y can effectively be used as randomization device.

these are characterized by a weighting of studentized signals that resembles a triangular kernel. □

Proposition 1. *In the running example, the following statements hold true.*

(i) If

$$C \|x_1 - x_0\| < \sqrt{\pi/2} \cdot \sigma_1, \tag{3.12}$$

then the following decision rule is uniquely (up to almost sure agreement) MMR optimal:

$$\begin{aligned} d_{m_0^*} &:= \mathbf{1}\{w_{m_0^*}^\top Y \geq 0\}, \\ w_{m_0^*}^\top &:= \left(1, \frac{\max\{m_0^* - C\|x_2 - x_0\|, 0\}/\sigma_2^2}{(m_0^* - C\|x_1 - x_0\|)/\sigma_1^2}, \ldots, \frac{\max\{m_0^* - C\|x_n - x_0\|, 0\}/\sigma_n^2}{(m_0^* - C\|x_1 - x_0\|)/\sigma_1^2}\right), \end{aligned}$$

where $m_0^* > C\|x_i - x_0\|$ solves a simple fixed point problem ((B.12) in Online Appendix B.2).

(ii) If

$$C \|x_1 - x_0\| = \sqrt{\pi/2} \cdot \sigma_1,$$

then $d^* := \mathbf{1}\{(1, 0, \ldots, 0)^\top Y \geq 0\}$ is MMR optimal.

(iii) If

$$C \|x_1 - x_0\| > \sqrt{\pi/2} \cdot \sigma_1, \tag{3.13}$$

then the rule $d_{RT}^*(\cdot)$ defined in (3.7) with

$$\tilde{\sigma} = \sqrt{2C^2 \|x_1 - x_0\|^2 / \pi - \sigma_1^2}, \quad w^* = (1, 0, \ldots, 0)^\top \tag{3.14}$$

is MMR optimal. So is the rule d_{linear}^* as defined in (3.10), where $\rho^* > 0$ is uniquely defined by $\rho^* = C\|x_1 - x_0\|(1 - 2\Phi(\rho^*/\sigma_1))$, as well as all convex combinations of these rules.

(iv) If Equation (3.13) holds and the nearest neighbor is unique (i.e., $\|x_1 - x_0\| < \|x_2 - x_0\|$), no linear threshold rule is MMR optimal.

Proof. See Online Appendix B.2. □

Remark 3. All decision rules above converge to the one from case (ii) as $C\|x_1 - x_0\| \to \sqrt{\pi/2} \cdot \sigma_1$.

Remark 4. This solution relates to the literature as follows. The problem is within the framework considered by Yata (2021), and his analysis applies; in particular, our linear index differs from his w^* only by a more explicit characterization. The alternative solutions in (iii), the uniqueness statement, and part (iv) are new. Ishihara and Kitagawa (2021) numerically find a solution within

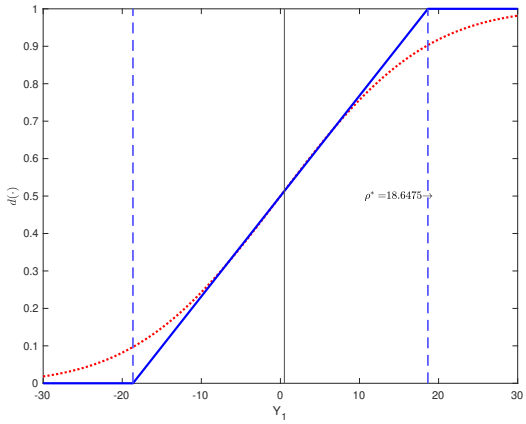

(a) d^*_{RT} (red) and d^*_{linear} (blue)

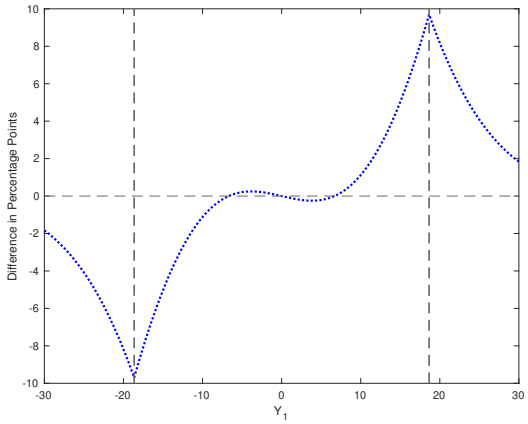

(b) $d^*_{\text{linear}} - d^*_{\text{RT}}$

Figure 1: Panel a): Visualization of rules (3.7) and (3.10) in the running example with estimates from two countries. (Setting $x_0 = 0$ yields $x_1 = -7.5$, $\sigma_1 = 3.9$, $x_2 = 7.9$, $\sigma_2 = 2.4$. We set $C = 2.5$.) Panel b): Difference between rules as a function of the nearest neighbor's outcome.

the class of symmetric threshold rules (i.e., rules of form $\mathbf{1}\{w^\top Y \geq 0\}$). This in principle recovers the global solution if $C \|x_1 - x_0\| \leq \sqrt{\pi/2} \cdot \sigma_1$ but will exclude all globally MMR optimal decision rules otherwise. That said, Ishihara and Kitagawa's (2021) solution approach applies considerably more generally.

Remark 5. If μ is exogenous and known, then the decision rule $d^*_{\text{known }\mu}(\mu) := \max\{\min\{\overline{I}(\mu)/(\overline{I}(\mu) - \underline{I}(\mu)), 1\}, 0\}$ uniquely attains MMR (Manski, 2007b). Only our new rule d^*_{linear} converges to $d^*_{\text{known }\mu}$ in certain special cases. Similarly, Ishihara and Kitagawa (2021, Section 3.3) discuss that an analogous convergence fails for any MMR rule that they propose. This may appear puzzling; however, any presumption that MMR rules "should" converge to such limits delicately depends on how one conceives the limit of the decision problem. Hence, it is not clear that we observe failure of any convergence that "should" have occurred.

We conclude that different MMR rules can lead to rather different policy actions for the same data. This difference is illustrated in Figure 1 for parameter values calibrated to Ishihara and Kitagawa's (2021) empirical example. How serious a challenge it is depends on one's view. If one truly thinks of MMR as encoding a decision maker's complete preferences, and hence of competing optimal rules as mutually indifferent, it is not much of a concern. However, it may make it harder to communicate MMR-based decisions to policy makers. In addition, the next section's results will show that one may plausibly have preferences among the different MMR rules.

4 Opportunities

In view of the preceding findings, we propose two different avenues for research: the judicious profiling out of parameters when evaluating risk functions, which we hope mitigates the thrust of some of the above results; and refinements of MMR, specifically toward least randomizing MMR, which gives rise to a unique solution in our setting.

4.1 Profiled Risk

Profiling out some parameters of expected welfare or regret may yield interesting insights. Beyond allowing to better visualize risk profiles of decision rules, it may render them comparable and may even exclude some of them as (in a redefined sense) inadmissible. We illustrate this idea by exploring a *profiled expected regret* criterion. While we could profile out any known function $h(m(\theta))$, we simplify exposition by restricting attention to linear indices. Thus, for a vector $w \in \mathbb{R}^n \backslash \{\mathbf{0}\}$, let

$$\Gamma_w := \left\{\gamma \in \mathbb{R} : w^\top m(\theta) = \gamma, \theta \in \Theta\right\}$$

be the image of the transformation $\theta \mapsto w^\top m(\theta)$. Then the worst-case expected regret of a rule d can be expressed as

$$\sup_{\theta\in\Theta} R(d,\theta) = \sup_{\gamma\in\Gamma_w}\left(\sup_{\theta\in\Theta: w^\top m(\theta)=\gamma} R(d,\theta)\right). \tag{4.1}$$

That is, we split (3.5) into an inner optimization problem with fixed level set $\{w^\top m(\theta) = \gamma\}$ and an outer optimization over γ. The value function of the inner problem may be of interest, thus define:

Definition 5 (w-Profiled Regret)**.** The w-profiled regret function $\overline{R}_w : \mathcal{D}_n \times \Gamma_w \to \mathbb{R}$ is given by

$$\overline{R}_w(d,\gamma) := \sup_{\theta\in\Theta: w^\top m(\theta)=\gamma} U(\theta)\left(\mathbf{1}\{U(\theta)\geq 0\} - \mathbb{E}_{m(\theta)}[d(Y)]\right), \tag{4.2}$$

where $Y \sim N(m(\theta), \Sigma)$.

In problems where MMR optimal rules depend on the data only through some linear combination $(w^*)^\top Y$, it seems reasonable to set w equal to w^*. However, there could be other vectors w of interest, as we illustrate when extrapolating Local Average Treatment effects in Section 5.1.

Remark 6 (Not all decision rules are admissible with respect to w-profiled regret). One could say

that a decision rule d is *w-profiled regret admissible* if there is no other rule d' for which

$$\overline{R}_w(d', \gamma) \leq \overline{R}_w(d, \gamma)$$

for every $\gamma \in \Gamma_w$, with strict inequality for some $\gamma \in \Gamma_w$. Even if Theorem 1 applies, a decision rule can fail to be w-profiled-regret admissible for some w. Indeed, the no-data rule $d_{\text{coin-flip}}(Y) = 1/2$ is w^*-profiled regret dominated by any MMR optimal rule in our running example.[10]

Remark 7 (Bayes rules with respect to w-profiled regret are Π^*-minimax under some conditions). Consider any decision rule that minimizes

$$\inf_{d \in \mathcal{D}_n} \int_{\Gamma_w} \overline{R}_w(d, \gamma) d\pi^*(\gamma), \tag{4.3}$$

the weighted average of w-profiled regret for some prior π^* over Γ_w. If π^* has full support and profiled regret is continuous in γ (for any d), then any solution to (4.3) is w-profiled-regret admissible; see Ferguson (1967, Theorem 3, Section 2, p. 62) and Berger (1985, Theorem 9, Section 4, p. 254).

We next provide sufficient conditions under which decision rules that solve (4.3) can be interpreted as Π^*-minimax decision rules in the sense of Berger (1985, Definition 13, p. 216), where Π^* is a class of priors over Θ.[11] Let D_w denote the class of decision rules that depend on the data only through $w^\top Y \sim N(w^\top m(\theta), w^\top \Sigma w)$. Consider the Π^*-minimax problem

$$\inf_{d \in \mathcal{D}_w} \left(\sup_{\pi \in \Pi^*} \int_\Theta R(d, \theta) d\pi(\theta) \right). \tag{4.4}$$

Let Π^* collect all priors over Θ for which $w^\top m(\theta) \sim \pi^*$, where π^* is the prior over Γ_w used in (4.3). This class of priors has recently been advocated by Giacomini and Kitagawa (2021). Their Theorem 2 establishes that, for any $d \in D_w$,

$$\sup_{\pi \in \Pi^*} \int_\Theta R(d, \theta) d\pi(\theta) = \int_{\Gamma_w} \overline{R}_w(d, \gamma) d\pi^*(\gamma). \tag{4.5}$$

Thus, the problem in (4.4) is equivalent to minimizing average w-profiled regret over decision rules that depend on the data only through $w^\top Y$. In contrast, we next show that decision rules minimizing average w-profiled regret cannot in general be computed by minimizing posterior regret.

Remark 8 (Bayes' rules with respect to w-profiled regret do not in general minimize Π^*-posterior

[10] Simple algebra shows that that rule's profiled regret is *minimized* at $\gamma = 0$, where it coincides with the *maximized* regret of any MMR optimal rule. This dominance can also be verified if $C \|x_1 - x_0\| \leq \sqrt{\pi/2} \cdot \sigma_1$. See Figure 2(a).

[11] To avoid notational conflict, we write Π^*-minimax instead of the more common Γ-minimax.

expected loss). Consider the loss function

$$L(a,\theta) := U(\theta)[\mathbf{1}\{U(\theta) \geq 0\} - a]$$

and note that $R(d,\theta) = \mathbb{E}_{m(\theta)}[L(d,\theta)]$. An alternative to Π^*-minimax decision rules are rules that minimize the worst-case posterior expected loss; see Berger (1985, Definition 10, p. 205) and Giacomini, Kitagawa, and Read (2021, Equation 2.11). To illustrate the difference, restrict attention to rules that depend on the data through $w^\top Y$. Then Π^*-posterior expected loss minimization solves

$$\inf_{a\in[0,1]} \sup_{\pi\in\Pi^*} \mathbb{E}_\pi[L(a,\theta) \mid w^\top Y]. \tag{4.6}$$

Theorem 2 in Giacomini and Kitagawa (2021) can be used to show that

$$\sup_{\pi\in\Pi^*} \mathbb{E}_\pi[L(a,\theta) \mid w^\top Y] = \int_{\Gamma_w} \left(\sup_{\theta\in\Theta: w^\top m(\theta)=\gamma} L(a,\theta) \right) d\pi^*(\gamma \mid w^\top Y). \tag{4.7}$$

Thus, the criterion in (4.6) is analogous to the one suggested by Christensen et al. (2022). One can easily construct examples in which these criteria disagree; we provide one after Figure 2. Indeed, recall that in terminology from axiomatic decision theory, Equation (4.7) uses full Bayesian updating of multiple prior preferences. This model is known to give rise to dynamically inconsistent behavior, meaning that a decision maker's preferred action after observing the data may not be what her preferred ex-ante decision rule recommends for that data realization.[12]

Running Example—Continued: Recall that, for some parameter values, there is an MMR rule depending on the data only through Y_1. Thus, consider w^*-profiled regret for $w^* = (1,0,\ldots,0)^\top$. Using (4.2), the w^*-profiled regret function equals

$$\overline{R}_{w^*}(d,\gamma) = \sup_{(\mu_0,\mu_1,\ldots,\mu_n)^\top\in\mathbb{R}^{n+1}:\mu_1=\gamma,\mu_0\in I(\mu),\mu\in M} \mu_0 \left(\mathbf{1}\{\mu_0 \geq 0\} - \mathbb{E}_\mu[d(Y)]\right),$$

and we can easily verify that $\Gamma_{w^*} = \mathbb{R}$. □

For the same parameters used in Figure 1, Figure 2(a) depicts w^*-profiled regret over the range $\gamma \in [-30, 30]$ for four decision rules. The blue (solid) line is d^*_{linear} (see Equation (3.10)); the red (dotted) line is d^*_{RT} (see Equation (3.7)); the black (bimodal, solid) line is the symmetric threshold

[12] See Hanany and Klibanoff (2007) for a detailed explanation and further references. To be clear, our point is that the criteria do not in general agree, not that this makes one of them better.

rule $d_0^*(Y) = \mathbf{1}\{Y_1 \geq 0\}$; and the green (dashed) line is $d_{\text{coin-flip}} = 1/2$.[13]

An immediate use of Figure 2(a) is to compare decision rules in terms of their worst-case regret. For example, consistent with Theorem 3-(ii), d_0^* is not MMR optimal. At the same time, d_0^* optimizes worst-case posterior expected loss under a symmetric prior on θ, providing the example alluded to in Remark 8 and illustrating that the current setting is not among the special cases in which minimax regret is dynamically consistent. For the specific parameter values used here, we can furthermore show that d_0^* is minimax regret optimal among linear threshold rules; thus, the black line also illustrates the minimax regret efficiency loss from restricting attention to linear threshold rules.

Figure 2(a) also reveals interesting differences among MMR optimal rules. In particular, d^*_{linear} has smaller w^*-profiled regret than d^*_{RT} virtually everywhere (though algebra shows that this dominance does not hold for small nonzero γ). Finally, Figure 2(a) illustrates that, as remarked earlier, $d_{\text{coin-flip}}$ is w^*-profiled regret inadmissible in this example. We hope that future research will yield more powerful results on profiled admissibility.

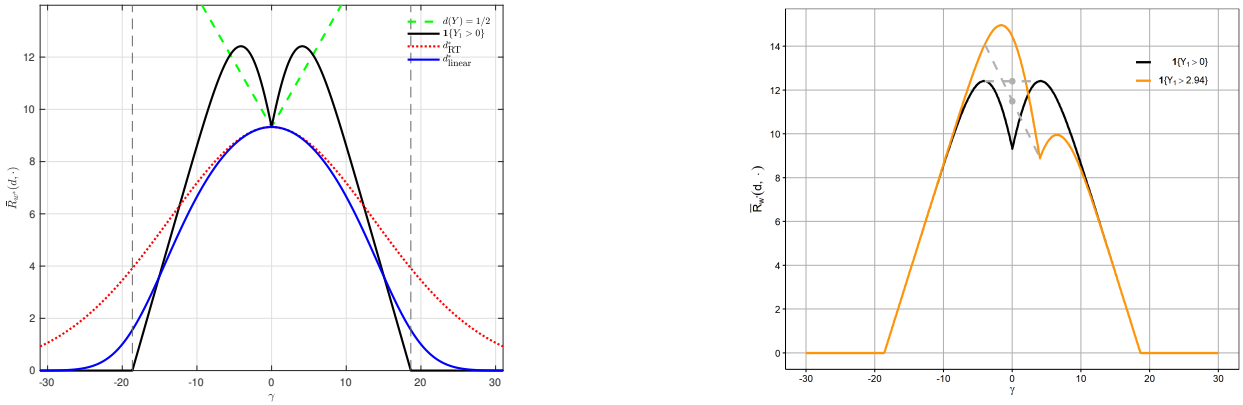

(a) w^*-profiled regret of four decision rules

(b) $\mathbf{1}\{Y_1 > 0\}$ v.s. $\mathbf{1}\{Y_1 > 2.94\}$.

Figure 2: w^*-profiled regret in the running example; parameter values are as in Figure 1.

We close with three brief remarks about profiled regret. First, the considerable difference between profiled regret functions in Figure 2(a) continues beyond the figure. Indeed, the ratio $\overline{R}_{w^*}(d^*_{\text{linear}}, \gamma)/\overline{R}_{w^*}(d^*_{\text{RT}}, \gamma)$ decays to zero at exponential rate as $\gamma \to \pm\infty$.

Second, consider a symmetric, uniform two-point prior on γ supported at the two points that maximize the black (solid) line in Figure 2(a). Among nonrandomized rules, the threshold rule $d_0^*(Y) = \mathbf{1}\{Y_1 > 0\}$ solves the corresponding Π^*-posterior loss minimization problem (4.6). However, d_0^* cannot solve the corresponding Π^*-minimax problem (4.4), even when we consider only

[13] Appendix B.3.1 presents algebraic and computational details that underlie this figure.

nonrandomized threshold decision rules. This is because, as visualized in Figure 2(b), the threshold rule $\mathbf{1}\{Y_1 > 2.94\}$ has smaller average w^*-profiled regret with respect to this prior.

Third, computation of profiled regret can frequently be simplified. Algebra shows that $\overline{R}_w(d, \gamma)$ equals the maximum between

$$\bar{k}_w^+(\gamma) := \sup_{\theta \in \Theta} U(\theta)(1 - \mathbb{E}_{m(\theta)}[d(Y)]) \quad \text{s.t. } w^\top m(\theta) = \gamma, \quad U(\theta) \geq 0, \tag{4.8}$$

and

$$\bar{k}_w^-(\gamma) := \sup_{\theta \in \Theta} -U(\theta)\mathbb{E}_{m(\theta)}[d(Y)] \quad \text{s.t. } w^\top m(\theta) = \gamma, \quad U(\theta) \leq 0. \tag{4.9}$$

In principle, these are the value functions of two infinite-dimensional, nonlinear optimization problems. However, they can be recast as finite dimensional. For example, $\bar{k}_w^+(\gamma)$ equals

$$\overline{I}_w^+(\gamma) := \sup_{\mu \in M} \overline{I}(\mu)(1 - \mathbb{E}_\mu[d(Y)])) \quad \text{s.t. } w^\top \mu = \gamma, \quad \overline{I}(\mu) \geq 0. \tag{4.10}$$

This problem has a scalar choice variable, one linear equality constraint, and one potentially nonlinear inequality constraint. The bottleneck is evaluation of $\overline{I}(\mu)$. The running example admits a closed-form solution for $\overline{I}(\mu)$, so that evaluating (4.10) is easy.

More generally, the computational cost of evaluating (4.10) can be reduced by imposing more structure on the parameter space Θ. For instance, when Θ is convex, the set M is convex as well; the optimization problem is then over a convex subset of $\mathbb{R}^n$. Moreover, if $m(\cdot)$ is linear and $U(\cdot)$ is concave, the function $\overline{I}(\mu)$ can be shown to be concave. This means that under Assumption 1, the optimization problem in (4.10) is convex.

4.2 Least Randomizing MMR Optimal Rules

We next argue that further refining the MMR criterion presents an interesting research opportunity and may even lead to unique recommendations. To this purpose, we propose consideration of, and characterize, the *least randomizing* MMR rule. Our main motivation is that, despite the wide adoption of randomized treatment allocations in economics and the social sciences, policy makers might shy away from exposing only a fraction of a population to the new policy. Thus, we attempt to recommend actions $a \in (0, 1)$ as infrequently as possible.[14]

[14] A different approach to deal with multiplicity of MMR optimal rules would be to directly build attitude to randomization into the welfare or risk functions. For example, consider the mean square regret criterion proposed

To formalize this, observe that all of d^*_{RT}, d^*_{linear}, and d^*_{limit} can be considered smoothed versions of d^*_0 in a sense that we now make precise and that resembles the progression from Manski (1975) to Horowitz (1992). Let $F : \mathbb{R} \to [0,1]$ be a c.d.f. and consider a decision rule of form $F \circ w^* := F((w^*)^\top Y) \in \mathcal{D}_n$; that is, the step function d^*_0 was smoothed into a c.d.f. We will restrict attention to c.d.f.'s that are *symmetric* (i.e., $F(-x) = 1 - F(x)$) and *unimodal* (i.e., $F(\cdot)$ is convex for $x \leq 0$ and concave otherwise). Let $\mathcal{F}$ be the set of all such c.d.f's and let

$$\tilde{\mathcal{D}}_n := \{F \circ w^* \in \mathcal{D}_n : F \in \mathcal{F}\}.$$

Note that each rule $F \circ w^* \in \tilde{\mathcal{D}}_n$ depends on the data only via $(w^*)^\top Y$ and is nondecreasing in $(w^*)^\top Y$. Moreover, for each $F \circ w^* \in \tilde{\mathcal{D}}_n$, the interval on which treatment assignment is randomized equals (up to closure)

$$V(F \circ w^*) := \big(\sup\{x \in \mathbb{R} : F(x) = 0\}, \inf\{x \in \mathbb{R} : F(x) = 1\}\big). \tag{4.11}$$

All MMR decision rules considered in this paper are in $\tilde{\mathcal{D}}_n$. We next show that d^*_{linear} is least randomizing among them and among all other MMR decision rules that might exist in this class.

Theorem 4. *Suppose all conditions of Theorem 3 hold. If $F \circ w^* \in \tilde{\mathcal{D}}_n$ is MMR optimal, then $V(d^*_{linear} \circ w^*) \subseteq V(F \circ w^*)$, with equality if and only if $F = d^*_{linear}$.*

Proof. See Appendix A.4. □

In words, any symmetric, weakly increasing and unimodal MMR optimal rule that depends on data only via $(w^*)^\top Y$ must have a randomization area that is wider than that of d^*_{linear}, strictly so if it is a meaningfully distinct rule. Thus, the least randomizing criterion provides a pragmatic and unique refinement among the set of known MMR optimal rules.

To establish Theorem 4, we first show that for any rule $F \circ w^* \in \tilde{\mathcal{D}}_n$, its expected regret at any parameter θ for which $(w^*)^\top m(\theta) = 0$ equals the MMR value of the problem. If $F \circ w^* \in \tilde{\mathcal{D}}_n$ is MMR optimal, its expected regret must therefore be maximized at $(w^*)^\top m(\theta) = 0$. For any symmetric and unimodal c.d.f. F with $V(d^*_{\text{linear}} \circ w^*) \not\subseteq V(F \circ w^*)$, we can show that a necessary condition for this maximization fails.

Running Example—Continued: Recall that, applied to the running example and for C large

in Kitagawa et al. (2022). In a stylized example of a one-dimensional signal, Theorem 3.1 in Kitagawa et al. (2023) implies, in general, a unique minimax-optimal rule based on mean square regret.

enough, d^*_{linear} can be expressed as

$$d^*_{\text{linear}}(Y_1) := \begin{cases} 0, & Y_1 < -\rho^*, \\ \frac{Y_1+\rho^*}{2\rho^*}, & -\rho^* \leq Y_1 \leq \rho^*, \\ 1, & Y_1 > \rho^*, \end{cases} \tag{4.12}$$

where $\rho^* \in (0, C\,\|x_1 - x_0\|)$ is uniquely defined by

$$\rho^* = C\,\|x_1 - x_0\|\,(1 - 2\Phi\,(\rho^*/\sigma_1)).$$

Of note, d^*_{linear} only randomizes when $|Y_1| < \rho^*$. This is of interest because, whenever this happens, the natural plug-in estimator of the identified set for μ_0 given Y_1, i.e. the interval

$$[\,Y_1 - C\,\|x_1 - x_0\|\;,\; Y_1 + C\,\|x_1 - x_0\|\,]\,, \tag{4.13}$$

contains 0. Conversely, $d^*_{\text{linear}}(Y_1)$ always (never) implements the new policy when (4.13) is to the right (left) of zero. In this sense, the estimated identified set is explicitly used for decision making. In contrast, d^*_{RT} always randomizes the policy recommendation, although for large Y_1 the fraction of population assigned to treatment will be large.

It would be interesting to compare $d^*_{\text{linear}}(Y_1)$ to a plug-in rule based on an estimated version of $I(\mu)$.[15] In the example, a natural such estimator would be $[\widehat{\underline{I}}, \widehat{\overline{I}}]$, where

$$\widehat{\overline{I}} = \min_{i=1,\ldots,n} \{\hat{\mu}_i + C\,\|x_i - x_0\|\}\,, \quad \widehat{\underline{I}} = \max_{i=1,\ldots,n} \{\hat{\mu}_i - C\,\|x_i - x_0\|\}\,,$$

and $(\hat{\mu}_1, \ldots, \hat{\mu}_n)$ is a constrained (to M) maximum likelihood estimator. Based on the estimated identified set $[\widehat{\underline{I}}, \widehat{\overline{I}}]$, a natural plug-in rule is then

$$d_{\text{plug-in}}(Y) := \begin{cases} 0, & \widehat{\overline{I}} < 0, \\ \frac{\widehat{\overline{I}}}{\widehat{\overline{I}} - \widehat{\underline{I}}}, & \text{otherwise}, \\ 1, & \widehat{\underline{I}} > 0. \end{cases} \tag{4.14}$$

In Appendix B.3.2, we plot and compare the w^*-profiled regrets of $d_{\text{plug-in}}$ and other rules, including

[15] Christensen et al. (2022) formally compare their proposed rule to "as-if" rules. However, Bayes optimal rules with respect to $w-$profiled regret generally disagree with those in Christensen et al. (2022) (see Remark 8), so their findings do not apply here.

d^*_{linear} and d^*_{RT}. Although $d_{\text{plug-in}}$ is not MMR optimal, its worst-case regret is only slightly larger than the MMR value, and its profiled regret is a bell-shaped curve similar to that of d^*_{linear}. □

5 Further Applications

5.1 Extrapolating Local Average Treatment Effects

We next apply our analysis to extrapolation of Local Average Treatment Effects (Mogstad et al., 2018; Mogstad and Torgovitsky, 2018). Let $Z \in \{0,1\}$ be a binary instrument, $D \in \{0,1\}$ a binary treatment assignment, and $(Y(1), Y(0))$ potential outcomes under treatment and control. As usual, the observed outcome is $Y = DY(1) + (1-D)Y(0)$. To simplify exposition, we assume that there are no covariates and that $Y(1), Y(0) \in \{0,1\}$. Following Heckman and Vytlacil (1999, 2005),[16] let $p(z) := P\{D = 1 \mid Z = z\}$ be the propensity score and write $D = \mathbf{1}\{V \leq p(Z)\}$, where $(V \mid Z = z) \sim \text{Unif}[0,1]$. The parameter space Θ contains all tuples $\theta := (p(1), p(0), \text{MTE}(\cdot))$, where $p(1) \in [0,1]$, $p(0) \in [0,1]$, $p(1) \geq p(0)$, and $\text{MTE}(\cdot)$ is the marginal treatment effect function

$$\text{MTE}(v) := \mathbb{E}[Y(1) - Y(0) \mid V = v].$$

The policy maker observes

$$\begin{pmatrix} \hat{m}_1 \\ \hat{m}_2 \end{pmatrix} \sim N\left(\begin{pmatrix} m_1(\theta) \\ m_2(\theta) \end{pmatrix}, \Sigma \right), \tag{5.1}$$

where

$$\begin{aligned} m_1(\theta) &:= \mathbb{E}[Y \mid Z = 1] - \mathbb{E}[Y \mid Z = 0] = \int_{p(0)}^{p(1)} \text{MTE}(v) dv \\ m_2(\theta) &:= \mathbb{E}[D \mid Z = 1] - \mathbb{E}[D \mid Z = 0] = p(1) - p(0) \end{aligned}$$

are the population reduced-form and first-stage coefficients and Σ is positive definite. We assume that the policy of interest would expand the complier subpopulation through an additive shift of size $\alpha > 0$ in the propensity score. Mogstad et al. (2018) show that the payoff relevant parameter then is the "policy-relevant treatment effect" (Heckman and Vytlacil, 2005)

$$\text{PRTE}(\alpha) = \mathbb{E}[Y(1) - Y(0) \mid V \in (p(0), p(1) + \alpha]].$$

[16] See, for example, Assumption I and Equation (2) in Mogstad and Torgovitsky (2018). Also see Imbens and Angrist (1994) for additional references.

If we normalize $W(0,\theta) = 0$, the welfare contrast $U(\theta)$ is equal to PRTE(α) and can be written as

$$U(\theta) := \text{PRTE}(\alpha) = \frac{m_1(\theta)}{\alpha + m_2(\theta)} + \frac{1}{\alpha + m_2(\theta)} \int_{p(1)}^{p(1)+\alpha} \text{MTE}(v)dv. \tag{5.2}$$

Hence, the decision maker wants to find an optimal treatment policy given partial identification of parameter (5.2) in model (5.1). In Online Appendix B.3.3, we verify that Theorem 1 applies. Therefore, any decision rule is admissible in this example. For example, if the PRTE(α) were the payoff relevant parameter, implementing a policy for large values of the IV estimator would be admissible, as would be the approach of Christensen et al. (2022), who discuss the same application.

It would be an interesting exercise to report the w-profiled regret of decision rules. In the context of this example, one might reasonably use

$$w = (1, -\beta_0)^\top / \sqrt{(1, -\beta_0)\Sigma(1, -\beta_0)^\top}$$

for some $\beta_0 \in \mathbb{R}$. For motivation, note that the square of

$$(m_1(\theta) - \beta_0 m_2(\theta)) / \sqrt{(1, -\beta_0)\Sigma(1, -\beta_0)^\top}$$

can be viewed as the population Anderson and Rubin (1949) statistic for the null hypothesis of β_0. Thus, the profiled regret function reports the worst-case regret as one keeps constant the population analog of that statistic.

5.2 Decision-theoretic Breakdown Analysis

Consider a policy maker who uses quasi-experimental data but is worried about confounding. More specifically, she assumes a constant treatment effect model and unconfoundedness given covariates (X, W), motivating the linear regression model

$$Y = \gamma_0 + \beta_{\text{long}} D + \gamma_1^\top X + \gamma_2^\top W + e,$$

where Y is observed outcome, D is the binary treatment, and e is a projection residual. The infeasible optimal treatment policy is $\mathbf{1}\{\beta_{\text{long}} \geq 0\}$. However, W is unobserved, so that the policy

maker can only estimate the "medium" regression

$$Y = \pi_0 + \beta_{\text{med}} D + \pi_1^\top X + u,$$

where u is a projection residual.[17] In general, if there exists selection on unobservables (that is, D is correlated with W), then β_{long} is only partially identified. Specifically, Diegert et al. (2022, Theorem 4) show that the identified set of β_{long} given β_{med} is

$$\beta_{\text{long}} \in [\beta_{\text{med}} - k, \beta_{\text{med}} + k],$$

where

$$k := \begin{cases} \sqrt{\frac{\text{var}\left(Y^{\perp D,X}\right)}{\text{var}\left(D^{\perp X}\right)} \frac{\overline{r}_D^2 R^2_{D\sim X}}{1 - R^2_{D\sim X} - \overline{r}_D^2}}, & \text{if } 0 \leq \overline{r}_D < \sqrt{1 - R^2_{D\sim X}}, \\ \infty, & \text{if } \overline{r}_D \geq \sqrt{1 - R^2_{D\sim X}}, \end{cases}$$

and where $\text{var}\left(Y^{\perp D,X}\right)$ is the variance of the residual from projecting Y onto $(1, D, X)$, $\text{var}\left(D^{\perp X}\right)$ is the variance of the residual from projecting D onto $(1, X)$, $R^2_{D\sim X}$ is the R^2 from projecting D onto $(1, X)$, and $\overline{r}_D \geq 0$ is a user-specified sensitivity parameter that measures the relative importance of selection on unobservables versus selection on observables.[18]

Diegert et al. (2022) use this result to ask: How strong does omitted variables bias have to be to potentially overturn findings based on β_{med}? At population level, the answer is that this can happen if $|k| > |\beta_{\text{med}}|$, a condition that can be related to primitive parameters through the above display and for which Diegert et al. (2022) provide estimation and inference theory.

Suppose now that there is an estimator $\hat{\beta}_{\text{med}} \sim N(\beta_{\text{med}}, \sigma^2)$. Then our results apply upon letting $\theta = (\beta_{\text{long}}, \beta_{\text{med}})^\top \in \mathbb{R}^2$, $U(\theta) = \beta_{\text{long}}$ and $m(\theta) = \beta_{\text{med}}$. In particular, when $k > \sqrt{\frac{\pi}{2}}\sigma$, there are infinitely many MMR optimal rules, with the least randomizing one among known ones being

$$d^*_{\text{linear}}(\hat{\beta}_{\text{med}}) := \begin{cases} 0, & \hat{\beta}_{\text{med}} < -\rho^*, \\ \frac{\hat{\beta}_{\text{med}} + \rho^*}{2\rho^*}, & -\rho^* \leq \hat{\beta}_{\text{med}} \leq \rho^*, \\ 1, & \hat{\beta}_{\text{med}} > \rho^*, \end{cases} \tag{5.3}$$

[17] We express all regressions as projections for alignment with the literature and because only projection algebra is used. However, motivating $\mathbf{1}\{\beta_{\text{long}} \geq 0\}$ as optimal usually requires causal interpretation and therefore slightly stronger assumptions on e; in other words, readers may want to think of the long regression as causal and the medium one as best linear prediction. See Hansen (2022, Chapter 2), whose notation we also borrow, for a lucid discussion.

[18] In practice, $\text{var}\left(Y^{\perp D,X}\right)$, $\text{var}\left(D^{\perp X}\right)$ and $R^2_{D\sim X}$ are unknown. However, as discussed in Section 2, we may treat the joint distribution of $\{Y, D, X\}$ as multivariate normal with known covariance matrix. As $\text{var}\left(Y^{\perp D,X}\right)$, $\text{var}\left(D^{\perp X}\right)$, and $R^2_{D\sim X}$ are functions of the covariance matrix, it is then coherent to take them to be known as well.

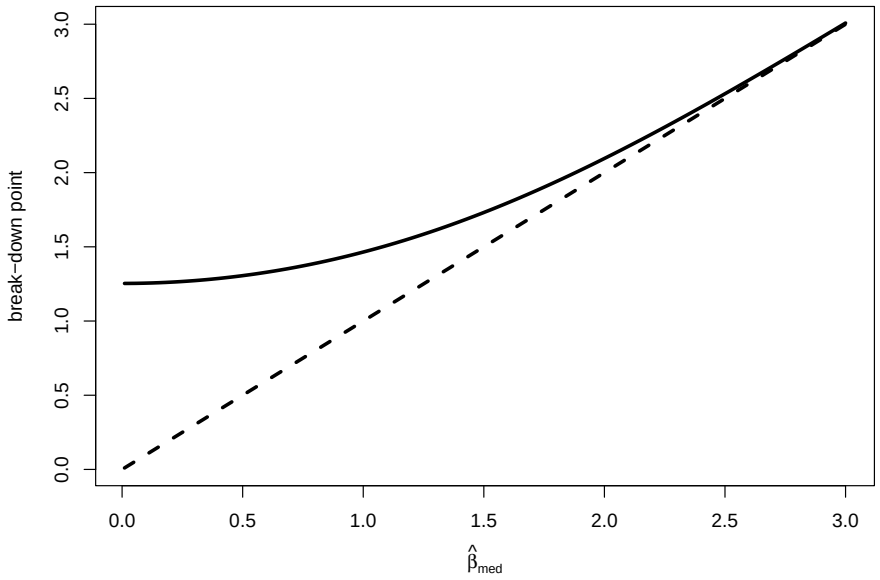

Figure 3: $\bar{k}$ (solid) and $\tilde{k}$ (dashed) as a function of $\hat{\beta}_{\text{med}}$

where $\rho^* > 0$ uniquely solves $\rho^* = k(1-2\Phi(-\rho^*/\sigma))$. When $k \leq \sqrt{\frac{\pi}{2}}\sigma$, we know from Stoye (2012a) that $d_0^*(\hat{\beta}_{\text{med}}) = \mathbf{1}\{\hat{\beta}_{\text{med}} \geq 0\}$ is essentially uniquely MMR optimal.

These results motivate a complementary breakdown analysis guided by statistical decision theory. For given $\hat{\beta}_{\text{med}} > 0$, we can ask: How large could k have to be so that the MMR optimality criterion still supports assigning the new policy without any hedging?[19] Due to its least randomizing property, d^*_{linear} implies the tightest possible answer to this question. Specifically, MMR supports non-randomized policy assignment up to the "decision theoretic breakdown point"

$$\begin{aligned} \bar{k}(\hat{\beta}_{\text{med}}) &:= \sup\{k > 0 : d^*(\hat{\beta}_{\text{med}}) = 1\} \\ d^*(\hat{\beta}_{\text{med}}) &:= \begin{cases} \mathbf{1}\{\hat{\beta}_{\text{med}} \geq 0\} & \text{if } k \leq \sqrt{\frac{\pi}{2}}\sigma \\ d^*_{\text{linear}}(\hat{\beta}_{\text{med}}) & \text{if } k > \sqrt{\frac{\pi}{2}}\sigma \end{cases}. \end{aligned}$$

Figure 3 displays both breakdown points as functions of $\hat{\beta}_{\text{med}}$ when $\sigma = 1$. It turns out that the decision theoretic breakdown point tolerates more ambiguity; this difference is salient for smaller values of $\hat{\beta}_{\text{med}}$ and vanishes as $\hat{\beta}_{\text{med}}$ diverges.

[19] Informal exploration of this question goes back at least to Stoye (2009b, see Table 3).

6 Conclusion

In this paper, we used statistical decision theory to argue that treatment choice problems with partial identification present important theoretical and practical challenges as well as interesting research opportunities. For a large and empirically relevant class of such problems, we show that every decision rule is admissible, that maximin welfare optimality criterion often select no-data decision rules, and that there are infinitely many minimax regret optimal rules, all of which randomize the policy action at least for some data realizations. These results stand in stark contrast with treatment choice problems with point-identified welfare.

We also discuss concrete ideas for overcoming these issues. First, judicious *profiling* of regret may help to visualize and summarize the risk function, suggest that some rules outperform others by commonsensical standards, and potentially give rise to an informative admissibility criterion. Second, we provide a decision rule that is *least randomizing* in a large class of MMR optimal rules including all known ones. We illustrate our results in three applications that arise in applied work: extrapolation of experimental estimates for policy adoption, policy-making with quasi-experimental data when omitted variable bias is a concern, and extrapolation of Local Average Treatment Effects.

We believe there are many interesting research directions for future work. For example, it would be worthwhile to investigate the profiled regret of different decision rules suggested in the recent literature. We would also be interested to learn about other refinement criteria. Finally, it would be interesting to investigate more in detail the properties of the "as if" or plug-in approach that uses information from an estimated identified set for decision making.

References

Adjaho, C. and T. Christensen (2022): "Externally Valid Treatment Choice," *arXiv preprint arXiv:2205.05561*.

Anderson, T. and H. Rubin (1949): "Estimation of the Parameters of a Single Equation in a Complete System of Stochastic Equations," *The Annals of Mathematical Statistics*, 20, 46–63.

Andrews, I. and A. Mikusheva (2022): "GMM is Inadmissible Under Weak Identification," *arXiv preprint arXiv:2204.12462*.

Aradillas Fernández, A., J. L. Montiel Olea, C. Qiu, J. Stoye, and S. Tinda (2024): "Robust Bayes Treatment Choice with Partial Identification," Discussion paper, Cornell University.

Armstrong, T. B. and M. Kolesár (2018): "Optimal inference in a class of regression models," *Econometrica*, 86, 655–683.

Athey, S. and S. Wager (2021): "Efficient policy learning with observational data," *Econometrica*, 89, 133–161.

Ben-Michael, E., D. J. Greiner, K. Imai, and Z. Jiang (2021): "Safe policy learning through extrapolation: Application to pre-trial risk assessment," *arXiv preprint arXiv:2109.11679*.

Ben-Michael, E., K. Imai, and Z. Jiang (2022): "Policy learning with asymmetric utilities," *arXiv preprint arXiv:2206.10479*.

Berger, J. (1985): *Statistical decision theory and Bayesian analysis*, Springer.

Bhattacharya, D. and P. Dupas (2012): "Inferring welfare maximizing treatment assignment under budget constraints," *Journal of Econometrics*, 167, 168–196.

Canner, P. L. (1970): "Selecting One of Two Treatments when the Responses are Dichotomous," *Journal of the American Statistical Association*, 65, 293–306.

Casella, G. and R. L. Berger (2002): *Statistical inference*, Thomson Learning.

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

——— (2011): "Bayesian Aspects of Treatment Choice," in *The Oxford Handbook of Bayesian Econometrics*, Oxford University Press.

Chen, H. and P. Guggenberger (2024): "A note on minimax regret rules with multiple treatments in finite samples," Discussion paper, The Pennsylvania State University.

Christensen, T., H. R. Moon, and F. Schorfheide (2022): "Optimal Discrete Decisions when Payoffs are Partially Identified," *arXiv preprint arXiv:2204.11748*.

Clarke, F. H. (1990): *Optimization and nonsmooth analysis*, SIAM.

D'Adamo, R. (2021): "Policy Learning Under Ambiguity," *arXiv preprint arXiv:2111.10904*.

Dehejia (2005): "Program evaluation as a decision problem," *Journal of Econometrics*, 125, 141–173.

Diegert, P., M. A. Masten, and A. Poirier (2022): "Assessing omitted variable bias when the controls are endogenous," *arXiv preprint arXiv:2206.02303*.

Donoho, D. L. (1994): "Statistical estimation and optimal recovery," *The Annals of Statistics*, 22, 238–270.

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

Giacomini, R. and T. Kitagawa (2021): "Robust Bayesian Inference for Set-Identified Models," *Econometrica*, 89, 1519–1556.

Giacomini, R., T. Kitagawa, and M. Read (2021): "Robust Bayesian Analysis for Econometrics," *CEPR Discussion Paper No. 16488*.

Hanany, E. and P. Klibanoff (2007): "Updating preferences with multiple priors," *Theoretical Economics*, 2, 261–298.

Hansen, B. E. (2022): *Econometrics*, Princeton University Press.

Heckman, J. J. and E. J. Vytlacil (1999): "Local instrumental variables and latent variable models for identifying and bounding treatment effects," *Proceedings of the National Academy of Sciences*, 96, 4730–4734.

——— (2005): "Structural equations, treatment effects, and econometric policy evaluation 1," *Econometrica*, 73, 669–738.

Hirano, K. and J. R. Porter (2009): "Asymptotics for statistical treatment rules," *Econometrica*, 77, 1683–1701.

——— (2020): "Asymptotic analysis of statistical decision rules in econometrics," in *Handbook of Econometrics, Volume 7A*, ed. by S. N. Durlauf, L. P. Hansen, J. J. Heckman, and R. L. Matzkin, Elsevier, vol. 7 of *Handbook of Econometrics*, 283–354.

Horowitz, J. L. (1992): "A Smoothed Maximum Score Estimator for the Binary Response Model," *Econometrica*, 60, 505–531.

Ida, T., T. Ishihara, K. Ito, D. Kido, T. Kitagawa, S. Sakaguchi, and S. Sasaki (2022): "Choosing Who Chooses: Selection-Driven Targeting in Energy Rebate Programs," Tech. rep., National Bureau of Economic Research.

Imbens, G. W. and J. D. Angrist (1994): "Identification and Estimation of Local Average Treatment Effects," *Econometrica*, 467–475.

Ishihara, T. and T. Kitagawa (2021): "Evidence Aggregation for Treatment Choice," ArXiv:2108.06473 [econ.EM], https://doi.org/10.48550/arXiv.2108.06473.

Kallus, N. and A. Zhou (2018): "Confounding-robust policy improvement," *Advances in neural information processing systems*, 31.

Karlin, S. and H. Rubin (1956): "The theory of decision procedures for distributions with monotone likelihood ratio," *The Annals of Mathematical Statistics*, 272–299.

Kido, D. (2022): "Distributionally Robust Policy Learning with Wasserstein Distance," *arXiv preprint arXiv:2205.04637*.

——— (2023): "Locally Asymptotically Minimax Statistical Treatment Rules Under Partial Identification," *arXiv preprint arXiv:2311.08958*.

Kitagawa, T., S. Lee, and C. Qiu (2022): "Treatment Choice with Nonlinear Regret," *arXiv preprint arXiv:2205.08586*.

——— (2023): "Treatment choice, mean square regret and partial identification," *The Japanese Economic Review*, 74, 573–602.

Kitagawa, T., S. Sakaguchi, and A. Tetenov (2021): "Constrained classification and policy learning," *arXiv preprint*.

Kitagawa, T. and A. Tetenov (2018): "Who should be treated? Empirical welfare maximization methods for treatment choice," *Econometrica*, 86, 591–616.

——— (2021): "Equality-Minded Treatment Choice," *Journal of Business Economics and Statistics*, 39, 561–574.

Kitagawa, T. and G. Wang (2020): "Who should get vaccinated? Individualized allocation of vaccines over SIR network," *arXiv preprint*.

Lehmann, E. L. and J. P. Romano (2005): *Testing statistical hypotheses*, vol. 3, Springer.

Lei, L., R. Sahoo, and S. Wager (2023): "Policy Learning under Biased Sample Selection," *arXiv preprint arXiv:2304.11735*.

Low, M. G. (1995): "Bias-variance tradeoffs in functional estimation problems," *The Annals of Statistics*, 23, 824–835.

Manski, C. F. (1975): "Maximum score estimation of the stochastic utility model of choice," *Journal of Econometrics*, 3, 205–228.

——— (2000): "Identification problems and decisions under ambiguity: empirical analysis of treatment response and normative analysis of treatment choice," *Journal of Econometrics*, 95, 415–442.

——— (2004): "Statistical treatment rules for heterogeneous populations," *Econometrica*, 72, 1221–1246.

——— (2005): *Social choice with partial knowledge of treatment response*, Princeton University Press.

——— (2007a): *Identification for prediction and decision*, Harvard University Press.

——— (2007b): "Minimax-regret treatment choice with missing outcome data," *Journal of Econometrics*, 139, 105–115.

——— (2020): "Towards credible patient-centered meta-analysis," *Epidemiology*.

——— (2022): "Identification and Statistical Decision Theory," *arXiv preprint arXiv:2204.11318*.

Manski, C. F. and A. Tetenov (2007): "Admissible treatment rules for a risk-averse planner with experimental data on an innovation," *Journal of Statistical Planning and Inference*, 137, 1998–2010.

Mbakop, E. and M. Tabord-Meehan (2021): "Model selection for treatment choice: Penalized welfare maximization," *Econometrica*, 89, 825–848.

Menzel, K. (2023): "Transfer Estimates for Causal Effects across Heterogeneous Sites," *arXiv preprint arXiv:2305.01435*.

Mogstad, M., A. Santos, and A. Torgovitsky (2018): "Using instrumental variables for inference about policy relevant treatment parameters," *Econometrica*, 86, 1589–1619.

Mogstad, M. and A. Torgovitsky (2018): "Identification and extrapolation of causal effects with instrumental variables," *Annual Review of Economics*, 10, 577–613.

Müller, U. K. (2011): "Efficient tests under a weak convergence assumption," *Econometrica*, 79, 395–435.

Parmigiani, G. (1992): "Minimax, information and ultrapessimism," *Theory and Decision*, 33, 241–252.

Rockafellar, R. T. (1997): *Convex analysis*, vol. 11, Princeton university press.

Sadler, E. (2015): "Minimax and the value of information," *Theory and Decision*, 78, 575–586.

Savage, L. (1951): "The theory of statistical decision," *Journal of the American Statistical Association*, 46, 55–67.

Schlag, K. H. (2006): "ELEVEN - Tests needed for a Recommendation," Tech. rep., European University Institute Working Paper, ECO No. 2006/2, https://cadmus.eui.eu/bitstream/handle/1814/3937/ECO2006-2.pdf.

Skorohod, A. V. (2012): *Integration in Hilbert space*, vol. 79, Springer Science & Business Media.

Stoye, J. (2007): "Minimax Regret Treatment Choice With Incomplete Data and many Treatments," *Econometric Theory*, 23, 190–199.

——— (2009a): "Minimax regret treatment choice with finite samples," *Journal of Econometrics*, 151, 70–81.

——— (2009b): "Partial identification and robust treatment choice: an application to young offenders," *Journal of Statistical Theory and Practice*, 3, 239–254.

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

——— (2012b): "New perspectives on statistical decisions under ambiguity," *Annu. Rev. Econ.*, 4, 257–282.

Tetenov, A. (2012): "Statistical treatment choice based on asymmetric minimax regret criteria," *Journal of Econometrics*, 166, 157–165.

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.

Yu, G. and P. Kouvelis (1995): "On min-max optimization of a collection of classical discrete

optimization problems," in *Minimax and applications*, Springer, 157–171.

A Proofs of Main Results

A.1 Proof of Theorem 1

Suppose by contradiction that some rule d is dominated. Then there exists d' such that

$$U(\theta)\mathbb{E}_{m(\theta)}\left[d'(Y)\right] \geq U(\theta)\mathbb{E}_{m(\theta)}\left[d(Y)\right] \tag{A.1}$$

for all $\theta \in \Theta$.

Step 1: We first show that (A.1) must hold with equality for any $\theta \in m^{-1}(\mathcal{S})$. Suppose not, then there exists $\theta^* \in m^{-1}(\mathcal{S})$ such that

$$U(\theta^*)\mathbb{E}_{m(\theta^*)}\left[d'(Y)\right] > U(\theta^*)\mathbb{E}_{m(\theta^*)}\left[d(Y)\right], \tag{A.2}$$

which implies that $U(\theta^*) \neq 0$. Without loss of generality, assume $U(\theta^*) > 0$. Then

$$\mathbb{E}_{m(\theta^*)}\left[d'(Y)\right] > \mathbb{E}_{m(\theta^*)}\left[d(Y)\right].$$

Define $\mu^* := m(\theta^*)$. Since $\mu^* \in \mathcal{S}$, there must exist $\theta_{\mu^*}, \tilde{\theta}_{\mu^*} \in \Theta$ such that $U(\theta_{\mu^*}) < 0$, $U(\tilde{\theta}_{\mu^*}) > 0$. Therefore, (A.2) implies

$$U(\theta_{\mu^*})\mathbb{E}_{m(\theta_{\mu^*})}\left[d'(Y)\right] < U(\theta_{\mu^*})\mathbb{E}_{m(\theta_{\mu^*})}\left[d(Y)\right], \tag{A.3}$$

contradicting (A.1). We conclude that

$$U(\theta)\mathbb{E}_{m(\theta)}\left[d'(Y)\right] = U(\theta)\mathbb{E}_{m(\theta)}\left[d(Y)\right]$$

for any $\theta \in m^{-1}(\mathcal{S})$.

Step 2: We next show that for any $\mu \in \mathcal{S}$, $\mathbb{E}_{\mu}\left[d'(Y) - d(Y)\right] = 0$. Because of nontrivial partial identification, for any $\mu \in \mathcal{S}$ there exists θ_{μ} such that $U(\theta_{\mu}) \neq 0$. Step 1 then implies that for any $\mu \in \mathcal{S}$,

$$U(\theta_{\mu})\mathbb{E}_{\mu}\left[d'(Y)\right] = U(\theta_{\mu})\mathbb{E}_{\mu}\left[d(Y)\right].$$

Since $U(\theta_\mu) \neq 0$, the desired result follows.

STEP 3: Consider now the family of distributions

$$Y \sim N(\mu, \Sigma), \quad \mu \in \mathcal{S}.$$

We will show that this collection of distributions is *complete* (Casella and Berger, 2002, Definition 6.2.21). Define the vector $\tilde{\mu} := \Sigma^{-1}\mu$ ranging over the set

$$\tilde{\mathcal{S}} := \{\tilde{\mu} \in \mathbb{R}^n \mid \tilde{\mu} = \Sigma^{-1}\mu, \quad \mu \in \mathcal{S}\}.$$

Note $\tilde{\mathcal{S}}$ is open under Definition 2. Furthermore, the pdf $p_{\tilde{\mu}}$ of Y given $\tilde{\mu}$ is of the form

$$p_{\tilde{\mu}}(Y) = h(Y)C(\tilde{\mu}) \exp\left[\tilde{\mu}^\top Y\right].$$

Thus, the family of distributions for Y is complete by Casella and Berger (2002, Theorem 6.2.25).

STEP 4: We use this completeness result to show that $d'(y) - d(y) = 0$ for almost all $y \in \mathbb{R}^n$. First, the definition of completeness readily implies this for all $y \in \mathbb{R}^n$ except for a set of realizations that have zero probability under any $N(\mu, \Sigma)$ as μ ranges over $\mathcal{S}$. By Skorohod (2012), the Gaussian measure in $\mathbb{R}^n$ is absolute continuous with respect to Lebesgue measure. This step's claim follows.

CONCLUSION: We find that $\mathbb{E}_{m(\theta)}[d'(Y)] = \mathbb{E}_{m(\theta)}[d(Y)]$ for all $\theta \in \Theta$, a contradiction.

A.2 Proof of Theorem 2

First, we can bound

$$\begin{aligned}\sup_{d\in\mathcal{D}_n} \inf_{\theta\in\Theta} \mathbb{E}_{m(\theta)}[W(d(Y),\theta)] &= \sup_{d\in\mathcal{D}_n} \inf_{\theta\in\Theta} \left[W(0,\theta) + U(\theta)\mathbb{E}_{m(\theta)}[d(Y)]\right] \\ &\leq \sup_{d\in\mathcal{D}_n} \inf_{\theta\in\Theta:U(\theta)\leq 0} \left[W(0,\theta) + U(\theta)\mathbb{E}_{m(\theta)}[d(Y)]\right] \leq \inf_{\theta\in\Theta:U(\theta)\leq 0} W(0,\theta),\end{aligned}$$

using that $\mathbb{E}_{m(\theta)}[d(Y)] \geq 0$. To see that this bound is tight, write

$$\inf_{\theta\in\Theta} \mathbb{E}_{m(\theta)}[W(d_{\text{no-data}},\theta)] = \inf_{\theta\in\Theta} \left[W(0,\theta) + U(\theta)\mathbb{E}_{m(\theta)}[d_{\text{no-data}}(Y)]\right] = \inf_{\theta\in\Theta} W(0,\theta).$$

and recall that $\inf_{\theta\in\Theta} W(0,\theta) = \inf_{\theta\in\Theta:U(\theta)\leq 0} W(0,\theta)$ by assumption.

A.3 Proof of Theorem 3

A.3.1 Proof of Part (i) of Theorem 3

Let $\mathbf{R}$ denote the minimax value of the policy maker's decision problem:

$$\mathbf{R} := \inf_{d \in \mathcal{D}_n} \sup_{\theta \in \Theta} \left\{ U(\theta) \left(\mathbf{1}\{U(\theta) \geq 0\} - \mathbb{E}_{m(\theta)}[d(Y)] \right) \right\}. \tag{A.4}$$

Step 1 (Minimax Regret Value): We first show that, if Assumption 1 holds and there exists an MMR optimal rule that depends on the data only through $(w^*)^\top Y$ and satisfies Equations (3.8) and (3.9), then

$$\mathbf{R} = (1/2) \cdot \bar{k}_{w^*}(0), \tag{A.5}$$

where

$$\bar{k}_{w^*}(0) := \sup_{\theta \in \Theta} U(\theta) \quad \text{s.t} \quad (w^*)^\top m(\theta) = 0. \tag{A.6}$$

Lemma B.1 in Appendix B shows that, under this step's premise, Equation (3.8) implies

$$(1/2) \cdot \bar{k}_{w^*}(0) \leq \mathbf{R}.$$

Since d^* is MMR optimal, Equations (3.8)-(3.9) and centrosymmetry of Θ imply

$$\mathbf{R} = \sup_{\theta \in \Theta, m(\theta) = \mathbf{0}} R(d^*, \theta) = \sup_{\theta \in \Theta, m(\theta) = \mathbf{0}} U(\theta) \left(\mathbf{1}\{U(\theta) \geq 0\} - \frac{1}{2} \right) = \frac{\overline{I}(\mathbf{0})}{2}.$$

By definition, $\overline{I}(\mathbf{0}) \leq \bar{k}_{w^*}(0)$. Thus, we have

$$(1/2) \cdot \bar{k}_{w^*}(0) \leq \mathbf{R} = (1/2) \cdot \overline{I}(\mathbf{0}) \leq (1/2) \cdot \bar{k}_{w^*}(0).$$

Step 2 (Upper bound for the worst-case regret of decision rules that depend on the data only through $(w^*)^\top Y$): We obtain an upper bound for the worst-case regret of such rules by linearizing the parameter space. We introduce some notation to formalize this step.

Let $\Gamma_{w^*} := \left\{ \gamma \in \mathbb{R} | (w^*)^\top m(\theta) = \gamma, \theta \in \Theta \right\}$ be the image of the transformation $\theta \mapsto (w^*)^\top m(\theta)$. We define the identified set for the welfare contrast $U(\theta)$ given $\gamma \in \Gamma_{w^*}$ as

$$ISU_{w^*}(\gamma) := \{u \in \mathbb{R} \mid U(\theta) = u, (w^*)^\top m(\theta) = \gamma, \theta \in \Theta\}. \tag{A.7}$$

Any decision rule that depends on the data only through $(w^*)^\top Y$ can be identified with a (measurable) function d from $\mathbb{R}$ to $[0,1]$. For future reference, let $\mathcal{D}$ collect all such functions. The worst-case expected regret of any $d \in \mathcal{D}$ can be expressed as

$$\sup_{\theta\in\Theta} \left(U(\theta)\left(\mathbf{1}\{U(\theta)\geq 0\} - \mathbb{E}_\gamma[d((w^*)^\top Y)]\right)\right) \tag{A.8}$$

$$= \sup_{\gamma\in\Gamma_{w^*}} \left(\sup_{\theta\in\Theta,\, (w^*)^\top m(\theta)=\gamma} U(\theta)\left(\mathbf{1}\{U(\theta)\geq 0\} - \mathbb{E}_\gamma[d((w^*)^\top Y)]\right)\right)$$

$$= \sup_{\gamma\in\Gamma_{w^*}} \left(\sup_{U^*\in ISU_{w^*}(\gamma)} U^*\left(\mathbf{1}\{U^*\geq 0\} - \mathbb{E}_\gamma[d((w^*)^\top Y)]\right)\right), \tag{A.9}$$

where the expectation $\mathbb{E}_\gamma[\cdot]$ is taken over $(w^*)^\top Y \sim N(\gamma, (w^*)^\top \Sigma w^*)$. For $\gamma \in \Gamma_{w^*}$, define

$$\underline{k}_{w^*}(\gamma) := \inf ISU_{w^*}(\gamma) = \inf\{U(\theta) : (w^*)^\top m(\theta) = \gamma, \theta \in \Theta\} \tag{A.10}$$

$$\bar{k}_{w^*}(\gamma) := \sup ISU_{w^*}(\gamma) = \sup\{U(\theta) : (w^*)^\top m(\theta) = \gamma, \theta \in \Theta\}. \tag{A.11}$$

By centrosymmetry of Θ and linearity of $U(\theta)$ and $m(\theta)$, we have that

$$\inf\{U(\theta) : (w^*)^\top m(\theta) = \gamma, \theta \in \Theta\} = -\sup\{U(\theta) : (w^*)^\top m(\theta) = -\gamma, \theta \in \Theta\}$$
$$\implies \underline{k}_{w^*}(\gamma) = -\bar{k}_{w^*}(-\gamma)$$

and therefore that $ISU_{w^*}(\gamma) \subseteq \left[-\bar{k}_{w^*}(-\gamma),\ \bar{k}_{w^*}(\gamma)\right], \forall \gamma \in \Gamma_{w^*}$. Equation (A.9) then implies that (A.8) is bounded above by

$$\sup_{\gamma\in\Gamma_{w^*}} \left(\sup_{U^*\in\left[-\bar{k}_{w^*}(-\gamma),\, \bar{k}_{w^*}(\gamma)\right]} U^*\left(\mathbf{1}\{U^*\geq 0\} - \mathbb{E}_\gamma[d((w^*)^\top Y)]\right)\right). \tag{A.12}$$

Lemma B.2 shows that, under Assumption 1, $\bar{k}_{w^*}(\gamma)$ is concave (and therefore $-\bar{k}_{w^*}(-\gamma)$ is convex). Furthermore, Lemma B.3 shows that, under Assumption 1, the superdifferential of the function $\bar{k}_{w^*}(\cdot)$ at $\gamma = 0$,

$$\partial\bar{k}_{w^*}(0) := \{s \in \mathbb{R} \mid \bar{k}_{w^*}(\gamma) \leq \bar{k}_{w^*}(0) + s\gamma, \quad \forall \gamma \in \Gamma_{w^*}\}, \tag{A.13}$$

(see p. 214-215 in Rockafellar (1997)) is nonempty, bounded, and closed.

Let $s_{w^*}(0)$ be the largest element of $\partial\bar{k}_{w^*}(0)$ and suppose without loss of generality that $s_{w^*}(0) \geq$

0. Step 3 in the proof of Lemma B.3 established that Γ_{w^*} is symmetric around 0. The definition of superdifferential then gives

$$s_{w^*}(0)\gamma - \bar{k}_{w^*}(0) \leq -\bar{k}_{w^*}(-\gamma) \quad \forall \gamma \in \Gamma_{w^*}. \tag{A.14}$$

It follows that

$$[-\bar{k}_{w^*}(-\gamma)\ ,\ \bar{k}_{w^*}(\gamma)] \subseteq [s_{w^*}(0)\gamma - \bar{k}_{w^*}(0)\ ,\ s_{w^*}(0)\gamma + \bar{k}_{w^*}(0)].$$

Substituting into (A.12) then implies that (A.8) is further bounded above by

$$\sup_{\gamma \in \Gamma_{w^*}} \left(\sup_{U^* \in \left[s_{w^*}(0)\gamma - \bar{k}_{w^*}(0)\ ,\ s_{w^*}(0)\gamma + \bar{k}_{w^*}(0)\right]} U^* \left(\mathbf{1}\{U^* \geq 0\} - \mathbb{E}_\gamma[d((w^*)^\top Y)] \right) \right). \tag{A.15}$$

We note that the choice set for U^* in this expression is linear in γ.

Step 3 ("Linear Embedding" Minimax Regret Problem): The previous step and the fact that $\Gamma_{w^*} \subseteq \mathbb{R}$ imply that (A.8) is bounded above by

$$\inf_{d \in \mathcal{D}} \sup_{\gamma \in \mathbb{R}} \left(\sup_{U^* \in \left[s_{w^*}(0)\gamma - \bar{k}_{w^*}(0)\ ,\ s_{w^*}(0)\gamma + \bar{k}_{w^*}(0)\right]} U^* \left(\mathbf{1}\{U^* \geq 0\} - \mathbb{E}_\gamma[d\left(\widehat{\gamma}\right)] \right) \right), \tag{A.16}$$

where

$$\widehat{\gamma} \sim N\left(\gamma\ ,\ (w^*)^\top \Sigma w^*\right), \quad \gamma \in \mathbb{R}.$$

Lemma B.4 in Appendix B shows that if $s_{w^*}(0) > 0$ and $\bar{k}_{w^*}(0) > \sqrt{\frac{\pi}{2}}\sqrt{(w^*)^\top \Sigma w^*} \cdot s_{w^*}(0)$, then (A.16) equals $\bar{k}_{w^*}(0)/2$ and there are infinitely many rules that give such value. In particular, a solution is given by any convex combination of the following rules:

$$d^*_{\text{RT}}(\widehat{\gamma}) \ := \ \Phi\left(\widehat{\gamma} / \sqrt{\frac{2 \cdot \bar{k}_{w^*}(0)^2}{\pi \cdot s_{w^*}(0)^2} - (w^*)^\top \Sigma w^*} \right) \tag{A.17}$$

$$d^*_{\text{linear}}(\widehat{\gamma}) \ := \ \begin{cases} 0, & \widehat{\gamma} < -\rho^*, \\ \frac{\widehat{\gamma} + \rho^*}{2\rho^*}, & -\rho^* \leq \widehat{\gamma} \leq \rho^*, \\ 1, & \widehat{\gamma} > \rho^*, \end{cases} \tag{A.18}$$

where $\rho^* \in \left(0, \frac{\bar{k}_{w^*}(0)}{s_{w^*}(0)}\right)$ is the unique solution to

$$\left(\frac{s_{w^*}(0)}{2 \cdot \bar{k}_{w^*}(0)}\right)\rho^* - \frac{1}{2} + \Phi\left(-\frac{\rho^*}{\sqrt{(w^*)^\top \Sigma w^*}}\right) = 0. \tag{A.19}$$

Step 4 (Rules that solve the "Linear Embedding" minimax regret problem also solve the original problem). If $s_{w^*}(0) > 0$, $\bar{k}_{w^*}(0) > \sqrt{\frac{\pi}{2}}\sqrt{(w^*)^\top \Sigma w^*} \cdot s_{w^*}(0)$, and $d^\bigstar \in D$ solves the linear embedding problem, then

$$d^\bigstar \circ w^*(Y) := d^\bigstar((w^*)^\top Y) \in \mathcal{D}_n \tag{A.20}$$

is MMR optimal in the original decision problem (A.4).

This is because Step 1 implies that

$$(1/2)\bar{k}_{w^*}(0) = \mathbf{R} \leq \sup_{\theta \in \Theta} \left\{U(\theta)\left(\mathbf{1}\{U(\theta) \geq 0\} - \mathbb{E}_{m(\theta)}[d^\bigstar((w^*)^\top Y)]\right)\right\}$$

and Step 2 implies that

$$\begin{aligned} & \sup_{\theta \in \Theta} \left\{U(\theta)\left(\mathbf{1}\{U(\theta) \geq 0\} - \mathbb{E}_{m(\theta)}[d^\bigstar((w^*)^\top Y)]\right)\right\} \\ \leq & \sup_{\gamma \in \mathbb{R}} \left(\sup_{U^* \in \left[s_{w^*}(0)\gamma - \bar{k}_{w^*}(0)\,,\, s_{w^*}(0)\gamma + \bar{k}_{w^*}(0)\right]} U^*\left(\mathbf{1}\{U^* \geq 0\} - \mathbb{E}_\gamma[d^\bigstar(\widehat{\gamma})]\right)\right) \\ = & \frac{\bar{k}_{w^*}(0)}{2}, \end{aligned}$$

where the last equality follows from Step 3. Consequently,

$$\sup_{\theta \in \Theta} \left\{U(\theta)\left(\mathbf{1}\{U(\theta) \geq 0\} - \mathbb{E}_{m(\theta)}[d^\bigstar((w^*)^\top Y)]\right)\right\} = \frac{\bar{k}_{w^*}(0)}{2}.$$

Step 5: Finally, we show that the assumptions of Theorem 3 imply

$$s_{w^*}(0) > 0$$

and

$$\bar{k}_{w^*}(0) > \sqrt{\pi/2} \cdot \sqrt{(w^*)^\top \Sigma w^*} \cdot s_{w^*}(0).$$

First, we show that $s_{w^*}(0) > 0$. The definitions of $\overline{I}(\cdot)$ and $\bar{k}_{w^*}(\cdot)$ imply

$$\overline{I}(\mu) \leq \bar{k}_{w^*}((w^*)^\top \mu), \text{ for all } \mu \in M.$$

As $s_{w^*}(0)$ is a supergradient of $\bar{k}_{w^*}(0)$,

$$\bar{k}_{w^*}((w^*)^\top \mu) \leq \bar{k}_{w^*}(0) + s_{w^*}(0)((w^*)^\top \mu), \text{ for all } \mu \in M.$$

Step 1 showed that $\overline{I}(\mathbf{0}) = \bar{k}_{w^*}(0)$. Hence, combining the above equations yields

$$\overline{I}(\mu) \leq \overline{I}(\mathbf{0}) + s_{w^*}(0)(w^*)^\top \mu, \text{ for all } \mu \in M. \tag{A.21}$$

If $s_{w^*}(0) = 0$, Equation (A.21) then implies $\overline{I}(\mu) \leq \overline{I}(\mathbf{0})$ for all $\mu \in M$, contradicting the assumption that there exists $\mu \in M$ such that $\overline{I}(\mu) > \overline{I}(\mathbf{0})$.

Second, since Step 1 showed that $\overline{I}(\mathbf{0}) = \bar{k}_{w^*}(0)$, then

$$\bar{k}_{w^*}(0) > \sqrt{\pi/2} \cdot \sqrt{(w^*)^\top \Sigma w^*} \cdot s_{w^*}(0)$$

holds when $\overline{I}(\mathbf{0})$ is large enough, in particular whenever

$$\overline{I}(\mathbf{0}) > \sqrt{\pi/2} \cdot \sqrt{(w^*)^\top \Sigma w^*} \cdot s_{w^*}(0).$$

Conclusion: Steps 1-5 imply there are infinitely many rules that solve the problem (A.4).

A.3.2 Proof of Part (ii) of Theorem 3

Consider any decision rule $d_{\mathrm{m}}(\cdot)$ that depends on the data only as nondecreasing function of $w^\top Y$ (for some $w \neq \mathbf{0}$ that is not necessarily w^*) and such that $d_{\mathrm{m}}(\cdot) \in \{0, 1\}$ for all data realizations. Then we must have $d_{\mathrm{m}}(w^\top Y) = \mathbf{1}\left\{w^\top Y \geq c\right\}$ for some $-\infty \leq c \leq \infty$. The worst-case expected

regret of such a rule satifies

$$\begin{aligned} & \sup_{\theta\in\Theta} U(\theta)\left(\mathbf{1}\{U(\theta)\geq 0\}-\mathbb{E}_{m(\theta)}[d_m(Y)]\right) && \text{(A.22)}\\ \geq\ & \sup_{\theta\in\Theta:m(\theta)=\mathbf{0}} U(\theta)\left(\mathbf{1}\{U(\theta)\geq 0\}-\mathbb{E}_{m(\theta)}[d_m(Y)]\right)\\ =\ & \max\left\{-\underline{I}(\mathbf{0})\mathbb{E}_0[d_m(Y)],\overline{I}(\mathbf{0})(1-\mathbb{E}_0[d_m(Y)])\right\}\\ \geq\ & \overline{I}(\mathbf{0})/2, && \text{(A.23)}\end{aligned}$$

where we used that $\underline{I}(\mathbf{0})=-\overline{I}(\mathbf{0})$ by centrosymmetry, and the last inequality is strict unless $c=0$. As $\overline{I}(\mathbf{0})/2$ is the MMR value of the problem, $d^*_{\mathrm{m}}(\cdot)$ cannot be MMR optimal if $c\neq 0$. For $w=w^*$, substantial additional algebra that we relegate to Lemma B.8 extends the result to $c=0$ (by showing that a first-order condition cannot hold at $(w^*)^\top Y=0$).

A.3.3 Proof of Part (iii) of Theorem 3

The preceding argument established the claim for rules of form $\mathbf{1}\{w^\top Y\geq c\}$, where $c\neq 0$. It remains to consider symmetric threshold rules $\mathbf{1}\{w^\top Y\geq 0\}$. To this end, bound the worst-case expected regret of such rules as follows:

$$\begin{aligned} & \sup_{\theta\in\Theta} U(\theta)\left(\mathbf{1}\{U(\theta)\geq 0\}-\mathbb{E}_{m(\theta)}[\mathbf{1}\{w^\top Y\geq 0\}]\right)\\ \geq\ & \sup_{\theta\in\Theta,m(\theta)=\mu,U(\theta)\geq 0,\mu\in M} U(\theta)\left(\mathbf{1}\{U(\theta)\geq 0\}-\mathbb{E}_{\mu}[\mathbf{1}\{w^\top Y\geq 0\}]\right)\\ =\ & \sup_{\mu\in M,\overline{I}(\mu)>0} \overline{I}(\mu)\left(1-\mathbb{E}_{\mu}[\mathbf{1}\{w^\top Y\geq 0\}]\right)\\ =\ & \sup_{\mu\in M,\overline{I}(\mu)>0} \overline{I}(\mu)\Phi\left(-\frac{w^\top\mu}{\sqrt{w^\top\Sigma w}}\right)\\ :=\ & \sup_{\mu\in M,\overline{I}(\mu)>0} \mathrm{g}_w(\mu).\end{aligned}$$

Note that $\mathrm{g}_w(\mathbf{0})=\overline{I}(\mathbf{0})/2$ is the MMR value of the problem, implying that $\mu=\mathbf{0}$ attains this value under any symmetric threshold rule. For such a rule to be MMR optimal, $\mu=\mathbf{0}$ must then be a local constrained maximum point of $\mathrm{g}_w(\mu)$. Because M contains an open set including $\mathbf{0}$ by Definition 2 and Assumption 1 and since we assumed differentiability of $\overline{I}(\mu)$ at $\mathbf{0}$, this requires a

first-order condition

$$\frac{\partial \mathrm{g}_w(\mu)}{\partial \mu_j} \mid_{\mu=\mathbf{0}} = \frac{1}{2}\frac{\partial \overline{I}(\mathbf{0})}{\partial \mu_j} - \frac{w_j}{\sqrt{w^\top \Sigma w}}\overline{I}(\mathbf{0})\phi(0) = 0, \; j = 1 \ldots n. \tag{A.24}$$

To simplify expressions, change co-ordinates (if necessary) so that $w^* = (1, 0, \ldots, 0)^\top$. Because w^* must fulfil (A.24), we have that $\frac{\partial \overline{I}(\mathbf{0})}{\partial \mu_j} = 0$ for $j = 2, \ldots, n$. But this, in turn, means that (A.24) requires $w_2 = \ldots = w_n = 0$. Next, noting that w in a symmetric threshold rule is determined only up to scale, we restrict attention to $w_1 \in \{-1, 0, 1\}$. But if $w^* = (1, 0, \ldots, 0)$ solves (A.24) for $j = 1$, then $-w^*$ cannot solve it because the sign change does not affect the denominator dividing w_1. Finally, $w = \mathbf{0}$, i.e. never adopting treatment, is excluded by part (ii) (it is the same as setting $c = \infty$ there) and is also easily seen directly to not be MMR optimal.

A.4 Proof of Theorem 4

STEP 1: If $F \circ w^* \in \tilde{\mathcal{D}}_n$ is MMR optimal, then $V(d^*_{\text{linear}} \circ w^*) \subseteq V(F \circ w^*)$.

To see this, pick any $F \circ w^* \in \tilde{\mathcal{D}}_n$ that is MMR optimal. Then we can write

$$F \circ w^*(Y) = F((w^*)^\top Y) = F(\widehat{\gamma}),$$

where $F \in \mathcal{F}$ is a symmetric and unimodal c.d.f (thus weakly increasing as well), $\widehat{\gamma} := (w^*)^\top Y \sim N(\gamma, \sigma^2)$, with $\gamma \in \Gamma_{w^*}$ defined in Step 2 of the proof for Theorem 3(i), and $\sigma^2 = (w^*)^\top \Sigma w^*$. The worst-case expected regret of rule $F \circ w^*$ is

$$\sup_{\theta \in \Theta} U(\theta)\left(\mathbf{1}\{U(\theta) \geq 0\} - \mathbb{E}_{m(\theta)}[F((w^*)^\top Y)]\right) = \sup_{\gamma \in \Gamma_{w^*} : \bar{k}_{w^*}(\gamma) > 0} \bar{k}_{w^*}(\gamma)\left(1 - \mathbb{E}_\gamma[F(\widehat{\gamma})]\right),$$

where $\bar{k}^*_w$ is defined in (A.11). Letting $g_F(\gamma) := \bar{k}_{w^*}(\gamma)\left(1 - \mathbb{E}_\gamma[F(\widehat{\gamma})]\right)$. Since $\widehat{\gamma} \sim N(\gamma, \sigma^2)$, we may further calculate (using integration by parts)

$$\begin{aligned}
\mathbb{E}_\gamma[F(\widehat{\gamma})] &= \int F(s) d\Phi(\frac{s-\gamma}{\sigma}) \\
&= \Phi\left(\frac{s-\gamma}{\sigma}\right) F(s)|_{-\infty}^{\infty} - \int \Phi\left(\frac{s-\gamma}{\sigma}\right) dF(s) \\
&= 1 - \int \Phi\left(\frac{s-\gamma}{\sigma}\right) dF(s).
\end{aligned}$$

Therefore, $g_F(\gamma) = \bar{k}_{w^*}(\gamma)\int \Phi(\frac{s-\gamma}{\sigma})dF(s)$. Note that $F(-x) = 1 - F(x)$ for all $x \in \mathbb{R}$; hence, $\int \Phi(\frac{s}{\sigma})dF(s) = \frac{1}{2}$ and therefore $g_F(0) = \frac{\bar{k}_{w^*}(0)}{2}$. By Step 1 for the proof of Theorem 3(i), $\frac{\bar{k}_{w^*}(0)}{2}$ is the MMR value of the problem. MMR optimality of $F \circ w^*$ implies

$$0 \in \arg \sup_{\gamma \in \Gamma_{w^*}, \bar{k}_{w^*}(\gamma) > 0} g_F(\gamma).$$

By Lemma B.3, 0 is an interior point of $\{\gamma \in \mathbb{R} : \bar{k}_{w^*}(\gamma) > 0, \gamma \in \Gamma_{w^*}\}$. Thus, let $\partial g_F(0)$ denote the generalized gradient of $g_F(\cdot)$ at 0.[20] Then $0 \in \partial g_F(0)$ is necessary for optimality. To show that it fails, compute the generalized gradient as[21]

$$\begin{aligned}&\tilde{s}_{w^*}(0)\int \Phi(\frac{s}{\sigma})dF(s) - \frac{\bar{k}_{w^*}(0)}{\sigma}\int \phi\left(\frac{s}{\sigma}\right)dF(s)\\ =&\frac{\tilde{s}_{w^*}(0)}{2} - \frac{\bar{k}_{w^*}(0)}{\sigma}\int \phi\left(\frac{s}{\sigma}\right)dF(s),\end{aligned}$$

where $\tilde{s}_{w^*}(0)$ is a supergradient of $\bar{k}_{w^*}(\gamma)$ at $\gamma = 0$. Therefore, we conclude

$$\frac{\tilde{s}_{w^*}(0)}{2} - \frac{\bar{k}_{w^*}(0)}{\sigma}\int \phi\left(\frac{s}{\sigma}\right)dF(s) = 0 \iff \int \phi\left(\frac{s}{\sigma}\right)dF(s) = \frac{\tilde{s}_{w^*}(0)\sigma}{2\bar{k}_{w^*}(0)}$$

for some $\tilde{s}_{w^*}(0) > 0$.

Next, $d^*_{\text{linear}} \in \mathcal{F}$ can be verified to solve the linear embedding problem (A.16) (Lemma B.4). In particular, evaluating $g^{(1)}_{\text{linear}}(\gamma)$ at $\gamma = 0$, where $g_{\text{linear}}(\gamma)$ is defined in Lemma B.7 with $k = \frac{\bar{k}_{w^*}(0)}{s_{w^*}(0)}$ and $\sigma^2 = (w^*)^\top \Sigma w^*$, one finds

$$\int \phi\left(\frac{s}{\sigma}\right)d(d^*_{\text{linear}}(s)) = \int_{-\rho^*}^{\rho^*} \phi\left(\frac{s}{\sigma}\right)\frac{1}{2\rho^*}ds = \frac{s_{w^*}(0)\sigma}{2\bar{k}_{w^*}(0)}, \tag{A.25}$$

where $s_{w^*}(0) > 0$ is the largest supergradient of $\bar{k}_{w^*}(\gamma)$ at $\gamma = 0$.

Since f is symmetric around 0, $V(F \circ w^*)$ is as well. Write $V(F \circ w^*) := (-a_F, a_F)$ for some a_F and $V(d^*_{\text{linear}} \circ w^*) := (-\rho^*, \rho^*)$. Suppose by contradiction that $V(d^*_{\text{linear}} \circ w^*) \not\subseteq V(F \circ w^*)$. Then

[20] The generalized gradient of $g : \mathbb{R} \to \mathbb{R}$ at x equals $\partial g(x) := \{\xi \in \mathbb{R} : \limsup_{y \to x, t \downarrow 0} \frac{f(y+tv)-f(y)}{t} \geq \xi v, \forall v \in \mathbb{R}\}$. See Clarke (1990, p. 27).

[21] We can verify, following the same steps in the proof for Lemma B.8, that conditions of Proposition 2.3.13 in Clarke (1990) are satisfied, so that the chain rule can be applied.

$a_F < \rho^*$, but that would imply

$$\int \phi(\frac{s}{\sigma})dF(s) = \int_{-a_F}^{a_F} \phi(\frac{s}{\sigma})dF(s) > \frac{s_{w^*}(0)\sigma}{2\bar{k}_{w^*}(0)} \geq \frac{\tilde{s}_{w^*}(0)\sigma}{2\bar{k}_{w^*}(0)}, \tag{A.26}$$

where the first inequality follows by the assumption that $dF(x) = 0$ for all $x \notin (-a_F, a_F)$ and $F(x)$ is symmetric and unimodal, and the second inequality follows as $s_{w^*}(0)$ is the largest supergradient of $\bar{k}_{w^*}(0)$. Thus, $0 \notin \partial g_F(0)$, a contradiction.

STEP 2: Next, $V(d^*_{\text{linear}} \circ w^*) = V(F \circ w^*)$ if and only if $F = d^*_{\text{linear}}$. The "if" direction is obvious. To see "only if," suppose by contradiction that there exists some $\tilde{F} \in \mathcal{F}$ not equal to d^*_{linear} but such that $V(d^*_{\text{linear}} \circ w^*) = V(F \circ w^*)$ and $\tilde{F} \circ w^*$ is MMR optimal. Then

$$\int \phi\left(\frac{s}{\sigma}\right) d\tilde{F}(s) = \int_{-\rho^*}^{\rho^*} \phi\left(\frac{s}{\sigma}\right) d\tilde{F}(s) > \int_{-\rho^*}^{\rho^*} \phi\left(\frac{s}{\sigma}\right) \frac{1}{2\rho^*} ds = \frac{s_{w^*}(0)\sigma}{2\bar{k}_{w^*}(0)} \geq \frac{\tilde{s}_{w^*}(0)\sigma}{2\bar{k}_{w^*}(0)},$$

where the first step uses $V(d^*_{\text{linear}} \circ w^*) = V(F \circ w^*)$, the second one that $\tilde{F}$ is symmetric and unimodal, the third one uses (A.25), and the last one that $s_{w^*}(0)$ is the largest supergradient of $\bar{k}_{w^*}(0)$. Thus, $\tilde{F} \circ w^*$ cannot be MMR optimal, a contradiction.

B Online Appendix

B.1 Lemmas for Theorem 3

Lemma B.1. Consider a treatment choice problem with payoff function (2.1) and statistical model (2.3) that exhibits nontrivial partial identification in the sense of Definition 2. Suppose that Assumption 1 holds and there exists an MMR optimal rule d^* depending on the data only through $(w^*)^\top Y$. If Equation (3.8) holds, then

$$\mathbf{R} := \sup_{\theta\in\Theta} U(\theta)\left(\mathbf{1}\{U(\theta)\geq 0\} - \mathbb{E}_{m(\theta)}[d^*((w^*)^\top Y)]\right) \geq \frac{\bar{k}_{w^*}(0)}{2},$$

where

$$\bar{k}_{w^*}(0) := \sup_{\theta\in\Theta} U(\theta) \quad \text{s.t} \quad (w^*)^\top m(\theta) = 0.$$

Proof. Since the distribution of ${w^*}^\top Y$ only depends on $m(\theta)$ through $(w^*)^\top m(\theta)$, we can write

$$\begin{aligned}
\mathbf{R} &= \sup_{\theta\in\Theta} U(\theta)\left(\mathbf{1}\{U(\theta)\geq 0\} - \mathbb{E}_{(w^*)^\top m(\theta)}[d^*((w^*)^\top Y)]\right) \\
&\geq \sup_{\theta\in\Theta:(w^*)^\top m(\theta)=0} U(\theta)\left(\mathbf{1}\{U(\theta)\geq 0\} - \mathbb{E}_0[d^*((w^*)^\top Y)]\right) \\
&\overset{(1)}{=} \sup_{\theta\in\Theta,(w^*)^\top m(\theta)=0} U(\theta)\left(\mathbf{1}\{U(\theta)\geq 0\} - \frac{1}{2}\right) \\
&= \frac{1}{2}\max\left\{\sup_{\theta\in\Theta:(w^*)^\top m(\theta)=0,U(\theta)\geq 0} U(\theta)\,,\ \sup_{\theta\in\Theta:(w^*)^\top m(\theta)=0,U(\theta)\leq 0} -U(\theta)\right\} \\
&\overset{(2)}{=} \frac{1}{2}\max\left\{\sup_{\theta\in\Theta:(w^*)^\top m(\theta)=0,U(\theta)\geq 0} U(\theta)\,,\ \sup_{\theta\in\Theta:(w^*)^\top m(\theta)=0,U(-\theta)\geq 0} U(-\theta)\right\} \\
&\overset{(3)}{=} \frac{1}{2}\sup_{\theta\in\Theta:(w^*)^\top m(\theta)=0,} U(\theta),
\end{aligned}$$

where (1) uses that $\mathbb{E}_0[d^*((w^*)^\top Y)] = 1/2$ by (3.8), (2) uses linearity of $U(\cdot)$, and (3) uses centrosymmetry of Θ and linearity of $m(\cdot)$ and $U(\cdot)$. □

Lemma B.2. Suppose Assumption 1 holds. Then $\bar{k}_{w^*}(\gamma)$ is concave in $\gamma \in \Gamma_{w^*}$.

Proof. Fix any $\gamma_1, \gamma_2 \in \Gamma_{w^*}$ and $\varepsilon > 0$. By the definition of $\bar{k}_{w^*}(\gamma_1)$ and $\bar{k}_{w^*}(\gamma_2)$, there exist

$\theta_{\gamma_1,\varepsilon}, \theta_{\gamma_2,\varepsilon} \in \Theta$ such that

$$(w^*)^\top m(\theta_{\gamma_1,\varepsilon}) = \gamma_1,\ U(\theta_{\gamma_1,\varepsilon}) > \bar{k}_{w^*}(\gamma_1) - \varepsilon,$$
$$(w^*)^\top m(\theta_{\gamma_2,\varepsilon}) = \gamma_2,\ U(\theta_{\gamma_2,\varepsilon}) > \bar{k}_{w^*}(\gamma_2) - \varepsilon,$$

By convexity of Θ, $t\theta_{\gamma_1,\varepsilon} + (1-t)\theta_{\gamma_2,\varepsilon} \in \Theta$ as well. Moreover, as $m(\cdot)$ is linear,

$$\begin{aligned}(w^*)^\top m\left(t\theta_{\gamma_1,\varepsilon} + (1-t)\theta_{\gamma_2,\varepsilon}\right) &= t(w^*)^\top m(\theta_{\gamma_1,\varepsilon}) + (1-t)(w^*)^\top m(\theta_{\gamma_2,\varepsilon}) \\ &= t\gamma_1 + (1-t)\gamma_2.\end{aligned}$$

Next, for $t \in [0,1]$, let $\tilde{\gamma}_t := t\gamma_1 + (1-t)\gamma_2$. Then

$$\begin{aligned}\bar{k}_{w^*}(\tilde{\gamma}_t) &= \sup_{(w^*)^\top m(\theta) = t\gamma_1 + (1-t)\gamma_2, \theta \in \Theta} U(\theta) \\ &\geq U(t\theta_{\gamma_1,\varepsilon} + (1-t)\theta_{\gamma_2,\varepsilon}) \\ &= tU(\theta_{\gamma_1,\varepsilon}) + (1-t)U(\theta_{\gamma_2,\varepsilon}) \\ &> t\left(\bar{k}_{w^*}(\gamma_1) - \varepsilon\right) + (1-t)\left(\bar{k}_{w^*}(\gamma_2) - \varepsilon\right) \\ &= t\bar{k}_{w^*}(\gamma_1) + (1-t)\bar{k}_{w^*}(\gamma_2) - \varepsilon,\end{aligned}$$

where the third relation follows from linearity of $U(\cdot)$. This means that

$$\bar{k}_{w^*}(\tilde{\gamma}_t) > t\bar{k}_{w^*}(\gamma_1) + (1-t)\bar{k}_{w^*}(\gamma_2) - \varepsilon,$$

for any $\varepsilon > 0$. We conclude that

$$\bar{k}_{w^*}(\tilde{\gamma}_t) \geq t\bar{k}_{w^*}(\gamma_1) + (1-t)\bar{k}_{w^*}(\gamma_2).$$

□

Lemma B.3. Consider a treatment choice problem with payoff function (2.1) and statistical model (2.3) that exhibits nontrivial partial identification in the sense of Definition 2. If Assumption 1 holds, then the superdifferential of $\bar{k}_{w^*}(\gamma)$ at $\gamma = 0$ is nonempty, bounded, and closed.

Proof. To see closure, let $s_n \to s^*$ be a converging sequence of elements in the superdifferential. By definition, for every $\gamma \in \Gamma_{w^*}$ we have

$$\bar{k}_{w^*}(\gamma) \leq \bar{k}_{w^*}(0) + s_n\gamma.$$

But then, for every $\gamma \in \Gamma_{w^*}$

$$\bar{k}_{w^*}(\gamma) \leq \bar{k}_{w^*}(0) + s^*\gamma.$$

Thus, $s^* \in \partial \bar{k}_{w^*}(0)$.

For nonemptiness and boundedness, by Rockafellar (1997, Theorem 23.4), it suffices to show that $0 \in \text{int}(\Gamma_{w^*})$.

Step 1: We first show that $0 \in \Gamma_{w^*}$. As Θ is centrosymmetric, convex and nonempty, $\overline{\mathbf{0}} \in \Theta$, where $\overline{\mathbf{0}}$ is the zero vector in Θ. As $m(\cdot)$ is linear, it follows $(w^*)^\top m(\overline{\mathbf{0}}) = (w^*)^\top \mathbf{0} = 0$, where $\mathbf{0} := 0_{n\times 1}$. That is, $0 \in \Gamma_{w^*}$.

Step 2: We show that there exists some $\gamma \neq 0$ such that $\gamma \in \Gamma_{w^*}$. Let $\mathcal{S}$ be an open set in $\mathbb{R}^n$ for which

$$\underline{I}(\mu) < 0 < \overline{I}(\mu), \quad \forall\, \mu \in \mathcal{S}.$$

Such a set exists because the problem exhibits nontrivial partial identification. Pick any $\theta \in \Theta$ such that $m(\theta) \in \mathcal{S}$. If $(w^*)^\top m(\theta) \neq 0$, then the initial claim of step 2 follows. If $(w^*)^\top m(\theta) = 0$, then note as $\mathcal{S}$ is open, we can pick some $\epsilon > 0$ small enough such that $m(\theta) + \epsilon w^* \in \mathcal{S}$. It follows then $(w^*)^\top (m(\theta) + \epsilon w^*) = \epsilon > 0$. Therefore, the claim of step 2 is verified.

Step 3: We show that Γ_{w^*} is symmetric around zero. Using the conclusion from Step 2, pick any $\gamma \in \Gamma_{w^*}$ and $\gamma \neq 0$. Then there exists $\theta \in \Theta$ such that $\gamma = (w^*)^\top m(\theta)$. Since Θ is centrosymmetric, $-\theta \in \Theta$ as well. As $(w^*)^\top m(-\theta) = -(w^*)^\top m(\theta) = -\gamma$ by linearity of $m(\cdot)$, one has $-\gamma \in \Gamma_{w^*}$. That is, M is symmetric around zero.

Steps 1-3 then imply that Γ_{w^*} is a symmetric interval around zero and that $0 \in \text{int}(\Gamma_{w^*})$. □

Lemma B.4. If $s_{w^*}(0) > 0$ and $\bar{k}_{w^*}(0) > \sqrt{\frac{\pi}{2}}\sqrt{(w^*)^\top \Sigma w^*} \cdot s_{w^*}(0)$, then the right-hand side of (A.16) equals $\bar{k}_{w^*}(0)/2$ and infinitely many rules attain this value. In particular, any convex combination of the rules in Equations (A.17)-(A.18) solves the problem in Expression (A.16).

Proof. As $s_{w^*}(0) > 0$, the value of the linear embedding problem in (A.16) equals $s_{w^*}(0)$ times

$$\inf_{d\in\mathcal{D}} \sup_{\gamma\in\mathbb{R}} \left(\sup_{\tilde{U}^* \in [-k+\gamma,\, k+\gamma]} \tilde{U}^* \left(\mathbf{1}\{\tilde{U}^* \geq 0\} - \mathbb{E}_\gamma[d(\widehat{\gamma})] \right) \right), \tag{B.1}$$

where

$$\widehat{\gamma} \sim N(\gamma, \sigma^2), \quad \sigma^2 = (w^*)^\top \Sigma w^*, \quad k := \frac{\bar{k}_{w^*}(0)}{s_{w^*}(0)}.$$

Stoye (2012a) shows that, if $k > \sqrt{\pi/2}\,\sigma$ (which holds if and only if $\bar{k}_{w^*}(0) > \sqrt{\frac{\pi}{2}}\sqrt{(w^*)^\top \Sigma w^*}$ ·

$s_{w^*}(0)$), then

$$d^*_{\text{Gaussian}} := \Phi\big(\hat{\gamma}/\sqrt{2k^2/\pi - \sigma^2}\big)$$

solves (B.1) and its worst-case regret is attained at $\gamma = 0$; that is, (B.1) equals

$$\sup_{\gamma\in\mathbb{R}}\left(\sup_{\tilde{U}^*\in[-k+\gamma,\,k+\gamma]} \tilde{U}^*\left(\mathbf{1}\{\tilde{U}^* \geq 0\} - \mathbb{E}_\gamma[d^*_{\text{Gaussian}}(\widehat{\gamma})]\right)\right),$$

and this expression equals

$$\sup_{\tilde{U}^*\in[-k,\,k]} \tilde{U}^*\left(\mathbf{1}\{\tilde{U}^* \geq 0\} - \mathbb{E}_\gamma[d^*_{\text{Gaussian}}(\widehat{\gamma})]\right).$$

Moreover, the MMR value (B.1) equals $k/2$. Since $d^*_{\text{RT}} = d^*_{\text{Gaussian}}$, this implies that d^*_{RT} solves the problem in Equation (A.16) and that this problem has value $\bar{k}_{w^*}(0)/2$.

Lemma B.5 establishes that d^*_{linear} and d^*_{mixture} equally attain MMR. Since the set of MMR optimal rules is closed under convex combination, this establishes the claim. □

Lemma B.5. If $k > \sqrt{\pi/2}\sigma$, then following rule solves the linear embedding minimax problem defined in (B.1):

$$d^*_{\text{linear}} := \begin{cases} 0, & \hat{\gamma} < -\rho^*, \\ \frac{\hat{\gamma}+\rho^*}{2\rho^*}, & -\rho^* \leq \hat{\gamma} \leq \rho^*, \\ 1, & \hat{\gamma} > \rho^*, \end{cases} \tag{B.2}$$

where $\rho^* \in (0, k)$ is the unique solution of

$$\left(\frac{\rho^*}{2\cdot k}\right) - \frac{1}{2} + \Phi\left(-\frac{\rho^*}{\sigma}\right) = 0. \tag{B.3}$$

Proof. Lemma B.6(i) shows that $\rho^* \in (0, k)$ exists and is unique. Recall again from Stoye (2012a) that, if $k > \sqrt{\pi/2}\,\sigma$, the MMR value of the problem is $k/2$, where

$$I(\gamma) := [-k+\gamma,\ k+\gamma], \quad R(d, \gamma, \gamma^*) := \gamma^*\left(\mathbf{1}\{\gamma^* \geq 0\} - \mathbb{E}_\gamma[d]\right). \tag{B.4}$$

Furthermore, by definition of the minimax problem,

$$\inf_{d\in\mathcal{D}} \sup_{\gamma^*\in I(\gamma),\gamma\in\mathbb{R}} R(d, \gamma, \gamma^*) \leq \sup_{\gamma^*\in I(\gamma),\gamma\in\mathbb{R}} R(d^*_{\text{linear}}, \gamma, \gamma^*).$$

And Lemma B.7 shows that, when $k > \sqrt{\pi/2}\sigma$,

$$\sup_{\gamma^* \in I(\gamma), \gamma \in \mathbb{R}} R(d^*_{\text{linear}}, \gamma, \gamma^*) = k/2$$

and

$$\sup_{\gamma^* \in I(0)} R(d^*_{\text{linear}}, 0, \gamma^*) = k/2$$

Hence, we conclude that

$$\inf_{d \in \mathcal{D}} \sup_{\gamma^* \in I(\gamma), \gamma \in \mathbb{R}} R(d, \gamma, \gamma^*) = \sup_{\gamma^* \in I(\gamma), \gamma \in \mathbb{R}} R(d^*_{\text{linear}}, \gamma, \gamma^*) = \frac{k}{2},$$

and thus d^*_{linear} is MMR optimal, and its worst-case regret is achieved at $\gamma = 0$. □

Lemma B.6. Consider ρ^* defined in Lemma B.5.

(i) $\rho^* \in (0, k)$ exists and is uniquely defined when $\frac{k}{\sigma} > \sqrt{\frac{\pi}{2}}$.

(ii) The value of ρ^* is strictly decreasing in σ. Moreover, $\rho^* \to k$ when $\sigma \to 0$.

(iii) The value of ρ^* is strictly increasing in k.

Proof. Note

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

is equivalent to

$$1 - \frac{k}{\rho^*}\left(1 - 2\Phi\left(-\frac{\rho^*}{\sigma}\right)\right) = 0. \tag{B.5}$$

Write $\mathbf{g}(\rho; k, \sigma) = 1 - \frac{k}{\rho}\left(1 - 2\Phi\left(-\frac{\rho}{\sigma}\right)\right)$.

To see (i), further write

$$\begin{aligned}
\frac{\partial}{\partial \rho}\mathbf{g} &= \frac{k}{\rho^2}\left(1 - 2\Phi\left(-\frac{\rho}{\sigma}\right)\right) - \frac{2k}{\rho\sigma}\phi\left(-\frac{\rho}{\sigma}\right) \\
&= \frac{k}{\rho^2}\left(1 - 2\left(\Phi\left(-\frac{\rho}{\sigma}\right) + \frac{\rho}{\sigma}\phi\left(-\frac{\rho}{\sigma}\right)\right)\right) \\
&= \frac{k}{\rho^2}\left(1 - 2\left(\Phi\left(-\frac{\rho}{\sigma}\right) - \phi'\left(-\frac{\rho}{\sigma}\right)\right)\right).
\end{aligned}$$

Note $\Phi(x) < \frac{1}{2}$ and $\phi'(x) > 0$ for all $x < 0$. Thus, $\Phi\left(-\frac{\rho}{\sigma}\right) - \phi'\left(-\frac{\rho}{\sigma}\right) < \frac{1}{2}$ for all $\rho > 0$. It follows

that $\frac{\partial}{\partial \rho}\mathbf{g} > 0$, i.e. $\mathbf{g}$ is strictly increasing in ρ for all $\rho > 0$. Furthermore, note

$$\lim_{\rho \to 0} \mathbf{g}(\rho, k, \sigma) = 1 - \frac{k}{\sigma}\sqrt{\frac{2}{\pi}} = 1 - \frac{k}{\sigma}\frac{1}{C}$$
$$\mathbf{g}(k, k, \sigma) = 2\Phi\left(-\frac{k}{\sigma}\right).$$

If $\frac{k}{\sigma} > C$, note $\lim \mathbf{g}_{\rho \to 0}(\rho, k, \sigma) < 0$, $\mathbf{g}(k, k, \sigma) > 0$ and $\mathbf{g}(\cdot; k, \sigma)$ is continuous and strictly increasing. Thus, there exists a unique ρ^* such that $\mathbf{g}(\rho^*; k, \sigma) = 0$.

To see (ii), note first that $\frac{\partial}{\partial \sigma}\mathbf{g} = \frac{2k}{\sigma^2}\phi\left(-\frac{\rho}{\sigma}\right) > 0$. Thus, viewing ρ^* as a function of k and σ, we can see

$$\frac{\partial \rho^*}{\partial \sigma} = -\frac{\frac{\partial}{\partial \sigma}\mathbf{g}(\rho^*; k, \sigma)}{\frac{\partial}{\partial \rho^*}\mathbf{g}(\rho^*; k, \sigma)} < 0.$$

Therefore, ρ^* is strictly decreasing in σ. When $\sigma \to 0$, $\Phi\left(-\frac{\rho}{\sigma}\right) \to 0$ for each fixed $\rho > 0$. Then, in the limit when $\sigma = 0$, $\mathbf{g}(\rho^*; k, 0) = 0$ is solved by setting $\rho^* = k$.

To see (iii), note first that $\frac{\partial}{\partial k}\mathbf{g}(\rho^*; k, \sigma) = -\left(1 - 2\Phi\left(-\frac{\rho^*}{\sigma}\right)\right)/\rho^* < 0$. The remaining proof mimics that of statement (ii). □

Lemma B.7. Suppose $k > \sqrt{\frac{\pi}{2}}\sigma$. Then

$$\sup_{\gamma^* \in I(\gamma), \gamma \in \mathbb{R}} R(d^*_{\text{linear}}, \gamma, \gamma^*) = k/2 \tag{B.6}$$

and

$$\sup_{\gamma^* \in I(0)} R(d^*_{\text{linear}}, 0, \gamma^*) = k/2, \tag{B.7}$$

where $I(\gamma)$ and $R(d, \gamma, \gamma^*)$ are defined in (B.4).

Proof. We may write the left-hand side of (B.6) as

$$\sup_{\gamma^* \in I(\gamma), \gamma \in \mathbb{R}} R(d^*_{\text{linear}}, \gamma, \gamma^*) = \sup_{\gamma + k \geq 0} (\gamma + k)\left(1 - \mathbb{E}_\gamma[d^*_{\text{linear}}]\right) = \sup_{\gamma + k \geq 0} g_{\text{linear}}(\mu) \tag{B.8}$$

$$g_{\text{linear}}(\gamma) := (\gamma + k)\int_0^1 \Phi\left(\frac{2\rho^* x - \rho^* - \gamma}{\sigma}\right) dx, \tag{B.9}$$

where the first equality follows from symmetry of the parameter space and the fact that $d^*_{\text{linear}}(-x) =$

$1 - d^*_{\text{linear}}(x)$ for all $x \in \mathbb{R}$, and the second equality follows by applying change-of-variable twice:

$$\begin{aligned}
&\mathbb{E}_\gamma[d^*_{\text{linear}}] \\
&= \int (d^*_{\text{linear}}(x)) d\Phi\left(\frac{x-\gamma}{\sigma}\right) \\
&= \int_{-\rho^*}^{\rho^*} \frac{x+\rho^*}{2\rho^*} d\Phi\left(\frac{x-\gamma}{\sigma}\right) + \int_{\rho^*}^{\infty} d\Phi\left(\frac{x-\gamma}{\sigma}\right) \\
&= \left[\frac{x+\rho^*}{2\rho^*}\Phi\left(\frac{x-\gamma}{\sigma}\right)\right]_{-\rho^*}^{\rho^*} - \int_{-\rho^*}^{\rho^*} \Phi\left(\frac{2\rho^*\left(\frac{x+\rho^*}{2\rho^*} - \frac{1}{2}\right) - \gamma}{\sigma}\right) d\left(\frac{x+\rho^*}{2\rho^*}\right) + 1 - \Phi\left(\frac{\rho^*-\gamma}{\sigma}\right) \\
&= 1 - \int_0^1 \Phi\left(\frac{2\rho^* x - \rho^* - \gamma}{\sigma}\right) dx.
\end{aligned}$$

Now, $\mathbb{E}_0[d^*_{\text{linear}}(\hat{\gamma})] = \frac{1}{2}$ and $g_{\text{linear}}(0) = \frac{k}{2}$ by construction. Below, we show that $g_{\text{linear}}(\gamma)$ is first increasing and then decreasing on $[-k, \infty)$ with unique maximum at $\gamma = 0$, establishing the claim. To see this, take first and second derivatives of $g_{\text{linear}}(\gamma)$:

$$\begin{aligned}
g^{(1)}_{\text{linear}}(\gamma) &= -\frac{\gamma+k}{\sigma}\int_0^1 \phi\left(\frac{2\rho^* x - \rho^* - \gamma}{\sigma}\right) dx + \int_0^1 \Phi\left(\frac{2\rho^* x - \rho^* - \gamma}{\sigma}\right) dx, \\
g^{(2)}_{\text{linear}}(\gamma) &= \frac{\gamma+k}{\sigma^2}\int_0^1 \phi^{(1)}\left(\frac{2\rho^* x - \rho^* - \gamma}{\sigma}\right) dx - \frac{2}{\sigma}\int_0^1 \phi\left(\frac{2\rho^* x - \rho^* - \gamma}{\sigma}\right) dx \\
&= \frac{\gamma+k}{\sigma^2}\int_0^1 \frac{\gamma+\rho^*-2\rho x}{\sigma}\phi\left(\frac{2\rho^* x - \rho^* - \gamma}{\sigma}\right) dx - \frac{2}{\sigma}\int_0^1 \phi\left(\frac{2\rho^* x - \rho^* - \gamma}{\sigma}\right) dx \\
&= \frac{\gamma+k}{2\rho^*\sigma}\int_{\frac{-\gamma-\rho^*}{\sigma}}^{\frac{-\gamma+\rho^*}{\sigma}} -t\phi(t)dt - \frac{1}{\rho^*}\int_{\frac{-\gamma-\rho^*}{\sigma}}^{\frac{-\gamma+\rho^*}{\sigma}} \phi(t)dt \\
&= \frac{\gamma+k}{2\rho^*\sigma}\int_{\frac{\gamma-\rho^*}{\sigma}}^{\frac{\gamma+\rho^*}{\sigma}} t\phi(t)dt - \frac{1}{\rho^*}\int_{\frac{\gamma-\rho^*}{\sigma}}^{\frac{\gamma+\rho^*}{\sigma}} \phi(t)dt \\
&= \underbrace{\frac{1}{2\rho^*}\int_{\frac{\gamma-\rho^*}{\sigma}}^{\frac{\gamma+\rho^*}{\sigma}} \phi(t)dt}_{:=A}\left(\frac{\gamma+k}{\sigma}\frac{\int_{\frac{\gamma-\rho^*}{\sigma}}^{\frac{\gamma+\rho^*}{\sigma}} t\phi(t)dt}{\int_{\frac{\gamma-\rho^*}{\sigma}}^{\frac{\gamma+\rho^*}{\sigma}} \phi(t)dt} - 2\right),
\end{aligned}$$

where the second equality for $g^{(2)}_{\text{linear}}(\gamma)$ uses that $\phi'(x) = -x\phi(x)$ for all $x \in \mathbb{R}$, the third one follows from integration by change-of-variable, and the fourth equality follows from change-of-variable again

and $\phi(x) = \phi(-x)$ for all $x \in \mathbb{R}$. As $A > 0$, the sign of $g^{(2)}_{\text{linear}}(\gamma)$ is determined by

$$g^*_{\text{linear}}(\gamma) := \frac{\gamma + k}{\sigma} \frac{\int_{\frac{\gamma-\rho^*}{\sigma}}^{\frac{\gamma+\rho^*}{\sigma}} t\phi(t)dt}{\int_{\frac{\gamma-\rho^*}{\sigma}}^{\frac{\gamma+\rho^*}{\sigma}} \phi(t)dt} - 2.$$

Furthermore, we can write

$$g^*_{\text{linear}}(\gamma) = \frac{\gamma + k}{\sigma} \mathbb{E}\left[Z \mid \frac{\gamma - \rho^*}{\sigma} \leq Z \leq \frac{\gamma + \rho^*}{\sigma}\right] - 2,$$

where $\mathbb{E}\left(Z \mid a \leq Z \leq b\right)$ denotes the conditional expectation of a standard normal random variable Z conditional on $a \leq Z \leq b$. We are only interested in $\gamma + k \geq 0$. Also, note $\mathbb{E}\left[Z \mid \frac{\gamma-\rho^*}{\sigma} \leq Z \leq \frac{\gamma+\rho^*}{\sigma}\right]$ strictly increases in γ and has the same sign as γ. Moreover, note $\mathbb{E}\left[Z \mid -\frac{\rho^*}{\sigma} \leq Z \leq \frac{\rho^*}{\sigma}\right] = 0$, implying $g^*_{\text{linear}}(0) = -2$. Thus, we conclude $g^{(2)}_{\text{linear}}(\gamma) < 0$ for all γ below some strictly positive threshold and $g^{(2)}_{\text{linear}}(\gamma) > 0$ for all larger γ. That is, $g_{\text{linear}}(\gamma)$ is first concave and then convex, with the inflexion occurring at a strictly positive point.

Since $g_{\text{linear}}(\gamma) \geq 0$ when $\gamma \geq -k$ and it can also be verified that $g_{\text{linear}}(-k) = 0$ and $\lim_{\gamma\to\infty} g_{\text{linear}}(\gamma) = 0$, we conclude that $g_{\text{linear}}(\cdot)$ is first strictly increasing and then strictly decreasing, with a unique maximum. Furthermore, note

$$\begin{aligned}
g^{(1)}_{\text{linear}}(0) &= -\frac{k}{\sigma}\int_0^1 \phi\left(\frac{2\rho^* x - \rho^*}{\sigma}\right) dx + \int_0^1 \Phi\left(\frac{2\rho^* x - \rho^*}{\sigma}\right) dx \\
&= -\frac{k}{2\rho^*}\int_{-\frac{\rho^*}{\sigma}}^{\frac{\rho^*}{\sigma}} \phi\left(t\right) dt + \frac{\sigma}{2\rho^*}\int_{-\frac{\rho^*}{\sigma}}^{\frac{\rho^*}{\sigma}} \Phi\left(t\right) dt \\
&= -\frac{k}{2\rho^*}\left(\Phi\left(\frac{\rho^*}{\sigma}\right) - \Phi\left(-\frac{\rho^*}{\sigma}\right)\right) + \frac{1}{2} \\
&= \frac{1}{2} - \frac{k}{2\rho^*}\left(1 - 2\Phi\left(-\frac{\rho^*}{\sigma}\right)\right) \\
&= 0,
\end{aligned}$$

where the second equality applies change-of-variable and the last equality follows from the definition of ρ^*, which exists and is unique when $k > C\sigma$ by Lemma B.6. Thus, $g_{\text{linear}}(\cdot)$ has a unique maximization point in $[-k, \infty)$ at $\gamma = 0$. □

Lemma B.8. Under assumptions made in Theorem 3,

$$d_0 := d_0((w^*)^\top Y) := \mathbf{1}\{(w^*)^\top Y \geq 0\}$$

is not MMR optimal.

Proof. By Step 1 in the proof of part (i) of Theorem 3, we know the MMR value of problem (A.4) equals $\bar{k}_{w^*}(0)/2$. In contrast, we will show that

$$R^*_{w^*,0} := \sup_{\theta\in\Theta} U(\theta)\left(\mathbf{1}\{U(\theta)\geq 0\} - \mathbb{E}_{m(\theta)}[d_0((w^*)^\top Y)]\right) > \bar{k}_{w^*}(0)/2.$$

Write

$$\begin{aligned} R^*_{w^*,0} &= \sup_{\gamma\in\Gamma_{w^*}} \left(\sup_{U^*\in(-\bar{k}_{w^*}(-\gamma),\, \bar{k}_{w^*}(\gamma))} U^*\left(\mathbf{1}\{U^*\geq 0\} - \mathbb{E}_\gamma[d_0((w^*)^\top Y)]\right)\right) \\ &= \sup_{\gamma\in\Gamma_{w^*}, \bar{k}_{w^*}(\gamma)>0} \bar{k}_{w^*}(\gamma)\left(1 - \mathbb{E}_\gamma[d_0((w^*)^\top Y)]\right) \\ &= \sup_{\gamma\in\Gamma_{w^*}, \bar{k}_{w^*}(\gamma)>0} \bar{k}_{w^*}(\gamma)\Phi\left(-\frac{\gamma}{\sigma}\right), \end{aligned}$$

using first centrosymmetry of Θ and then $(w^*)^\top Y \sim N(\gamma,\sigma^2)$, where $\sigma^2 := (w^*)^\top \Sigma w^*$. Write $g(\gamma) := \bar{k}_{w^*}(\gamma)\Phi(-\frac{\gamma}{\sigma})$ for $\gamma\in\Gamma_{w^*}$ such that $\bar{k}_{w^*}(\gamma) > 0$. By definition

$$g(0) = \bar{k}_{w^*}(0)/2.$$

Let $\partial g(\cdot)$ be the generalized gradient of $g(\cdot)$. In the following, we show that $0 \notin \partial g(0)$. By Clarke (1990, Proposition 2.3.2), 0 is then not a local maximum or minimum; hence, $R^*_{w^*,0} > \bar{k}_{w^*}(0)/2$.

Note $\Phi(\cdot)$ is strictly differentiable and thus Lipschitz near 0 (Clarke, 1990, Proposition 2.2.4). Also, as $\bar{k}_{w^*}(0)$ is concave and bounded from below near 0, $\bar{k}_{w^*}(\cdot)$ must be Lipschitz near 0 (Clarke, 1990, Proposition 2.2.6). Moreover, both $\bar{k}_{w^*}$ and Φ are regular at 0 (Clarke, 1990, Proposition 2.3.6) as well as positive. By Clarke (1990, Proposition 2.3.13), the (appropriately generalized) chain rule can be applied to g to characterize $\partial g(0)$.

As Φ is strictly differentiable, its generalized gradient coincides with the unique derivative (Clarke, 1990, Proposition 2.2.4)). As $\bar{k}_{w^*}$ is concave and Lipschitz near 0, its generalized gradient coincides with its superdifferential (Clarke, 1990, Proposition 2.2.7)). Hence, let $\tilde{s}_{w^*}(0)$ be a

supergradient of $\bar{k}_{w^*}(0)$. We may calculate the generalized gradient of $g(\gamma)$ at $\gamma = 0$ as

$$\begin{aligned}
&\frac{\tilde{s}_{w^*}(0)}{2} - \frac{\bar{k}_{w^*}(0)\phi(0)}{\sigma} \\
=&\frac{\phi(0)}{\sigma}\left(\sqrt{\frac{\pi}{2}}\sigma \cdot \tilde{s}_{w^*}(0) - \bar{k}_{w^*}(0)\right) \\
\leq&\frac{\phi(0)}{\sigma}\left(\sqrt{\frac{\pi}{2}}\sigma \cdot s_{w^*}(0) - \bar{k}_{w^*}(0)\right) < 0,
\end{aligned}$$

where the first inequality follows as $s_{w^*}(0)$ is the largest supergradient so $s_{w^*}(0) \geq \tilde{s}_{w^*}(0)$, and the second inequality follows from noting $\overline{I}(\mathbf{0}) = \bar{k}^*_w(0)$ under stated assumptions (by Step 1 in the proof of Theorem 3(i)) and by picking $\overline{I}(\mathbf{0})$ large enough so that $\overline{I}(\mathbf{0}) > \sqrt{\pi/2} \cdot \sigma \cdot s_{w^*}(0)$. Thus, we have shown that $\partial g(0) < 0$, and therefore 0 is not a local maximum or minimum. □

B.2 Proof of Proposition 1

Statements (i)-(ii). In the running example, the expected regret of a rule $d(\cdot)$ can be written as

$$R(d, \mu, \mu_0) = \mu_0 \left(\mathbf{1}\{\mu_0 \geq 0\} - \mathbb{E}_\mu[d(Y)]\right), \quad \mu \in M, \mu_0 \in I(\mu).$$

where $Y \sim N(\mu, \Sigma)$; for future use, we state its likelihood

$$f(y \mid \mu) = \frac{1}{\sqrt{(2\pi)^n \Pi_{j=1}^n \sigma_j^2}} \exp\left(-\frac{1}{2}\sum_{i=1}^n \frac{(y_i - \mu_i)^2}{\sigma_i^2}\right).$$

Consider the class of decision rules (parameterized by scalar $m_0 \geq C\|x_1 - x_0\|$)

$$\begin{aligned}
d_{m_0} &:= \mathbf{1}\{w_{m_0}^\top Y \geq 0\} \\
w_{m_0}^\top &:= \left(1, \frac{\max\{m_0 - C\|x_2 - x_0\|, 0\}/\sigma_2^2}{(m_0 - C\|x_1 - x_0\|)/\sigma_1^2}, \ldots, \frac{\max\{m_0 - C\|x_n - x_0\|, 0\}/\sigma_n^2}{(m_0 - C\|x_1 - x_0\|)/\sigma_1^2}\right),
\end{aligned}$$

with the understanding that, for $m_0 = C\|x_1 - x_0\|$, we have $w_{m_0}^\top = (1, 0, \ldots, 0)$. Consider also the class of priors π_{m_0} that randomize evenly over $\left\{\left(\mu_0, \mu^\top\right)^\top, \left(-\mu_0, -\mu^\top\right)^\top\right\}$, where

$$\begin{aligned}
\mu_0 &= m_0 \\
\mu_j &= \max\{m_0 - C\|x_j - x_0\|, 0\}, \quad j = 1, \ldots, n
\end{aligned} \tag{B.10}$$

for some $m_0 \geq C\|x_1 - x_0\|$. We will show that i) for any prior π_{m_0}, rule d_{m_0} is a corresponding Bayes rule, uniquely so if $m_0 > C\|x_1 - x_0\|$, ii) for any decision rule d_{m_0}, a prior $\pi_{\tilde{m}_0}$ (note $\tilde{m}_0 \neq m_0$ in general) is least favorable, and finally that iii) the resulting best-response mapping has a fixed point m_0^*. This fixed point defines the MMR rule from the proposition, which is furthermore unique whenever it is uniquely Bayes against $\pi_{m_0^*}$.[22]

Regarding step i), if $m_0 > C\|x_1 - x_0\|$, the unique Bayes response to π_{m_0} equals

$$\begin{aligned} & \mathbf{1}\{\mathbb{E}[\mu_0 \mid Y] \geq 0\} \\ = \; & \mathbf{1}\{f(Y|\mu) - f(Y|-\mu) \geq 0\} \\ = \; & d_{m_0}(Y), \end{aligned}$$

where the last step uses familiar normal likelihood algebra.[23] Any decision rule is Bayes against π_{m_0} if $m_0 = C\|x_1 - x_0\|$.

Regarding ii), observe that expected regret of d_{m_0} depends on θ only through (μ, μ_0) and that maximizing it amounts to solving

$$\begin{aligned} & \sup_{\mu \in M, \mu_0 \in I(\mu)} \mu_0 \left(\mathbf{1}\{\mu_0 \geq 0\} - \Phi\left(\frac{w_{m_0}^\top \mu}{\sqrt{w_{m_0}^\top \Sigma w_{m_0}}} \right) \right) \\ = & \sup_{\mu \in M, \mu_0 \in I(\mu), \mu_0 \geq 0} \mu_0 \Phi\left(-\frac{w_{m_0}^\top \mu}{\sqrt{w_{m_0}^\top \Sigma w_{m_0}}} \right), \end{aligned}$$

where we used centrosymmetry of Θ and where

$$\begin{aligned} M &= \{\mu \in \mathbb{R}^n : |\mu_i - \mu_j| \leq C\|x_i - x_j\|, i, j = 1 \ldots n, i \neq j\}. \\ I(\mu) &= \{u \in \mathbb{R} : |\mu_i - u| \leq C\|x_i - x_0\|, \quad i = 1, \ldots . n\}. \end{aligned}$$

Let $j_{m_0}^*$ be the highest index j for which $w_{m_0,j}$ is not 0. Then, $\mu_0 \Phi\left(-\frac{w_{m_0}^\top \mu}{\sqrt{w_{m_0}^\top \Sigma w_{m_0}}} \right)$ does not depend on $(\mu_{j_{m_0}^*+1}, \ldots, \mu_n)$ and decreases in $(\mu_1, \ldots, \mu_{j_{m_0}^*})$. It follows that some prior $\pi_{\tilde{m}_0}$ is least favorable,

[22] Thus, we use the game theoretic characterization of maximin-type decision rules (e.g., Berger (1985, Section 5)).

[23] This may be easier to see upon multiplying w_{m_0} through by $m_0 - C\|x_1 - x_0\|)/\sigma_1^2$. Our notation is meant to clarify continuity and convergence to $(1, 0, \ldots, 0)$.

where furthermore $\tilde{m}_0$ is the optimal argument μ_0 in

$$\sup_{\mu\in M,\mu_0\in I(\mu),\mu_0\geq 0} \mu_0\Phi\left(-\frac{w_{m_0}^\top\mu}{\sqrt{w_{m_0}^\top\Sigma w_{m_0}}}\right) = \sup_{\mu_0\geq 0} g(\mu_0,m_0),$$
$$g(\mu_0,m_0) = \mu_0\Phi\left(-\frac{\sum_{j=1}^n \frac{\max\{m_0-C\|x_j-x_0\|,0\}(\mu_0-C\|x_j-x_0\|)}{\sigma_j^2}}{\sqrt{\sum_{j=1}^n \frac{\max^2\{m_0-C\|x_j-x_0\|,0\}}{\sigma_j^2}}}\right).$$

Regarding iii), the best-response mapping $\psi : [C\|x_1-x_0\|,\infty) \rightrightarrows [0,\infty)$ defined by[24]

$$\psi(m_0) := \arg\sup_{\mu_0\geq 0} g(\mu_0,m_0) \tag{B.11}$$

has a fixed point $m_0^* \in \psi(m_0^*)$. To see this, first compactify the domain of μ_0 in the above definition by noting that $\tilde{m}_0$ can be universally bounded from above. This is because, for any m_0 under consideration, one has $g(C\|x_1-x_0\|,m_0) = C\|x_1-x_0\|/2$ but also

$$g(\mu_0,m_0) = \mu_0(1-\Pr(w_{m_0}^\top Y\geq 0)) \leq \mu_0(1-\Pr(Y\geq \mathbf{0})) = \mu_0\left(1-\prod_{j=1}^n \Phi(-\mu_j/\sigma_j)\right).$$

Using (B.10), this upper bound is seen to vanish as $m_0\to\infty$. Hence, it is w.l.o.g. to change the constraint set in (B.11) to $0\leq\mu_0\leq\overline{\mu}_0$ for $\overline{\mu}_0$ large enough (but independent of m_0). Given this compactification, continuity of $g(\cdot)$ implies nonemptiness and upper hemicontinuity of $\psi(\cdot)$. Next, by algebra resembling Proposition 7 in Stoye (2012a, see also Lemma B.7 above), for any fixed m_0 the function $g(\mu_0,m_0)$ is first concave then convex in μ_0 and converges to 0 $[-\infty]$ as $\mu_0\to\infty$ $[\mu_0\to-\infty]$. Hence, $\psi(\cdot)$ is interval-valued. These observations jointly imply that the graph of $\psi(\cdot)$ is path-connected. We show below that (with slight abuse of notation for set-valued mappings) $\psi(C\|x_1-x_0\|)\geq C\|x_1-x_0\|$, and we already know that $\psi(m_0)<m_0$ for $m_0>\overline{\mu}_0$. This establishes existence of m_0^*.

We next show that $\psi(C\|x_1-x_0\|)\geq C\|x_1-x_0\|$, strictly so if $C\|x_1-x_0\|<\sqrt{\pi/2}\cdot\sigma_1$. This is because for $m_0=C\|x_1-x_0\|$, we have

$$\begin{aligned} g(\mu_0,m_0) &= \mu_0\Phi\left(\frac{C\|x_1-x_0\|-\mu_0}{\sigma_1}\right) \\ \Longrightarrow \left.\frac{\partial g(\mu_0,m_0)}{\partial\mu_0}\right|_{\mu_0=C\|x_1-x_0\|} &= -\frac{C\|x_1-x_0\|\,\phi(0)}{\sigma_1}+\frac{1}{2}. \end{aligned}$$

[24]The mapping is in fact a function, but establishing that would be unnecessary work.

After substituting in for $\phi(0)$ and simplifying, the above partial derivative is seen to have the same sign as $\sqrt{\pi/2}\cdot\sigma_1 - C\|x_1 - x_0\|$. This establishes the claim and also proves statement (ii) because the fixed point $m_0^* = C\|x_1 - x_0\|$ has been discovered for that statement's case.

This concludes the proof. For ease of computation, we note that, after tedious algebra, m_0^* can be uniquely characterized by

$$\frac{\Phi\left(-\sqrt{\sum_{j=1}^n \frac{\max^2\left\{m_0^* - C\|x_j - x_0\|, 0\right\}}{\sigma_j^2}}\right)}{\phi\left(-\sqrt{\sum_{j=1}^n \frac{\max^2\left\{m_0^* - C\|x_j - x_0\|, 0\right\}}{\sigma_j^2}}\right)} = m_0^* \frac{\sum_{j=1}^n \frac{\max\left\{m_0^* - C\|x_j - x_0\|, 0\right\}}{\sigma_j^2}}{\sqrt{\sum_{j=1}^n \frac{\max^2\left\{m_0^* - C\|x_j - x_0\|, 0\right\}}{\sigma_j^2}}}. \tag{B.12}$$

Statement (iii). Below, we show $d^*_{\mathrm{RT}}((w^*)^\top Y) = \Phi\big((w^*)^\top Y/\tilde{\sigma}\big)$, where $\tilde{\sigma} = \sqrt{2C^2\|x_1 - x_0\|^2/\pi - \sigma_1^2}$ and $w^* = (1, 0, \ldots, 0)^\top$ is MMR optimal and satisfies (3.8) and (3.9). Therefore, results in Theorem 3 apply. In particular, Lemma B.4 implies that d^*_{linear} defined in (3.10) with ρ such that $\frac{C\|x_1 - x_0\| - \rho^*}{2C\|x_1 - x_0\|} = \Phi\left(-\rho^*/\sigma_1\right)$, as well as any convex combination of d^*_{RT} and d^*_{linear}, are also MMR optimal. Let

$$\mathbf{R} := \min_{d\in\mathcal{D}_n} \sup_{\theta\in\Theta} U(\theta)\left(\mathbf{1}\{U(\theta) \geq 0\} - \mathbb{E}_{m(\theta)}[d(Y)]\right)$$

be the minimax value of the problem in the running example.

Step 1: We show $\mathbf{R} \leq C\|x_1 - x_0\|/2$ because this bound is attained by $d^*_{\mathrm{RT}}(\cdot)$. To see this, use Step 2 of the proof of Theorem 3 to write

$$\begin{aligned}
\mathbf{R} &\leq \sup_{\theta\in\Theta} \theta_0 \left(\mathbf{1}\{\theta_0 \geq 0\} - \mathbb{E}_{m(\theta)}[d^*_{\mathrm{RT}}((w^*)^\top Y)]\right) && \text{(B.13)}\\
&= \sup_{\gamma\in\mathbb{R}} \left(\sup_{U^*\in\left[-\bar{k}_{w^*}(-\gamma), \bar{k}_{w^*}(\gamma)\right]} U^* \left(\mathbf{1}\{U^* \geq 0\} - \mathbb{E}_\gamma[d^*_{\mathrm{RT}}((w^*)^\top Y)]\right)\right), && \text{(B.14)}
\end{aligned}$$

where $w^* = (1, 0, \ldots, 0)^\top$ and $\bar{k}_{w^*}(\gamma)$ is defined in (A.11) and calculated as

$$\begin{aligned}
\sup\quad & \theta_0 \\
\text{s.t.}\quad & |\theta_i - \theta_j| \leq C\|x_i - x_j\|, \quad i, j = 0, ..., n, \\
& \theta_1 = \gamma.
\end{aligned} \tag{B.15}$$

As $\|x_1 - x_0\| \leq \|x_j - x_0\|$ for all $j = 1, ..., n$, the linear program (B.15) admits a simple solution $\bar{k}_{w^*}(\gamma) = \gamma + C\|x_0 - x_1\|$. Therefore, (B.14) can be further written as

$$\sup_{\gamma \in \mathbb{R}} \left(\sup_{U^* \in [\gamma - C\|x_0 - x_1\|,\, \gamma + C\|x_0 - x_1\|]} U^* \left(\mathbf{1}\{U^* \geq 0\} - \mathbb{E}_\gamma[d^*_{\mathrm{RT}}((w^*)^\top Y)] \right) \right), \tag{B.16}$$

where $(w^*)^\top Y \sim N(\gamma, \sigma_1^2)$. Applying Stoye (2012a) with $k = C\|x_0 - x_1\|$ and $\sigma = \sigma_1$, we conclude that when $C\|x_1 - x_0\| > \sqrt{\pi/2} \cdot \sigma_1$, $d^*_{\mathrm{RT}}((w^*)^\top Y)$ solves the minimax problem

$$\min_{d \in \mathcal{D}} \sup_{\gamma \in \mathbb{R}} \left(\sup_{U^* \in [\gamma - C\|x_0 - x_1\|,\, \gamma + C\|x_0 - x_1\|]} U^* \left(\mathbf{1}\{U^* \geq 0\} - \mathbb{E}_\gamma[d((w^*)^\top Y)] \right) \right)$$

and the value of (B.16) equals $C\|x_0 - x_1\|/2$.

STEP 2: $\mathbf{R} \geq C\|x_1 - x_0\|/2$ because this value is attained by setting $(\theta_1, \ldots, \theta_n) = \mathbf{0}$:

$$\begin{aligned} \mathbf{R} &\geq \min_{d \in \mathcal{D}_n} \sup_{\theta \in \Theta : \theta = (\theta_0, \mathbf{0})} \theta_0 \left(\mathbf{1}\{\theta_0 \geq 0\} - \mathbb{E}_{\mathbf{0}}[d(Y)] \right) \\ &= \min_{d \in \mathcal{D}_n} \max \left\{ \overline{I}(\mathbf{0}) \left(1 - \mathbb{E}_0[d(Y)] \right), \overline{I}(\mathbf{0}) \mathbb{E}_{\mathbf{0}}[d(Y)] \right\}, \end{aligned}$$

where the last line used $\underline{I}(\mathbf{0}) = -\overline{I}(\mathbf{0})$. This minimum is attained by any rule with $\mathbb{E}_{\mathbf{0}}[d(Y)] = 1/2$, and its value equals $\overline{I}(\mathbf{0})/2 = C\|x_1 - x_0\|/2$.

As $\mathbb{E}_{\mathbf{0}}[d^*_{\mathrm{RT}}((w^*)^\top Y)] = 1/2$, the last step above also verifies (3.8). Optimality of d^*_{RT} in Step 1 verified (3.9).

Statement (iv). Using that x_1 is a unique nearest neighbor, we have that, for μ close to $\mathbf{0}$, $\overline{I}(\mu) = \mu_1 + C\|x_1 - x_0\|$. This is clearly differentiable at $\mu = \mathbf{0}$ and Theorem 3(iii) therefore applies.

B.3 Additional Results

B.3.1 Profiled Regret in the Running Example

In the running example, consider w^*-profiled regret, recalling that $w^* = (1, 0)^\top$. By the definition in (4.2), the w^*-profiled regret function of a rule $d \in \mathcal{D}_2$ equals

$$\overline{R}_{w^*}(d, \gamma) = \sup_{\theta \in \Theta \text{ s.t. } m_1(\theta) = \gamma} U(\theta) \left(\mathbf{1}\{U(\theta) \geq 0\} - \mathbb{E}_{m(\theta)}[d(Y)] \right).$$

Specifically, we look at three rules:

$$\begin{aligned} d_0((w^*)^\top Y) &= d_0(Y_1) = \mathbf{1}\{Y_1 \geq 0\} \\ d^*_{\mathrm{RT}}((w^*)^\top Y) &= d^*_{\mathrm{RT}}(Y_1) \\ d^*_{\mathrm{linear}}((w^*)^\top Y) &= d^*_{\mathrm{linear}}(Y_1) \end{aligned}$$

defined in (3.7)-(3.11).

As these rules depend on data only via $(w^*)^\top Y$, for $d \in \{d_0, d^*_{\mathrm{RT}}, d^*_{\mathrm{linear}}\}$ we can write

$$\begin{aligned} \overline{R}_{w^*}(d,\gamma) &= \sup_{\theta\in\Theta \text{ s.t. } m_1(\theta)=\gamma} U(\theta)\left(\mathbf{1}\{U(\theta)\geq 0\} - \mathbb{E}_{m_1(\theta)}[d(Y_1)]\right) \\ &= \sup_{U^*\in\left[-\bar{k}_{w^*}(-\gamma),\, \bar{k}_{w^*}(\gamma)\right]} U^*\left(\mathbf{1}\{U^*\geq 0\} - \mathbb{E}_\gamma[d(Y_1)]\right), \end{aligned}$$

where $\bar{k}_{w^*}(\gamma)$ is defined in (A.11), and (B.15) solves for $\bar{k}_{w^*}(\gamma) = \gamma + k$, where $k = C\|x_0 - x_1\|$. Therefore, we may further calculate

$$\begin{aligned} &\overline{R}_{w^*}(d,\gamma) \\ &= \sup_{U^*\in[\gamma-k,\,\gamma+k]} U^*\left(\mathbf{1}\{U^*\geq 0\} - \mathbb{E}_\gamma[d(Y_1)]\right) \\ &= \max\left\{ \sup_{U^*\in[\gamma-k,\,\gamma+k],U^*\geq 0} U^*\left(1-\mathbb{E}_\gamma[d(Y_1)]\right), \sup_{U^*\in[\gamma-k,\,\gamma+k],U^*\leq 0} -U^*\mathbb{E}_\gamma[d(Y_1)] \right\} \\ &= \begin{cases} (-\gamma+k)\mathbb{E}_\gamma[d(Y_1)], & \text{if } \gamma < -k, \\ \max\left\{(\gamma+k)\left(1-\mathbb{E}_\gamma[d(Y_1)]\right), (-\gamma+k)\mathbb{E}_\gamma[d(Y_1)]\right\}, & \text{if } -k\leq\gamma\leq k, \\ (\gamma+k)(1-\mathbb{E}_\gamma[d(Y_1)]), & \text{if } \gamma > k. \end{cases} \end{aligned}$$

As $Y_1 \sim N(\gamma,\sigma_1)$, algebra shows $\mathbb{E}_\gamma[d_0(Y_1)] = \Phi\left(\frac{\gamma}{\sigma_1}\right)$, $\mathbb{E}_\gamma[d^*_{\mathrm{RT}}(Y_1)] = \Phi\left(\sqrt{\frac{\pi}{2}}\frac{\gamma}{k}\right)$, and

$$\begin{aligned} \mathbb{E}_\gamma[d^*_{\mathrm{linear}}(Y_1)] = \Phi\left(\frac{\gamma-\rho^*}{\sigma_1}\right) &+ \frac{\sigma_1}{2\rho^*}\frac{1}{\sqrt{2\pi}}\left[e^{-\frac{1}{2}\left(\frac{\rho^*+\gamma}{\sigma_1}\right)^2} - e^{-\frac{1}{2}\left(\frac{\rho^*-\gamma}{\sigma_1}\right)^2}\right] \\ &+ \frac{\gamma+\rho^*}{2\rho^*}\left[\Phi\left(\frac{\rho^*-\gamma}{\sigma_1}\right) - \Phi\left(\frac{-\rho^*-\gamma}{\sigma_1}\right)\right]. \end{aligned}$$

The w^*-profiled regret of the three rules can then be calculated easily following our characterizations above.

B.3.2 Plot of Profiled Regret for the Plug-in Rule

We plot and compare the w^*-profiled regrets of $d_{\text{plug-in}}$ defined in (4.14) and those of other rules, including d^*_{linear} and d^*_{RT}. See Figure 4 for details. Note $d_{\text{plug-in}}$ is not MMR optimal: its w^*-profiled regret at $\gamma = 0$ is slightly larger than the MMR value of the problem. Moreover, the w^*-profiled regret curve of $d_{\text{plug-in}}$ is bell-shaped and similar to that of d^*_{linear}.

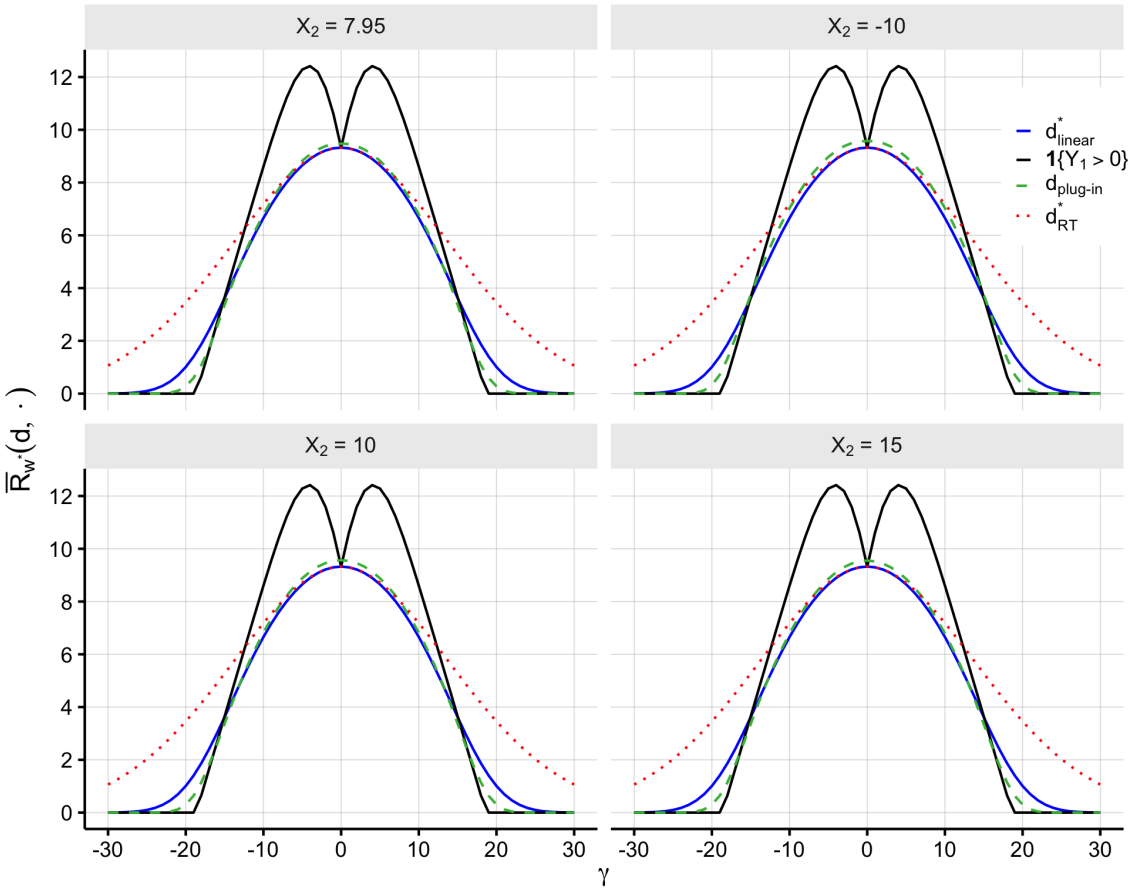

Figure 4: w^*-profiled regrets of $d_{\text{plug-in}}$ in (4.14) and other rules. The top left plot uses parameters of the running examples from Figure 1; the rest of the plots change the value of x_2 while keeping other parameters the same.

B.3.3 Verification of Theorem 1 for the Example in Section 5.1

By Theorem 1, it suffices to show that the statistical model in (5.1) and the welfare contrast in (5.2) display nontrivial partial identification in the sense of Definition 2. Note the image of

$(m_1(\theta), m_2(\theta))^\top$ in this example is defined as

$$M = \left\{ (\mu_1, \mu_2)^\top \in \mathbb{R}^2 | m_1(\theta) = \mu_1, m_2(\theta) = \mu_2, \theta \in \Theta \right\}.$$

The restrictions on $p(0)$ and $p(1)$ are: $p(1) \in [0,1]$, $p(0) \in [0,1]$, $p(1) \geq p(0)$ and $p(1) + \alpha \leq 1$. Without any additional shape restrictions on $\text{MTE}(\cdot)$, any functions from $[0,1]$ to $[-1,1]$ are compatible with the model. Hence, we can deduce that

$$M = \left\{ (\mu_1, \mu_2)^\top \in \mathbb{R}^2 | \mu_1 \in [-\mu_2, \mu_2], 0 \leq \mu_2 \leq 1 - \alpha \right\}.$$

The identified set of $U(\theta)$ is then

$$I(\mu_1, \mu_2) = \left\{ u \in \mathbb{R} | U(\theta) = u, m_1(\theta) = \mu_1, m_2(\theta) = \mu_2, \theta \in \Theta \right\}, \text{for all } (\mu_1, \mu_2)^\top \in M.$$

with extrema

$$\begin{aligned}
\overline{I}(\mu_1, \mu_2) &= \sup_{m_1(\theta)=\mu_1, m_2(\theta)=\mu_2, \theta \in \Theta} \left\{ \frac{m_1(\theta)}{\alpha + m_2(\theta)} + \frac{1}{\alpha + m_2(\theta)} \int_{p(1)}^{p(1)+\alpha} \text{MTE}(v) dv \right\} \\
&= \frac{\mu_1}{\alpha + \mu_2} + \frac{1}{\alpha + \mu_2} \sup \left\{ \int_{p(1)}^{p(1)+\alpha} \text{MTE}(v) dv \right\} \\
&= \frac{\mu_1}{\alpha + \mu_2} + \frac{\alpha}{\alpha + \mu_2}
\end{aligned}$$

and similarly

$$\underline{I}(\mu) = \frac{\mu_1}{\alpha + \mu_2} - \frac{\alpha}{\alpha + \mu_2}.$$

Let

$$\mathcal{S} = \left\{ (\mu_1, \mu_2)^\top \in \mathbb{R}^2 : -\min(\alpha, 1-\alpha) < \mu_1 < \min(\alpha, 1-\alpha), 0 \leq \mu_2 \leq 1 - \alpha \right\} \subseteq \mathbb{R}^2.$$

We can verify that for any $(\mu_1, \mu_2)^\top \in \mathcal{S}$,

$$\underline{I}(\mu_1, \mu_2) = \frac{\mu_1}{\alpha + \mu_2} - \frac{\alpha}{\alpha + \mu_2} < 0 < \frac{\mu_1}{\alpha + \mu_2} + \frac{\alpha}{\alpha + \mu_2} = \overline{I}(\mu_1, \mu_2).$$

Therefore, the statistical model in (5.1) and the welfare contrast in (5.2) exhibit nontrivial partial identification as defined in 2, and Theorem 1 applies to the example in Section 5.1.

Externally Valid Selection of Experimental Sites via the k-Median Problem*

José Luis Montiel Olea† Brenda Prallon† Chen Qiu†
Jörg Stoye† Yiwei Sun†

August 2024

Abstract

We present a decision-theoretic justification for viewing the question of how to best choose *where* to experiment in order to optimize external validity as a *k-median (clustering) problem*, a popular problem in computer science and operations research. We present conditions under which minimizing the *worst-case*, welfare-based regret among *all* nonrandom schemes that select k sites to experiment is approximately equal—and sometimes exactly equal—to finding the k most central vectors of baseline site-level covariates. The k-median problem can be formulated as a linear integer program. Two empirical applications illustrate the theoretical and computational benefits of the suggested procedure.

1 Introduction

A common concern in randomized evaluations of new policies is their *external validity*; that is, whether the estimated effects of a policy intervention carry over to new samples or populations (Vivalt, 2020; Gechter, 2024; Duflo, Glennerster, and Kremer, 2007, Chapter 8). There is a large body of literature arguing that the external validity of randomized evaluations can be improved by explicitly incorporating this goal into the experimental design by, for example, carefully deciding where to experiment; see Degtiar and Rose (2023) for an overview and references, and also the recent work of Chassang and Kapon (2022). For instance, if a researcher has access to multiple

*We would like to thank Isaiah Andrews, Tim Armstrong, Michèle Belot, Michael Gechter, Jesse Goodman, Kei Hirano, Yiqi Liu, Francesca Molinari, Charles Manski, Guillaume Pouliot, David Shmoys, and Davide Viviano for helpful feedback, comments, and suggestions. We would also like to thank Eddie Ramirez Saquic for excellent research assistance. We gratefully acknowledge financial support from the NSF under grant SES-2315600.

†Cornell University. Department of Economics. Corresponding author: *montiel.olea@gmail.com*

sites to experimentally evaluate a new policy, it is possible to use the information available before the evaluation (such as site-level characteristics) to *nonrandomly* select one site—or more—that are "representative" of the populations of interest. Recently, both Egami and Lee (2024) and Gechter, Hirano, Lee, Mahmud, Mondal, Morduch, Ravindran, and Shonchoy (2024) have shown that a research design that nonrandomly selects experimental sites—referred to broadly in the literature as *purposive sampling* (Cook, Campbell, and Shadish, 2002, p. 511)—is useful to improve the external validity of randomized evaluations.

This paper presents a theoretical justification for viewing the question of how to best design a purposive sampling scheme—or, equivalently, how to best nonrandomly select k sites to experiment—with the goal of optimizing external validity as a *k-median problem*. This is a classical problem in computer science and operations research (Williamson and Shmoys, 2011; Cohen-Addad, Esfandiari, Mirrokni, and Narayanan, 2022). Broadly speaking, in a k-median problem, there is a set of *facilities* and a set of *clients*; the goal is to open at most k facilities and *connect* each client to at least one facility at minimal total connection cost. We next provide more specific details on how these problems relate.

Statistical Decision Theory: Under conditions that will be explained clearly in Sections 2 and 3, Theorem 1 in this paper shows that the *worst-case* (welfare-based) *regret* of *any* purposive sampling scheme that selects at most k sites to experiment is approximately equal—and *sometimes exactly equal*—to the objective function of a k-median problem with the following features. First, the sites available for experimentation are treated as facilities. Second, the sites where the policy maker would like to implement the new policy are treated as clients. Third, the connection cost between clients and facilities is proportional to the Euclidean distance between the available site-level covariates. Importantly, we show that the solution of the k-median problem is exactly "minimax-regret" optimal (among all purposive sampling schemes) when i) the candidate sites for experimentation and the policy-relevant sites are *disjoint* (Remark 4) and ii) the treatment effect heterogeneity across sites accommodated in the parameter space is *substantial* (in a sense Theorem 1 makes precise). When ii) holds but i) does not, we show that the optimized value of the k-median problem only approximates the minimax-regret value, and that the approximation error improves as either the number of policy-relevant sites increases or the experiments conducted become more precise (Remark 5). The link with the k-median problem established in this paper thus shows that selecting the k sites that have the most central vector of covariates tends to optimize external validity (in a minimax-regret sense).

In order to formalize this connection, we leverage recent developments in the literature of treatment choice problems with partial identification; for example Yata (2021); Ishihara and Kitagawa

(2021); Montiel Olea, Qiu, and Stoye (2023). Note first that in the k-median problem, each client is typically connected to only one facility (otherwise, the connection cost would not be minimized). In the context of the site selection problem, a connection between a site i (where an experiment was conducted) and a site j (where no experimentation occurred) means that the estimated effects obtained in site i are used to inform policy decisions in site j. This means that even when experimental outcomes in $k > 1$ sites are available, the solution to the k-median problem would typically prescribe the policy maker to only use the information from the site with the smallest connection cost: the *nearest neighbor*. But when is it decision-theoretically optimal for a policy maker to behave in this way? After k sites have been selected for experimentation, the policy maker faces the "evidence aggregation" problem introduced in Ishihara and Kitagawa (2021) and recently discussed in Yata (2021), Christensen, Moon, and Schorfheide (2022) and Montiel Olea et al. (2023). That is, the policy maker must decide how to use the estimated treatment effects in different—and potentially heterogeneous—sites to make a treatment choice (implement the new policy or preserve the status quo) in the site of interest. It is known that there are conditions under which it is minimax-regret optimal for the policy maker to base decisions on the nearest neighbor (the site with the most similar site-level characteristics), provided the true treatment effects are allowed to vary substantially as a function of site-level covariates; see, for example, Proposition 1 in Montiel Olea et al. (2023). Moreover, in this case, the optimized worst-case regret is proportional to the distance between the baseline covariates of the site of interest and those of its nearest neighbor. Thus, these results on treatment choice problems with partial identification are fundamental to provide a decision-theoretic justification for the use of the k-median (clustering) problem to optimize external validity.

Algorithms: The connection with the k-median problem clarifies the problem's difficulty but also suggests efficient algorithms. To see the need for those, recall that any purposive sampling scheme optimizes over "n choose k" potential site combinations, where n is the total number of sites available for experimentation. Optimization over purposive sampling schemes also requires the evaluation of some measure of performance that depends on the dimension of the site-level covariates, d. Thus, optimally choosing a purposive sampling scheme by simply evaluating the performance of each combination is costly when "n choose k" or d is large.

Conceptually, the connection with the k-median problem allows us to understand the computational complexity of finding a minimax-regret optimal purposive sampling scheme under the conditions of Theorem 1. Since the k-median problem is known to be NP-hard (Kariv and Hakimi, 1979; Megiddo and Supowit, 1984; Cohen-Addad, De Mesmay, Rotenberg, and Roytman, 2018), there is no algorithm whose computational time scales polynomially in the problem's inputs (d, k, n).

However, from a practical perspective, it is also known that the k-median problem admits a linear integer program formulation (Williamson and Shmoys, 2011, Chapter 7.7, p. 185), and is routinely solved with off-the-shelf algorithms such as different versions of the branch-and-bound method (Bertsimas and Weismantel, 2005, Chapter 11.1). In addition, different branch-and-bound algorithms either find a solution with provable optimality or, if stopped early, generate a report on the suboptimality of the solution found (Bertsimas, King, and Mazumder, 2016). As we explain later, in our application with $n = 41$, $d = 13$, any problem for $k \in \{1, \ldots, 10\}$ can be solved to provable optimality in just a few seconds using a personal laptop (see Figure 5 in Section 5).

Related Literature: Our results build on two recent papers that presented novel purposive sampling strategies to select experimental sites so as to optimize external validity. Gechter et al. (2024) present an elegant decision-theoretic approach that frames external validity as a policy problem and—under the assumption that the policy maker has a priori information about the effects of the new policy across sites—recommend a Bayesian approach for choosing where to experiment. Egami and Lee (2024) use the principle behind synthetic control (Abadie, Diamond, and Hainmueller, 2010) to recommend what they call the *synthetic purposive sampling* of sites; specifically, they select good "donors," defined as sites whose weighted average of observed baseline characteristics are close to the characteristics of the sites of interest. It is important to note that our decision-theoretic framework leans heavily on Gechter et al. (2024), but also that the sites selected through the k-median problem can be interpreted as a degenerate synthetic purposive sampling strategy, whereby each unit's associated synthetic unit is just its nearest neighbor.

This paper also contributes to the literature arguing that a "modern, decision-theoretic framework can help clarify important practical questions of experimental design" (Banerjee, Chassang, and Snowberg, 2017). Although decision-theoretic approaches to external validity are recent, a large body of work used statistical decision theory to analyze other aspects of experimental design, such as sample size determination (Raiffa and Schlaifer, 1961; Manski and Tetenov, 2016, 2019; Azevedo, Deng, Montiel Olea, Rao, and Weyl, 2020; Azevedo, Mao, Montiel Olea, and Velez, 2023; Hu, Zhu, Brunskill, and Wager, 2024). Finally, our notion of external validity is conceptually related to several areas of research in econometrics, machine learning, and statistics such as domain adaptation (Mansour, Mohri, and Rostamizadeh, 2009; Ben-David, Blitzer, Crammer, Kulesza, Pereira, and Vaughan, 2010), distributional shifts (Duchi and Namkoong, 2021; Sugiyama, Krauledat, and Müller, 2007; Adjaho and Christensen, 2022), learning under biased sampling (Sahoo, Lei, and Wager, 2022), and cross-domain transfer estimation and performance (Andrews, Fudenberg, Liang, and Wu, 2022; Menzel, 2023). To the best of our knowledge, none of these papers contain decision-theoretic analyses of *where* to experiment.

OUTLINE: This paper is organized as follows. Section 2 introduces the formal framework. Section 3 presents the main result of this paper (Theorem 1) and explains how the k-median problem arises naturally when trying to incorporate external validity concerns into the decision problem of where to experiment. Section 4 presents the linear integer program formulation of the k-median problem. Section 5 presents two illustrative empirical applications. Section 6 considers extensions of the baseline model. Section 7 concludes. Proofs of the main results can be found in Appendix A. Additional results are collected in the Supplementary Appendix.

2 Setting up the Decision Problem

2.1 Notation

A policy maker considers a set of $S \in \mathbb{N}$ candidate sites to evaluate and, eventually, implement a new policy of interest. The sites are indexed by $s \in \mathcal{S} \equiv \{1, \ldots, S\}$. In order to accommodate situations in which the policy maker is not necessarily able to experiment in all the candidate sites, we assume there is a nonempty subset $\mathcal{S}_E \subseteq \mathcal{S}$ of what we term *experimental* sites.[1] Throughout the paper, and to avoid a trivial instance of the site selection problem, we assume that there are at least two experimental sites (i.e., $\text{card}(\mathcal{S}_E) \geq 2$). It is also possible that institutional restrictions preclude the eventual implementation of the policy of interest in all of the candidate sites. Thus, it will be convenient to denote by $\mathcal{S}_P \subseteq \mathcal{S}$ the nonempty set of *policy* or *policy-relevant* sites.

In principle—and in order to accommodate different scenarios that could arise in empirical work—we allow for the case in which there is overlap between experimental and policy sites (i.e., $\mathcal{S}_E \cap \mathcal{S}_P \neq \emptyset$). We note, however, that the case in which these two sets are disjoint (i.e., $\mathcal{S}_E \cap \mathcal{S}_P = \emptyset$) will be of particular interest. This is because when there is no overlap between experimental and policy sites, the extrapolation problem faced by the policy maker becomes evident: a policy decision needs to be made in sites where there has never been a previous experimental evaluation of the policy of interest.

For each site, the policy maker observes a vector of *site characteristics* $X_s \in \mathbb{R}^d$ that may affect the treatment effect. Thus, we allow for *treatment effect heterogeneity* across sites, but we restrict this heterogeneity by assuming that it depends on observable characteristics. Specifically, let the function $\tau : \mathbb{R}^d \to \mathbb{R}$ define conditional (on X) average treatment effects. We posit that any pair

[1] As discussed in Allcott (2015), there are often systematic reasons determining the eligibility of certain sites for experimentation. For example, in microfinance RCT studies, experiments often require large sample sizes and well-managed microfinance institutions (MFIs), characteristics more commonly found in older and larger institutions. To qualify for clinical trials involving a new surgical procedure, hospitals and surgeons need to have both experience in the procedure and a history of low mortality rates.

of sites with "similar" observed characteristics also have "similar" treatment effects, formally by assuming that τ is a Lipschitz function (with respect to the Euclidean norm) with known constant C. We formally state this assumption in Section 2.4. It is not innocuous, but we will argue that it can be replaced by other continuity-like conditions (such as Hölder continuity) while maintaining the main message of the paper.[2] We let $\text{Lip}_C(\mathbb{R}^d)$ denote the space of all Lipschitz functions from $\mathbb{R}^d$ to $\mathbb{R}$ with constant C.

2.2 Statistical Model for the Site Selection Problem

As in Gechter et al. (2024), the policy maker must choose a strict subset of experimental sites $\mathscr{S} \subset \mathcal{S}_E$.[3] In this paper, we focus on the case in which there is a restriction on the total number of experimental sites that the policy maker can select. That is, there is an integer $k \in \mathbb{N}$, $k < \text{card}(\mathcal{S}_E)$, such that $\mathscr{S}$ must belong to the set

$$\mathcal{A}(k) := \{\mathscr{S} \subset \mathcal{S}_E \mid \text{card}(\mathscr{S}) \leq k\}. \tag{1}$$

What we have in mind with this restriction is that the policy maker may not have enough administrative resources to run more than k experiments to evaluate the policy of interest. Our notation also allows for the possibility that the policy maker does not want to experiment at all.

If the policy maker decides to experiment in a nonempty set $\mathscr{S} \in \mathcal{A}(k)$ of cardinality $\text{card}(\mathscr{S}) \leq k$, then she will observe $\text{card}(\mathscr{S})$ treatment effect estimates. We collect these estimates in a vector of dimension $\text{card}(\mathscr{S})$. In a slight abuse of notation, let $\mathscr{S}_1 < \mathscr{S}_2 < \ldots < \mathscr{S}_{\text{card}(\mathscr{S})}$ denote the indices of the $\text{card}(\mathscr{S})$ experimental sites in $\mathscr{S}$. Letting $\widehat{\tau}_s$ denote the estimated treatment effect in site s, we can define the vector

$$\widehat{\tau}_{\mathscr{S}} := (\hat{\tau}_{\mathscr{S}_1}, \ldots, \hat{\tau}_{\mathscr{S}_{\text{card}(\mathscr{S})}})^\top. \tag{2}$$

Analogously, we can denote the vector of true treatment effects for the experimental sites in $\mathscr{S}$ as

$$\tau_{\mathscr{S}} := (\tau(X_{\mathscr{S}_1}), \ldots, \tau(X_{\mathscr{S}_{\text{card}(\mathscr{S})}}))^\top. \tag{3}$$

[2] In Section 6.3, we discuss how our results extend to other functional classes that could be used to restrict treatment effect heterogeneity, such as collections of functions that belong to a convex and centrosymmetric space. Such restrictions were recently used in the econometrics literature to analyze estimation, inference, and other decision problems that arise in a nonparametric regression setup (Yata, 2021; Armstrong and Kolesár, 2018).

[3] We require $\mathscr{S}$ to be a strict subset of $\mathcal{S}_E$ because if we allow the policy maker to experiment in all sites, and there is no cost of experimentation that varies at the site level, then there is no site selection problem. We consider the case in which $\mathscr{S}$ is allowed to equal $\mathcal{S}_E$ in Section 6.1.

We assume that the treatment effect estimators obtained in each site are normally (and independently) distributed around the vector of true effects:

$$\widehat{\tau}_{\mathscr{S}} \sim \mathcal{N}_{\text{card}(\mathscr{S})}\left(\tau_{\mathscr{S}}\,,\ \Sigma_{\mathscr{S}}\right),\ \text{where } \Sigma_{\mathscr{S}} := \text{diag}(\sigma^2_{\mathscr{S}_1}, \ldots, \sigma^2_{\mathscr{S}_{\text{card}(\mathscr{S})}}). \tag{4}$$

Following Gechter et al. (2024), treat $\Sigma_{\mathscr{S}}$ as known.

The normality assumption in (4) is unlikely to hold exactly; however, it is common to assume that treatment effect estimates from randomized controlled trials are asymptotically normal with asymptotic variances that can be estimated consistently. Treating the limiting normal model as an exact finite-sample statistical model then eases exposition and allows us to focus on the core features of the site selection problem. Indeed, working directly with such a limiting model is common in applications of statistical decision theory to econometrics; see Müller (2011) and the references therein for theoretical support and applications in the context of testing problems and Ishihara and Kitagawa (2021), Stoye (2012), or Tetenov (2012) for precedents in closely related work. Gechter et al. (2024) use the same statistical model, but our parameter space has a more specific form as treatment effects are controlled by the Lipschitz function τ.

After observing $\widehat{\tau}_{\mathscr{S}}$, the policy maker chooses an *action* $a_s \in [0,1]$ at each policy-relevant site $s \in \mathcal{S}_P$. We interpret this action as the proportion of a population in the site that will be randomly assigned to the new policy. Thus, $a_s = 1$ means that everyone in site s is exposed to the new policy, and $a_s = 0$ means that the status quo at the site is preserved. Under this interpretation, $a_s = .5$ means that 50% of the population at site s will be exposed at random to the new policy; however, the formal development equally applies to either individual or population-level randomization. Our interpretation abstracts from integer issues arising with small populations.

Thus, we can define a *treatment rule* T as a (measurable) function $T : \mathbb{R}^{\text{card}(\mathscr{S})} \to [0,1]^{\text{card}(\mathcal{S}_P)}$ that maps experimental outcomes to (possibly) randomized policy actions in each of the policy-relevant sites. It will be sometimes convenient to use $T_s(\cdot)$ to denote the specific treatment rule for site $s \in \mathcal{S}_P$ implied by T and $\mathcal{T}_{\mathscr{S}}$ to denote the set of all treatment rules. Note that we index the treatment rules by the selected experimental sites, $\mathscr{S}$, to be explicit about the fact that the data used to inform policy will vary depending on the choice of $\mathscr{S}$. We call $T \in \mathcal{T}_{\mathscr{S}}$ *nonrandomized* if for every $s \in \mathcal{S}_P$ we have $T_s(z) \in \{0,1\}$ for (Lebesgue) almost every $z \in \mathbb{R}^{\text{card}(\mathscr{S})}$. Otherwise, we say that the rule is *randomized*. We will focus on decision rules that belong to the set

$$\mathcal{T}^{1/2}_{\mathscr{S}} := \left\{T \in \mathcal{T}_{\mathscr{S}} \mid \text{ for any diagonal matrix } \Sigma, \text{ and } \forall s \in \mathcal{S}_P, \mathbb{E}[T_s(U)] = 1/2, U \sim \mathcal{N}_{\text{card}(\mathscr{S})}(\mathbf{0}, \Sigma)\right\}.$$

These are treatment rules in which the ex-ante probability of implementing the policy is 50%

whenever the true treatment effects at the sites experimented on are zero. For the moment, and for the sake of exposition, we assume that there is no cost of experimentation. While this assumption is clearly unrealistic, we later show that the main conclusions of our analysis are robust to adding fixed costs to the objective function of the k-median problem.

2.3 Welfare and Regret

Let $\#\mathcal{S}_P := \text{card}(\mathcal{S}_P)$. We assume that the *welfare* of a decision rule T, given that sites $\mathscr{S}$ are selected for experimentation, corresponds to the average expected welfare across policy-relevant sites:

$$\mathcal{W}(T, \mathscr{S}, \tau) := \frac{1}{\#\mathcal{S}_P} \sum_{s \in \mathcal{S}_P} \tau(X_s) \mathbb{E}_{\tau_{\mathscr{S}}}[T_s(\widehat{\tau}_{\mathscr{S}})], \tag{5}$$

where $\mathbb{E}_{\tau_{\mathscr{S}}}[T_s(\widehat{\tau}_{\mathscr{S}})]$ means that the expectation is taken assuming $\widehat{\tau}_{\mathscr{S}} \sim \mathcal{N}_{\text{card}(\mathscr{S})}(\tau_{\mathscr{S}}, \Sigma_{\mathscr{S}})$, and $\tau(X_s)$ is the true treatment effect at site s.

The *regret* of policy $(T, \mathscr{S})$ can then be shown to equal

$$\mathcal{R}(T, \mathscr{S}, \tau) := \frac{1}{\#\mathcal{S}_P} \sum_{s \in \mathcal{S}_P} \tau(X_s) \left(\mathbf{1}\{\tau(X_s) \geq 0\} - \mathbb{E}_{\tau_{\mathscr{S}}}[T_s(\widehat{\tau}_{\mathscr{S}})] \right). \tag{6}$$

Our focus will be on finding the purposive sampling scheme that minimizes worst-case regret.

Definition 1 (*MMR optimal purposive sampling scheme and treatment rule*)**.** The pair $(T^*, \mathscr{S}^*) \in \mathcal{T}_{\mathscr{S}}^{1/2} \times \mathcal{A}(k)$ is minimax-regret (MMR) optimal among all purposive sampling schemes and treatment rules if

$$\sup_{\tau \in \text{Lip}_C(\mathbb{R}^d)} \mathcal{R}(T^*, \mathscr{S}^*, \tau) = \inf_{\mathscr{S} \in \mathcal{A}(k), T \in \mathcal{T}_{\mathscr{S}}^{1/2}} \sup_{\tau \in \text{Lip}_C(\mathbb{R}^d)} \mathcal{R}(T, \mathscr{S}, \tau). \tag{7}$$

Remark 1. In the standard definition of minimax-regret optimality, it is common to allow the decision maker to select *randomized* decision rules. Definition 1 implies an asymmetric treatment of randomization: While we allow the policy maker to randomize policy implementation choices, we are restricting her to pick the experimental sites in a deterministic fashion. The nonrandom selection of experimental sites—referred to broadly in the literature as *purposive sampling* (Cook et al., 2002, p. 511)—is common in practice; see the discussion about purposive sampling in Egami and Lee (2024). In Section 6, we discuss the challenges we encountered in trying to allow for the random selection of experimental sites. □

Remark 2. It will sometimes be convenient to rewrite the right-hand side of (7) as

$$\inf_{\mathscr{S}\in\mathcal{A}(k)}\left(\inf_{T\in\mathcal{T}_{\mathscr{S}}^{1/2}}\sup_{\tau\in\mathrm{Lip}_C(\mathbb{R}^d)}\mathcal{R}(T,\mathscr{S},\tau)\right).$$

This suggests that, conceptually, the solutions of the minimax-regret problem can be done in two steps. First, one can analyze the problem of policy implementation given the experimental outcomes at sites $\mathscr{S}$. Then, one can optimize over the sites where to experiment.

This distinction is helpful because the first step is related to the "evidence aggregation" problem analyzed in Ishihara and Kitagawa (2021). As mentioned in the Introduction, we can leverage recent results in Yata (2021) and Montiel Olea et al. (2023) to better understand the minimax-regret solutions to this problem. A key difference from Ishihara and Kitagawa (2021) is that we typically have more than one policy-relevant site. □

2.4 Main Assumptions

We make the following assumptions. To begin, treatment effects vary continuously with the site-level covariates:

Assumption 1. τ is a Lipschitz function (with respect to the Euclidean norm) with known constant C. That is, for any $x, x' \in \mathbb{R}^d$, $|\tau(x) - \tau(x')| \leq C\|x - x'\|$, where $\|\cdot\|$ denotes the Euclidean norm.

We also impose a regularity condition on site-level covariates. That is, we assume that all observed covariates are different:

Assumption 2. $X_s \neq X_{s'}\ \forall\ s, s' \in \mathcal{S}$.

Even if this were not the case in raw data, one would presumably want to induce it by adding site fixed effects.[4]

Finally, for each policy site $s \in \mathcal{S}_P$ and every $\mathscr{S} \in \mathcal{A}(k)$, denote by $N_{\mathscr{S}}(s) \in \mathscr{S}$ its *nearest neighbor* in $\mathscr{S}$ (or the nearest neighbor with the smallest index in case of multiplicity). That is, for every $s \in \mathcal{S}_P$:

$$\|X_s - X_{N_{\mathscr{S}}(s)}\| \leq \|X_s - X_{s'}\|\ \ \forall s' \in \mathscr{S}.$$

[4]In section 6.3, we discuss how Assumption 1 can be modified to accommodate other general distance measures, as well as some forms of unobserved treatment heterogeneity (for which Assumption 2 may be dropped).

3 Main Result

This section presents our main result: The problem of finding the purposive sampling scheme that minimizes the worst-case regret is approximately equal—and sometimes equal—to solving a k-median problem.

In order to derive this, first recall the definition of the k-median problem.

Definition 2 (*k-median problem*)**.** We say that a purposive sampling scheme, $\mathscr{S} \in \mathcal{A}(k)$, solves the k-median problem if it solves

$$\inf_{\mathscr{S}\in\mathcal{A}(k)} \sum_{s\in\mathcal{S}_P} \|X_s - X_{N_{\mathscr{S}}(s)}\|. \tag{8}$$

In Definition 2, the policy-relevant sites (with indexes in $\mathcal{S}_P$) are *clients* and the experimental sites (with indexes in $\mathcal{S}_E$) are *facilities*. The connection cost between a facility i and a client j is the Euclidean distance $\|X_j - X_i\|$. Since the goal of the k-median problem is to choose the k facilities that minimize connection cost, each client $s \in \mathcal{S}_P$ gets connected to the facility that is closest among those in $\mathscr{S}$. Hence, the term $X_{N_{\mathscr{S}}(s)}$ appears in Definition 2. Finally, Equation (8) can be also written as

$$\inf_{\mathscr{S}\in\mathcal{A}(k)} \left(\sum_{s\in\mathcal{S}_P\cap\mathscr{S}} \|X_s - X_{N_{\mathscr{S}}(s)}\| + \sum_{s\in\mathcal{S}_P\setminus\mathscr{S}} \|X_s - X_{N_{\mathscr{S}}(s)}\| \right) = \inf_{\mathscr{S}\in\mathcal{A}(k)} \sum_{s\in\mathcal{S}_P\setminus\mathscr{S}} \|X_s - X_{N_{\mathscr{S}}(s)}\|, \tag{9}$$

where the equality follows from the fact that for each facility $s \in \mathcal{S}_P$ that is also a client $s \in \mathscr{S}$, connection cost becomes zero.

There is no scope for interesting analysis, or indeed for justification of experiments, unless noise in the signal is bounded. Therefore, define $C^* := \max_{\mathscr{S}\in\mathcal{A}(k), s\in\mathcal{S}_P\setminus\mathscr{S}} \left\{ \sqrt{\frac{\pi}{2}} \frac{\sigma_{N_{\mathscr{S}}(s)}}{\left\|X_s - X_{N_{\mathscr{S}}(s)}\right\|} \right\} < \infty$. Let σ_E denote the largest standard deviation among the potential experimental sites. Our main theorem is the following:

Theorem 1. Suppose Assumptions 1-2 hold. If $C > C^*$, then for any $\mathscr{S} \in \mathcal{A}(k)$:

$$\left| \left(\inf_{T\in\mathcal{T}_{\mathscr{S}}^{1/2}} \sup_{\tau\in \mathrm{Lip}_C(\mathbb{R}^d)} \mathcal{R}(T, \mathscr{S}, \tau) \right) - \frac{C}{2} \frac{1}{\#\mathcal{S}_P} \cdot \left(\sum_{s\in\mathcal{S}_P\setminus\mathscr{S}} \|X_s - X_{N_{\mathscr{S}}(s)}\| \right) \right| \tag{10}$$

is at most

$$B \cdot \sigma_E \cdot \frac{\min\{\#\left(\mathcal{S}_E \cap \mathcal{S}_P\right), k\}}{\#\mathcal{S}_P}, \tag{11}$$

where $B \equiv \arg\max_{z \geq 0} z\Phi(-z)$.

Proof. Follows directly from Lemma 1 and 2 below. □

Remark 3. The left-hand expression in (10), i.e., $\inf_{T \in \mathcal{T}_{\mathscr{S}}^{1/2}} \sup_{\tau \in \text{Lip}_C(\mathbb{R}^d)} \mathcal{R}(T, \mathscr{S}, \tau)$, is the *worst-case regret* of the purposive sampling scheme $\mathscr{S}$ assuming optimality of subsequent treatment choice $T_{\mathscr{S}}$. Theorem 1 thus shows that the (optimized) worst-case regret of any purposive sampling scheme can be uniformly approximated (after scaling by the factor $C/(2\#\mathcal{S}_P)$) by the objective function of the *k-median* problem in Definition 2—which is equivalent to the objective function of the problem in the right-hand side of Equation (9). □

Remark 4. When the candidate sites for experimentation ($\mathcal{S}_E$) and the policy-relevant sites ($\mathcal{S}_P$) are disjoint, then for any $\mathscr{S} \in \mathcal{A}(k)$:

$$\left(\inf_{T \in \mathcal{T}_{\mathscr{S}}^{1/2}} \sup_{\tau \in \text{Lip}_C(\mathbb{R}^d)} \mathcal{R}(T, \mathscr{S}, \tau) \right) = \frac{C}{2} \frac{1}{\#\mathcal{S}_P} \cdot \left(\sum_{s \in \mathcal{S}_P \setminus \mathscr{S}} \|X_s - X_{N_{\mathscr{S}}(s)}\| \right). \tag{12}$$

Consequently, any solution of the k-median problem is *a fortiori* an *exact* minimax-regret solution for the site selection problem. This follows from the fact that if a purposive sampling scheme, $\mathscr{S}^*$, minimizes the right-hand side of Equation (12), then it will automatically minimize the left-hand side, which, by Remark 2, defines a minimax-regret optimal purposive sampling scheme. □

Remark 5. Theorem 1 implies that even in the case where $\mathcal{S}_E \cap \mathcal{S}_P \neq \emptyset$,

$$\left| \inf_{\mathscr{S} \in \mathcal{A}(k)} \left(\inf_{T \in \mathcal{T}_{\mathscr{S}}^{1/2}} \sup_{\tau \in \text{Lip}_C(\mathbb{R}^d)} \mathcal{R}(T, \mathscr{S}, \tau) \right) - \frac{C}{2} \frac{1}{\#\mathcal{S}_P} \cdot \inf_{\mathscr{S} \in \mathcal{A}(k)} \left(\sum_{s \in \mathcal{S}_P \setminus \mathscr{S}} \|X_s - X_{N_{\mathscr{S}}(s)}\| \right) \right| \tag{13}$$

is at most equal to (11). Thus, the quality of this approximation improves as either the number of policy-relevant sites increases or the potential experiments conducted become more precise (σ_E becomes smaller). The approximation deteriorates as k increases. This happens because in approximating the worst-case regret of a purposive sampling scheme, we ignored the regret associated with sites selected for experimentation; we further discuss this below when we explain Lemma 1. This means that, as k increases, more sites will be ignored in our approximation and the upper bound will become looser. □

Remark 6. The tight connection to the k-median problem relies on the Lipschitz constant to be large enough. In other words, the k-median problem only has a decision-theoretic justification when the treatment effect heterogeneity across sites accommodated in the parameter space is substantial. □

The proof of Theorem 1 follows from the two Lemmas below, which bound the site selection problem's minimax-regret value from above and below, respectively.

Lemma 1 (Upper bound). Suppose Assumptions 1-2 hold. For every $\mathscr{S} \in \mathcal{A}(k)$, there exists a constant $C(\mathscr{S})$ such that, if $C > C(\mathscr{S})$, then

$$\inf_{T \in \mathcal{T}_{\mathscr{S}}^{1/2}} \sup_{\tau \in \mathrm{Lip}_C(\mathbb{R}^d)} \mathcal{R}(T, \mathscr{S}, \tau) \leq \left(\frac{B}{\#\mathcal{S}_P} \sum_{s \in \mathscr{S} \cap \mathcal{S}_P} \sigma_s \right) + \left(\frac{C}{2} \frac{1}{\#\mathcal{S}_P} \sum_{s \in \mathcal{S}_P \setminus \mathscr{S}} \|X_s - X_{N_{\mathscr{S}}(s)}\| \right), \tag{14}$$

where $B \equiv \arg\max_{z \geq 0} z\Phi(-z)$.

Proof. See Appendix A.1. □

Remark 7. Lemma 1 implies that the minimax-regret value of the site selection problem in Definition 1 can be upper bounded by the solution of the *uncapacitated k-facility location problem*.[5] Just as in the k-median problem, there is a set of facilities $\mathcal{F}$ and a set of clients (or cities) $\mathcal{C}$. Now we assume that there is an opening cost $c_i \in \mathbb{R}_+$ associated with each facility $i \in \mathcal{F}$. The connection cost between facilities and clients is as before. The goal is to open at most k facilities and connect each client to one facility so that total cost is minimized. Thus, the problem

$$\inf_{\mathscr{S} \in \mathcal{A}(k)} \left(\left(\frac{B}{\#\mathcal{S}_P} \sum_{s \in \mathscr{S} \cap \mathcal{S}_P} \sigma_s \right) + \left(\frac{C}{2} \frac{1}{\#\mathcal{S}_P} \sum_{s \in \mathcal{S}_P \setminus \mathscr{S}} \|X_s - X_{N_{\mathscr{S}}(s)}\| \right) \right) \tag{15}$$

can be viewed as a k-facility location problem. Just as before, the set of facilities is $\mathcal{S}_E$ and the set of clients is $\mathcal{S}_P$. The connection cost between sites $s \in \mathcal{S}_E$ and $s' \in \mathcal{S}_p$ is $(C/(2\#\mathcal{S}_P))\|X_s - X_{s'}\|$. The opening cost for any facility $s \in \mathcal{S}_E$ is $(B/\#\mathcal{S}_P)\sigma_s$. Minimizing the upper bound in Lemma 1 is thus equivalent to solving the k-facility location problem. As shown in Appendix A.1, the upper bound arises by bounding the worst-case sum of regrets across sites by the corresponding sum of worst-case regrets. For selected sites that are both facilities and clients, the worst-case regret is obtained by solving a point-identified treatment choice problem. For policy-relevant sites where no experiment was conducted, it is obtained by solving the partially-identified treatment choice problem in Ishihara and Kitagawa (2021). When $\mathcal{S}_E$ and $\mathcal{S}_P$ are disjoint, the first component of the upper bound vanishes, and the bound becomes proportional to the solution of the k-median problem. □

[5] Uncapacitated means that there is no capacity constraint on the number of clients that each facility can accommodate. See (Williamson and Shmoys, 2011, Chapter 4.5) and Zhang (2007).

Lemma 2 (Lower bound). Suppose Assumptions 1-2 hold. For every $\mathscr{S} \in \mathcal{A}(k)$:

$$\inf_{T \in \mathcal{T}_{\mathscr{S}}^{1/2}} \sup_{\tau \in \mathrm{Lip}_C(\mathbb{R}^d)} \mathcal{R}(T, \mathscr{S}, \tau) \geq \left(\frac{C}{2} \frac{1}{\#\mathcal{S}_P} \sum_{s \in \mathcal{S}_P \setminus \mathscr{S}} \|X_s - X_{N_{\mathscr{S}}(s)}\| \right). \tag{16}$$

Proof. See Appendix A.2. □

Remark 8. Lemma 2 implies that the minimax-regret value of the site selection problem in Definition 1 can be lower bounded by the solution of the *k-median problem*:

$$\inf_{\mathscr{S} \in \mathcal{A}(k)} \left(\frac{C}{2} \frac{1}{\#\mathcal{S}_P} \sum_{s \in \mathcal{S}_P \setminus \mathscr{S}} \|X_s - X_{N_{\mathscr{S}}(s)}\| \right). \tag{17}$$

Theorem 1 thus follows by noting that the upper and lower bound match up to the *opening costs* of the facilities. As noted before, the lower and upper bound match when $\mathcal{S}_E \cap \mathcal{S}_P = \emptyset$. □

Remark 9. While the computation of an exact minimax-regret sampling scheme (whether purposive or randomized) appears intractable, our ability to provide nontrivial upper bounds on worst-case regret does not hinge on C being large enough. Indeed, one could always compute such a bound by (i) forcing the treatment decision T_s to only depend on site $N_{\mathscr{S}}(s)$ and (ii) computing worst-case expected regret separately across sites, ignoring that the regret-maximizing parameter configurations may be mutually inconsistent across sites. Both manipulations increase regret and therefore define an upper bound. This bound is easy to compute for any given sampling scheme because the induced minimax-regret treatment choice problem was solved in Stoye (2012). For C large enough, it will coincide with the bound in Theorem 1. The caveat is that this bound is nonlinear; hence, while it is easy to compute once, it does not generate a tight connection to the k-median problem and will be hard to optimize for large instances of the site selection problem. □

4 Integer Programming and the k-median Problem

It is well known (Williamson and Shmoys, 2011, Chapter 7.7, p. 185) that the k-median problem in (17) can be formulated as the following linear integer program:

$$\min_{\{y_i, x_{i,j}\}_{i \in \mathcal{S}_E, j \in \mathcal{S}_P}} \sum_{i \in \mathcal{S}_E, j \in \mathcal{S}_P} x_{i,j} \cdot c(j,i) \tag{18}$$

$$\text{such that } \sum_{i \in \mathcal{S}_E} x_{i,j} = 1, \quad \forall j \in \mathcal{S}_P, \tag{19}$$

$$\sum_{i \in \mathcal{S}_E} y_i \leq k, \tag{20}$$

$$0 \leq x_{i,j} \leq y_i, \quad i \in \mathcal{S}_E, j \in \mathcal{S}_P, \tag{21}$$

$$y_i \in \{0,1\}, x_{i,j} \in \{0,1\}, \quad i \in \mathcal{S}_E, j \in \mathcal{S}_P. \tag{22}$$

In this formulation, the choice variables are the $\{0,1\}$-valued variables y_i and $x_{i,j}$, for $i \in \mathcal{S}_E, j \in \mathcal{S}_P$. The variable y_i indicates whether site i is chosen for experimentation, and $x_{i,j}$ indicates whether experimental site i is the nearest neighbor of policy site j. The total number of sites chosen for experimentation is restricted to be no greater than k, and site i can only be the nearest neighbor of policy site j if site i is chosen for experimentation. The connection cost between a facility i and a client j—denoted $c(j,i)$—is given by $\|X_i - X_j\|$.[6] It is worth highlighting that the program above is linear in the choice variables $x_{i,j}$ and y_i.

A major advantage of the integer programming formulation of the k-median problem is that most ecosystems for scientific computing offer different solvers for linear and nonlinear integer programs. For the applications in this paper, we use the `MIP` solver in `Gurobi` (Gurobi Optimization, LLC, 2023) through the Python package `gurobipy`. The `Gurobi` software is highly optimized, especially for large-scale problems, and it offers an academic license. Even though the scale of the applications presented in Section 5 is not large enough for the efficiency advantages of `Gurobi` to become salient over other solvers, we wanted to showcase its ease of use. It also integrates seamlessly with Google Colab, providing us a way to build self-contained and reproducible examples.[7]

The `MIP` solver uses a version of the branch-and-bound algorithm; see Bertsimas and Weismantel (2005), Chapter 11.1, for a general description of this algorithm. Broadly speaking, this algorithm works by first finding the solution to the linear programming (LP) relaxation of the original integer

[6] The connection cost in the objective function of the integer program differs from the connection cost in the k-median problem in (17) by a constant factor $C/(2\#\mathcal{S}_P)$, which does not affect the solution of the k-median problem. Solving the linear integer program described above is equivalent to solving (17).

[7] https://colab.research.google.com/

problem. This is known as the *relaxation* step. Then, the problem is split into two sub-problems (*branching*), according to the integer constraints on the solution. This method then gives bounds on the integer solution, and the algorithm is applied recursively until the lower bound and the upper bound converge up to a tolerance parameter. The recursion creates *nodes*, and there are some strategies to determine which nodes should be explored further; for example, nodes that have integer solutions do not require any more branching.

`Gurobi` implements several additional steps that help the branch-and-bound algorithm be more efficient.[8] The *presolve* step reduces the number of effective constraints of the problem by checking if the integer requirement can eliminate some of them. As the name suggests, it is performed before the start of the branch-and-bound algorithm. *Cutting planes* tightens the feasible region by adding linear inequalities to eliminate fractional solutions; it is performed during the branch-and-bound process, and for this reason the algorithm used by `Gurobi` can be referred to as a version of a "branch-and-cut" algorithm. Finally, the `MIP` solver implements several *heuristics*, for example by rounding the component of a solution that is closest to an integer, fixing it, and hoping the other components will converge to integers quicker.

In principle, one can always solve the k-median problem in (17) by enumerating all possible size k subsets of $\mathcal{S}_E$. Such an algorithm runs in time proportional to

$$\binom{\#\mathcal{S}_E}{k} \cdot \#\mathcal{S}_P \cdot k \cdot d$$

time and therefore scales poorly when the term "$\#\mathcal{S}_E$ choose k" is large. However, we are able to evaluate the performance of our preferred solver by brute-force solving smaller but not trivial instances of the problem, notably the entire application in Section 5.2.

Figure 9 in Appendix B.1 shows an example of the output obtained after using the `MIP` solver in `Gurobi` to solve the linear integer program for the application in Section 5.2. In this example, the scale of the problem is given by $\#\mathcal{S}_E = \#\mathcal{S}_P = 15, k = 6$, and $d = 8$. We defer the details of the application to Section 5.2.

While we do not use it in this paper, we finally note that there is a large literature studying efficient (polynomial-time) *approximation* algorithms for the k-median problem, going back to seminal work of Charikar, Guha, Tardos, and Shmoys (2002). A basic idea in these algorithms is to consider a linear programming relaxation of the integer program associated with the k-median problem (Williamson and Shmoys, 2011, Chapter 7.7). Even though the scale of the problems analyzed in this paper does not require the implementation of any of these algorithms, there are several pa-

[8] See https://www.gurobi.com/resources/mixed-integer-programming-mip-a-primer-on-the-basics/.

pers that present theoretical performance guarantees for them; see for example Cohen-Addad et al. (2022) and also the references in Cohen-Addad et al. (2018). Finally, it is worth mentioning that when k is fixed (and not viewed as part of the problem's input), there is an approximate algorithm that runs in time proportional to $n \cdot d$; see Kumar, Sabharwal, and Sen (2010). Such an algorithm could be useful when n and d are large and k is small.

5 Applications

5.1 Mobile Financial Services in Bangladesh

Lee, Morduch, Ravindran, Shonchoy, and Zaman (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 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. 61 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.

We ask: Do the findings in Lee et al. (2021) generalize to other migration corridors, i.e., combinations of origin and destination districts, in Bangladesh? Is the corridor selected by Lee et al. (2021) a good choice for a researcher who is concerned about external validity? 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". Figure 1 displays this corridor along with other common ones. The 41 migration corridors analyzed in Gechter et al. (2024) are depicted as dotted lines connecting an origin and a destination.[9]

[9]We thank Michael Gechter for gracefully sharing part of the dataset used in Gechter et al. (2024).

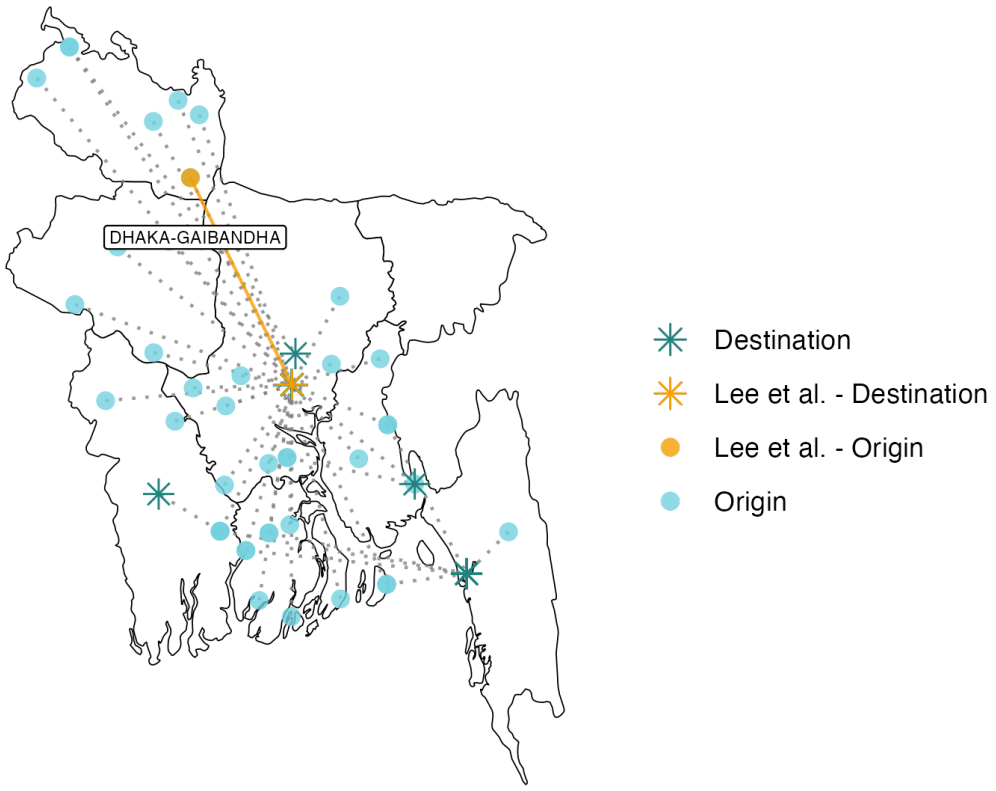

Figure 1: Bangladesh Migration Corridors

Notes: The map of Bangladesh with the origins of the migration corridors marked as light blue dots and the destinations marked as dark blue stars. Following the terminology used in Gechter et al. (2024), *origin* refers to a worker's home, and destination refers to where the worker migrates for work. The corridor where the experiment was originally implemented in Lee et al. (2021), Dhaka-Gaibandha, is highlighted in yellow.

In Lee et al.'s 2021 words, "[t]he particular nature of our sample potentially limits the external validity" of the analysis. In short, migration corridors differ in characteristics ranging from the distance between the origin and the destination to the cost of living and the average wages. Figure 2 presents a box plot of $d = 13$ (standardized) characteristics that Gechter et al. (2024) collected for each of the 41 migration corridors. We take these corridors to be our potential experimental and policy-relevant sites. That is, $\mathcal{S}_E = \mathcal{S}_P$, and $\text{card}(\mathcal{S}_E) = \text{card}(\mathcal{S}_P) = 41$. Below we present results for $k \in \{1, 2\}$.

Selected Site when $k = 1$: When $k = 1$, the selected site based on the k-median problem is Dhaka-Pabna. This is also the solution obtained by using the synthetic purposive sampling

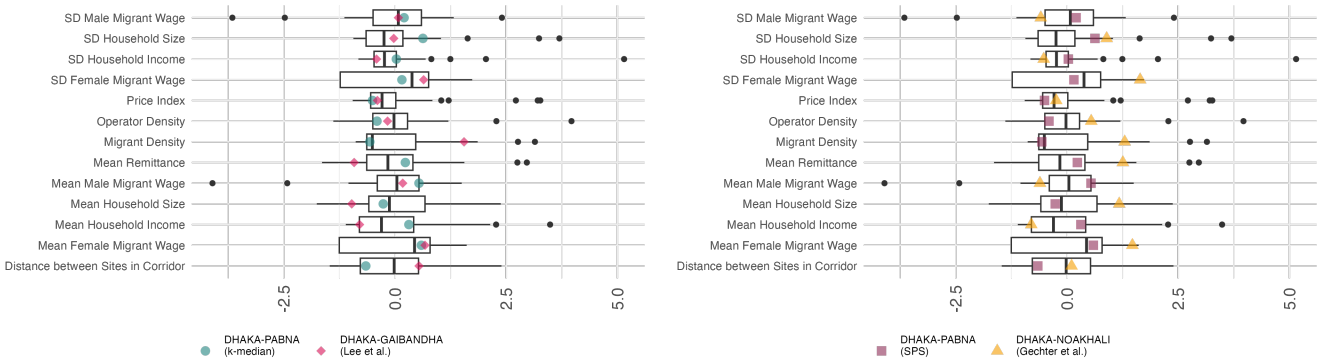

(a) k-median (b) Egami and Lee (2024) and Gechter et al. (2024)

Figure 2: Solution of the k-median Problem ($k = 1$) and Distribution of Site-Level Covariates

Notes: For each covariate, the box represents the interquartile range (IQR), the vertical black line represents the median, and the horizontal line shows the "theoretical minimum" (defined as $Q_1 - 1.5\text{IQR}$) and "theoretical maximum" (defined as $Q_3 + 1.5\text{IQR}$). Black dots are outliers, defined as observations that fall beyond the theoretical minimum and maximum. Each panel depicts the sites selected by the different approaches when $k = 1$.

approach (henceforth, SPS) in Egami and Lee (2024).

Figure 2a presents a visual comparison of the covariates of Dhaka-Pabna (blue circles) and Dhaka-Gaibandha (pink diamonds), the original site selected by Lee et al. (2021). The covariate values for Dhaka-Gaibandha are slightly outside the interquartile range for at least three covariates: migrant density, mean remittance, and mean household size. In comparison, all but one covariate value for Dhaka-Pabna are within the interquartile range. The figure also shows that two key covariates of Dhaka-Gaibandha are right at the edges of the interquartile range: distance between sites in the corridor (3rd quartile) and mean household income (1st quartile). One might conjecture that the effects of adopting a mobile banking technology to transfer money particularly depend on distance and household incomes, suggesting that the Dhaka-Gaibandha corridor may not be the most representative.[10] Interestingly, Dhaka-Pabna has opposite features: the distance between destination and origin in this corridor is short (1st quartile) and households are relatively better off in terms of income (3rd quartile). The use of the minimax criterion might explain why a corridor with these characteristics may be a good choice for extrapolating experimental results.

Figure 2b also presents the covariates of Dhaka-Noakhali (yellow triangles), the migration corridor selected by the Bayesian approach of Gechter et al. (2024).[11] We note that for 10 out of the

[10]Gechter et al. (2024) suggest that these qualities could explain the large treatment effects found by Lee et al. (2021).

[11]By construction, the solution of Gechter et al. (2024) depends on the choice of prior. The results herein reported are based on their preferred prior specification; see Section 5.3 p.p.23 in Gechter et al. (2024).

13 variables that control treatment effect heterogeneity, Dhaka-Noakhali has covariates that are typically outside the interquartile range.

Selected Sites when $k = 2$: When $k = 2$, the k-median solution is to pick Dhaka-Pabna and Dhaka-Pirojpur. The former also solved the k-median problem for $k = 1$.

Figure 3a presents the covariates of Dhaka-Pabna (filled, blue circle) and Dhaka-Pirojpur (hollow, blue circle). While four covariate values of Dhaka-Pirojpur (standard deviation of male migrant wage, standard deviation of household income, price index, and migrant density) are outside the interquartile range, they still appear more central than Dhaka-Gaibandha and Dhaka-Noakhali. However, relative to Dhaka-Pabna, the solution for $k = 2$ adds a considerably less central site; Figure 4 illustrates that this site is a good nearest neighbor for some sites that would not otherwise be well matched.

Figure 3b presents the solutions of Egami and Lee (2024) (squares) and Gechter et al. (2024) (triangles).[12] The solution of the k-median problem and synthetic purposive sampling are no longer the same. We also note that the two sites selected by synthetic purposive sampling differ from the site selected by this procedure when $k = 1$.

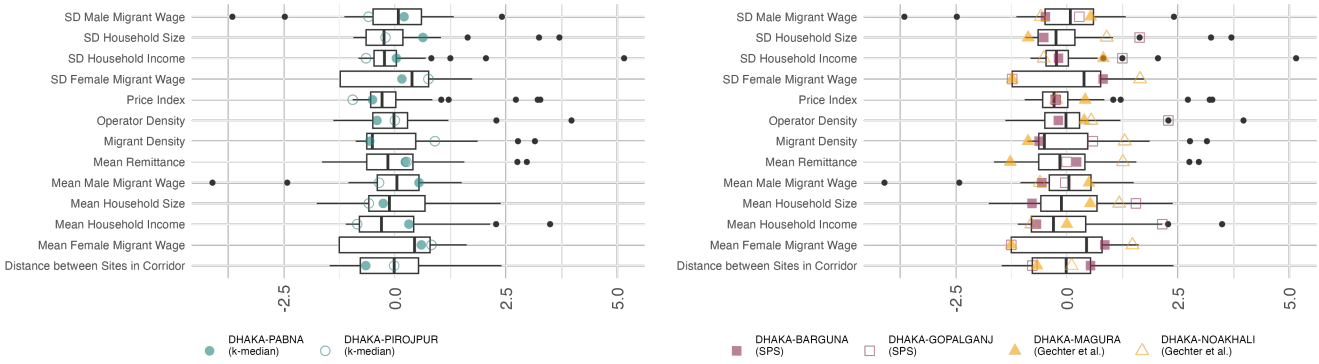

(a) k-median (b) Egami and Lee (2024) and Gechter et al. (2024)

Figure 3: Solution of the k-median Problem ($k = 2$) and Distribution of Site-level Covariates

Notes: Box plots for the distribution of covariates among migration corridors, constructed as explained in Figure 2. Each panel depicts the site selected by the different approaches when $k = 2$.

Figure 4 presents a simple visualization of optimal connections between sites under different optimality criteria. Note that the underlying scatter plots in both panels are the same; they visualize the location of corridors in (distance, household income)-space. However, Panel 4a indicates which

[12]Gechter et al. (2024) impose an additional constraint on purposive sampling schemes: they require the two migration corridors selected for experimentation to have origins in different divisions. We note that both the solutions of Egami and Lee (2024) and the k-median problem satisfy this constraint as well.

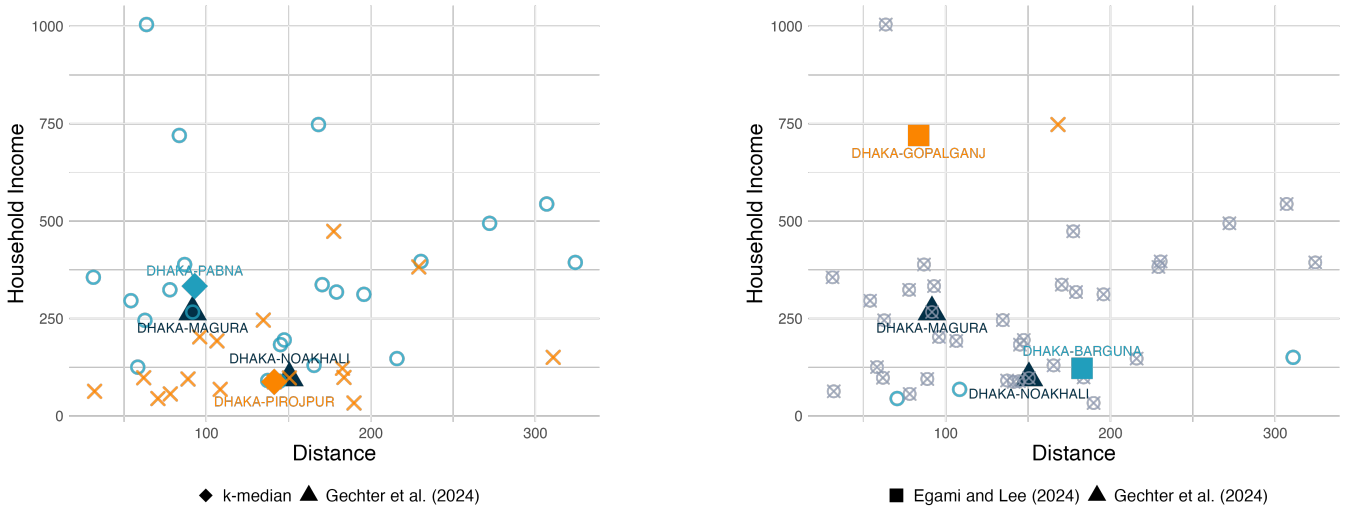

(a) k-median (b) Egami and Lee (2024)

Figure 4: Optimal Experimental Sites for $k = 2$ on a Two-dimensional Covariate Plane

Notes: Each point represents a destination-origin migration corridor (site) on the two-dimensional covariate plane, using two covariates: distance between the two ends of each corridor and average household income in the home district (where the average is taken over households with a reported migrant). The solutions of each method are indicated by the shapes of the solid dots. In Panel 4a, blue circles denote policy sites with Dhaka-Pabna as the nearest neighbor, and orange crosses represent those with Dhaka-Pirojpur as the nearest neighbor. Similarly, in Panel 4b, blue circles and orange crosses follow the same coding, indicating that these policy sites rely on a single experimental site for constructing their synthetic control. Additionally, gray cross-circles represent policy sites that use both experimental sites for their synthetic control.

sites were selected through solving the k-median problem and which of these sites any other site was matched to, where blue circles represent matches with Dhaka-Pabna and orange crosses represent matches with Dhaka-Pirojpur. Figure 4a also presents the sites selected by Gechter et al. (2024), i.e. Dhaka-Magura and Dhaka-Noakhali, both marked with dark blue triangles. They appear close to the sites selected by the k-median problem, but in terms of that problem's criterion function, they only rank 455 among 820 candidate solutions and miss the problem's optimal value by 17%. Panel 4b similarly visualizes the sites selected by Egami and Lee's (2024) synthetic purposive sampling approach, i.e. Dhaka-Gopalganj (orange square) and Dhaka-Barguna (blue square).[13] We also visualize how policy sites are matched with experimental sites: Orange crosses correspond to migration corridors that assign all of their weight to Dhaka-Gopalganj; blue circles correspond to migration corridors that assign all of their weight to Dhaka-Barguna; all other sites (gray crossed circles) assign strictly positive weights to both donor sites. This illustrates that these sampling schemes

[13]We generated this figure by using Egami and Lee's (2024) code on Gechter et al.'s (2024) data.

meaningfully differ. This would even be true if the synthetic purposive sampling approach were implemented but forcing degenerate (single donor) matches, because Egami and Lee's (2024) approach would then reduce to a k-mean problem, i.e. using squared Euclidean distance as connection cost.

Computational Costs: The above examples are small enough so that the k-median problem could easily be solved by brute-force enumeration of 41 and 820 candidate solutions, respectively. Needless to say, such an approach would not scale—for example, in this same application, $k = 10$ induces $1,121,099,408$ candidate solutions.

Indeed, Figure 5a compares time to solve the k-median problem for $k \in \{1, \ldots, 10\}$ using a) the integer program formulation of the k-median problem in Section 4 and b) brute-force enumeration. The `MIP` solver in `Gurobi` solves all instances of the problem to provable optimality in less than one second each. In contrast, brute-force enumeration takes approximately 5 hours for $k = 10$.[14]

One potential benefit of brute-force enumeration is that one can check for multiple solutions, which actually occurred at $k = 6$. In `Gurobi`, an ad hoc search could be conducted by modifying the random number seed or using the concurrent optimizer, but discovery would not be guaranteed.

Figure 5b reports the time needed to implement the synthetic purposive sampling approach of Egami and Lee (2024). This was done by using the `spsR` package with the option to use the `Gurobi` solver on the background. We consider it to be fast, taking less than 40 seconds for $k = 10$. However, the synthetic purposive sampling problem can be formulated as a quadratic mixed integer program, while k-median is linear. This is evident from the graph; the k-median problem is solved almost instantly for every k up to 10.

5.2 Multi-Country Survey Experiments in Europe

Our second application revisits a multi-country survey experiment originally conducted and analyzed in Naumann, F. Stöetzer, and Pietrantuono (2018) and discussed in Egami and Lee (2024). The question of interest is whether native-born inhabitants of a particular country are more supportive of immigration depending on whether the potential migrants are high-skilled or low-skilled. Naumann et al. (2018) carried out a survey experiment in 15 European countries: Austria, Belgium, Czech Republic, Denmark, Finland, France, Germany, Ireland, Netherlands, Norway, Slovenia, Spain, Sweden, Switzerland, and the United Kingdom. Respondents were native-born individuals and were randomly assigned to report their attitudes towards either high-skilled ("treatment") or

[14]This was run in a Windows XPS with 10 cores and 32GB of RAM using R. The code is parallelized, and to avoid memory issues when $k > 8$, it runs a C++ function in the background to evaluate one combination at a time. This ensures that we are giving brute-force enumeration the best chance of success.

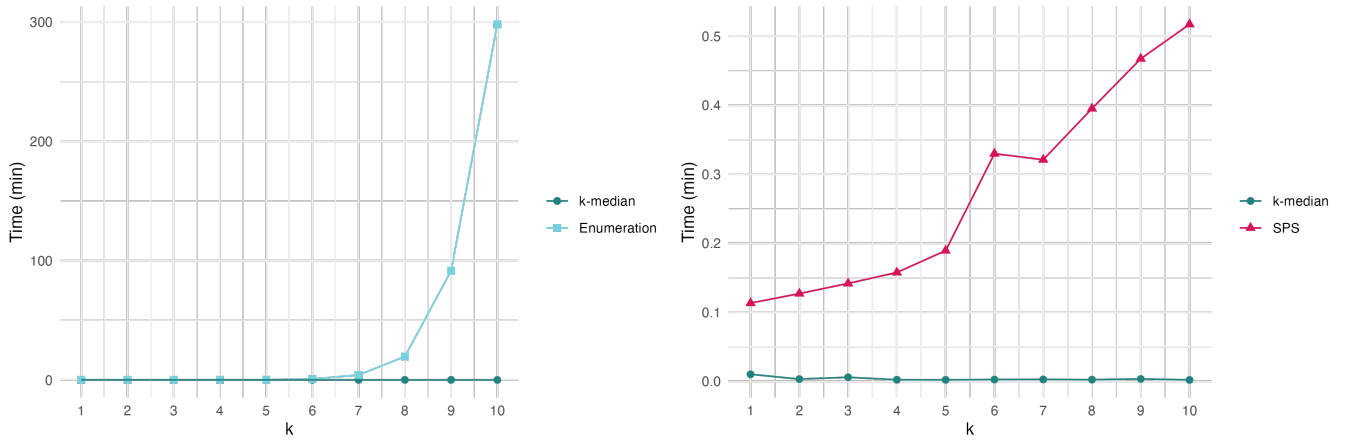

(a) Integer Program vs. Enumeration

(b) Integer Program vs. SPS

Figure 5: Time Needed to Solve the k-median Problem $k \in \{1, \ldots 10\}$

Notes: Time comparison of different purposive sampling approaches. Time (vertical axis) is in minutes. The dark, blue line with circles represents the time it takes to solve the integer program in Section 4 using the `MIP` solver in `Gurobi` to provable optimality (Gurobi gives the solution in less than a second). The light, blue line with squares in Panel a) represents the time it takes to solve the k-median problem using brute-force enumeration. The red line with triangles in Panel b) represents the time needed to implement the synthetic purposive sampling approach using the `spsR` package.

low-skilled ("control") immigrants.

While the experiments have already been conducted and outcomes of each experiment are available for all countries, we consider the situation of a researcher that can only conduct $k = 6$ experiments; here the choice of $k = 6$ mirrors Egami and Lee (2024). We let all 15 countries be both potential experimental and policy-relevant sites. That is, $\mathcal{S}_E = \mathcal{S}_P$, and $\text{card}(\mathcal{S}_E) = \text{card}(\mathcal{S}_P) = 15$. [15]

The only data needed to solve the k-median problem are site-level covariates of both experimental and policy-relevant sites. We use the same covariates as Egami and Lee (2024).[16] They are in

[15] A potential situation we have in mind is one of a researcher who would like to give policy advice to the policy makers in these countries on whether to initiate an immigration reform that could favor either high-skilled or low-skilled immigrants. The researcher knows that policy makers are interested in voters' attitudes towards these types of reforms. We assume that the researcher is only able to experiment in a subset of countries due to administrative or budget constraints, and that he/she needs to extrapolate voters' attitudes of the other policy-relevant sites based on the experimental estimates. The researcher needs to decide whether to recommend the implementation of an immigration reform directed to either high-skilled or low-skilled immigrants.

[16] These are: Gross Domestic Product (GDP), size of migrant population, unemployment rate, proportion of females, mean age, mean education, baseline level of support for immigration by the general public, and a categorical variable that indicates the subregions in Europe (i.e., South, North, East, and West). The covariate data can be accessed in the open-source software R package, `spsR`.

different scales: For example, GDP is measured in 2015 U.S. dollars, while the unemployment rate is reported in percentage points. As is commonly recommended in clustering problems and also done in Egami and Lee (2024), we standardize all of them.

When $k = 6$, the k-median problem is solved by the Czech Republic, Denmark, France, Ireland, Spain, and Switzerland. Figure 6 visualizes the distribution of standardized covariates, along with the sites selected by both the k-median and the synthetic purposive sampling approach. Four of the six selected sites are common to both approaches (Czech Republic, Denmark, Spain, Switzerland), but synthetic purposive sampling chooses Germany and the Netherlands instead of France and Ireland.

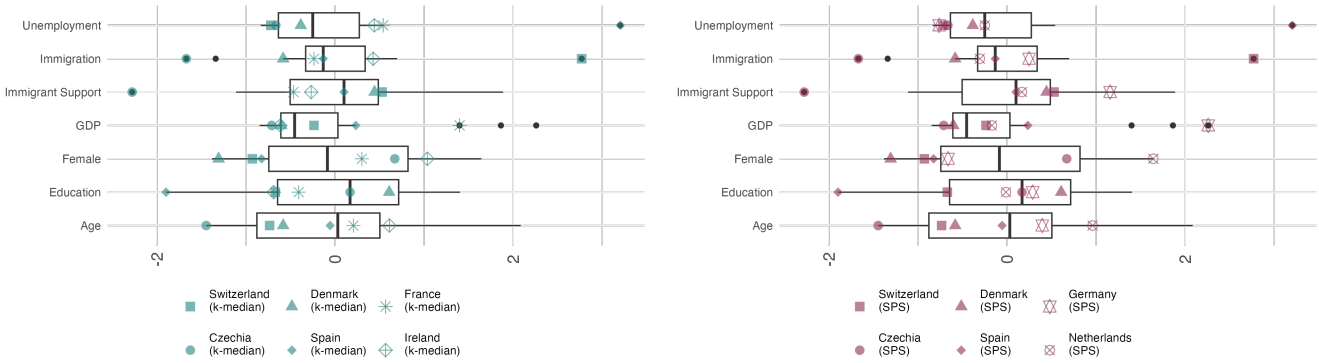

(a) k-median (b) Egami and Lee (2024)

Figure 6: Distribution of Site-level Covariates in The Multi-Country Survey Experiment ($k = 6$)

Notes: Box plots showing the distribution of covariates among the fifteen countries. The box plots are constructed as explained in Figure 2. Each panel depicts the site selected by the different approaches when $k = 6$. Solid shapes indicate sites that are in both the solution of the k-median problem and the synthetic purposive sampling approach. Hollow shapes indicate solutions that differ across methods.

Figure 7 visualizes the connection networks induced by the different solutions. In both panels, red circles represent countries that are selected for experimentation. The gray lines in Figure 7 indicate the connections between the experimental sites and the policy-relevant sites that were not selected for experimentation (blue circles). For example, we can see that each policy-relevant site is connected to exactly one experimental site, i.e. its nearest neighbor; for example, the United Kingdom uses only the information of France. The connection network in Figure 7b is considerably more dense, with all policy-relevant sites connected to more than one experimental site. For instance, the synthetic experiment for United Kingdom assigns positive weights to the Czech Republic, Netherlands, and Germany. To further aid the visual interpretation of the connection network, we color each connection differently to capture its strength. For example, for the United

Kingdom, the strongest connection is to the Netherlands (0.74), whereas the weakest connection of the United Kingdom is to the Czech Republic (0.01).

Figure 7a also shows that three selected countries in the k-median problem (Switzerland, Czech Republic, and Spain) are not connected to any of the other policy-relevant sites. This suggests that these countries are selected because no other country provides a close enough match for them. In contrast, in Figure 7b, these countries (Switzerland, Czech Republic, and Spain) receive positive weights from at least five other countries.

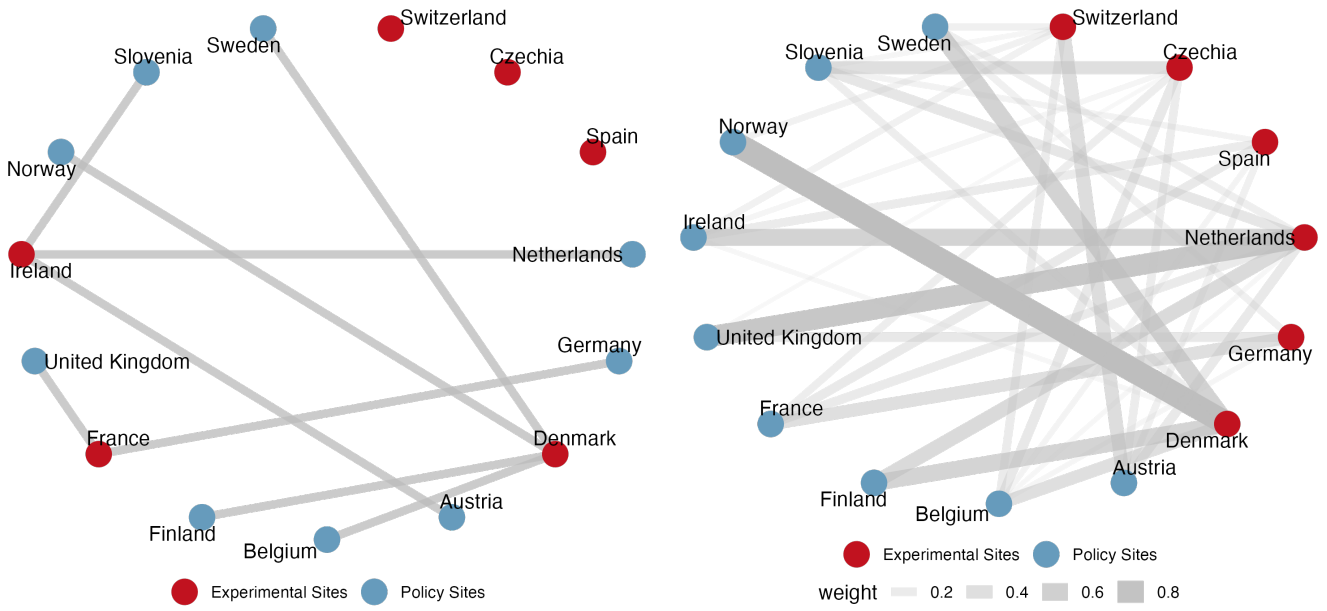

(a) k-median (b) Egami and Lee (2024)

Figure 7: Connection Network for $k = 6$

Notes: In Panel (7a), a connection between a blue dot and a red dot indicates that the corresponding experimental site is the nearest neighbor of the policy site. In Panel (7b), each blue dot may be connected to one or multiple red dots, indicating that the corresponding policy site uses the weighted average of one or more experimental sites to construct its synthetic control. The width and the transparency of the connection line indicate the weight that each policy site puts on the connecting experimental site.

Approximation Error in Theorem 1: Because experimental estimates and corresponding standard errors are available for all 15 countries, we can compute the approximation error in Theorem 1. Define the *relative* approximation error of the k-median approximation to minimax regret as

$$\frac{\inf_{\mathscr{S}\in\mathcal{A}(k)}\left(\inf_{T\in\mathcal{T}_{\mathscr{S}}^{1/2}}\sup_{\tau\in\mathrm{Lip}_C(\mathbb{R}^d)}\mathcal{R}(T,\mathscr{S},\tau)\right)}{\frac{C}{2}\frac{1}{\#\mathcal{S}_P}\inf_{\mathscr{S}\in\mathcal{A}(k)}\left(\sum_{s\in\mathcal{S}_P\setminus\mathscr{S}}\|X_s-X_{N_{\mathscr{S}}(s)}\|\right)}. \tag{23}$$

Algebraic manipulations shows that this expression is bounded from below by

$$\max\left\{0, 1 - \frac{B \cdot \sigma_E \cdot \frac{\min\{\#(\mathcal{S}_E \cap \mathcal{S}_P), k\}}{\#\mathcal{S}_P}}{\frac{C}{2}\frac{1}{\#\mathcal{S}_P}\inf_{\mathscr{S}\in\mathcal{A}(k)}\left(\sum_{s\in\mathcal{S}_P\setminus\mathscr{S}}\|X_s - X_{N_{\mathscr{S}}(s)}\|\right)}\right\}, \tag{24}$$

and from above by

$$1 + \frac{B \cdot \sigma_E \cdot \frac{\min\{\#(\mathcal{S}_E \cap \mathcal{S}_P), k\}}{\#\mathcal{S}_P}}{\frac{C}{2}\frac{1}{\#\mathcal{S}_P}\inf_{\mathscr{S}\in\mathcal{A}(k)}\left(\sum_{s\in\mathcal{S}_P\setminus\mathscr{S}}\|X_s - X_{N_{\mathscr{S}}(s)}\|\right)}, \tag{25}$$

where $B \equiv \arg\max_{z\geq 0} z\Phi(-z)$. The closer the lower and upper bounds are, the better the k-median approximation is.

Figure 8 displays these bounds for values of $k \in \{1, \ldots, 10\}$. As k increases, the approximation becomes worse; as mentioned before, this is driven by ignoring the experimental sites themselves in bounding regret. In this application, the k-median solution still works relatively well when $k = 6$, with a relative error of about $\pm 13\%$. As k increases to 10, the relative error is about $\pm 50\%$.[17]

6 Extensions

6.1 Fixed Costs of Experimentation

We next allow for the possibility that running an experiment in a site $s \in \mathcal{S}_E$ has a fixed cost c_s. This means that the welfare of a decision rule T, given that sites $\mathscr{S}$ are selected for experimentation, corresponds to

$$\mathcal{W}_c(T, \mathscr{S}, \tau) := \frac{1}{\#\mathcal{S}_P}\left(\sum_{s\in\mathcal{S}_P}\tau(X_s)\mathbb{E}_{\tau_{\mathscr{S}}}[T_s(\widehat{\tau}_{\mathscr{S}})] - \sum_{s\in\mathscr{S}} c_s\right). \tag{27}$$

Based on this welfare function, the oracle action for the policy maker is to implement the policy in any policy-relevant site s for which $\tau(X_s) \geq 0$. Expected regret of $(T, \mathscr{S})$ becomes

$$\mathcal{R}_c(T, \mathscr{S}, \tau) := \frac{1}{\#\mathcal{S}_P}\sum_{s\in\mathscr{S}} c_s + \frac{1}{\#\mathcal{S}_P}\sum_{s\in\mathcal{S}_P}\tau(X_s)\left(\mathbf{1}\{\tau(X_s) \geq 0\} - \mathbb{E}_{\tau_{\mathscr{S}}}[T_s(\widehat{\tau}_{\mathscr{S}})]\right). \tag{28}$$

[17] For these computations, we picked C to be the smallest Lipschitz constant needed to capture the heterogeneity of estimated treatment effects in the data. That is, we pick C as

$$\max_{s,s'\in\mathcal{S}_E\cup\mathcal{S}_P}\frac{|\hat{\tau}_s - \hat{\tau}_{s'}|}{\|X_s - X_{s'}\|}, \tag{26}$$

where $\hat{\tau}_s$ denote the estimated treatment effects for site s. It turns out that this C is comfortably "large" in the sense of Theorem 1. We provide additional details in Appendix B.3.

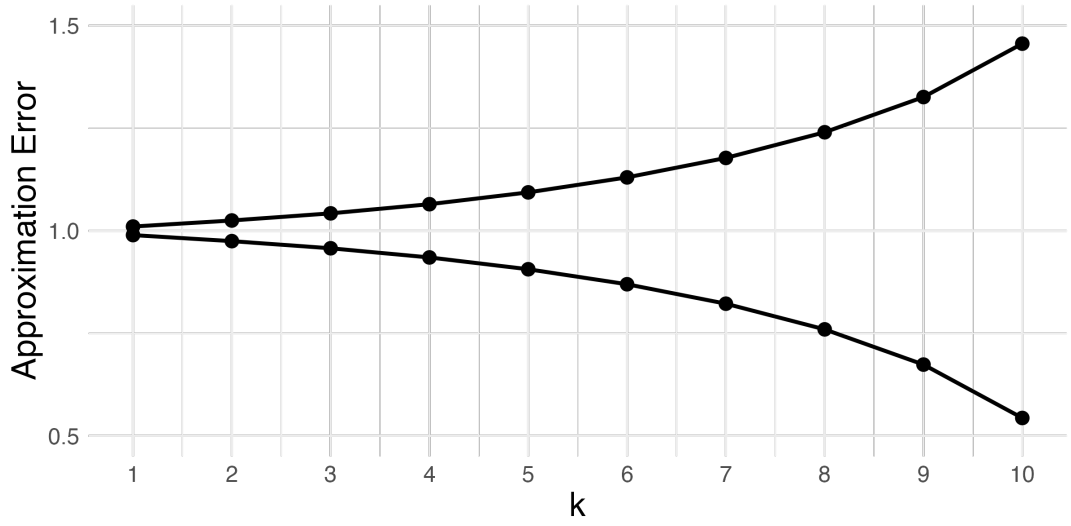

Figure 8: Approximation Error of The k-median Solutions

Notes: Approximation error of the k-median problem in terms of lower and upper bound on the fraction of the true minimax-regret solution over the k-median solution; cf. (23). The Lipschitz constant is chosen to be the smallest Lipschitz constant that is needed to explain the data; cf. (26).

Under the assumptions of Theorem 1, it is possible to show (by extending Lemma 1 and 2 to account for the fixed costs of experimentation) that

$$\inf_{\mathscr{S}\in\mathcal{A}(k)}\left(\inf_{T\in\mathcal{T}_{\mathscr{S}}^{1/2}}\sup_{\tau\in\mathrm{Lip}_C(\mathbb{R}^d)}\mathcal{R}_c(T,\mathscr{S},\tau)\right)$$

can be approximated by $(C/2\#\mathcal{S}_P)$ times

$$\inf_{\mathscr{S}\in\mathcal{A}(k)}\left(\sum_{s\in\mathscr{S}}\left(\frac{2}{C}\right)c_s+\sum_{s\in\mathcal{S}_P\setminus\mathscr{S}}\|X_s-X_{N_{\mathscr{S}}(s)}\|\right), \tag{29}$$

and that the approximation error is the same as the one given in Theorem 1.

The problem in (29) is the *metric uncapacitated k-facility location problem* that was discussed in Remark 7. This is a common extension of the k-median problem where there is a fixed cost of

opening each facility. The problem can also be formulated as a linear integer program, namely

$$\min_{\{y_i, x_{i,j}\}_{i \in \mathcal{S}_E, j \in \mathcal{S}_P}} \left(\sum_{i \in \mathcal{S}_E} y_i c_i^* + \sum_{i \in \mathcal{S}_E, j \in \mathcal{S}_P} x_{i,j} \cdot c(j,i) \right) \tag{30}$$

$$\text{such that } \sum_{i \in \mathcal{S}_E} x_{i,j} = 1, \quad \forall j \in \mathcal{S}_P, \tag{31}$$

$$\sum_{i \in \mathcal{S}_E} y_i \leq k, \tag{32}$$

$$0 \leq x_{i,j} \leq y_i, \quad i \in \mathcal{S}_E, j \in \mathcal{S}_P, \tag{33}$$

$$y_i \in \{0,1\}, x_{i,j} \in \{0,1\}, \quad i \in \mathcal{S}_E, j \in \mathcal{S}_P. \tag{34}$$

Just as before, the choice variable y_i indicates whether facility i is open, and $x_{i,j}$ indicates whether client j is assigned to facility i. Constraints are imposed to ensure that each client j is assigned to at least one facility, that there are no more than k facilities in total, and that a client can only be assigned to an open facility. The connection cost between a facility i and a client j is still given by $\|X_i - X_j\|$, but now there is a fixed cost of opening a facility i given by $c_i^* \equiv 2c_i/C$. This means that—in contrast to the case in which there is no fixed cost—choosing the optimal sites to maximize external validity now requires knowledge of the Lipschitz constant C (as this constant appears explicitly on the fixed cost c_i^*). This is simply because, as we have shown, the scale of regret absent costs of experimentation depends on C. Consequently, every time a new site is considered for experimentation, there is a trade-off between its contribution to reduce regret and the fixed cost of experimentation. We note that when C is large, the fixed costs in the objective function become negligible, and the solution of the uncapacitated facility location problem can be approximated by the solution of the k-median problem.

6.2 Random Selection of Experimental Sites

So far, our analysis focused on minimizing the worst-case (welfare-based) regret among all *purposive* sampling schemes that select at most k sites. That is, we excluded randomized (including nonuniformly randomized) sampling schemes. This is an important limitation—a widespread view among experimenters is 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" (Duflo et al., 2007, p. 3953).

We now extend our baseline framework (which excludes fixed costs of experimentation) to allow for randomized site selection. First, we note that whether randomization is potentially desired

delicately depends on the decision-theoretic setup. For example, the optimal sampling scheme in the Bayesian setting of Gechter et al. (2024) will typically be purposive. Whether randomization improves minimax regret depends on how exactly the decision problem is formulated, which can be related to what we refer to as the *timing* assumptions in an implicit game that the decision maker plays against a malicious "nature". We will first clarify this observation and then provide a brief illustration of randomized solutions.

To formalize the discussion, let M denote the cardinality of $\mathcal{A}(k)$. Let $\Delta(\mathcal{A}(k))$ denote the set of all probability distributions over the M elements of $\mathcal{A}(k)$. We define a randomized site selection as a probability distribution $p := (p_1, \ldots, p_M) \in \Delta(\mathcal{A}(k))$. The econometrician will pick a subset from the experimental sites by drawing one realization of a distribution $p^* \in \Delta(\mathcal{A}(k))$ that she specified. Denote the randomly selected sites by $\mathscr{S}^*$.

In our setting, whether the decision maker will *want to* randomize crucially depends on timing assumptions; that is, the moment in the game in which nature can move to harm the decision maker. Consider first the case in which an adversarial nature may harm the policy maker by choosing τ only after seeing the realization of $\mathscr{S}^*$ (and knowing $\mathcal{T}_{\mathscr{S}^*}$). Then, the risk of using treatment rule $T \in \mathcal{T}_{\mathscr{S}^*}$ is

$$\mathcal{R}(T, \mathscr{S}^*, \tau),$$

and the worst-case payoff becomes

$$\sup_{\tau \in \mathrm{Lip}_C(\mathbb{R}^d)} \mathcal{R}(T, \mathscr{S}^*, \tau).$$

The minimax problem faced by the econometrician after $\mathscr{S}^*$ has been realized is:

$$V(\mathscr{S}^*) := \inf_{T \in \mathcal{T}_{\mathscr{S}^*}} \sup_{\tau \in \mathrm{Lip}_C(\mathbb{R}^d)} \mathcal{R}(T, \mathscr{S}^*, \tau).$$

With slight abuse of notation, let $p(\mathscr{S}^*)$ denote the probability of choosing $\mathscr{S}^*$ at random under $p \in \Delta(\mathcal{A}(k))$. The (ex ante) expected payoff of any randomized site selection is

$$\sum_{\mathscr{S}^* \in \mathcal{A}(k)} p(\mathscr{S}^*) \cdot V(\mathscr{S}^*),$$

and the optimal randomized site selection solves

$$\inf_{p \in \Delta(\mathcal{A}(k))} \left(\sum_{\mathscr{S}^* \in \mathcal{A}(k)} p(\mathscr{S}^*) \cdot V(\mathscr{S}^*) \right).$$

But this problem is solved by any p supported on $\arg\min_{\mathscr{S}} V(\mathscr{S})$, the set of purposive sampling schemes that solve the site selection problem. If that problem's solution is unique, the policy maker will *never* randomize under this timing of the game. Furthermore, this "sequential" timing may feel natural in applications that we have in mind.

That said, the timing that seems more in line with Wald's (1950) original application of the minimax principle is likely one in which the policy maker commits to both a randomized sampling scheme and a set of contingent (on sampling scheme) decision rules and nature adversarially picks τ before any randomization was realized. We next briefly discuss this possibility.

To see that randomization might strictly speaking be optimal, consider a stylized example where $k = 1$ and the covariates of each site are equal to its index: $S_E = \{1, 4\}$, and $S_P = \{2, 3\}$. For simplicity, suppose furthermore that $\hat{\tau}_s = \tau_s$, i.e. there is no sampling uncertainty in the treatment effects. This example can be solved for those combinations of sampling scheme and treatment assignment rule that achieve exact MMR. As we formally show in Appendix B.2, the exact MMR attainable by purposive sampling equals $3C/4$, whereas the exact MMR with randomized sampling equals $C/2$.[18] Thus, in principle there can be a gain to randomized sampling.

To see that solving this problem can quickly become very hard, consider now the same example, except that the experimental sites coincide with the policy sites at $\{1, 2\}$. Then we can find the MMR optimal combination of purposive sampling scheme and treatment assignment rule, and we can also verify that randomized site selection will strictly reduce worst-case regret. However, we are unable to characterize the exact solution for this, still extremely structured, case. We leave further exploration of randomized site selection, including the potentially fruitful investigation of "good enough" (in terms of regret) such schemes, for future research.

6.3 Other Restrictions on Treatment Heterogeneity

With Assumption 1, the effect of the policy of interest is a Lipschitz function of observed covariates with respect to the Euclidean distance. This means that the effect of the policy of interest in two sites with the same covariates is assumed to be the same (in other words, there is no *unobserved treatment effect heterogeneity*). In this section, we show that our main result can also accommodate other general measures of the distance between covariates (thus allowing for other functional classes), as well as some forms of unobserved treatment effect heterogeneity (provided $\mathcal{S}_E \cap \mathcal{S}_P = \emptyset$).

[18] The assumption of perfect signals is for simplicity. The example is rigged to resemble cases analyzed in Montiel Olea et al. (2023) and Stoye (2012) and we would accordingly be able to generalize it, but for our present purpose, solving for arbitrary sampling variances would only add tedium.

6.3.1 Other Distance Measure Based on Observed Covariates

To see how our results can allow for other types of distances, let $m(x, x') : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}_+$ be a measure of distance (metric) between $x, x' \in \mathbb{R}^d$.[19] Then, Assumption 1 may be modified as:

Assumption 3. τ is a Lipschitz function (with respect to metric $m(\cdot, \cdot)$) with known constant C. That is, for any $x, x' \in \mathbb{R}^d$, $|\tau(x) - \tau(x')| \leq Cm(x, x')$.

For example, $m(x, x') = \left((x - x')^\top (x - x')\right)^{1/2}$ is the Euclidean distance in Assumption 1 and $m(x, x') = \left((x - x')^\top W (x - x')\right)^{1/2}$ for some positive definite matrix W is a W-weighted Euclidean distance. Choosing $m(x, x') = [\tilde{m}(x, x')]^\alpha$ for some $\alpha \in (0, 1)$ and some distance measure $\tilde{m}(\cdot, \cdot)$ also effectively allows us to model τ as a Hölder continuous function of order α (Dudley, 2002, p. 56). With Assumptions 2 and 3, $N_{\mathscr{S}}(s)$ is understood to be a nearest neighbor in $\mathscr{S}$ measured in terms of metric $m(\cdot, \cdot)$ that appears in Assumption 3. Moreover, $\text{Lip}_C(\mathbb{R}^d)$ now stands for the space of all Lipschitz functions from $\mathbb{R}^d$ to $\mathbb{R}$ with constant C, but in terms of metric $m(\cdot, \cdot)$. We note that we are slightly abusing notation, because we are not using m to index the nearest-neighbor function $N_{\mathscr{S}}(s)$ or the Lipschitz functional class $\text{Lip}_C(\mathbb{R}^d)$, although their respective definitions explicitly on the choice of metric $m(\cdot, \cdot)$. Under a general metric $m(\cdot, \cdot)$, the k-median problem is simply modified as:

Definition 3. We say that a purposive sampling scheme, $\mathscr{S} \in \mathcal{A}(k)$, solves the k-median problem with a metric cost function $m(\cdot, \cdot)$ if it solves

$$\inf_{\mathscr{S} \in \mathcal{A}(k)} \sum_{s \in \mathcal{S}_P} m(X_s, X_{N_{\mathscr{S}}(s)}). \tag{35}$$

Since the proof of Theorem 1 does not rely on any specific property of the Euclidean distance, the results of Theorem 1 still hold under a general metric $m(\cdot, \cdot)$, by replacing the Euclidean distance $\|\cdot - \cdot\|$ with $m(\cdot, \cdot)$.

6.3.2 Unobserved Treatment Heterogeneity

Now, we show how our assumptions can be modified to accommodate some forms of unobserved treatment heterogeneity. Instead of viewing policy effects as a function of observed covariates X only, we now model the policy effects at the site level. Specifically, for each $s \in \mathcal{S}$, denote by $\tau_s \in \mathbb{R}$ the policy effect in site s. Then, instead of viewing τ as a function, we will let $\tau := (\tau_1, \tau_2, ..., \tau_S)^\top \in$

[19] That is, we assume that $m(x, x') > 0$ for any $x \neq x'$, and that for any x we have $m(x, x) = 0$. We also assume that the function is symmetric, in that $m(x, x') = m(x', x)$. And finally, we assume that m satisfies the triangle inequality $m(x, x') \leq m(x, x'') + m(x', x'')$ for any x, x', x''. Also see Dudley (2002, p.20).

$\mathbb{R}^S$ be a vector of dimension S that represents the policy effects of all sites in $\mathcal{S}$. Then, Assumptions 1 and 2 can be replaced with the following:

Assumption 4. For any $s, s' \in \mathcal{S}$, $s \neq s'$, τ satisfies:

$$|\tau_s - \tau_{s'}| \leq Cm(X_s, X_{s'}) + c, \tag{36}$$

where $C > 0$ and $c > 0$ are both known, and $m(\cdot, \cdot)$ is a metric.

Assumption 4 allows sites with the same covariates to have different policy effects. The difference, however, is assumed to be at most c. We motivate Assumption 4 using a simple linear regression model in Appendix B.4. Denote by $\mathcal{F}_{C,c} \subset \mathbb{R}^S$ the collection of all vectors of dimension $\mathbb{R}^S$ satisfying (36). For this parameter class, the minimax-regret optimal purposive sampling scheme and treatment rule are redefined to solve

$$\inf_{\mathscr{S} \in \mathcal{A}(k), T \in \mathcal{T}_{\mathscr{S}}^{1/2}} \sup_{\tau \in \mathcal{F}_{C,c}} \mathcal{R}(T, \mathscr{S}, \tau). \tag{37}$$

It turns out that the purposive sampling scheme that solves (37) can still be approximated by the solution of the k-median problem in Definition 3 with a different cost function.

To see this, note that the functional class $\mathcal{F}_{C,c}$ is still convex and centrosymmetric. Thus, we can show that, for the evidence aggregation problem studied in Montiel Olea et al. (2023), the conclusion of their Proposition 1(iii) extends to $\mathcal{F}_{C,c}$ as well. Then, under Assumption 4, Lemmas 1 and 2 continue to hold with minor modifications. In particular, when C is large enough (with a threshold that can be exactly characterized), we have that

$$\inf_{T \in \mathcal{T}_{\mathscr{S}}^{1/2}} \sup_{\tau \in \mathcal{F}_{C,c}} \mathcal{R}(T, \mathscr{S}, \tau) \leq \left(\frac{B}{\#\mathcal{S}_P} \sum_{s \in \mathscr{S} \cap \mathcal{S}_P} \sigma_s \right) + \left(\frac{1}{\#\mathcal{S}_P} \sum_{s \in \mathcal{S}_P \setminus \mathscr{S}} \frac{Cm(X_s, X_{N_{\mathscr{S}}(s)}) + c}{2} \right).$$

In addition, the following lower bound can also be verified to hold

$$\inf_{T \in \mathcal{T}_{\mathscr{S}}^{1/2}} \sup_{\tau \in \mathcal{F}_{C,c}} \mathcal{R}(T, \mathscr{S}, \tau) \geq \left(\frac{1}{\#\mathcal{S}_P} \sum_{s \in \mathcal{S}_P \setminus \mathscr{S}} \frac{Cm(X_s, X_{N_{\mathscr{S}}(s)}) + c}{2} \right).$$

The two bounds above imply that, even when treatment effect heterogeneity is characterized by Assumption 4, the optimized value of the minimax problem (37) can still be uniformly approximated by the following objective function

$$\inf_{\mathscr{S}\in\mathcal{A}(k)} \left(\frac{1}{\#\mathcal{S}_P} \sum_{s\in\mathcal{S}_P\setminus\mathscr{S}} \left(\frac{Cm(X_s, X_{N_{\mathscr{S}}(s)}) + c}{2} \right) \right).$$

When $\mathcal{S}_E \cap \mathcal{S}_P = \emptyset$, the solution to the problem above is the same as

$$\inf_{\mathscr{S}\in\mathcal{A}(k)} \left(\frac{1}{\#\mathcal{S}_P} \sum_{s\in\mathcal{S}_P} \left(\frac{Cm(X_s, X_{N_{\mathscr{S}}(s)}) + c}{2} \right) \right).$$

This problem is equivalent to the k-median problem in Definition 3, where the connection cost between sites s and s' is given by $m(X_s, X_{s'})$ (neither C nor c are required as inputs).

7 Conclusion

This paper presented a decision-theoretic justification for viewing the question of how to best choose *where* to experiment in order to optimize external validity as a *k-median (clustering) problem*. More concretely, we presented conditions under which minimizing the *worst-case*, welfare-based regret among all purposive (nonrandomized) schemes that select k sites is approximately equal, and can be exactly equal, to finding the k most central vectors of baseline site-level covariates.

We believe there are many interesting directions for future work. For example, while we focused on purposive sampling schemes, it would be interesting to better understand the value of randomized sampling schemes and whether site-level covariates can be used to design such randomized selection with an eye on external validity. We also think that discussions around the relation between the k-median problem and synthetic purposive sampling of Egami and Lee (2024) open interesting research directions to provide a decision-theoretic justification for the use the synthetic control of Abadie et al. (2010). Finally, it would be interesting to explore the use of (mixed) integer programming techniques to solve for the purposive sampling schemes suggested in Egami and Lee (2024) and Gechter et al. (2024).

A Proofs of Main Results

A.1 Proof of Lemma 1

Fix the selected sites $\mathscr{S}$, and denote the cardinality of $\mathscr{S}$ as $\text{card}(\mathscr{S})$. Let $\mathscr{S}_1 < \mathscr{S}_2 < \ldots < \mathscr{S}_{\text{card}(\mathscr{S})}$ denote the indices of the $\text{card}(\mathscr{S})$ experimental sites in $\mathscr{S}$. For a given experimental site $s \in \mathscr{S}$, let $\widehat{\tau}_s$ denote its corresponding treatment effect estimate. Let

$$\widehat{\tau}_{\mathscr{S}} := (\hat{\tau}_{\mathscr{S}_1}, \ldots, \hat{\tau}_{\mathscr{S}_{\text{card}(\mathscr{S})}})^\top$$

denote the vector containing the estimates for each experimental site.

For each policy-relevant site, $s \in \mathcal{S}_P$, recall that $N_{\mathscr{S}}(s)$ denotes its nearest neighbor among the sites $\mathscr{S}$ (or the nearest neighbor with the smallest index in case of multiplicity). Partition the policy-relevant sites as

$$\mathcal{S}_P = (\mathscr{S} \cap \mathcal{S}_P) \cup (\mathcal{S}_P \backslash \mathscr{S}).$$

Consider the decision rule, $T^* \in \mathcal{T}_{\mathscr{S}}^{1/2}$, that recommends, for each policy-relevant site $s \in \mathcal{S}_P$ and given data $\widehat{\tau}_{\mathscr{S}}$, the following action:

i) For $s \in \mathcal{S}_P \cap \mathscr{S}$,

$$T_s^*(\widehat{\tau}_{\mathscr{S}}) := \mathbf{1}\{\widehat{\tau}_s \geq 0\}.$$

ii) For $s \in \mathcal{S}_P \backslash \mathscr{S}$, set

$$T_s^*(\widehat{\tau}_{\mathscr{S}}) := \Phi\left(\frac{\widehat{\tau}_{N_{\mathscr{S}}(s)}}{\tilde{\sigma}_s}\right),$$

where

$$\tilde{\sigma}_s := \sqrt{\left(\frac{C\|X_s - X_{N_{\mathscr{S}}(s)}\|}{\sqrt{\pi/2}}\right)^2 - \sigma_s^2}.$$

Note that the expression in ii) above is well-defined for every $s \in \mathcal{S}_P$ when C is large enough. This decision rule is the minimax-regret optimal rule (provided C is large enough) for the evidence aggregation framework discussed in Yata (2021) and Montiel Olea et al. (2023), both of which build upon Stoye (2012), and is indeed in $\mathcal{T}_{\mathscr{S}}^{1/2}$.

Define $C(\mathscr{S})$ to be any value of C for which $\tilde{\sigma}_s > 0$ for every $s \in \mathcal{S}_P$. Note that by definition of infimum,

$$\begin{aligned}\inf_{T\in\mathcal{T}_{\mathscr{S}}} \sup_{\tau\in \mathrm{Lip}_C(\mathbb{R}^d)} \mathcal{R}(T,\mathscr{S},\tau) &\leq \sup_{\tau\in \mathrm{Lip}_C(\mathbb{R}^d)} \mathcal{R}(T^*,\mathscr{S},\tau), \\ &\leq \frac{1}{\#\mathcal{S}_P}\sum_{s\in\mathcal{S}_P} \sup_{\tau\in \mathrm{Lip}_C(\mathbb{R}^d)} \left(\tau(X_s)\left(\mathbf{1}\{\tau(X_s)\geq 0\} - \mathbb{E}_{\tau_{\mathscr{S}}}[T_s^*(\widehat{\tau}_{\mathscr{S}})]\right)\right),\end{aligned}$$

where the second inequality follows from the fact that

$$\mathcal{R}(T,\mathscr{S},\tau) := \frac{1}{\#\mathcal{S}_P}\sum_{s\in\mathcal{S}_P} \tau(X_s)\left(\mathbf{1}\{\tau(X_s)\geq 0\} - \mathbb{E}_{\tau_{\mathscr{S}}}[T_s(\widehat{\tau}_{\mathscr{S}})]\right).$$

It is a well-known result (and can be verified by algebra) that for any $s \in \mathscr{S} \cap \mathcal{S}_P$,

$$\sup_{\tau\in \mathrm{Lip}_C(\mathbb{R}^d)} \left(\tau(X_s)\left(\mathbf{1}\{\tau(X_s)\geq 0\} - \mathbb{E}_{\tau_{\mathscr{S}}}[T_s^*(\widehat{\tau}_{\mathscr{S}})]\right)\right) = B\sigma_s, \tag{38}$$

where $B \equiv \arg\max_{z\geq 0} z\Phi(-z)$.

Moreover, it follows from (Montiel Olea et al., 2023, Proposition 1 (iii) and its proof) that if $C > \max_{s\in\mathcal{S}_P\setminus\mathscr{S}} \left\{\sqrt{\frac{\pi}{2}}\frac{\sigma_{N_{\mathscr{S}}(s)}}{\left\|X_s - X_{N_{\mathscr{S}}(s)}\right\|}\right\} := C(\mathscr{S})$, where $\sigma_{N_{\mathscr{S}}(s)}$ denotes the standard deviation of the nearest neighbor estimate in $\mathscr{S}$ for site $s \in \mathcal{S}_P\setminus\mathscr{S}$, it holds that for any $s \in \mathcal{S}_P\setminus\mathscr{S}$:

$$\sup_{\tau\in \mathrm{Lip}_C(\mathbb{R}^d)} \left(\tau(X_s)\left(\mathbf{1}\{\tau(X_s)\geq 0\} - \mathbb{E}_{\tau_{\mathscr{S}}}[T_s^*(\widehat{\tau}_{\mathscr{S}})]\right)\right) = \frac{C}{2}\|X_s - X_{N_{\mathscr{S}}(s)}\|. \tag{39}$$

Equations (38)-(39) imply:

$$\inf_{T\in\mathcal{T}_{\mathscr{S}}^{1/2}} \sup_{\tau\in \mathrm{Lip}_C(\mathbb{R}^d)} \mathcal{R}(T,\mathscr{S},\tau) \leq \left(\frac{B}{\#\mathcal{S}_P}\sum_{s\in\mathscr{S}\cap\mathcal{S}_P}\sigma_s\right) + \left(\frac{C}{2}\frac{1}{\#\mathcal{S}_P}\sum_{s\in\mathcal{S}_P\setminus\mathscr{S}}\|X_s - X_{N_{\mathscr{S}}(s)}\|\right). \tag{40}$$

A.2 Proof of Lemma 2

Fix the selected sites $\mathscr{S}$, and denote the cardinality of $\mathscr{S}$ as $\mathrm{card}(\mathscr{S})$. Let $\mathscr{S}_1 < \mathscr{S}_2 < \ldots < \mathscr{S}_{\mathrm{card}(\mathscr{S})}$ denote the indices of the $\mathrm{card}(\mathscr{S})$ experimental sites in $\mathscr{S}$. For any $X \in \mathbb{R}^d$, define $X_{\mathcal{N}_{\mathscr{S}}}$ to be the element in $\{X_{\mathscr{S}_1}, \ldots, X_{\mathscr{S}_{\mathrm{card}(\mathscr{S})}}\}$ that is closest to X in terms of $\|\cdot\|$ (if there is more than one closest element, pick the X associated to the smallest index). The proof has three parts.

PART I: Consider the function $\tau^* : \mathbb{R}^d \to \mathbb{R}$ such that:

$$\tau^*(X) = C\|X - X_{\mathcal{N}_\mathscr{S}}\|.$$

We start by showing that the function τ^* is Lipschitz with constant C. To see this, consider three cases:

Case 1: Suppose first that $X, X' \in \{X_{\mathscr{S}_1}, \ldots, X_{\mathscr{S}_{\text{card}(\mathscr{S})}}\}$. In this case, we trivially have

$$|\tau^*(X) - \tau^*(X')| = 0 \leq C\|X - X'\|.$$

Case 2: Suppose now that $X \notin \{X_{\mathscr{S}_1}, \ldots, X_{\mathscr{S}_{\text{card}(\mathscr{S})}}\}$, but $X' \in \{X_{\mathscr{S}_1}, \ldots, X_{\mathscr{S}_{\text{card}(\mathscr{S})}}\}$. In this case,

$$|\tau^*(X) - \tau^*(X')| = C\|X - X_{\mathcal{N}_\mathscr{S}}\| \leq C\|X - X'\|,$$

where the last inequality follows by the definition of $X_{\mathcal{N}_\mathscr{S}}$ and the fact that $X' \in \{X_{\mathscr{S}_1}, \ldots, X_{\mathscr{S}_{\text{card}(\mathscr{S})}}\}$.

Case 3: Finally, take $X, X' \notin \{X_{\mathscr{S}_1}, \ldots, X_{\mathscr{S}_{\text{card}(\mathscr{S})}}\}$. In this case,

$$|\tau^*(X) - \tau^*(X')| = C|\|X - X_{\mathcal{N}_\mathscr{S}}\| - \|X' - X'_{\mathcal{N}_\mathscr{S}}\||.$$

Without loss of generality, assume that $\|X - X_{\mathcal{N}_\mathscr{S}}\| \geq \|X' - X'_{\mathcal{N}_\mathscr{S}}\|$. Then,

$$\begin{aligned}
|\tau^*(X) - \tau^*(X')| &= C\left(\|X - X_{\mathcal{N}_\mathscr{S}}\| - \|X' - X'_{\mathcal{N}_\mathscr{S}}\|\right) \\
&\leq C\left(\|X - X'_{\mathcal{N}_\mathscr{S}}\| - \|X' - X'_{\mathcal{N}_\mathscr{S}}\|\right) \\
&= C\left(\|X - X' + X' - X'_{\mathcal{N}_\mathscr{S}}\| - \|X' - X'_{\mathcal{N}_\mathscr{S}}\|\right) \\
&\leq C\|X - X'\|,
\end{aligned}$$

where the first inequality uses the definition of $X_{\mathcal{N}_\mathscr{S}}$, and the last display uses the triangle inequality. We conclude that τ^* is a Lipschitz function with constant C, which means it is included in our parameter space.

PART II: Since $\tau^* \in \text{Lip}_C(\mathbb{R}^d)$, then for any treatment rule T:

$$\sup_{\tau \in \text{Lip}_C(\mathbb{R}^d)} \mathcal{R}(T, \mathscr{S}, \tau) \geq \mathcal{R}(T, \mathscr{S}, \tau^*). \tag{41}$$

Moreover, by definition,

$$\mathcal{R}(T, \mathscr{S}, \tau) := \frac{1}{\#\mathcal{S}_P} \sum_{s \in \mathcal{S}_P} \tau(X_s) \left(\mathbf{1}\{\tau(X_s) \geq 0\} - \mathbb{E}_{\tau_{\mathscr{S}}}[T_s(\widehat{\tau}_{\mathscr{S}})] \right).$$

Since $\tau^*(X_s) = 0$ for all $s \in \mathcal{S}_P \cap \mathscr{S}$, we have that

$$\begin{aligned} \mathcal{R}(T, \mathscr{S}, \tau^*) &= \frac{1}{\#\mathcal{S}_P} \sum_{s \in \mathcal{S}_P \setminus \mathscr{S}} \tau^*(X_s) \left(\mathbf{1}\{\tau^*(X_s) \geq 0\} - \mathbb{E}_{\tau^*_{\mathscr{S}}}[T_s(\widehat{\tau}_{\mathscr{S}})] \right), \\ &= \frac{1}{\#\mathcal{S}_P} \sum_{s \in \mathcal{S}_P \setminus \mathscr{S}} C \|X_s - X_{N_{\mathscr{S}}(s)}\| \left(1 - \mathbb{E}_{\mathbf{0}}[T_s(\widehat{\tau}_{\mathscr{S}})]\right), \end{aligned} \tag{42}$$

where the last line uses the definition of τ^*. Equations (41) and (42) thus imply that for any $T \in \mathcal{T}^{1/2}_{\mathscr{S}}$:

$$\sup_{\tau \in \mathrm{Lip}_C(\mathbb{R}^d)} \mathcal{R}(T, \mathscr{S}, \tau) \geq \left(\frac{C}{2} \frac{1}{\#\mathcal{S}_P} \sum_{s \in \mathcal{S}_P \setminus \mathscr{S}} \|X_s - X_{N_{\mathscr{S}}(s)}\| \right). \tag{43}$$

Part III: Equation (43) implies

$$\inf_{T \in \mathcal{T}^{1/2}_{\mathscr{S}}} \sup_{\tau \in \mathrm{Lip}_C(\mathbb{R}^d)} \mathcal{R}(T, \mathscr{S}, \tau) \geq \left(\frac{C}{2} \frac{1}{\#\mathcal{S}_P} \sum_{s \in \mathcal{S}_P \setminus \mathscr{S}} \|X_s - X_{N_{\mathscr{S}}(s)}\| \right).$$

References

Abadie, A., A. Diamond, and J. Hainmueller (2010): "Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program," *Journal of the American statistical Association*, 105, 493–505.

Adjaho, C. and T. Christensen (2022): "Externally Valid Treatment Choice," *arXiv preprint arXiv:2205.05561*.

Allcott, H. (2015): "Site selection bias in program evaluation," *The Quarterly Journal of Economics*, 130, 1117–1165.

Andrews, I., D. Fudenberg, A. Liang, and C. Wu (2022): "The Transfer Performance of Economic Models," *arXiv preprint arXiv:2202.04796*.

Armstrong, T. B. and M. Kolesár (2018): "Optimal inference in a class of regression models," *Econometrica*, 86, 655–683.

Azevedo, E. M., A. Deng, J. L. Montiel Olea, J. Rao, and E. G. Weyl (2020): "A/b testing with fat tails," *Journal of Political Economy*, 128, 4614–000.

Azevedo, E. M., D. Mao, J. L. Montiel Olea, and A. Velez (2023): "The A/B testing problem with Gaussian priors," *Journal of Economic Theory*, 210, 105646.

Banerjee, A. V., S. Chassang, and E. Snowberg (2017): "Decision theoretic approaches to experiment design and external validity," in *Handbook of Economic Field Experiments*, Elsevier, vol. 1, 141–174.

Ben-David, S., J. Blitzer, K. Crammer, A. Kulesza, F. Pereira, and J. W. Vaughan (2010): "A theory of learning from different domains," *Machine learning*, 79, 151–175.

Bertsimas, D., A. King, and R. Mazumder (2016): "Best subset selection via a modern optimization lens," *The Annals of Statistics*, 44, 813 – 852.

Bertsimas, D. and R. Weismantel (2005): *Optimization Over Integers*, Dynamic Ideas.

Charikar, M., S. Guha, É. Tardos, and D. B. Shmoys (2002): "A Constant-Factor Approximation Algorithm for the k-Median Problem," *Journal of Computer and System Sciences*, 65, 129–149.

Chassang, S. and S. Kapon (2022): "Designing randomized controlled trials with external validity in mind," Tech. rep., National Bureau of Economic Research.

Christensen, T., H. R. Moon, and F. Schorfheide (2022): "Optimal decision rules when payoffs are partially identified," *arXiv preprint arXiv:2204.11748*.

Cohen-Addad, V., A. De Mesmay, E. Rotenberg, and A. Roytman (2018): "The bane of low-dimensionality clustering," in *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, SIAM, 441–456.

Cohen-Addad, V., H. Esfandiari, V. Mirrokni, and S. Narayanan (2022): "Improved approximations for Euclidean k-means and k-median, via nested quasi-independent sets," in *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, 1621–1628.

Cook, T. D., D. T. Campbell, and W. Shadish (2002): *Experimental and quasi-experimental designs for generalized causal inference*, vol. 1195, Houghton Mifflin Boston, MA.

Degtiar, I. and S. Rose (2023): "A review of generalizability and transportability," *Annual Review of Statistics and Its Application*, 10, 501–524.

Duchi, J. C. and H. Namkoong (2021): "Learning models with uniform performance via distributionally robust optimization," *The Annals of Statistics*, 49, 1378–1406.

Dudley, R. (2002): *Real Analysis and Probability*, vol. 74, Cambridge University Press.

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," .

Gechter, M. (2024): "Generalizing the Results from Social Experiments: Theory and Evidence from India," *Journal of Business & Economic Statistics*, 42, 801–811.

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," .

Gurobi Optimization, LLC (2023): "Gurobi Optimizer Reference Manual," .

Hu, Y., H. Zhu, E. Brunskill, and S. Wager (2024): "Minimax-Regret Sample Selection in Randomized Experiments," *arXiv preprint arXiv:2403.01386*.

Ishihara, T. and T. Kitagawa (2021): "Evidence Aggregation for Treatment Choice," ArXiv:2108.06473 [econ.EM], `https://doi.org/10.48550/arXiv.2108.06473`.

Kariv, O. and S. L. Hakimi (1979): "An algorithmic approach to network location problems. I: The p-centers," *SIAM journal on applied mathematics*, 37, 513–538.

Kumar, A., Y. Sabharwal, and S. Sen (2010): "Linear-time approximation schemes for clustering problems in any dimensions," *Journal of the ACM (JACM)*, 57, 1–32.

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.

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

——— (2019): "Trial Size for Near-Optimal Choice Between Surveillance and Aggressive Treatment: Reconsidering MSLT-II," *The American Statistician*, 73, 305–311.

Mansour, Y., M. Mohri, and A. Rostamizadeh (2009): "Domain Adaptation: Learning Bounds and Algorithms," in *Proceedings of The 22nd Annual Conference on Learning Theory (COLT 2009)*, Montréal, Canada.

Megiddo, N. and K. J. Supowit (1984): "On the complexity of some common geometric location problems," *SIAM journal on computing*, 13, 182–196.

Menzel, K. (2023): "Transfer Estimates for Causal Effects across Heterogeneous Sites," *arXiv preprint arXiv:2305.01435*.

Montiel Olea, J. L., C. Qiu, and J. Stoye (2023): "Decision Theory for Treatment Choice Problems with Partial Identification," *arXiv preprint arXiv:2312.17623*.

Müller, U. K. (2011): "Efficient tests under a weak convergence assumption," *Econometrica*, 79, 395–435.

Naumann, E., L. F. Stöetzer, and G. Pietrantuono (2018): "Attitudes towards highly skilled and low-skilled immigration in Europe: A survey experiment in 15 European countries," *European Journal of Political Research*, 57, 1009–1030.

Raiffa, H. and R. Schlaifer (1961): "Applied statistical decision theory." .

Sahoo, R., L. Lei, and S. Wager (2022): "Learning from a biased sample," *arXiv preprint arXiv:2209.01754*.

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

Sugiyama, M., M. Krauledat, and K.-R. Müller (2007): "Covariate shift adaptation by importance weighted cross validation." *Journal of Machine Learning Research*, 8.

Tetenov, A. (2012): "Statistical treatment choice based on asymmetric minimax regret criteria," *Journal of Econometrics*, 166, 157–165.

Vivalt, E. (2020): "How much can we generalize from impact evaluations?" *Journal of the European Economic Association*, 18, 3045–3089.

Wald, A. (1950): *Statistical Decision Functions*, Wiley publications in statistics, New York; Chapman & Hall: London.

Williamson, D. P. and D. B. Shmoys (2011): *The design of approximation algorithms*, Cambridge university press.

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

Zhang, P. (2007): "A new approximation algorithm for the k-facility location problem," *Theoretical Computer Science*, 384, 126–135.

B Supplemental Appendix

B.1 Gurobi's Output

In this section, we discuss some parameters in the `Gurobi` optimizer that can be tuned to improve its performance.[20] First, `MIPFocus` specifies, on a high level, whether the solution should prioritize speed or optimality. The default value of this parameter is 0, which achieves a balance between searching for new solutions and proving optimality of the current solution. We set this parameter to 2, prioritizing finding good-quality and optimal solutions. A second set of important parameters affect how the solver terminates. There are two termination choices for `MIP` models: 1) a restriction on the runtime of the algorithm, such as using `TimeLimit` to limit the wall-clock runtime, and 2) controlling the *optimality* gap, by setting a parameter `MIPGap` that stops the algorithm when the relative gap between the best-known solution and the best known bound on the solution objective is less than the specified value. In this application, we let the runtime to be the default value (infinity) and set the tolerance level to 10^{-6}. In Figure 9, the `Gurobi` solver outputs not only the solution and the optimal value of the problem but also the output gap. In this example, we can see that the gap between bounds is 0, demonstrating that the k-median problem has been solved to *provable optimality*. An interesting feature of using the integer programming formulation of the k-median problem is that even when we are not able to solve the problem to provable optimality, the optimality gap provides a *suboptimality* guarantee for the obtained solution. Finally, the output reports the algorithm's runtime, which was only 0.05 seconds for this example. There are more parameters one can tune to control the root relaxation, the aggressiveness of the cutting plane strategies and the level of presolve, tolerance parameters for primal feasibility, the integer feasibility, and more. For the applications in this paper, we let the rest of the parameters be the default.

B.2 Exact Statement and Proof of Result Discussed in Section 6.2

Consider the following specialization of our setting: $d = 1$ and therefore $X \in \mathbb{R}$, $\mathcal{S}_E = \{1, 4\}$, $\mathcal{S}_P = \{2, 3\}$, and for all sites we just have $X_s = s$. (Intuitively, we consider 4 sites lined up equidistantly on a straight line, where the middle two sites are the policy sites.) Suppose furthermore that sampling distributions of signals are degenerate; formally, $\hat{\tau}_s = \tau_s := \tau(X_s)$ for all s.

Lemma 3. Under this section's additional assumptions (see paragraph immediately above):

1. If sampling schemes may randomize, the lowest regret achievable by any combination of sampling scheme and treatment rule equals $C/2$.

[20] For more specifics, see `https://www.gurobi.com/documentation/current/refman/mip_models.html`.

```
Set parameter MIPGap to value 1e-06
Set parameter MIPFocus to value 2
Gurobi Optimizer version 11.0.2 build v11.0.2rc0 (linux64 - "Ubuntu 22.04.3 LTS")

CPU model: Intel(R) Xeon(R) CPU @ 2.20GHz, instruction set [SSE2|AVX|AVX2]
Thread count: 1 physical cores, 2 logical processors, using up to 2 threads

Optimize a model with 241 rows, 240 columns and 690 nonzeros
Model fingerprint: 0x4ef0ef14
Variable types: 0 continuous, 240 integer (240 binary)
Coefficient statistics:
  Matrix range     [1e+00, 1e+00]
  Objective range  [9e-02, 4e-01]
  Bounds range     [1e+00, 1e+00]
  RHS range        [1e+00, 6e+00]
Found heuristic solution: objective 3.9100279
Presolve time: 0.00s
Presolved: 241 rows, 240 columns, 690 nonzeros
Variable types: 0 continuous, 240 integer (240 binary)
Found heuristic solution: objective 2.5710277
Root relaxation presolved: 241 rows, 240 columns, 690 nonzeros


Root relaxation: objective 1.404464e+00, 38 iterations, 0.00 seconds (0.00 work units)

    Nodes    |    Current Node    |     Objective Bounds      |     Work
 Expl Unexpl |  Obj  Depth IntInf | Incumbent    BestBd   Gap | It/Node Time

*    0     0               0       1.4044643    1.40446  0.00%     -    0s

Explored 1 nodes (38 simplex iterations) in 0.05 seconds (0.00 work units)
Thread count was 2 (of 2 available processors)

Solution count 3: 1.40446 2.57103 3.91003

Optimal solution found (tolerance 1.00e-06)
Best objective 1.404464273486e+00, best bound 1.404464273486e+00, gap 0.0000%
Runtime: 0.05749797821044922 seconds
Value of k-median problem when k = 6 is 1.404464273486204
The optimal experimental sites are ['Switzerland', 'Czechia', 'Denmark', 'Spain', 'France', 'Ireland']
```

Figure 9: Gurobi Output

Notes: An example output from the `MIP` solver in `Gurobi` using the multi-country experiment for $k = 6$. See details of the application in Section 5.2.

2. If sampling has to be purposive (nonrandomized), the lowest regret achievable in combination with any treatment rule equals $3C/4$.

Proof. To see claim 1, consider the distribution under which $(\tau_1, \tau_2, \tau_3, \tau_4)$ is uniformly distributed over $\{(0, C, C, 0), (0, -C, -C, 0)\}$. Then the experimental sites do not yield any information, and no sampling and decision rule can improve on tossing a coin, in which case expected regret equals $C/2$. We conclude that the MMR value of this decision problem is at least $C/2$.

Next, suppose the policy maker uniformly randomizes over $\mathscr{S} \subset \{1, 4\}$ for experimentation and then implements the new policy with probability $a_2 = a_3 = [(C + \hat{\tau}_{\mathscr{S}})/2C]_0^1$; here, the notation $[X]_0^1 := \min\{\max\{X, 0\}, 1\}$ indicates clamping of X to $[0, 1]$. We will show that the worst-case expected regret under this scheme is $C/2$, which therefore is the problem's MMR value and is

attained by this rule.

Careful book-keeping reveals that the policy maker's worst-case expected regret equals

$$\max_{\substack{(\tau_1,\tau_2,\tau_3,\tau_4):\\ |\tau_s-\tau_t|\leq C|s-t|}} \frac{1}{2}(\tau_2^+ + \tau_3^+)\cdot\frac{1}{2}\left(\left[\frac{C-\tau_1}{2C}\right]_0^1 + \left[\frac{C-\tau_4}{2C}\right]_0^1\right) + \frac{1}{2}(\tau_2^- + \tau_3^-)\cdot\frac{1}{2}\left(\left[\frac{C+\tau_1}{2C}\right]_0^1 + \left[\frac{C+\tau_4}{2C}\right]_0^1\right).$$

We will solve this by considering subcases. Suppose first that τ_2 and τ_3 have different signs. Since both objective and constraints are invariant under multiplying $(\tau_1,\tau_2,\tau_3,\tau_4)$ by -1, suppose without further loss of generality that $\tau_2 \geq 0 \geq \tau_3$. The optimization problem now simplifies to

$$\max_{\substack{(\tau_1,\tau_2,\tau_3,\tau_4):\\ |\tau_s-\tau_t|\leq C|s-t|}} \frac{|\tau_2|}{4}\underbrace{\left(\left[\frac{C-\tau_1}{2C}\right]_0^1 + \left[\frac{C-\tau_4}{2C}\right]_0^1\right)}_{\equiv B\in[0,2]} + \frac{|\tau_3|}{4}\underbrace{\left(\left[\frac{C+\tau_1}{2C}\right]_0^1 + \left[\frac{C+\tau_4}{2C}\right]_0^1\right)}_{=2-B} \leq \frac{\max\{|\tau_2|,|\tau_3|\}}{2} \leq \frac{C}{2},$$

where the first inequality is justified in the display and the second one follows because τ_2 and τ_3 have different signs but differ by at most C.

Now let τ_2 and τ_3 have the same sign, which we take without further loss of generality to be positive. Then we initially observe simplification to

$$\max_{\substack{(\tau_1,\tau_2,\tau_3,\tau_4):\\ |\tau_s-\tau_t|\leq C|s-t|}} \frac{1}{4}(\tau_2+\tau_3)\left(\left[\frac{C-\tau_1}{2C}\right]_0^1 + \left[\frac{C-\tau_4}{2C}\right]_0^1\right),$$

and we can furthermore concentrate out $\tau_1 = \tau_2 - C$, $\tau_4 = \tau_3 - C$ to get

$$\max_{\substack{(\tau_2,\tau_3):\\ |\tau_2-\tau_3|\leq C}} \frac{1}{4}(\tau_2+\tau_3)\left(\left[\frac{2C-\tau_2}{2C}\right]_0^1 + \left[\frac{2C-\tau_3}{2C}\right]_0^1\right).$$

Clamping of expressions at 1 cannot bind because τ_2 and τ_3 are positive. If clamping at 0 binds for both fractions, then the objective equals 0. Suppose clamping at 0 binds for one expression, say (without further loss of generality) because $\tau_2 > 2C$, then we have simplification to

$$\max_{\substack{(\tau_2,\tau_3):\\ |\tau_2-\tau_3|\leq C}} \frac{1}{4}(\tau_2+\tau_3)\frac{2C-\tau_3}{2C}.$$

Keeping in mind that $\tau_2 > 2C$ and therefore also $\tau_3 > C$ in this subcase, evaluation of derivatives shows that this expression decreases in τ_3; hence, $\tau_3 = \tau_2 - C$. Substituting this in, one can further verify the expression to be decreasing in τ_2; therefore, the maximal value in this subcase is attained

at a boundary point also covered by the next case (and, though not essential for the argument, this value can be verified to be $3C/8$).

Finally, if no clamping binds, we can reduce the problem to

$$\max_{\substack{(\tau_2,\tau_3):\\ |\tau_2-\tau_3|\leq C}} \frac{1}{4}(\tau_2+\tau_3)\frac{4C-\tau_2-\tau_3}{2C} = \frac{C}{2},$$

where the maximum is attained by setting $\tau_2+\tau_3 = 2C$.

Regarding claim 2, by the decision problem's symmetry, it is without further loss of generality to assume that site 1 is being sampled. MMR is then at least $3C/4$ because this value is achieved if the true parameter values $(\tau_1,\tau_2,\tau_3,\tau_4)$ are equally likely to be $(0,C,2C,3C)$ or $-(0,C,2C,3C)$.

We next show that this value is attained by uniformly assigning treatment with probability $a_2 = a_3 = [(3C-2\hat{\tau}_1)/6C]_0^1$. Indeed, worst case regret of this decision rule equals

$$\max_{\substack{(\tau_1,\tau_2,\tau_3,\tau_4):\\ |\tau_s-\tau_t|\leq C|s-t|}} \frac{1}{2}(\tau_2^+ + \tau_3^+)\left[\frac{3C-2\tau_1}{6C}\right]_0^1 + \frac{1}{2}(\tau_2^- + \tau_3^-)\left[\frac{3C+2\tau_1}{6C}\right]_0^1.$$

If τ_2 and τ_3 have different signs, we can bound this value by $C/2$ just as before. Suppose they have the same sign, which we take to be positive without further loss of generality, then the problem simplifies to

$$\max_{\substack{(\tau_1,\tau_2,\tau_3,\tau_4):\\ |\tau_s-\tau_t|\leq C|s-t|}} \frac{1}{2}(\tau_2+\tau_3)\left[\frac{3C-2\tau_1}{6C}\right]_0^1 = \max_{\tau_1}\frac{1}{2}(2\tau_1+3C)\left[\frac{3C-2\tau_1}{6C}\right]_0^1 = \frac{3C}{4}.$$

Here, the first equality concentrates out $\tau_2 = \tau_1 + C$ and $\tau_3 = \tau_2 + C$; the second equality uses that clamping cannot bind (clamping at 0 would set the expression to 0, clamping at 1 would imply that $\tau_2 < 0$), after which the problem is straightforwardly solved by $\tau_1 = 0$. □

B.3 Additional Analysis of the Survey Experiment

As mentioned in Section 5, we can obtain experimental estimates of all fifteen countries using the original experiments conducted in Naumann et al. (2018). Generally, experimental estimates of all policy sites are unknown and unattainable in most real-world applications; otherwise, there is no need to solve the site selection problem. However, as an illustrative example, we will leverage the information in these experiments to quantify the magnitude of the Lipschitz constant C needed to explain the treatment heterogeneity in the data and the constant $C(\mathscr{S})$ in the assumption of

Lemma 1 that gives the result of this paper.

The outcome of interest from the survey is a categorical variable indicating survey respondents' attitudes towards immigrants: 1 for Allow None; 2 for Allow A Few; 3 for Allow Some; 4 for Allow Many. For a more straightforward interpretation, we redefine the outcome variable to be binary: we let the outcome of the survey respondent be 1 if she answers 3 or 4, indicating "support"; otherwise, we let her outcome be 0, indicating "oppose." The treatment is also a binary variable, which equals 1 (0) if the survey is about high-skilled (low-skilled) immigrants. We use a simple difference-in-means estimator to estimate the treatment effect of each country. The table below shows the point estimates and their standard errors. The point estimates speak to the difference between the percentage of people who support high-skilled immigrants and the percentage of people who support low-skilled immigrants. For example, 25.8% more survey respondents in Austria are more supportive of high-skilled immigrants, compared to low-skilled immigrants.

Table 1: Experimental Estimates of Each Country

Country	Estimate	Standard Error
Austria	0.258906	0.026490
Belgium	0.232145	0.024740
Switzerland	0.285371	0.027139
Czechia	0.222865	0.019601
Germany	0.339650	0.016396
Denmark	0.293745	0.025156
Spain	0.265763	0.022948
Finland	0.403363	0.020222
France	0.275320	0.022045
United Kingdom	0.407362	0.022651
Ireland	0.238961	0.020889
Netherlands	0.301243	0.022853
Norway	0.262747	0.025613
Sweden	0.149249	0.019271
Slovenia	0.301862	0.025517

Notes: This table presents the experimental estimates of the treatment effect of the policy in each country. The original outcome from the survey is a categorical variable indicating survey respondents' attitudes towards immigrants: 1 for Allow None; 2 for Allow A Few; 3 for Allow Some; 4 for Allow Many. We redefine the outcome variable to be binary: we let the outcome of the survey respondent be 1 if she answers 3 or 4, indicating "support"; otherwise, we let her outcome be 0, indicating "oppose." The treatment is also a binary variable, which equals 1 (0) if the survey is about high-skilled (low-skilled) immigrants. We use a simple difference-in-means estimator to estimate the treatment effect of each country and their standard error.

In Figure 10a, each point represents a pair of two countries, and the slope from the origin to each point represents the smallest Lipschitz constant needed to explain the data observed for these two countries. Hence, the Lipschitz constant C that is able to explain the treatment effect heterogeneity

in the data for all countries is at least

$$\max\left\{\frac{|\hat{\tau}(X_i)-\hat{\tau}(X_j)|}{\|X_i-X_j\|}\right\}, \quad \forall i,j \in \mathcal{S}_E \cup \mathcal{S}_P, \tag{44}$$

which equals 0.0892 and corresponds to the slope of the red dashed line in Figure 10. The pair of countries that give the maximum of equation 44 is Finland and Sweden. Additionally, Montiel Olea et al. (2023) show that, for each possible set of experimental sites $\mathscr{S}$, the $C(\mathscr{S})$ that gives the solution in Lemma 1 is defined as

$$C(\mathscr{S}) := \max_{s\in\mathcal{S}_P\setminus\mathscr{S}}\left\{\sqrt{\frac{\pi}{2}}\frac{\sigma_{N_{\mathscr{S}(s)}}}{\left\|X_s - X_{N_{\mathscr{S}(s)}}\right\|}\right\}. \tag{45}$$

Replacing $\sigma_{N_{\mathscr{S}(s)}}$ with the corresponding estimated standard errors, we get $C(\mathscr{S})$ equals 0.0233, which corresponds to the slope of the blue dashed line in figure 10a. Both the numbers and the plot indicate the smallest Lipschitz constant C needed to explain the data is bigger than the largest lower bound of $C(\mathscr{S})$ that gives the nearest neighbor result. Additionally, Figure 10b shows a histogram of $C(\mathscr{S})$ for all possible $\mathscr{S}$, and the red dashed line is the smallest Lipschitz constant compatible with data. By visual inspection, we can conclude that, in this application, the assumption $C > C(\mathscr{S})$ is likely to hold. It is worth pointing out that this assumption is in general not testable because the experimental estimates of all policy sites are unknown and we cannot compute $C(\mathscr{S})$.

B.4 A Simple Example that Motivates Assumption 4

In this section, we provide a simple linear regression example to motivate Assumption 4. Suppose the effect of the status-quo is known in all sites and normalized to zero. For each site $s \in \mathcal{S}$, we have a random sample of n_s units in the experiment. Let $Y_{i,s}$ be the outcome under the policy of interest for unit $i \in \{1,...,n_s\}$. We assume that $Y_{i,s}$ is generated as follows:

$$Y_{i,s} = \beta X_{i,s} + \gamma Z_{i,s} + \varepsilon_{i,s},$$

where $X_{i,s} \in \mathbb{R}$ is the observed unit-level covariate for individual i in site s, $Z_{i,s} \in \mathbb{R}$ is the *unobserved* unit-level covariate, $\varepsilon_{i,s} \sim \mathcal{N}(0,\sigma^2_{\varepsilon,s})$ is an error term with $\sigma^2_{\varepsilon,s} > 0$ and is also independent of $(X_{i,s}, Z_{i,s})$, and $\beta,\gamma \in \mathbb{R}$ are the same across different sites. For simplicity, suppose that $X_{i,s}$ and $Z_{i,s}$ are jointly normal: $(X_{i,s}, Z_{i,s})^\top \sim \mathcal{N}(\mu_s, \Sigma_s)$, where $\mu_s \in \mathbb{R}^2$ and Σ_s is a 2×2 positive-definite covariance matrix.

Let $\bar{Y}_s := \frac{1}{n_s}\sum_{i=1}^{n_s} Y_{i,s}$ be the sample average of the observed outcome at site s, and $\bar{X}_s :=$

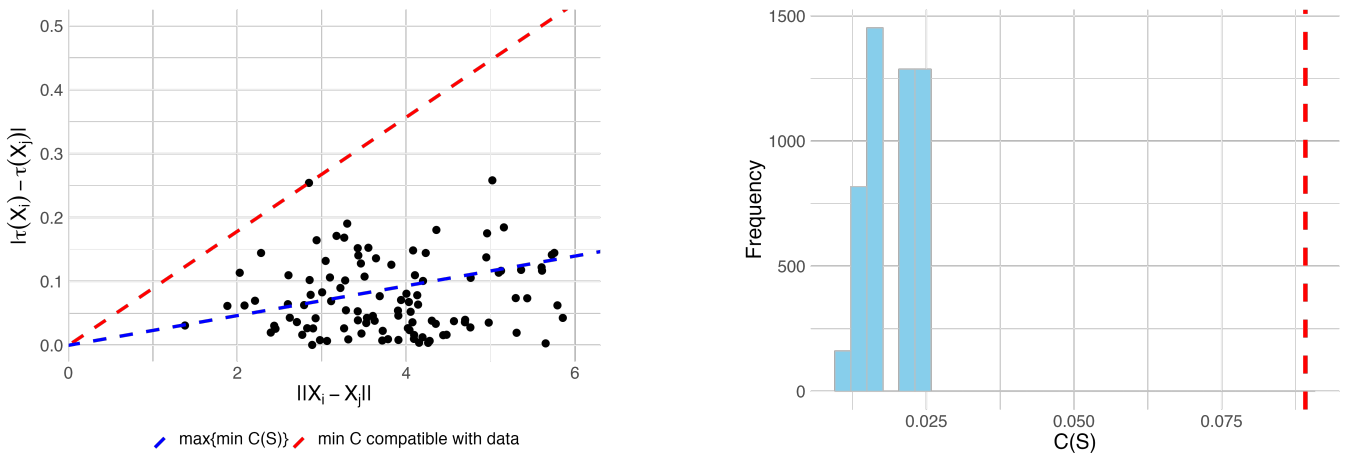

(a) The C Compatible with Data and $C(\mathscr{S})$

(b) Histogram of $C(\mathscr{S})$

Figure 10: Lipschitz Constant C and $C(\mathscr{S})$

Notes: The dots in Panel 10a represent all possible pairs among the 15 countries. The slope of each point connecting to the origin represents the value of the Lipschitz constant needed to explain the estimated treatment effects and site-level covariates observed for that pair of countries. The red dash line is the smallest C that is needed to explain all the data, computed using equation 44. The slope of the blue dashed line represents $C(\mathscr{S})$, computed using 45. Panel 10b presents a histogram of $C(\mathscr{S})$ for all possible choices of experimental sites $\mathscr{S}$. The red dashed line is the smallest Lipschitz constant needed to explain the data, corresponding to the same red dashed line in panel 10a.

$\frac{1}{n_s}\sum_{i=1}^{n_s} X_{i,s}, \bar{Z}_s := \frac{1}{n_s}\sum_{i=1}^{n_s} Z_{i,s}$ be the observed and *unobserved* site-level aggregate covariates, respectively. Under the above assumptions and conditional on observed site-level aggregate covariate $\bar{X}_s$, we have

$$\bar{Y}_s \mid \bar{X}_s \sim \mathcal{N}(\beta \bar{X}_s + \gamma \mathbb{E}[\bar{Z}_s|\bar{X}_s], \sigma_s^2), \tag{46}$$

where $\mathbb{E}[\bar{Z}_s|\bar{X}_s]$ is the expectation of $\bar{Z}_s$ conditional $\bar{X}_s$, and $\sigma_s^2 > 0$. Then, Assumption 4 models a case for which the policy effect of interest is

$$\tau_s := \beta \bar{X}_s + \gamma \mathbb{E}[\bar{Z}_s|\bar{X}_s], \forall s \in \mathcal{S}.$$

Furthermore, (46) implies that, conditional on $\bar{X}_s$, we have an unbiased and normal estimator for τ_s. We may also calculate that $\mathbb{E}[\bar{Z}_s|\bar{X}_s] = \alpha_s$ for some α_s that depends on each site s. Then, for $s, s' \in \mathcal{S}$, $s \neq s'$, we have:

$$\tau_s - \tau_{s'} = \beta\left(\bar{X}_s - \bar{X}_{s'}\right) + \gamma\left(\alpha_s - \alpha_{s'}\right), \tag{47}$$

implying that $|\tau_s - \tau_{s'}|$ can be bounded as in Assumption 4 for some positive C and c, as long as

we are willing to assume that β, γ are bounded, and that $|\alpha_s - \alpha_{s'}|$ are bounded uniformly among all $s, s' \in \mathcal{S}$.

Robust Bayes Treatment Choice with Partial Identification*

Andrés Aradillas Fernández† José Luis Montiel Olea‡ Chen Qiu§
Jörg Stoye¶ Serdil Tinda‖

August 2024

Abstract

We study a class of binary treatment choice problems with partial identification, through the lens of robust (multiple prior) Bayesian analysis. We use a convenient set of prior distributions to derive ex-ante and ex-post robust Bayes decision rules, both for decision makers who can randomize and for decision makers who cannot.

Our main messages are as follows: First, ex-ante and ex-post robust Bayes decision rules do not tend to agree in general, whether or not randomized rules are allowed. Second, randomized treatment assignment for some data realizations can be optimal in both ex-ante and, perhaps more surprisingly, ex-post problems. Therefore, it is usually with loss of generality to exclude randomized rules from consideration, even when regret is evaluated ex-post.

We apply our results to a stylized problem where a policy maker uses experimental data to choose whether to implement a new policy in a population of interest, but is concerned about the external validity of the experiment at hand (Stoye, 2012a); and to the aggregation of data generated by multiple randomized control trials in different sites to make a policy choice in a population for which no experimental data are available (Manski, 2020; Ishihara and Kitagawa, 2021).

KEYWORDS: treatment choice, partial identification, robust Bayes, Γ-minimax, posterior robustness

*We would like to thank Karun Adusumilli, Tim Armstrong, Tim Christensen, Patrik Guggenberger, Kei Hirano and Toru Kitagawa for helpful feedback, comments, and suggestions. We gratefully acknowledge financial support from NSF grant SES-2315600.

†Department of Economics, Cornell University, aa2679@cornell.edu
‡Department of Economics, Cornell University, jlo67@cornell.edu
§Department of Economics, Cornell University, cq62@cornell.edu
¶Department of Economics, Cornell University, stoye@cornell.edu
‖Department of Economics, Stanford University, serdil@stanford.edu

1 Introduction

A policy maker must decide between implementing a new policy or preserving the status quo. Her data provide information about the potential benefits of these two options. Unfortunately, these data only *partially identify* payoff-relevant parameters and may therefore not reveal, even in large samples, the correct course of action. Such *treatment choice problems with partial identification* have recently received growing interest; for example, see D'Adamo (2021), Ishihara and Kitagawa (2021), Yata (2021), Christensen, Moon, and Schorfheide (2022), Kido (2022) or Manski (2024). Several interesting problems that arise in empirical research can be recast using this framework. See Montiel Olea, Qiu, and Stoye (2023) and the end of this section for references.

This paper applies the *robust Bayes* approach—which interpolates between Bayesian and agnostic minimax analyses by evaluating minimax risk over a set of priors—to a class of treatment choice problems with partial identification. The use of the robust Bayes approach has drawn recent attention in problems that feature partial identification (Giacomini and Kitagawa, 2021; Giacomini, Kitagawa, and Read, 2021; Christensen et al., 2022). Indeed, due to partial identification, integrating Bayesian and minimax elements into decision making can be particularly attractive (see Poirier 1998; Moon and Schorfheide 2012 and references therein).

The robust Bayes approach can be applied *ex-ante* or *ex-post*, depending on whether the multiple priors are used to evaluate payoffs before or after seeing the data.[1] These concepts represent two different ways of resolving model ambiguity and sampling uncertainty, and both have been proposed to improve Bayesian robustness in decision problems.

By the well-known *dynamic consistency* of Bayesian decision making, ex-ante and ex-post robust Bayes coincide with each other—and with standard Bayes optimality—if the set of priors is a singleton. Indeed, this equivalence is used to calculate Bayes optimal decisions in practice because ex-post rules are usually computed but ex-ante Bayes optimality is claimed. As pointed out for the present context by Giacomini et al. (2021), they do not in general agree otherwise.[2] This raises the question: What are the qualitative and quantitative relationships between ex-ante and ex-post robust Bayes criteria when applied to treatment choice problems with partial identification? In this paper, we provide an answer to these questions in a general class of treatment choice problems, using a convenient class of priors.

[1] Giacomini et al. (2021) discuss both notions; they refer to the ex-ante and ex-post problems as "Gamma-minimax" and "Conditional Gamma-minimax", respectively. Christensen et al. (2022) focus on the ex-post problem for treatment choice problems with partial identification in a restricted class of decision rules.

[2] For example, for an estimation problem with a quadratic loss, Kitagawa (2012, Appendix B) derives the ex-post Γ-minimax action and shows that it is not ex-ante Γ-minimax optimal. This example does not apply to the treatment choice problem with partially-identified welfare, which is the main focus of this paper.

We use the same framework as Yata (2021) and Montiel Olea et al. (2023) but impose a simple instance of Giacomini and Kitagawa's (2021) set of priors, namely a symmetric and uniform two-point prior for reduced-form parameters and no restriction at all on unidentified parameters given reduced-form parameters. Working with regret, we formally define ex-ante and ex-post robust Bayes and, following Berger (1985), label them as "Γ-minimax regret" (Γ-MMR) and "Γ-posterior expected regret" (Γ-PER).[3] We then precisely characterize when these notions coincide and when they disagree. The main qualitative insights are as follows:

- Ex-ante Γ-MMR and ex-post Γ-PER criteria do not, in general, agree (although they may sometimes coincide). This conclusion does not hinge on whether or not randomization rules are allowed. If randomization is permitted, the criteria differ whenever updating the prior for the reduced-form parameter does not (ex-ante almost surely) resolve ambiguity regarding the sign of the optimal action. Even if randomization is not allowed, the criteria will disagree whenever the *identification power* of the model (in a sense we make precise) is sufficiently small compared to the *informativeness of the data*.

- Randomized rules are general solution concepts in both Γ-MMR and Γ-PER problems. That is, optimal rules often can or even must randomize even if regret is evaluated ex-post. For the ex-ante Γ MMR problem, whenever the identification power of the model is sufficiently small compared to the informativeness of the data, (1) infinitely many optimal rules exist, and (2) many of them are randomized rules. For the ex-post Γ-PER criterion, one might conjecture by analogy to single-prior Bayes inference that the optimal rules do not randomize. However, as long as (1) there exists ambiguity regarding the sign of the optimal action at the location of the two-point prior for the reduced-form parameter and (2) randomized rules are allowed, the optimal rule is unique and is randomized. Hence, exclusion of randomized rules is with loss of generality and cannot be justified by simply evaluating regret ex-post—it must come from other considerations, be the second-order preferences or logistical or convenience concerns.

Our point is not to advocate for either notion of robust Bayes criterion. We also do not aim to solve for robust Bayes criteria for more general sets of priors as this would get much more involved but (we suspect) not much more instructive to illustrate the points discussed above. What we hope to illustrate is when, and how Γ-MMR and Γ-PER criteria differ. We will also relate these results to timing assumptions in a fictitious game between the policy maker and an adversarial Nature.

[3] Among others, see Savage (1951), Manski (2004), Stoye (2012b), and Montiel Olea et al. (2023) for justifications of focusing on regret in treatment choice problems. In particular, while minimax loss can be an attractive alternative to minimax regret, the aforecited papers find that it leads to trivial recommendations in treatment choice settings; these findings are easily extended to our example.

Several auxiliary findings might be of independent interest. First, for Γ-MMR, even if we restrict the set of decision rules to be a class of non-randomized threshold rules based on the "efficient" linear index, the optimal threshold is not always zero. Given the apparent symmetry of the problem, we find this feature rather surprising. Second, whenever the dimension of the signal is larger than 1, there always exist (regardless of the parameter space and the variance of the signals) non-randomized linear-index threshold rules (with a threshold equal to zero) that are Γ-MMR optimal (among all decision rules). This is in stark contrast to Montiel Olea et al. (2023), in which no linear index rule is globally minimax regret optimal if the degree of partial identification is severe. The intuition is that the prior much reduces the state space; the signal space then becomes so rich relative to the state space that even linear threshold rules can effectively mimic randomization.

The literature on treatment choice with partially identified parameter has been growing since Manski (2004) and Dehejia (2005). For partial identification with known distribution of data, Manski (2000a, 2005, 2007a) and Stoye (2007) find minimax regret optimal treatment rules. For finite-sample minimax regret results with model ambiguity and sampling uncertainty, see Stoye (2012a,b), Yata (2021), Ishihara and Kitagawa (2021) and Montiel Olea et al. (2023); Kido's (2023) analysis is asymptotic. Bayes and robust Bayes approaches are analyzed by Chamberlain (2011), Giacomini and Kitagawa (2021), Giacomini et al. (2021), Christensen et al. (2022), among others. Earlier investigations of ex-ante and ex-post Γ-minimax estimators include DasGupta and Studden (1989); Betrò and Ruggeri (1992). Also see Vidakovic (2000) for a review. For different settings with point-identified welfare, finite- and large-sample results on optimal treatment choice rules were derived by Canner (1970), Chen and Guggenberger (2024), Hirano and Porter (2009, 2020), Kitagawa, Lee, and Qiu (2022), Schlag (2006), Stoye (2009), and Tetenov (2012b). There is also a large literature on optimal policy learning with covariates containing results with point identified (Bhattacharya and Dupas, 2012; Kitagawa and Tetenov, 2018, 2021; Mbakop and Tabord-Meehan, 2021; Kitagawa and Wang, 2023; Athey and Wager, 2021; Kitagawa, Sakaguchi, and Tetenov, 2021; Ida, Ishihara, Ito, Kido, Kitagawa, Sakaguchi, and Sasaki, 2022) as well as partially identified (Kallus and Zhou, 2018; Ben-Michael, Greiner, Imai, and Jiang, 2021; Ben-Michael, Imai, and Jiang, 2022; D'Adamo, 2021; Christensen et al., 2022; Adjaho and Christensen, 2022; Kido, 2022; Lei, Sahoo, and Wager, 2023) parameters. Guggenberger, Mehta, and Pavlov (2024) and Manski and Tetenov (2023) analyze related problems but focus on quantile, as opposed to expected, loss; Song (2014) considers partial identification but mean squared error regret loss.

The rest of this paper is organized as follows. Section 2 sets up the problem, provides examples, and defines both versions of robust Bayes optimality. Section 3 contains complete solutions for all aforementioned scenarios, Section 4 relates them to the "fictitious game" interpretation of minimax

theory, and Section 5 concludes. Proofs and auxiliary results are collected in the Appendix.

2 Framework

2.1 Actions, Payoffs, Statistical Model, and Decisions

Our setup follows Montiel Olea et al. (2023), who in turn follow Ferguson (1967) and others. Consider a policy maker who needs to choose an *action* $a \in [0,1]$ interpreted as probability of assigning treatment in the target population.[4] Her payoff when taking action $a \in [0,1]$ is captured by the welfare function

$$W(a,\theta) := aW(1,\theta) + (1-a)W(0,\theta), \tag{2.1}$$

where $\theta \in \Theta$ is an unknown state of the world or *parameter* and the functions $W(1,\cdot\,) : \Theta \to \mathbb{R}$ and $W(0,\cdot\,) : \Theta \to \mathbb{R}$ are known. Here, we may interpret $W(1,\theta)$ and $W(0,\theta)$ as the welfare of actions $a = 1$ (treating everyone in the population) and action $a = 0$ (treating no one in the population). Therefore, (2.1) implies that welfare is linear in actions, a standard assumption in the literature. Denote by $U(\theta) := W(1,\theta) - W(0,\theta)$ the *welfare contrast* at θ. If $U(\theta)$ were known to the policy maker, her optimal action would simply be

$$\mathbf{1}\left\{U(\theta) \geq 0\right\}. \tag{2.2}$$

Since $U(\theta)$ is in fact unknown, the policy maker gathers some empirical evidence to learn about θ. We assume that she observes a random vector $Y \in \mathbb{R}^n$ that follows a multivariate normal distribution:

$$Y \sim N(m(\theta), \Sigma), \tag{2.3}$$

where the function $m(\cdot\,) : \Theta \to \mathbb{R}^n$ and the positive definite matrix Σ are known.

Our focus is on the case when the data is not entirely informative about the sign of $U(\theta)$: Even if the policy maker perfectly learned $m(\theta)$, she could not (necessarily) pin down the sign of $U(\theta)$. To formally model such *treatment choice problems with (decision-relevant) partial identification*, let

$$M := \{\mu \in \mathbb{R}^n : m(\theta) = \mu, \theta \in \Theta\} \tag{2.4}$$

[4]Randomization could be i.i.d. across future potential treatment recipients, *fractional* in the sense of randomly assigning a certain fraction of the treatment population (in this sense, $a \in [0,1]$ can also be interpreted as the fraction of the population receiving the treatment), or an "all or nothing" randomization for the entire treatment population. While these might not be practically equivalent in all applications, they are in the current decision theoretic framework. See Manski and Tetenov (2007) for an exception in the related literature.

collect all the means of Y that can be generated as θ ranges over Θ. We refer to elements $\mu \in M$ as *reduced-form* parameters because they are identified in the statistical model (2.3) without further assumptions. Define the *identified set* for the welfare contrast given μ as

$$I(\mu) := \{u \in \mathbb{R} : U(\theta) = u, m(\theta) = \mu, \theta \in \Theta\} \tag{2.5}$$

and the corresponding upper and lower bounds as

$$\overline{I}(\mu) := \sup I(\mu), \quad \underline{I}(\mu) := \inf I(\mu). \tag{2.6}$$

Henceforth, when we refer to a treatment choice problem with partial identification, we mean there exists some open set in M such that for all μ in that open set, $\underline{I}(\mu) < 0 < \overline{I}(\mu)$.

A *decision rule* $d : \mathbb{R}^n \to [0, 1]$ is a (measurable) mapping from data Y to the unit interval $[0, 1]$. We call d *non-randomized* if it (almost surely, a.s.) maps into $\{0, 1\}$; otherwise, we call d *randomized*, including if it randomizes for some but not all data realizations. We use $\mathcal{D}_n$ to denote the set of all decision rules, and we consider decision rules the same if they a.s. agree. As a result, a rule is *unique* only up to a.s. agreement. The *oracle policy* $d^*_{\text{oracle}} := \mathbf{1}\{U(\theta) \geq 0\}$ is of special interest and for any given θ is contained in $\mathcal{D}_n$, but is not feasible in the statistical sense because $U(\theta)$ is not known.

In general, there will not be an unambiguously best feasible decision rule, a problem that gave rise to a large literature on different optimality criteria and their implementation. Before introducing the robust Bayes approach, we give two examples that fit into our general framework.

2.2 Examples

Example 1 (Stoye, 2012a)**.** This example is the one-dimensional version of the general setup and has been frequently analyzed before (Manski, 2000b; Brock, 2006; Stoye, 2012a; Tetenov, 2012a; Kitagawa, Lee, and Qiu, 2023). As we will explain below, one motivation for this example is to think of a policy maker that uses experimental data to choose whether to implement a new policy in a population of interest, but is concerned about the external validity of the experiment at hand. The treatment effect of action $a = 1$ is $\mu^* \in \mathbb{R}$, while the effect of action $a = 0$ is normalized to 0; thus, the policy maker's expected payoff equals $W(a, \mu^*) := a \cdot \mu^*$. The policy maker observes a realization of the one-dimensional statistic

$$\hat{\mu} \sim N(\mu, \sigma^2), \tag{2.7}$$

where $\sigma > 0$ is known and where $\mu \in \mathbb{R}$ is an *observable treatment effect*, in the sense that μ can be perfectly pinned down if infinite amount of data is provided. In this example, $\theta = (\mu, \mu^*)^\top$, $\Theta \subseteq \mathbb{R}^2$, $m(\theta) = \mu$, and $U(\theta) = \mu^*$.

Since the target population and the population from which data (2.7) is collected can be different, partial identification naturally arises. We assume that the observable and true treatment effects are constrained by $|\mu^* - \mu| \leq k$ for some known $k \geq 0$, implying

$$I(\mu) = [\mu - k, \mu + k], \quad \overline{I}(\mu) = \mu + k, \quad \underline{I}(\mu) = \mu - k, \quad \forall \mu \in \mathbb{R}.$$

The goal of the planner is to choose a statistical decision rule $d \in \mathcal{D}_1 : \mathbb{R} \to [0, 1]$ that maps observed data (2.7) to an action $a \in [0, 1]$ to be applied to the target population. Taking the normality assumption as an approximation, this stylized example could reflect model uncertainty (e.g., a treatment effect is estimated in a possibly somewhat misspecified model), external validity concerns (e.g., a randomized clinical trial was performed on volunteers), or a shift in the environment (e.g., we transfer estimates from study populations to treatment populations with slightly different covariates or are concerned about distributional drift over time). □

Example 2 (Ishihara and Kitagawa, 2021)**.** This example is taken from Ishihara and Kitagawa's (2021; see also Manski (2020)) "evidence aggregation" framework. A policy maker is interested in implementing a new policy in country $i = 0$ and observes estimates of the policy's effect for countries $i = 1, ..., n$. Let $Y = (Y_1, ..., Y_n)^\top \in \mathbb{R}^n$ denote these estimates and let $(x_0, \ldots, x_n)$ be nonrandom, d-dimensional baseline covariates. The policy maker is willing to extrapolate from her data by assuming that the welfare contrast of interest equals $U(\theta) = \theta(x_0)$ and that

$$Y = \begin{pmatrix} Y_1 \\ \vdots \\ Y_n \end{pmatrix} \sim N(m(\theta), \Sigma), \quad m(\theta) = \begin{pmatrix} \theta(x_1) \\ \vdots \\ \theta(x_n) \end{pmatrix}, \quad \Sigma = \operatorname{diag}(\sigma_1^2, \ldots, \sigma_n^2),$$

where $\theta : \mathbb{R}^d \to \mathbb{R}$ is an unknown Lipschitz function with known constant C. For notational simplicity, we can further write μ_i for $\theta(x_i)$. Thus, $Y_i \sim N(\mu_i, \sigma_i^2)$, $Y \sim N(\mu, \Sigma)$ and $U(\theta) = \mu_0$. The policy question is: Given data Y, what proportion of the population in country $i = 0$ should be assigned the new policy?

Let $\|\beta\| := \sqrt{\beta^\top \beta}$ be the Euclidean norm of a vector β. In this example, the identified set for the welfare contrast μ_0 is

$$I(\mu) = \{u \in \mathbb{R} : |\mu_i - u| \leq C \|x_i - x_0\|, \quad i = 1, \ldots, n\}.$$

The lower and upper bounds of $I(\mu)$ are simply intersection bounds:

$$\underline{I}(\mu) = \max_{i=1,\ldots,n} \{\mu_i - C\,\|x_i - x_0\|\}, \quad \overline{I}(\mu) = \min_{i=1,\ldots,n} \{\mu_i + C\,\|x_i - x_0\|\}\,.$$

□

2.3 Robust Bayes Optimality

Our setting up to here is as in Montiel Olea et al. (2023). We now connect it to the robust Bayes literature by imposing a *set of priors* Γ on θ. Specifically, for each prior π on θ, let π_μ be the corresponding prior on the point-identified reduced-form parameter $\mu \in M$, obtained as the push-forward measure induced by π and $m(\cdot)$, i.e.,

$$\pi_\mu(A) = \pi(m^{-1}(A)), \forall A \in \mathcal{F}_M,$$

where $m^{-1}(A) := \{\theta \in \Theta : m(\theta) \in A\}$ and $\mathcal{F}_M$ denotes the $\sigma-$algebra of M. Following Giacomini and Kitagawa (2021), we choose a particular single proper prior π_μ for $\mu \in M$ but leave the conditional prior of θ given μ, denoted as $\pi_{\theta|\mu}$, unrestricted except for

$$\pi_{\theta|\mu}\{U(\theta) \in I(\mu)\} = 1, \quad \pi_\mu\text{-a.s.} \tag{2.8}$$

Then, the class of priors Γ consists of all priors on θ induced by the single prior π_μ and any conditional prior $\pi_{\theta|\mu}$ that meets (2.8). Intuitively, we choose a single prior on the point-identified parameter and place no new restriction on the partially identified parameter $U(\theta)$. One can pick any proper prior π_μ; for tractability, we let π_μ to be supported on two symmetric points $\{\bar{\mu}, -\bar{\mu}\}$ with equal probability, where $\bar{\mu} \in \mathbb{R}^n$ is chosen by the decision maker. Henceforth, Γ is understood to refer to the implied set of priors:

$$\Gamma := \left\{ \pi_\theta = \int \pi_{\theta|\mu} d\pi_\mu : \pi_\mu \sim \text{unif}(\{-\bar{\mu}, \bar{\mu}\}), \pi_{\theta|\mu} \text{ satisfies (2.8)} \right\}. \tag{2.9}$$

While the set of priors Γ broadly puts us into the "robust Bayes" territory, it still does not pin down a uniquely best decision rule because the sign of $U(\theta)$ can remain ambiguous. Let

$$L(a, \theta) := \sup_{a' \in [0,1]} W(a', \theta) - W(a, \theta) = U(\theta)\,\{\mathbf{1}\{U(\theta) \geq 0\} - a\}$$

be the *regret* of action $a \in [0, 1]$. We will evaluate decision rules d by their *expected regret*, defined

as

$$\begin{aligned} R(d,\theta) &:= \mathbb{E}_{m(\theta)}[L(d(Y),\theta)] \\ &= U(\theta)\left\{\mathbf{1}\{U(\theta)\geq 0\} - \mathbb{E}_{m(\theta)}[d(Y)]\right\}, \end{aligned} \tag{2.10}$$

where for any $x \in \mathbb{R}^n$, $\mathbb{E}_x[\cdot]$ denotes expectation taken over $Y \sim N(x,\Sigma)$.

Even with the set of priors Γ given and the notion of expected regret chosen as the decision criterion, the robust Bayes literature contains multiple decision criteria that do not in general agree. The difference lies in when (and how) the expectations regarding the unknown parameter $\theta \in \Theta$ are taken. For a decision rule $d \in \mathcal{D}_n$, let

$$r(d,\pi) := \int_{\theta\in\Theta} R(d,\theta)d\pi(\theta)$$

be its Bayes expected regret of under a prior $\pi \in \Gamma$. Following Berger (1985, Definition 12, p. 216), we introduce the first robust Bayes optimality notion.

Definition 1 (*Ex-ante* Γ-Minimax Regret)**.** A decision rule $d^* \in \mathcal{D}_n$ is Γ-*minimax regret* (henceforth Γ-MMR) optimal if

$$\sup_{\pi\in\Gamma} r(d^*,\pi) = \inf_{d\in\mathcal{D}_n} \sup_{\pi\in\Gamma} r(d,\pi). \tag{2.11}$$

An alternative "posterior" robustness notion is also common in Bayesian analysis. For each action $a \in [0,1]$, define *posterior expected regret* under prior π (Berger, 1985, Definition 8, p. 159)

$$\rho(a,\pi_{\theta|Y}) := \int_{\tilde{\theta}\in\Theta} L(a,\tilde{\theta})d\pi_{\theta|Y}(\tilde{\theta}),$$

where $\pi_{\theta|Y}$ is the posterior distribution of θ given prior π and data Y.[5] Then we have the following, alternative optimality criterion (Berger, 1985, Definition 10, p. 205):

Definition 2 (*Ex-post* Γ-Posterior Expected Regret)**.** A decision rule $d^* \in \mathcal{D}_n$ is Γ-*posterior expected regret* (henceforth Γ-PER) optimal if

$$\sup_{\pi\in\Gamma} \rho(d^*(Y),\pi_{\theta|Y}) = \inf_{a\in[0,1]} \sup_{\pi\in\Gamma} \rho(a,\pi_{\theta|Y}), \quad \forall Y \in \mathbb{R}^n. \tag{2.12}$$

[5] In our setting, information from data Y does not revise the conditional prior $\pi_{\theta|\mu}$ (Giacomini and Kitagawa, 2021). For any event A in the σ-algebra of Θ, we therefore have $\pi_{\theta|Y}(A) = \int \pi_{\theta|\mu}(A)d\pi_{\mu|Y}$, where $\pi_{\mu|Y}$ is the posterior distribution of μ given Y.

If the decision maker's action space is restricted to $\{0,1\}$, i.e., randomization is not allowed, the above definition is revised by replacing $[0,1]$ with $\{0,1\}$.

The labelling of Γ-MMR as "ex-ante" versus Γ-PER as "ex-post" can be related to the timing of a fictitious game against an adversarial Nature; see Section 4 for additional discussion. If Γ were a singleton, the criteria would coincide and would also agree with (single-prior) Bayes optimality. They do not in general agree otherwise. The Γ-PER criterion is often easier to solve and is routinely employed to quantify posterior robustness of statistical decisions. However, the term "Gamma minimax" overwhelmingly (and even "robust Bayes" more often than not) refers to Γ-MMR; for example, see Berger (1985, Section 4.7.6).[6]

3 Main Results

In this section, we explicitly solve for the two versions of robust Bayes optimality under the set of priors (2.9). Following Yata (2021), we assume:

Assumption 1. (i) Θ is convex, centrosymmetric (i.e., $\theta \in \Theta$ implies $-\theta \in \Theta$) and nonempty.

(ii) $m(\cdot)$ and $U(\cdot)$ are linear.

These conditions are restrictive but turn out to encompass many examples of empirical relevance; see this paper's introduction and Montiel Olea et al. (2023); Yata (2021) for examples. Under Assumption 1, we let

$$\overline{I}(\bar{\mu}) + \underline{I}(\bar{\mu}) > 0$$

as a normalization.

We say a rule is a *linear-index threshold rule* if it has the form $\mathbf{1}\left\{\beta^\top Y \geq c\right\}$ for some $\beta \in \mathbb{R}^n$ and $c \in \mathbb{R}$. Linear-index threshold rules are nonrandomized. They are of particular interest because they form a complete class when $U(\theta)$ is point-identified (Karlin and Rubin, 1956) and have received particular attention in the recent literature (Ishihara and Kitagawa, 2021; Montiel Olea et al., 2023). For both $\Gamma-$MMR and Γ-PER, we will clarify when linear-index threshold rules are optimal and when they are not. For reasons that will become obvious, the following linear-index threshold rule is of particular interest:

$$d^*_{w,0} := d^*_{w,0}(Y) := \mathbf{1}\left\{w^\top Y \geq 0\right\}, \; w := \Sigma^{-1}\bar{\mu}. \tag{3.1}$$

[6] Giacomini, Kitagawa, and Read (2021) discuss both criteria for general loss functions and refer to Definitions 1 and 2 as the "unconditional Γ-minimax" and "conditional Γ-minimax" problems, respectively. For treatment choice problems with partial identification, Christensen et al. (2022) optimize the Γ-PER criterion, restricting the action space to be $\{0,1\}$.

For each vector $\beta \in \mathbb{R}^n$, let $\|\beta\|_\Sigma := \sqrt{\beta^\top \Sigma \beta}$. Thus, $\|w\|_\Sigma = \sqrt{w^\top \Sigma w} = \sqrt{\bar{\mu}^\top \Sigma^{-1} \bar{\mu}}$. Denote by $\Phi(\cdot)$ the standard normal c.d.f. and by $\Phi^{-1}(\cdot)$ its inverse, i.e. the corresponding quantile function.

3.1 Ex-ante Robust Bayes Optimality

Theorem 1 (Γ-MMR optimal decisions)**.** *Consider a treatment choice problem with welfare function* (2.1)*, statistical model* (2.3) *and a set of priors* (2.9)*, that satisfies Assumption 1.*

(i) If

$$\frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})} \geq \Phi(\|w\|_\Sigma), \tag{3.2}$$

then $d^*_{w,0}$ is uniquely Γ-MMR optimal.

(ii) If

$$\frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})} < \Phi(\|w\|_\Sigma), \tag{3.3}$$

then a rule $d \in \mathcal{D}_n$ attains Γ-MMR if, and only if, it implies

$$\mathbb{E}_{-\bar{\mu}}[d^*(Y)] = \frac{-\underline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}, \tag{3.4}$$

$$\mathbb{E}_{\bar{\mu}}[d^*(Y)] = \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}. \tag{3.5}$$

In particular, any rule of the form $\mathbf{1}\{w^\top Y \geq c\}$ for some $c \in \mathbb{R}$ is not *Γ-MMR optimal. The following rules, and any convex combination of them, are all Γ-MMR optimal:*

$$d^*_{RT} := \Phi\left(\frac{w^\top Y}{\tilde{\sigma}}\right), \tag{3.6}$$

$$d^*_{linear} := \begin{cases} 0, & w^\top Y < -\rho^*, \\ \frac{w^\top Y + \rho^*}{2\rho^*}, & -\rho^* \leq w^\top Y \leq \rho^*, \\ 1, & w^\top Y > \rho^*, \end{cases} \tag{3.7}$$

$$d^*_{step} := \begin{cases} \frac{1}{2} - \beta^*, & w^\top Y < 0, \\ \frac{1}{2} + \beta^*, & w^\top Y \geq 0, \end{cases} \tag{3.8}$$

where

$$\tilde{\sigma} = \sqrt{\left[\frac{\|w\|_\Sigma^2}{\Phi^{-1}\left(\frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})}\right)}\right]^2 - \|w\|_\Sigma^2},$$

$\rho^* > 0$ *is unique and solves* $\int_0^1 \Phi\left(\frac{2\rho^* x - \rho^* - \|w\|_\Sigma^2}{\|w\|_\Sigma}\right) dx = \frac{-\underline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})}$, *and*

$$\beta^* = \frac{\frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})} - \frac{1}{2}}{2\Phi\left(\|w\|_\Sigma\right) - 1} \in \left(0, \frac{1}{2}\right).$$

(iii) In case (3.3), among linear-index threshold rules of the form $\mathbf{1}\left\{w^\top Y \geq c\right\}$, *the optimal thresholds are* $\pm c^*$, *where*

$$c^* := \|w\|_\Sigma^2 - \|w\|_\Sigma \Phi^{-1}\left(\frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}\right) > 0.$$

(iv) In case (3.3) and if $n > 1$, *the following linear-index threshold rule is also* Γ*-MMR optimal:*

$$d^*_{w_{t^*},0}(Y) := \mathbf{1}\left\{w_{t^*}^\top Y \geq 0\right\},$$

where $w_{t^*} = \Sigma^{-1}\left(t^*\bar{\mu} + (1-t^*)\dot{\mu}\right)$, $\dot{\mu} \neq 0$ *is such that* $\dot{\mu}^\top \Sigma^{-1} \bar{\mu} = 0$,

$$t^* := \frac{1}{1 \pm \sqrt{\frac{(1-s^*)}{s^*}\frac{\|w\|_\Sigma^2}{\|\Sigma^{-1}\dot{\mu}\|_\Sigma^2}}}, \tag{3.9}$$

and

$$s^* := \frac{\left[\Phi^{-1}\left(\frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})}\right)\right]^2}{\|w\|_\Sigma^2} \in (0,1).$$

Thus, $d^*_{w_{t^*},0}$ *is also optimal among all linear threshold rules.*

Theorem 1 reveals that the Γ-MMR rules qualitatively change depending on whether condition (3.2) is met or not. This condition admits an intuitive interpretation: Up to clamping to the unit interval (i.e., values outside this interval are mapped to its edges), the left-hand side of (3.2) equals the unique minimax regret optimal rule for *known* μ (Manski, 2007b); therefore, it arguably measures the model's identification strength. The right-hand side of (3.2) can be interpreted as the informativeness of the data about which of $\{-\bar{\mu}, \bar{\mu}\}$ obtained. Therefore, Theorem 1 says that if the model's identification power is sufficiently large compared to the informativeness of the data, then, the unique Γ-MMR optimal rule is $d^*_{w,0}$, a non-randomized linear index rule with threshold

0 that effectively ignores the partial identification issue. In contrast, if the model's identification power is small compared to the informativeness of data, there are infinitely many Γ-MMR optimal rules, all of which satisfy (3.4) and (3.5). Examples include suitably smoothed versions of $d^*_{w,0}$ like d^*_{RT} and d^*_{linear} (the functional forms of which showed up in Montiel Olea et al. 2023) as well as d^*_{step} (which is a new result).

The contrast between these regimes becomes more intuitive upon inspecting the corresponding least favorable priors. When (3.2) holds, this prior is supported at two points ($\mu = \bar{\mu}$, $U(\theta) = \overline{I}(\bar{\mu})$ versus $\mu = -\bar{\mu}$, $U(\theta) = \underline{I}(-\bar{\mu})$) with equal probability, leading to unique optimal response $d^*_{w,0}$. In words, treatment should be assigned to all iff $\mu = \bar{\mu}$; inference toward μ and toward $U(\theta)$ effectively coincide. When (3.3) holds, the least favorable prior is supported on four points ($\mu = \bar{\mu}$, $U(\theta) = \overline{I}(\bar{\mu})$; $\mu = \bar{\mu}$, $U(\theta) = \underline{I}(\bar{\mu})$; $\mu = -\bar{\mu}$, $U(\theta) = \overline{I}(-\bar{\mu})$ and $\mu = -\bar{\mu}$, $U(\theta) = \underline{I}(-\bar{\mu})$), with a probability profile such that the data are actually uninformative about $U(\theta)$. This explains nonuniqueness and randomization in optimal decision rules.[7]

In the case when the model's identification power is small, we encounter several surprising findings. First, despite the problem's apparent symmetry, $d^*_{w,0}$ is *not* optimal even among linear-index threshold rules that use the index $w^\top Y$. Instead, Theorem 1(iii) characterizes exactly two optimal thresholds, one positive and one negative. Second, while the particular linear-index threshold rule $d^*_{w,0}$ is not Γ-MMR optimal, in higher dimensional problems ($n > 1$), there do exist linear-index threshold rules that are. This finding is in stark contrast to Montiel Olea et al. (2023), who find that, in a large class of special cases, *no* linear index threshold rule is MMR optimal. The crucial difference in settings is that the set of priors much constrains the decision theoretic problem's state space; as a result, the signal space is much richer than the state space, and this can be exploited to mimic randomization without nominally randomizing. Compare Manski's (2024) abstract observation that, if a policy maker is not allowed to explicitly randomize, sampling uncertainty can be beneficial by providing an implicit randomization device. We note that this phenomenon is reminiscent of classic "purification" results in game theory (Dvoretzky, Wald, and Wolfowitz, 1951; Khan, Rath, and Sun, 2006).[8] In contrast, it is not deeply related to the general intuition that "Bayesians don't randomize."

We next apply Theorem 1 to Example 1 and immediately get the following results.

Corollary 1 (Γ-MMR decisions in Example 1)**.** *In Example 1, the following statements are true:*

[7] To be clear, this does not mean that every decision rule becomes optimal. Among other things, rules are still constrained by what one may want to think of as indifference conditions required for mixed strategy equilibrium in a fictitious game; see again Section 4.

[8] We thank Elliot Lipnowski for reminding us of this literature.

(i) If

$$\frac{\bar{\mu}+k}{2k} \geq \Phi\left(\bar{\mu}/\sigma\right), \tag{3.10}$$

then

$$d_0^*(\cdot) := \mathbf{1}\{\hat{\mu} \geq 0\}$$

is the unique Γ*-MMR optimal decision rule.*

(ii) If (3.10) *fails, then a rule* $d \in \mathcal{D}_1$ *attains* Γ*-MMR if, and only if, it implies*

$$\mathbb{E}[d(\hat{\mu}) \mid \mu = -\bar{\mu}] = \frac{-\bar{\mu}+k}{2k}, \tag{3.11}$$

$$\mathbb{E}[d(\hat{\mu}) \mid \mu = \bar{\mu}] = \frac{\bar{\mu}+k}{2k}. \tag{3.12}$$

In particular, no linear threshold rule is Γ*-MMR optimal. The following rules, and any convex combination of them, are all* $\Gamma-$*MMR optimal:*

$$d_{RT}^* := \Phi\left(\frac{\hat{\mu}}{\tilde{\sigma}}\right), \tag{3.13}$$

$$d_{linear}^* := \begin{cases} 0, & \hat{\mu} < -\frac{\sigma^2\rho^*}{\bar{\mu}}, \\ \frac{\bar{\mu}\hat{\mu}+\sigma^2\rho^*}{2\sigma^2\rho^*}, & -\frac{\sigma^2\rho^*}{\bar{\mu}} \leq \hat{\mu} \leq \frac{\sigma^2\rho^*}{\bar{\mu}}, \\ 1, & \hat{\mu} > \frac{\sigma^2\rho^*}{\bar{\mu}}, \end{cases} \tag{3.14}$$

$$d_{step}^* := \begin{cases} \frac{1}{2} - \frac{\bar{\mu}}{2k\left(2\Phi\left(\frac{\bar{\mu}}{\sigma}\right)-1\right)}, & \hat{\mu} < 0, \\ \frac{1}{2} + \frac{\bar{\mu}}{2k\left(2\Phi\left(\frac{\bar{\mu}}{\sigma}\right)-1\right)}, & \hat{\mu} \geq 0, \end{cases} \tag{3.15}$$

where

$$\tilde{\sigma} = \sigma\sqrt{\left[\frac{\bar{\mu}}{\sigma\Phi^{-1}\left(\frac{\bar{\mu}+k}{2k}\right)}\right]^2 - 1},$$

and $\rho^* > 0$ *is unique and solves* $\int_0^1 \Phi\left(\frac{2\rho^* x - \rho^* - (\bar{\mu}/\sigma)^2}{\bar{\mu}/\sigma}\right) dx = \frac{-\bar{\mu}+k}{2k}$.

(iii) In case (ii), the best linear threshold rules in terms of Γ*-minimax regret are* $\mathbf{1}\{\hat{\mu} \geq \pm c^*\}$*, where*

$$c^* = \bar{\mu} - \sigma\Phi^{-1}\left(\frac{\bar{\mu}+k}{2k}\right).$$

Note that there is no analog to Theorem 1's case (iv); indeed, no linear threshold rule is optimal in case (ii). This is because the scalar nature of the signal Y shuts down the aforementioned purification

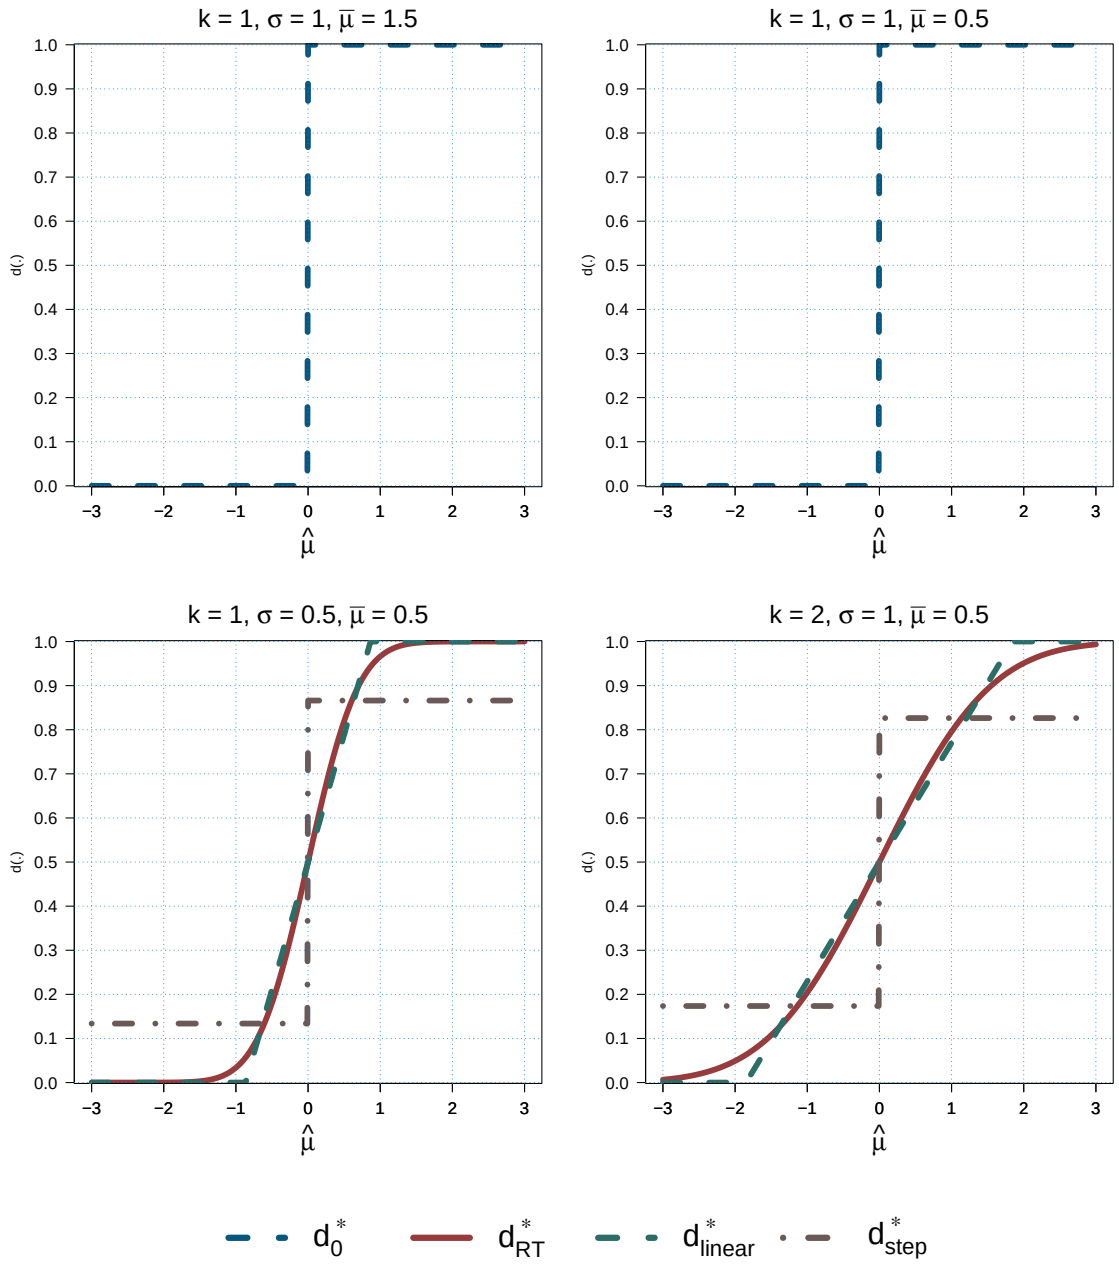

Figure 1: Γ-MMR optimal rules in Example 1

Notes: This figure reports the Γ-MMR optimal rules in Example 1 for various combinations of parameter values. In the top two panels, the combinations of parameter values satisfy (3.10). Therefore, the unique Γ-MMR optimal rule is d_0^*. In the bottom two panels, (3.10) fails. As a result, d_0^* is no longer Γ-MMR optimal. Instead, d_{RT}^*, d_{linear}^* and d_{step}^* are all Γ-MMR optimal.

mechanism. See Figure 1 for an illustration of different Γ-MMR optimal rules in Example 1 for selected parameter values of k, σ and $\bar{\mu}$.

3.2 Ex-post Robust Bayes Optimality

Theorem 2 (Γ-PER optimal decisions)**.** *Suppose all conditions in Theorem 1 hold true. Then:*

(i) If $\underline{I}(\bar{\mu}) < 0 < \overline{I}(\bar{\mu})$*, the unique* Γ*-PER optimal rule is*

$$d^*_{\text{PER}}(Y) = \begin{cases} \frac{-\underline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})}, & \text{if } w^\top Y < 0, \\ \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})}, & \text{if } w^\top Y \geq 0. \end{cases}$$

Otherwise, $d^*_{w,0}$ *is* Γ*-PER optimal.*

(ii) $d^*_{w,0}$ *is always the* Γ*-PER optimal non-randomized rule.*

The results of Theorem 2 offer some important clarifications regarding Γ-PER in treatment choice problems with partial identification. First, even if regret is evaluated according to the posterior distribution, it is not necessarily true that optimal rules are non-randomized. In fact, whenever there is model ambiguity regarding the sign of $U(\theta)$ (i.e., $\underline{I}(\bar{\mu}) < 0 < \overline{I}(\bar{\mu})$), the unique Γ-PER optimal rule is randomized. Therefore, restricting the action space to $\{0, 1\}$ in such problem is *not* without loss of generality even under the Γ-PER criterion. Comparing results with Theorem 1 also allows for instructive observations on when Γ-PER and Γ-MMR optimal rules agree or disagree; we elaborate these in Corollary 3. Applying Theorem 2 to Example 1, we finally obtain:

Corollary 2 (Γ-PER rules in Example 1)**.** *In Example 1, the following statements are true:*

(i) If $\bar{\mu} < k$*, then* Γ*-PER is uniquely minimized by*

$$d^*_{\text{PER}}(\hat{\mu}) = \begin{cases} \frac{k+\bar{\mu}}{2k} & \text{if } \hat{\mu} \geq 0 \\ \frac{k-\bar{\mu}}{2k} & \text{if } \hat{\mu} < 0 \end{cases}.$$

(ii) If $\bar{\mu} \geq k$*, then* d^*_0 *is* Γ*-PER optimal. Furthermore, it is always the* Γ*-PER optimal threshold rule.*

In Figure 2, we depict Γ-PER optimal rules for Example 1 with the same parameter values considered in Figure 1. We see clearly that Γ-MMR and -PER optimal rules coincide (if randomization is allowed) only in the special case when $k \leq \bar{\mu}$, an observation we generalize in Corollary

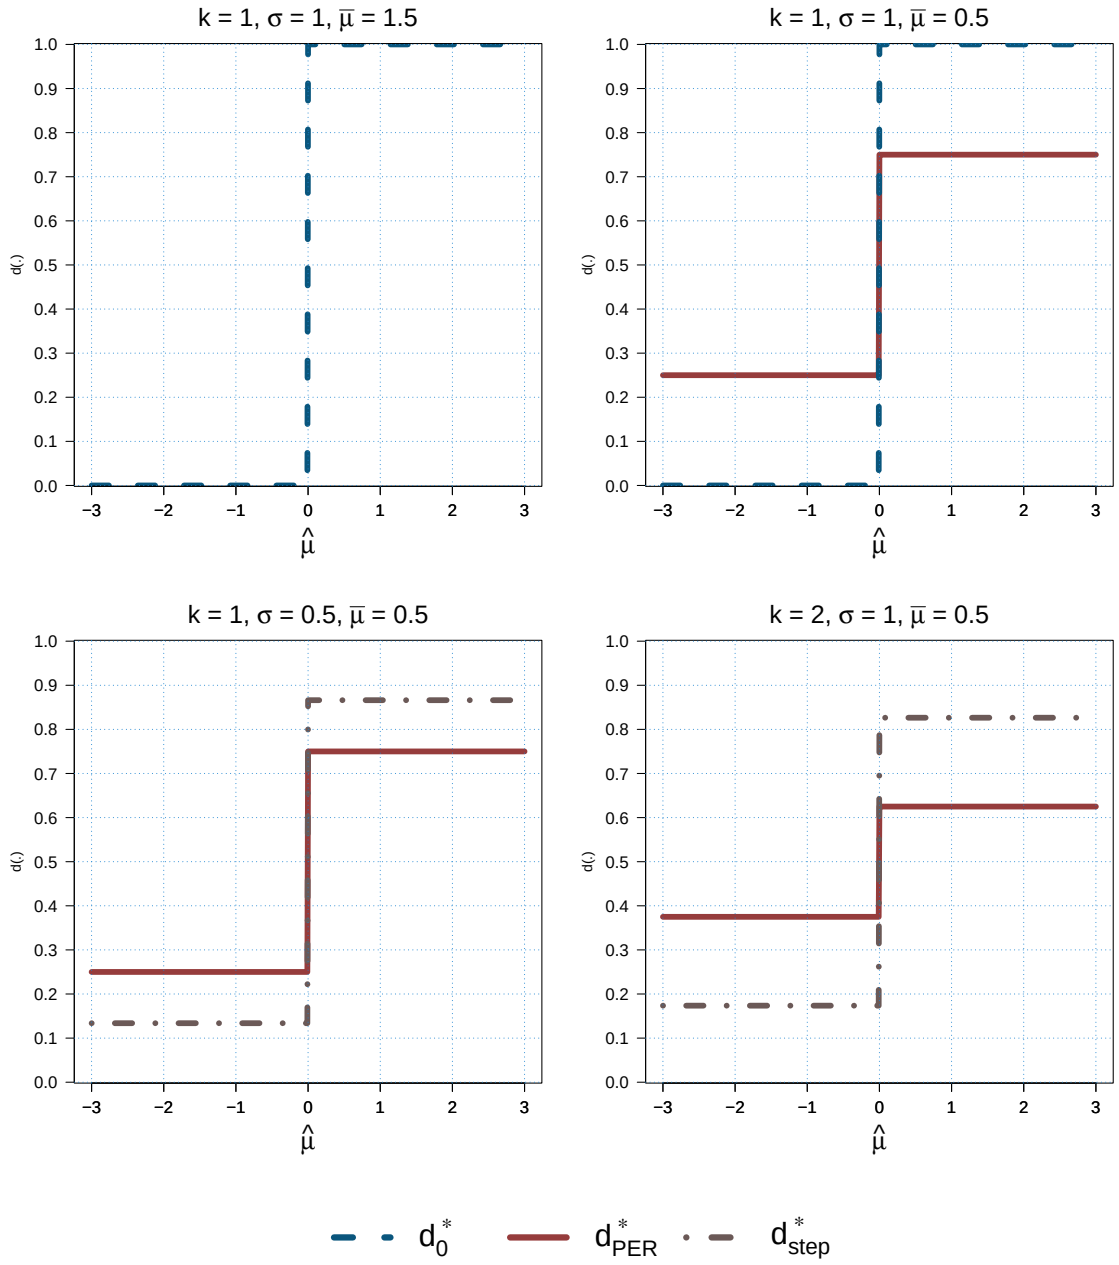

Figure 2: Γ-PER optimal rules in Example 1

Notes: This figure reports the Γ-PER optimal rules for the same parameter values considered in Figure 1. In the top left panel, Γ-PER and Γ-MMR optimal rules coincide and are both d_0^*. For the rest of the panels, the unique Γ-PER optimal rules are all d_{PER}^* and are different from any Γ-MMR optimal rules.

3. More specifically, in the top left panel of Figure 2, as $k \leq \bar{\mu}$, Γ-MMR and -PER coincide and are the non-randomized threshold rule d_0^*. In the top right panel, $k > \bar{\mu}$ and the Γ-PER optimal rule becomes d_{PER}^*. However, since (3.10) still holds, d_0^* is still Γ-MMR optimal. For the bottom two panels, as it still holds $k > \bar{\mu}$, d_{PER}^* is still Γ-PER optimal. However, the associated parameter values imply (3.10) fails. As a result, d_0^* is no longer Γ-MMR optimal and many Γ-MMR rules exist. But even in theses cases, Γ-MMR and -PER rules differ, as among the class of step function rules (which contain d_{PER}^*), only d_{step}^* is Γ-MMR optimal, still different from d_{PER}^*.

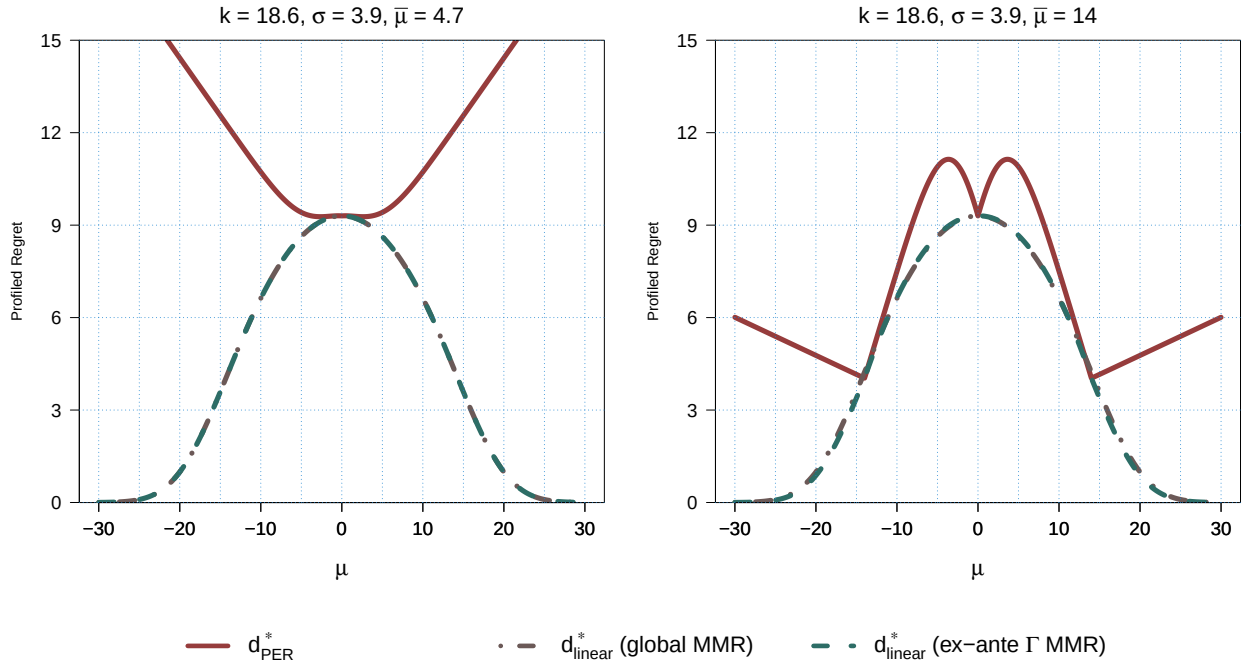

Figure 3: Profiled regret of Γ-PER and other rules in Example 1

Notes: This figure reports the (frequentist) profiled regrets, as a function of the true mean μ of data $\hat{\mu}$, of Γ-PER rule, Γ-MMR linear rule, and the least randomizing global MMR optimal rule (Montiel Olea et al. 2023) for the parameter values considered in Montiel Olea et al. (Figure 2, 2023). In the left plot, we see Γ-PER rule is visually dominated. In both plots, the profiled regret of Γ-MMR linear rule and the least randomizing global MMR optimal rule look very similar and are essentially overlapping with each other.

Applying Montiel Olea et al. (Theorem 1, 2023) to the current setting, we can conclude that all rules, including both Γ-MMR and Γ-PER optimal rules, are at least admissible. Therefore, it might be interesting and more useful to compare the (frequentist) profiled regrets (Montiel Olea et al. 2023) of ex-post Γ-PER and other rules as a function of the true but unknown mean μ of data $\hat{\mu}$.

In the context of Example 1, we can report $\bar{R}(d,\mu) := \sup_{\mu^* \in [\mu-k,\mu+k]} R(d,\mu,\mu^*)$ as μ varies for each rule d, where $R(d,\mu,\mu^*)$ is the expected regret of rule d as a function of μ and μ^*. Figure 3 gives a visual illustration. Surprisingly, when $\bar{\mu}$ is small and k large (left plot of Figure 3), the Γ-PER seems to be dominated in terms of the profiled regret. In fact, we can show that this is indeed the case whenever $\bar{\mu}$ is sufficiently small and k is sufficiently large (See Lemma D.2 in Appendix D for an exact statement). Intuitively, the Γ-PER rule mixes between a coin flip rule and the naive threshold rule (d_0^*). When $\bar{\mu}$ is small, Γ-PER rule is more analogous to the coin flip rule, which can be shown to be dominated in terms of profiled regret in Montiel Olea et al. (2023). Also note that the profiled regret of Γ-PER rule does not diminish to zero as μ goes to infinity, which may potentially be an undesirable property. Furthermore, Figure 3 also reveals that the $\Gamma-$MMR rule d^*_{linear}, though not globally MMR optimal, is numerically equivalent to the least randomizing global MMR optimal rule in some cases, especially when $\bar{\mu}$ is small. This interestingly provides a Robust Bayes interpretation of the least randomizing global MMR optimal rule derived by Montiel Olea et al. (2023).

3.3 When Do Ex-ante and Ex-post Agree?

Theorems 1 and 2 clarify to what extent Γ-MMR and Γ-PER criteria yield the same or different optimal rules. We spell this out in Corollary 3, which considers both cases when randomization is allowed and not allowed. The bottom line is that the assessment is sensitive to how the problem is set up. If underlying parameters lead to sufficiently small identification power, the criteria disagree.

Corollary 3. *Consider a treatment choice problem with welfare function* (2.1) *and statistical model* (2.3) *that satisfies Assumption 1.*

(i) Suppose randomization is allowed. Then the Γ*-MMR and* Γ*-PER optimal rules coincide if, and only if,* $\underline{I}(\bar{\mu}) \geq 0$.

(ii) Suppose randomization is not *allowed. Then the* Γ*-MMR and* Γ*-PER optimal rules coincide if, and only if,*

$$\frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})} \geq \Phi(\|w\|_\Sigma). \tag{3.16}$$

Thus, if randomization is allowed, Γ-MMR and Γ-PER optimal rules coincide only in the somewhat trivial case in which we a priori know that $U(\theta)$ and μ have the same sign, so that optimal treatment choice reduces to Bayesian inference on the point identified μ. They disagree in all other cases, including in settings where there are infinitely many $\Gamma-$MMR optimal rules.

One might conjecture that the two criteria agree more once randomization is excluded; after all, this leads to a much simpler action space. Part (ii) shows that there is some truth to this: The condition for agreement changes from $\underline{I}(\bar{\mu}) \geq 0$ to the strictly weaker (3.16). However, the criteria continue to disagree in many cases. It may be instructive to think of these cases in terms of "comparative statics." For example, consider holding all parameters of the problem fixed but scaling the signal variance Σ by a positive scalar, for example to reflect a change in sample size. Then (3.16) will hold if, and only if, Σ is *large* enough; hence, as long as $\underline{I}(\bar{\mu}) \geq 0$, increasing sample size will eventually cause disagreement between Γ-MMR and -PER even if randomization is excluded. Similarly, for fixed Σ, Γ-MMR and -PER rules will always disagree if the model's identification power is *small* enough.

4 Game Theoretic Interpretation

As Wald (1940) was first to discover, it can be instructive to think of minimax decision rules as equilibrium rules in a fictitious, zero-sum game against Nature. In particular, any equilibrium of the game in which the decision maker picks a decision rule and an adversarial Nature picks a prior characterizes a minimax (regret) decision rule and corresponding prior; furthermore, known features of zero-sum games may allow one to claim uniqueness of the solution and the like.[9] While usually stated for "agnostic" minimax rules, it is easy to see that these facts extend to our setting. They are liberally used in this paper's proofs; we here use them to briefly discuss some intuitions and contrast ex-ante versus ex-post robust Bayes, but also randomization versus no randomization, in terms of the fictitious game's timing.

First, this perspective gives further intuition for why MMR rules tend to randomize. Indeed, nontrivial simultaneous move zero-sum games frequently fail to have pure strategy equilibria, and that is precisely what happens here. Furthermore, in those cases where optimal decision rules randomize, the equilibria are "pooling" equilibria in which data are noninformative. This explains why these rules also tend to be nonunique; what disciplines them is Nature's best-response condition, not learning from data in equilibrium. This is also true for some nonrandomized threshold rules from Theorem 1(iv); these avoid randomization because uninformative features of the data can be used as randomization device, not because any updating occurs in equilibrium.

Second, ex-ante and ex-post rules can be distinguished in terms of the game being played. In the ex-ante (and original Waldian) perspective, this game is simultaneous move; in particular, Nature

[9] See Lehmann and Casella (1998, Theorem 1.4, Chap. 5) for a formal statement. The first explicit use of this technique in the literature on decisions under partial identification is in Stoye (2007).

moves before data Y are realized. In the ex-post perspective, Nature sees Y before choosing a prior. It is immediately clear that this may be easier to solve because it allows for backward induction. There is also an immediate sense that solutions might not agree, as we indeed found.

Finally, whether the decision maker is allowed to (or at least wants to) randomize or not can also be thought of as changing the timing of the fictitious game. Specifically, the more standard perspective is that the decision maker may randomize over decision rules and Nature must move before learning the outcome of this randomization. This setup will often yield randomized solutions. In contrast, if Nature is allowed to move after the decision maker's randomization realized, then any incentive to randomize is gone and we may as well restrict the action space to $\{0, 1\}$. Again, there is no reason to expect that solutions agree.

5 Conclusion

We studied treatment choice problems that display partial identification through the lens of the robust Bayes criteria. To do so, we take the general framework of Yata (2021) and others and embed in it a simple example of the set of priors advocated by Giacomini and Kitagawa (2021). We describe and contrast (ex-ante) Γ-minimax regret and (ex-post) Γ-posterior expected regret and analytically derive optimal solutions with and without randomization.

Our results contain two key messages that we think are valuable to the literature. First, with partial identification and multiple priors, ex-ante and ex-post assessments do *not* agree in general, whether or not randomized rules are allowed. This may at first seem expected due to dynamic inconsistency of multiple prior Bayes criteria, but is actually not obvious in view of the set of prior's specific structure. Second, randomization can be optimal in both ex-ante and ex-post problems—it is with loss of generality to exclude them even when regret is evaluated ex-post. We conjecture, but leave to future research, that our findings qualitatively extend to more general sets of priors. The contrast between the results also illustrates a need to better understand the comparative advantages—whether from a theoretical or practical perspective—of using one criterion over the other.

References

Adjaho, C. and T. Christensen (2022): "Externally Valid Treatment Choice," *arXiv preprint arXiv:2205.05561*.

Athey, S. and S. Wager (2021): "Efficient policy learning with observational data," *Econometrica*, 89, 133–161.

Ben-Michael, E., D. J. Greiner, K. Imai, and Z. Jiang (2021): "Safe policy learning through extrapolation: Application to pre-trial risk assessment," *arXiv preprint arXiv:2109.11679*.

Ben-Michael, E., K. Imai, and Z. Jiang (2022): "Policy learning with asymmetric utilities," *arXiv preprint arXiv:2206.10479*.

Berger, J. (1985): *Statistical decision theory and Bayesian analysis*, Springer.

Betrò, B. and F. Ruggeri (1992): "Conditional¡-minimax actions under convex losses," *Communications in Statistics-Theory and Methods*, 21, 1051–1066.

Bhattacharya, D. and P. Dupas (2012): "Inferring welfare maximizing treatment assignment under budget constraints," *Journal of Econometrics*, 167, 168–196.

Brock, W. A. (2006): "Profiling problems with partially identified structure," *The Economic Journal*, 116, F427–F440.

Canner, P. L. (1970): "Selecting One of Two Treatments when the Responses are Dichotomous," *Journal of the American Statistical Association*, 65, 293–306.

Chamberlain, G. (2011): "Bayesian Aspects of Treatment Choice," in *The Oxford Handbook of Bayesian Econometrics*, Oxford University Press.

Chen, H. and P. Guggenberger (2024): "A note on minimax regret rules with multiple treatments in finite samples," Tech. rep., The Pennsylvania State University.

Christensen, T., H. R. Moon, and F. Schorfheide (2022): "Optimal Discrete Decisions when Payoffs are Partially Identified," *arXiv preprint arXiv:2204.11748*.

D'Adamo, R. (2021): "Policy Learning Under Ambiguity," *arXiv preprint arXiv:2111.10904*.

DasGupta, A. and W. J. Studden (1989): "Frequentist behavior of robust Bayes estimates of normal means," *Statistics & Risk Modeling*, 7, 333–362.

Dehejia (2005): "Program evaluation as a decision problem," *Journal of Econometrics*, 125, 141–173.

Dvoretzky, A., A. Wald, and J. Wolfowitz (1951): "Elimination of Randomization in Certain Statistical Decision Procedures and Zero-Sum Two-Person Games," *The Annals of Mathematical Statistics*, 22, 1 – 21.

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

Giacomini, R. and T. Kitagawa (2021): "Robust Bayesian Inference for Set-Identified Models," *Econometrica*, 89, 1519–1556.

Giacomini, R., T. Kitagawa, and M. Read (2021): "Robust Bayesian Analysis for Econometrics," *CEPR Discussion Paper No. 16488*.

Guggenberger, P., N. Mehta, and N. Pavlov (2024): "Minimax regret treatment rules with finite samples when a quantile is the object of interest," Tech. rep., The Pennsylvania State University.

Hirano, K. and J. R. Porter (2009): "Asymptotics for statistical treatment rules," *Econometrica*, 77, 1683–1701.

——— (2020): "Asymptotic analysis of statistical decision rules in econometrics," in *Handbook of Econometrics, Volume 7A*, ed. by S. N. Durlauf, L. P. Hansen, J. J. Heckman, and R. L. Matzkin, Elsevier, vol. 7 of *Handbook of Econometrics*, 283–354.

Ida, T., T. Ishihara, K. Ito, D. Kido, T. Kitagawa, S. Sakaguchi, and S. Sasaki (2022): "Choosing Who Chooses: Selection-Driven Targeting in Energy Rebate Programs," Tech. rep., National Bureau of Economic Research.

Ishihara, T. and T. Kitagawa (2021): "Evidence Aggregation for Treatment Choice," ArXiv:2108.06473 [econ.EM], https://doi.org/10.48550/arXiv.2108.06473.

Kallus, N. and A. Zhou (2018): "Confounding-robust policy improvement," *Advances in neural information processing systems*, 31.

Karlin, S. and H. Rubin (1956): "The theory of decision procedures for distributions with monotone likelihood ratio," *The Annals of Mathematical Statistics*, 272–299.

Khan, M. A., K. P. Rath, and Y. Sun (2006): "The Dvoretzky-Wald-Wolfowitz Theorem and Purification in Atomless Finite-action Games," *International Journal of Game Theory*, 34, 91–104.

Kido, D. (2022): "Distributionally Robust Policy Learning with Wasserstein Distance," *arXiv preprint arXiv:2205.04637*.

——— (2023): "Locally Asymptotically Minimax Statistical Treatment Rules Under Partial Identification," *arXiv preprint arXiv:2311.08958*.

Kitagawa, T. (2012): "Estimation and inference for set-identified parameters using posterior lower probability," *Manuscript, UCL*.

Kitagawa, T., S. Lee, and C. Qiu (2022): "Treatment Choice with Nonlinear Regret," *arXiv preprint arXiv:2205.08586*.

——— (2023): "Treatment choice, mean square regret and partial identification," *The Japanese Economic Review*, 74, 573–602.

Kitagawa, T., S. Sakaguchi, and A. Tetenov (2021): "Constrained classification and policy learning," *arXiv preprint*.

Kitagawa, T. and A. Tetenov (2018): "Who should be treated? Empirical welfare maximization methods for treatment choice," *Econometrica*, 86, 591–616.

——— (2021): "Equality-Minded Treatment Choice," *Journal of Business Economics and Statistics*, 39, 561–574.

Kitagawa, T. and G. Wang (2023): "Who should get vaccinated? Individualized allocation of vaccines over SIR network," *Journal of Econometrics*, 232, 109–131.

Lehmann, E. L. and G. Casella (1998): *Theory of point estimation*, Springer Science & Business Media, second ed.

Lei, L., R. Sahoo, and S. Wager (2023): "Policy Learning under Biased Sample Selection," *arXiv preprint arXiv:2304.11735*.

Manski, C. F. (2000a): "Identification problems and decisions under ambiguity: empirical analysis of treatment response and normative analysis of treatment choice," *Journal of Econometrics*, 95, 415–442.

——— (2000b): "Identification problems and decisions under ambiguity: Empirical analysis of treatment response and normative analysis of treatment choice," *Journal of Econometrics*, 95, 415–442.

——— (2004): "Statistical treatment rules for heterogeneous populations," *Econometrica*, 72, 1221–1246.

——— (2005): *Social choice with partial knowledge of treatment response*, Princeton University Press.

——— (2007a): *Identification for prediction and decision*, Harvard University Press.

——— (2007b): "Minimax-regret treatment choice with missing outcome data," *Journal of Econometrics*, 139, 105–115.

——— (2020): "Towards credible patient-centered meta-analysis," *Epidemiology*, 31, 345–352.

——— (2024): "Identification and Statistical Decision Theory," *Econometric Theory*, in press.

Manski, C. F. and A. Tetenov (2007): "Admissible treatment rules for a risk-averse planner with experimental data on an innovation," *Journal of Statistical Planning and Inference*, 137, 1998–2010.

——— (2023): "Statistical decision theory respecting stochastic dominance," *The Japanese Economic Review*, 74, 447–469.

Mbakop, E. and M. Tabord-Meehan (2021): "Model selection for treatment choice: Penalized welfare maximization," *Econometrica*, 89, 825–848.

Montiel Olea, J. L., C. Qiu, and J. Stoye (2023): "Decision Theory for Treatment Choice Problems with Partial Identification," ArXiv: 2312.17623.

Moon, H. R. and F. Schorfheide (2012): "Bayesian and Frequentist Inference in Partially Identified Models," *Econometrica*, 80, 755–782.

Poirier, D. J. (1998): "Revising Beliefs in Nonidentified Models," *Econometric Theory*, 14, 483–509.

Savage, L. (1951): "The theory of statistical decision," *Journal of the American Statistical Association*, 46, 55–67.

Schlag, K. H. (2006): "ELEVEN - Tests needed for a Recommendation," Tech. rep., European University Institute Working Paper, ECO No. 2006/2, https://cadmus.eui.eu/bitstream/handle/1814/3937/ECO2006-2.pdf.

Song, K. (2014): "Point Decisions for Interval-Identified Parameters," *Econometric Theory*, 30, 334–356.

Stoye, J. (2007): "Minimax Regret Treatment Choice With Incomplete Data and many Treatments," *Econometric Theory*, 23, 190–199.

——— (2009): "Minimax regret treatment choice with finite samples," *Journal of Econometrics*, 151, 70–81.

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

——— (2012b): "New perspectives on statistical decisions under ambiguity," *Annu. Rev. Econ.*, 4, 257–282.

Tetenov, A. (2012a): "Measuring precision of statistical inference on partially identified parameters," *Discuss. Pap., Coll. Carlo Alberto, Torino*.

——— (2012b): "Statistical treatment choice based on asymmetric minimax regret criteria," *Journal of Econometrics*, 166, 157–165.

Vidakovic, B. (2000): "Γ-minimax: a paradigm for conservative robust Bayesians," *Robust bayesian analysis*, 241–259.

Wald, A. (1940): "The Fitting of Straight Lines if Both Variables are Subject to Error," *The Annals of Mathematical Statistics*, 11, 284–300.

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

We liberally invoke the game theoretic characterization alluded to in Section 4. Thus, proofs will frequently claim and verify equilibria of the fictitious game. Recall that, from basic facts about zero-sum games, if a decision rule *uniquely* best responds to some least favorable prior, it must be the *unique* equilibrium rule.

A.1 Proof of Theorem 1

A.1.1 Statement (i)

The expected regret of decision rule d is

$$R(d,\theta) = U(\theta)\left(\mathbf{1}\{U(\theta)\geq 0\} - \mathbb{E}_{m(\theta)}[d(Y)]\right), \quad \theta\in\Theta.$$

Recall that, by the definition of Γ, we have $\pi_\mu \sim \text{unif}(\{-\bar\mu,\bar\mu\})$ and, given $\mu = \pm\bar\mu$, $\pi_{\theta|\mu}(U(\theta)\in I(\mu)) = 1$. Therefore, the Bayes expected regret of d under prior $\pi\in\Gamma$ equals

$$\begin{aligned} r(d,\pi) = \frac{1}{2}\cdot\left[\int_{\tilde\theta\in\Theta} U(\tilde\theta)\left(\mathbf{1}\{U(\tilde\theta)\geq 0\} - \mathbb{E}_{\bar\mu}[d]\right) d\pi_{\theta|\bar\mu}(\tilde\theta)\right] \\ + \frac{1}{2}\cdot\left[\int_{\tilde\theta\in\Theta} U(\tilde\theta)\left(\mathbf{1}\{U(\tilde\theta)\geq 0\} - \mathbb{E}_{-\bar\mu}[d]\right) d\pi_{\theta|-\bar\mu}(\tilde\theta)\right], \end{aligned}$$

where $\mathbb{E}_{\bar\mu}[d] := \mathbb{E}_{\bar\mu}[d(Y)]$, $\mathbb{E}_{-\bar\mu}[d] := \mathbb{E}_{-\bar\mu}[d(Y)]$. One can easily solve for

$$\begin{aligned} \sup_{\pi\in\Gamma} r(d,\pi) = \frac{1}{2}\max\left\{\overline{I}(\bar\mu)(1-\mathbb{E}_{\bar\mu}[d]), -\underline{I}(\bar\mu)\mathbb{E}_{\bar\mu}[d]\right\} \\ + \frac{1}{2}\max\left\{\overline{I}(-\bar\mu)(1-\mathbb{E}_{-\bar\mu}[d]), -\underline{I}(-\bar\mu)\mathbb{E}_{-\bar\mu}[d]\right\}. \end{aligned}$$

The least favorable prior π^* equals

$$\pi^* = \left\{\pi^*_\mu \sim \text{unif}(\{-\bar\mu,\bar\mu\}), \pi^*_{\theta|\bar\mu}(U(\theta) = \overline{I}(\bar\mu)) = 1, \pi^*_{\theta|-\bar\mu}(U(\theta) = \underline{I}(-\bar\mu)) = 1\right\}. \tag{A.1}$$

Lemma B.1 shows that the unique Bayes rule against π^* is

$$d^*_{w,0} = \mathbf{1}\{w^\top Y\geq 0\}, \quad \text{where } w = \Sigma^{-1}\bar\mu.$$

Lemma B.2 establishes that $\sup_{\pi\in\Gamma} r(d^*_{w,0},\pi) = r(d^*_{w,0},\pi^*)$ as long as (3.2) holds true. This establishes the claim.

A.1.2 Statement (ii)

Step 1 We show that when (3.3) holds, any rule $d \in \mathcal{D}_n$ is Γ-MMR optimal if

$$\mathbb{E}_{\bar\mu}[d] = \frac{\overline{I}(\bar\mu)}{\overline{I}(\bar\mu) - \underline{I}(\bar\mu)}, \quad \mathbb{E}_{-\bar\mu}[d] = \frac{-\underline{I}(\bar\mu)}{\overline{I}(\bar\mu) - \underline{I}(\bar\mu)}. \tag{A.2}$$

The least favorable prior π^* is such that $\mu \sim \text{unif}(\{-\bar\mu,\bar\mu\})$, and when $\mu = \bar\mu$,

$$U(\theta) = \begin{cases} \overline{I}(\bar\mu), & \text{with probability (w.p.) } p_1, \\ \underline{I}(\bar\mu), & \text{w.p. } 1-p_1, \end{cases}$$

where $p_1 > 0$ is such that $p_1\overline{I}(\bar\mu) + (1-p_1)\underline{I}(\bar\mu) = 0$, and when $\mu = -\bar\mu$,

$$U(\theta) = \begin{cases} \overline{I}(-\bar\mu), & \text{w.p. } p_2, \\ \underline{I}(-\bar\mu), & \text{w.p. } 1-p_2, \end{cases}$$

where p_2 is such that $p_2\overline{I}(-\bar\mu) + (1-p_2)\underline{I}(-\bar\mu) = 0$. Lemma B.3 establishes that any decision rule is Bayes against this prior (intuitively because the data are uninformative), and Lemma B.4 further shows that, for any rule d that satisfies (A.2), $\sup_{\pi\in\Gamma} r(d,\pi) = r(d,\pi^*)$ obtains. This establishes the claim.

Step 2 We next verify the "only if" statement. Recall that the least favorable prior π^* must best respond to any optimal decision rule, i.e., for any MMR optimal rule d, π^* must solve

$$\begin{aligned}\sup_{\pi\in\Gamma} r(d,\pi) = &\frac{1}{2}\max\left\{\overline{I}(\bar\mu)(1-\mathbb{E}_{\bar\mu}[d]), -\underline{I}(\bar\mu)\mathbb{E}_{\bar\mu}[d]\right\} \\ &+ \frac{1}{2}\max\left\{\overline{I}(-\bar\mu)(1-\mathbb{E}_{-\bar\mu}[d]), -\underline{I}(-\bar\mu)\mathbb{E}_{-\bar\mu}[d]\right\}.\end{aligned}$$

This, however, requires that Nature is indifferent between $\overline{I}(\bar\mu)$ and $\underline{I}(\bar\mu)$ when $\mu = \bar\mu$ and similarly between $\overline{I}(-\bar\mu)$ and $\underline{I}(-\bar\mu)$ when $\mu = -\bar\mu$. That is, we must have

$$\overline{I}(\bar\mu)(1-\mathbb{E}_{\bar\mu}[d]) = -\underline{I}(\bar\mu)\mathbb{E}_{\bar\mu}[d], \quad \overline{I}(-\bar\mu)(1-\mathbb{E}_{-\bar\mu}[d]) = -\underline{I}(-\bar\mu)\mathbb{E}_{-\bar\mu}[d],$$

which is equivalent to (A.2).

Step 3 We show a rule of form $\mathbf{1}\left\{w^\top Y \geq c\right\}$ for some $c \in \mathbb{R}$ cannot be $\Gamma-$MMR optimal when (3.3) holds. Note $w = \Sigma^{-1}\bar{\mu}$. Thus,

$$\mathbb{E}_\mu[\mathbf{1}\{w^\top Y \geq c\}] = 1 - \Phi\left(\frac{c - w^\top \mu}{\sqrt{w^\top \Sigma w}}\right).$$

Suppose by contradiction that a rule $\mathbf{1}\left\{w^\top Y \geq c\right\}$ is optimal, then by statement (ii) we have

$$1 - \Phi\left(\frac{c - w^\top \bar{\mu}}{\sqrt{w^\top \Sigma w}}\right) = \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})} \tag{A.3}$$

$$1 - \Phi\left(\frac{c + w^\top \bar{\mu}}{\sqrt{w^\top \Sigma w}}\right) = \frac{-\underline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}. \tag{A.4}$$

By symmetry, (A.3) and (A.4) can both hold only if $c = 0$. But (3.3) then implies that

$$\Phi\left(\frac{w^\top \bar{\mu}}{\sqrt{w^\top \Sigma w}}\right) = \Phi(\|w\|_\Sigma) > \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})},$$

so that (A.3) cannot in fact hold when $c = 0$, a contradiction.

Step 4 We verify that d^*_{RT}, d^*_{linear} and d^*_{step} are all $\Gamma-$MMR optimal. Due to symmetry, it suffices to show that

$$\mathbb{E}_{\bar{\mu}}[d^*_{\mathrm{RT}}] = \mathbb{E}_{\bar{\mu}}[d^*_{\mathrm{linear}}] = \mathbb{E}_{\bar{\mu}}[d^*_{\mathrm{step}}] = \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}.$$

To see $\mathbb{E}_{\bar{\mu}}[d^*_{\mathrm{RT}}] = \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}$, consider the random threshold rule $\mathbf{1}\left\{w^\top Y \geq \xi\right\}$, where $\xi \sim N(0, \tilde{\sigma}^2)$ is independent of Y. As $w^\top Y - \xi \sim N\left(\|w\|^2_\Sigma, \|w\|^2_\Sigma + \tilde{\sigma}^2\right)$, algebra shows

$$\mathbb{E}_{\bar{\mu}}\left[\mathbf{1}\left\{w^\top Y \geq \xi\right\}\right] = \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}$$

as required. To see $\mathbb{E}_{\bar{\mu}}[d^*_{\mathrm{linear}}] = \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}$, let $\rho > 0$ and

$$d_{\mathrm{linear},\rho} := \begin{cases} 0, & w^\top Y < -\rho, \\ \frac{w^\top Y + \rho^*}{2\rho^*}, & -\rho \leq w^\top Y \leq \rho, \\ 1, & w^\top Y > \rho. \end{cases}$$

Applying Lemma B.7 in Montiel Olea et al. (2023) and integration by parts yield

$$f(\rho) := \mathbb{E}_{\bar{\mu}}[d_{\text{linear},\rho}] = 1 - \int_0^1 \Phi\left(\frac{2\rho x - \rho - \|w\|_\Sigma^2}{\|w\|_\Sigma}\right) dx = 1 - \frac{\|w\|_\Sigma}{2\rho} \int_{\frac{-\rho-\|w\|_\Sigma^2}{\|w\|_\Sigma}}^{\frac{\rho-\|w\|_\Sigma^2}{\|w\|_\Sigma}} \Phi(t)\, dt.$$

Note that $\lim_{\rho\downarrow 0} f(\rho) = 1 - \Phi(-\|w\|_\Sigma) = \Phi(\|w\|_\Sigma) > \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})}$, while L'Hopital's rule implies

$$\lim_{\rho\to\infty} f(\rho) = 1 - \frac{1}{2}\lim_{\rho\to\infty}\left\{\Phi\left(\frac{\rho - \|w\|_\Sigma^2}{\|w\|_\Sigma}\right) - \Phi\left(\frac{-\rho - \|w\|_\Sigma^2}{\|w\|_\Sigma}\right)\right\} = \frac{1}{2} < \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})},$$

where the last last inequality follows from $\overline{I}(\bar{\mu}) + \underline{I}(\bar{\mu}) > 0$ and $1 > \Phi(\|w\|_\Sigma) > \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})}$. Furthermore, by applying the chain rule, $\frac{\partial f(\rho)}{\partial \rho} < 0$. Therefore, $f(\cdot)$ is strictly decreasing in $(0,\infty)$. We conclude that there must exist some unique $\rho^* > 0$ such that $f(\rho^*) = \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})}$, or equivalently, $\int_0^1 \Phi\left(\frac{2\rho^* x - \rho^* - \|w\|_\Sigma^2}{\|w\|_\Sigma}\right) dx = \frac{-\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})}$, which implies $\mathbb{E}_{\bar{\mu}}[d^*_{\text{linear}}] = \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})}$.

Finally, we verify that $\mathbb{E}_{\bar{\mu}}[d^*_{\text{step}}] = \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})}$. For any $\beta \in \left(0, \frac{1}{2}\right)$, consider the following step function rule:

$$d_{\text{step},\beta} := \left(\frac{1}{2} - \beta\right)\mathbf{1}\left\{w^\top Y < 0\right\} + \left(\frac{1}{2} + \beta\right)\mathbf{1}\left\{w^\top Y \geq 0\right\}.$$

One then has

$$\begin{aligned}\mathbb{E}_{\bar{\mu}}[d_{\text{step},\beta}] =& \left(\frac{1}{2} - \beta\right)\Phi\left(-\|w\|_\Sigma\right) + \left(\frac{1}{2} + \beta\right)\left(1 - \Phi\left(-\|w\|_\Sigma\right)\right) \\ =& \frac{1}{2} + \beta\left(2\Phi\left(\|w\|_\Sigma\right) - 1\right).\end{aligned}$$

Setting $\mathbb{E}_{\bar{\mu}}[d_{\text{step},\beta^*}] = \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})}$ yields $\beta^* = \frac{\frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})} - \frac{1}{2}}{2\Phi(\|w\|_\Sigma)-1}$. As $\overline{I}(\bar{\mu})+\underline{I}(\bar{\mu}) > 0$ and (3.3) holds, $\frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})} > \frac{1}{2}$ and therefore $\beta^* > 0$. Furthermore, $\beta^* < \frac{1}{2}$ holds due to (3.3) as well. Since $d^*_{\text{step}} = d_{\text{step},\beta^*}$, we conclude that $\mathbb{E}_{\bar{\mu}}[d^*_{\text{step}}] = \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})}$.

A.1.3 Statement (iii)

For any rule of form $d_{w,c}(Y) := \mathbf{1}\{w^\top Y \geq c\}$ where $w = \Sigma^{-1}\bar{\mu}$ and $c \in \mathbb{R}$, we may calculate

$$\mathbb{E}_{\bar{\mu}}[d_{w,c}(Y)] = 1 - \Phi\left(\frac{c - w^\top\bar{\mu}}{\sqrt{w^\top\Sigma w}}\right)$$

and (due to Lemma D.1, recalling $\|w\|_\Sigma^2 = w^\top \Sigma w = \bar{\mu}^\top \Sigma^{-1} \bar{\mu}$)

$$g(c) := \sup_{\pi \in \Gamma} r(d_{w,c}, \pi) = \frac{1}{2} \max \left\{ \overline{I}(\bar{\mu}) \Phi \left(\frac{-\|w\|_\Sigma^2 + c}{\|w\|_\Sigma} \right), -\underline{I}(\bar{\mu}) \Phi \left(\frac{\|w\|_\Sigma^2 - c}{\|w\|_\Sigma} \right) \right\} + \frac{1}{2} \max \left\{ -\underline{I}(\bar{\mu}) \Phi \left(\frac{\|w\|_\Sigma^2 + c}{\|w\|_\Sigma} \right), \overline{I}(\bar{\mu}) \Phi \left(\frac{-\|w\|_\Sigma^2 - c}{\|w\|_\Sigma} \right) \right\}.$$

Lemma B.5 shows that g is decreasing on $[0, c^*]$ and increasing on $[c^*, \infty)$, implying that the optimal threshold rule is d_{w,c^*} when $c \in [0, \infty)$. By symmetry, $d_{w,-c^*}$ is optimal when $c \in (-\infty, 0]$, and $d_{w,-c^*}$ and d_{w,c^*} share the same worst-case expected regret.

A.1.4 Statement (iv)

In case (3.3) and when $n > 1$, there exists $\dot{\mu} \neq \mathbf{0}$ is such that $\dot{\mu}^\top \Sigma^{-1} \bar{\mu} = 0$, i.e., $\dot{\mu}$ is orthogonal to $\Sigma^{-1} \bar{\mu}$. For any $t \in \mathbb{R}$, let

$$d_{w_t,0}(Y) = \mathbf{1}\{w_t^\top Y \geq 0\}, \text{ where } w_t = \Sigma^{-1}(t\bar{\mu} + (1-t)\dot{\mu}). \tag{A.5}$$

Lemma B.6 shows that when $t = t^*$, $\mathbb{E}_{\bar{\mu}}[d_{w_{t^*},0}(Y)] = \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}$. By symmetry, one then also has $\mathbb{E}_{-\bar{\mu}}[d_{w_{t^*},c}(Y)] = \frac{-\underline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}$. Applying statement (ii) yields that $d_{w_{t^*},0}$ is MMR optimal.

A.2 Proof of Corollary 1

In Example 1, $\overline{I}(\bar{\mu}) = \bar{\mu} + k, \underline{I}(\bar{\mu}) = \bar{\mu} - k$, $\Sigma = \sigma^2$. The results of the corollary follow directly from Theorem 1(i)-(iii).

A.3 Proof of Theorem 2

Recall that the Γ-PER optimal rule solves

$$\inf_{a \in [0,1]} \sup_{\pi \in \Gamma} \int_{\tilde{\theta} \in \Theta} L(a, \tilde{\theta}) d\pi_{\theta|Y}(\tilde{\theta}), \quad \forall Y \in \mathbb{R}^n,$$

where $L(a, \theta) = U(\theta)(\mathbf{1}\{U(\theta) \geq 0\} - a)$, and $\pi_{\theta|Y}$ is the posterior distribution of θ given Y. If randomization is not allowed, the $\Gamma-$PER optimal rule solves

$$\inf_{a \in \{0,1\}} \sup_{\pi \in \Gamma} \int_{\tilde{\theta} \in \Theta} L(a, \tilde{\theta}) d\pi_{\theta|Y}(\tilde{\theta}), \quad \forall Y \in \mathbb{R}^n.$$

Statement (i) then follows from Lemma C.1; statement (ii) follows from Lemma C.2.

A.4 Proof of Corollary 2

Directly follows from Theorem 2.

A.5 Proof of Corollary 3

(i) "If": If $\underline{I}(\bar{\mu}) \geq 0$, then Theorem 1(i) and Theorem 2(i) apply and establish that $d^*_{w,0}$ is both the unique Γ-MMR and the unique $\Gamma-$PER optimal rule.

"Only if": If $\underline{I}(\bar{\mu}) < 0$, then d^*_{PER} is uniquely $\Gamma-$PER optimal by Theorem 2(i). If condition (3.2) holds as well, then Theorem 1(i) implies that $d^*_{w,0}$ is uniquely $\Gamma-$MMR optimal; hence, $\Gamma-$MMR and $\Gamma-$PER optimal rules disagree. If condition (3.3) applies, then, by Theorem 1(ii), a Γ-MMR optimal rule must satisfy (3.4)-(3.5). But d^*_{PER} can be written as

$$d^*_{\text{PER}} = d_{\text{step},\beta_{\text{PER}}} = \begin{cases} \frac{1}{2} - \beta_{\text{PER}}, & \text{if } w^\top Y < 0, \\ \frac{1}{2} + \beta_{\text{PER}}, & \text{if } w^\top Y \geq 0, \end{cases}$$

where $\beta_{\text{PER}} = \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})} - \frac{1}{2}$. Step 4 for the proof of Theorem 1(ii) shows that the only $\Gamma-$MMR optimal rule of form $d_{\text{step},\beta}$ is $d^*_{\text{step}} \neq d^*_{\text{PER}}$.

(ii) "If": When randomization is not allowed, Theorem 2(ii) shows that $d^*_{w,0}$ is always $\Gamma-$PER optimal. Under condition (3.16), Theorem 1(i) establishes that $d^*_{w,0}$ is $\Gamma-$MMR optimal as well.

"Only if": Again, $d^*_{w,0}$ is always $\Gamma-$PER optimal when randomization is not allowed. If (3.16) fails (i.e., (3.3) holds) and $n > 1$, it follows by Theorem 1(iv) that the randomized rule $d^*_{w_{t^*},0}$ is Γ-MMR optimal. If $n = 1$, Theorem 1(iii) implies that $d^*_{w,0}$ is not optimal even among linear threshold rules. Therefore, when (3.3) holds, Γ-MMR and Γ-PER optimal rules disagree with and without randomization.

B Technical Lemmas Supporting Ex-ante Analysis

Lemma B.1. *The Bayes rule supported by prior (A.1) is $d^*_{w,0}$.*

Proof. Note $\overline{I}(\bar{\mu}) > 0$ due to $\overline{I}(\bar{\mu}) + \underline{I}(\bar{\mu}) > 0$ and $\underline{I}(-\bar{\mu}) = -\overline{I}(\bar{\mu})$ by Lemma D.1. Given π^*, the

Bayes optimal rule must solve the posterior problem

$$\begin{aligned}\min_{a\in[0,1]} \quad & \int_{\tilde{\theta}\in\Theta} L(a,\tilde{\theta})d\pi^*_{\theta|Y}(\tilde{\theta}),\\ & \int_{\tilde{\theta}\in\Theta} L(a,\tilde{\theta})d\pi^*_{\theta|Y}(\tilde{\theta}) \propto \overline{I}(\bar{\mu})(1-a)\cdot\frac{1}{2}\cdot f(Y|\bar{\mu}) + \underline{I}(-\bar{\mu})(-a)\cdot\frac{1}{2}\cdot f(Y|-\bar{\mu}),\end{aligned}$$

where $f(Y|\bar{\mu})$ and $f(Y|-\bar{\mu})$ are the likelihood of Y at $\bar{\mu}$ and $-\bar{\mu}$. This problem is equivalent to

$$\min_{a\in[0,1]} \overline{I}(\bar{\mu})f(Y|\bar{\mu}) + a\underbrace{\overline{I}(\bar{\mu})}_{>0}(f(Y|-\bar{\mu}) - f(Y|\bar{\mu})).$$

Since $\overline{I}(\bar{\mu}) > 0$, the unique Bayes optimal rule is $\mathbf{1}\{f(Y|-\bar{\mu}) - f(Y|\bar{\mu}) \le 0\}$, which is equivalent to $d^*_{w,0}$ after further algebra. □

Lemma B.2. *Consider decision rule $d^*_{w,0}$ and prior π^* defined in (A.1). If (3.2) holds, then*

$$\sup_{\pi\in\Gamma} r(d^*_{w,0},\pi) = r(d^*_{w,0},\pi^*). \tag{B.1}$$

Proof. As $w^\top Y \sim N(w^\top\mu, w^\top\Sigma w)$, algebra shows

$$\mathbb{E}_\mu[d^*_{w,0}] = \Phi\left(\frac{w^\top\mu}{\sqrt{w^\top\Sigma w}}\right)$$

for all $\mu \in M$. In particular,

$$\mathbb{E}_{\bar{\mu}}[d^*_{w,0}] = \Phi\left(\frac{w^\top\bar{\mu}}{\sqrt{w^\top\Sigma w}}\right), \quad \mathbb{E}_{-\bar{\mu}}[d^*_{w,0}] = \Phi\left(-\frac{w^\top\bar{\mu}}{\sqrt{w^\top\Sigma w}}\right).$$

It follows that

$$\begin{aligned}\sup_{\pi\in\Gamma} r(d^*_{w,0},\pi) = {} & \frac{1}{2}\max\left\{\overline{I}(\bar{\mu})\Phi\left(-\frac{w^\top\bar{\mu}}{\sqrt{w^\top\Sigma w}}\right), -\underline{I}(\bar{\mu})\Phi\left(\frac{w^\top\bar{\mu}}{\sqrt{w^\top\Sigma w}}\right)\right\}\\ & + \frac{1}{2}\max\left\{\overline{I}(-\bar{\mu})\Phi\left(\frac{w^\top\bar{\mu}}{\sqrt{w^\top\Sigma w}}\right), -\underline{I}(-\bar{\mu})\Phi\left(-\frac{w^\top\bar{\mu}}{\sqrt{w^\top\Sigma w}}\right)\right\},\end{aligned}$$

and

$$r(d^*_{w,0},\pi^*) = \frac{1}{2}\overline{I}(\bar{\mu})\Phi\left(-\frac{w^\top\bar{\mu}}{\sqrt{w^\top\Sigma w}}\right) - \frac{1}{2}\underline{I}(-\bar{\mu})\Phi\left(-\frac{w^\top\bar{\mu}}{\sqrt{w^\top\Sigma w}}\right).$$

So, (B.1) holds as long as

(i) $\overline{I}(\bar{\mu})\Phi\left(-\frac{w^\top\bar{\mu}}{\sqrt{w^\top\Sigma w}}\right) \geq -\underline{I}(\bar{\mu})\Phi\left(\frac{w^\top\bar{\mu}}{\sqrt{w^\top\Sigma w}}\right)$,

(ii) $\overline{I}(-\bar{\mu})\Phi\left(\frac{w^\top\bar{\mu}}{\sqrt{w^\top\Sigma w}}\right) \leq -\underline{I}(-\bar{\mu})\Phi\left(-\frac{w^\top\bar{\mu}}{\sqrt{w^\top\Sigma w}}\right)$,

both of which are the same as (3.2) after further algebra, recalling that $\overline{I}(-\bar{\mu}) = -\underline{I}(\bar{\mu})$ and $-\underline{I}(-\bar{\mu}) = \overline{I}(\bar{\mu})$ by Lemma D.1, $w = \Sigma^{-1}\bar{\mu}$, and $\|w\|_\Sigma = \sqrt{\bar{\mu}^\top\Sigma^{-1}\bar{\mu}}$. □

Lemma B.3. *Consider the prior π^* such that $\mu \sim \text{unif}(\{-\bar{\mu}, \bar{\mu}\})$, and when $\mu = \bar{\mu}$,*

$$U(\theta) = \begin{cases} \overline{I}(\bar{\mu}), & w.p.\ p_1, \\ \underline{I}(\bar{\mu}), & w.p.\ 1-p_1, \end{cases}$$

where p_1 is such that $p_1\overline{I}(\bar{\mu}) + (1-p_1)\underline{I}(\bar{\mu}) = 0$, and when $\mu = -\bar{\mu}$,

$$U(\theta) = \begin{cases} \overline{I}(-\bar{\mu}), & w.p.\ p_2, \\ \underline{I}(-\bar{\mu}), & w.p.\ 1-p_2, \end{cases}$$

where p_2 is such that $p_2\overline{I}(-\bar{\mu}) + (1-p_2)\underline{I}(-\bar{\mu}) = 0$. Given this prior, any decision rule is Bayes optimal under (3.3).

Proof. As (3.3) holds, we have $\underline{I}(\bar{\mu}) < 0 < \overline{I}(\bar{\mu})$. Analogous to Lemma B.1, a Bayes rule must solve

$$\begin{aligned} \min_{a\in[0,1]} \quad & \frac{1}{2}\cdot f(Y \mid \bar{\mu})\left[p_1\overline{I}(\bar{\mu})(1-a) + (1-p_1)(-\underline{I}(\bar{\mu}))a\right] \\ & + \frac{1}{2}\cdot f(Y \mid -\bar{\mu})\left[p_2\overline{I}(-\bar{\mu})(1-a) + (1-p_2)(-\underline{I}(-\bar{\mu}))a\right]. \end{aligned}$$

Since $p_1\overline{I}(\bar{\mu}) + (1-p_1)\underline{I}(\bar{\mu}) = 0$ and $p_2\overline{I}(-\bar{\mu}) + (1-p_2)\underline{I}(-\bar{\mu}) = 0$, the objective is constant in a, hence the claim. □

Lemma B.4. *Consider the prior π^* such that $\mu \sim \text{unif}(\{-\bar{\mu}, \bar{\mu}\})$, when $\mu = \bar{\mu}$,*

$$U(\theta) = \begin{cases} \overline{I}(\bar{\mu}), & w.p.\ p_1, \\ \underline{I}(\bar{\mu}), & w.p.\ 1-p_1, \end{cases}$$

where p_1 is such that $p_1\overline{I}(\bar{\mu}) + (1-p_1)\underline{I}(\bar{\mu}) = 0$, and when $\mu = -\bar{\mu}$,

$$U(\theta) = \begin{cases} \overline{I}(-\bar{\mu}), & w.p.\ p_2, \\ \underline{I}(-\bar{\mu}), & w.p.\ 1-p_2, \end{cases}$$

where p_2 is such that $p_2\overline{I}(-\bar{\mu}) + (1-p_2)\underline{I}(-\bar{\mu}) = 0$. Then, when (3.3) holds, we have

$$\sup_{\pi\in\Gamma} r(d,\pi) = r(d,\pi^*)$$

for any decision rule $d \in \mathcal{D}_n$ such that (A.2) is true.

Proof. Note again that (3.3) implies $\underline{I}(\bar{\mu}) < 0$ and $\overline{I}(\bar{\mu}) > 0$. For any decision rule $d \in \mathcal{D}_n$ such that (A.2) is true, algebra shows

$$\overline{I}(\bar{\mu})(1 - \mathbb{E}_{\bar{\mu}}[d]) = -\underline{I}(\bar{\mu})\mathbb{E}_{\bar{\mu}}[d] = -\frac{\overline{I}(\bar{\mu})\underline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})},$$

$$\overline{I}(-\bar{\mu})(1 - \mathbb{E}_{-\bar{\mu}}[d]) = -\underline{I}(-\bar{\mu})\mathbb{E}_{-\bar{\mu}}[d] = -\frac{\overline{I}(\bar{\mu})\underline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})},$$

implying $\sup_{\pi\in\Gamma} r(d,\pi) = -\frac{\overline{I}(\bar{\mu})\underline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})}$. Meanwhile, for any $d \in \mathcal{D}_n$ such that (A.2) is true, we may also calculate the following that yields the desired conclusion:

$$r(d,\pi^*) = \frac{1}{2}p_1\overline{I}(\bar{\mu}) + \frac{1}{2}p_2\overline{I}(-\bar{\mu}) = -\frac{\overline{I}(\bar{\mu})\underline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})} = \sup_{\pi\in\Gamma} r(d,\pi).$$

□

Lemma B.5. *In case (3.3), the function*

$$\begin{aligned} g(c) = &\frac{1}{2}\max\left\{\overline{I}(\bar{\mu})\Phi\left(\frac{-\|w\|_\Sigma^2+c}{\|w\|_\Sigma}\right), -\underline{I}(\bar{\mu})\Phi\left(\frac{\|w\|_\Sigma^2-c}{\|w\|_\Sigma}\right)\right\} \\ &+ \frac{1}{2}\max\left\{-\underline{I}(\bar{\mu})\Phi\left(\frac{\|w\|_\Sigma^2+c}{\|w\|_\Sigma}\right), \overline{I}(\bar{\mu})\Phi\left(\frac{-\|w\|_\Sigma^2-c}{\|w\|_\Sigma}\right)\right\} \end{aligned}$$

is decreasing in $[0, c^*]$ and increasing in $[c^*, \infty)$.

Proof. Note when $c \in [0, c^*)$ and (3.3) holds, we have $-\underline{I}(\bar{\mu}) > 0$,

$$\begin{aligned} -\underline{I}(\bar{\mu})\Phi\left(\frac{\|w\|_\Sigma^2-c}{\|w\|_\Sigma}\right) &> \overline{I}(\bar{\mu})\Phi\left(\frac{-\|w\|_\Sigma^2+c}{\|w\|_\Sigma}\right), \\ -\underline{I}(\bar{\mu})\Phi\left(\frac{\|w\|_\Sigma^2+c}{\|w\|_\Sigma}\right) &> \overline{I}(\bar{\mu})\Phi\left(\frac{-\|w\|_\Sigma^2-c}{\|w\|_\Sigma}\right), \end{aligned}$$

and therefore,

$$g(c) = -\frac{\underline{I}(\bar{\mu})}{2}\left[\Phi\left(\frac{\|w\|_\Sigma^2-c}{\|w\|_\Sigma}\right) + \Phi\left(\frac{\|w\|_\Sigma^2+c}{\|w\|_\Sigma}\right)\right],$$

and

$$\frac{\partial g(c)}{\partial c} = -\frac{\underline{I}(\bar{\mu})}{2\|w\|_\Sigma}\left[\phi\left(\frac{\|w\|_\Sigma^2+c}{\|w\|_\Sigma}\right) - \phi\left(\frac{\|w\|_\Sigma^2-c}{\|w\|_\Sigma}\right)\right] < 0.$$

When $c = c^*$, note $\overline{I}(\bar{\mu})\Phi\left(\frac{-\|w\|_\Sigma^2+c^*}{\|w\|_\Sigma}\right) = -\underline{I}(\bar{\mu})\Phi\left(\frac{\|w\|_\Sigma^2-c^*}{\|w\|_\Sigma}\right)$ and

$$g(c^*) = -\frac{\underline{I}(\bar{\mu})}{2}\left[\Phi\left(\frac{\|w\|_\Sigma^2-c^*}{\|w\|_\Sigma}\right) + \Phi\left(\frac{\|w\|_\Sigma^2+c^*}{\|w\|_\Sigma}\right)\right].$$

When $c \in (c^*, \infty)$, note

$$-\underline{I}(\bar{\mu})\Phi\left(\frac{\|w\|_\Sigma^2-c}{\|w\|_\Sigma}\right) < \overline{I}(\bar{\mu})\Phi\left(\frac{-\|w\|_\Sigma^2+c}{\|w\|_\Sigma}\right),$$
$$-\underline{I}(\bar{\mu})\Phi\left(\frac{\|w\|_\Sigma^2+c}{\|w\|_\Sigma}\right) > \overline{I}(\bar{\mu})\Phi\left(\frac{-\|w\|_\Sigma^2-c}{\|w\|_\Sigma}\right),$$

so that

$$g(c) = \frac{1}{2}\left\{\overline{I}(\bar{\mu})\Phi\left(\frac{-\|w\|_\Sigma^2+c}{\|w\|_\Sigma}\right) - \underline{I}(\bar{\mu})\Phi\left(\frac{\|w\|_\Sigma^2+c}{\|w\|_\Sigma}\right)\right\},$$

and

$$\frac{\partial g(c)}{\partial c} = \frac{1}{2\|w\|_\Sigma}\left\{\overline{I}(\bar{\mu})\phi\left(\frac{-\|w\|_\Sigma^2+c}{\|w\|_\Sigma}\right) - \underline{I}(\bar{\mu})\phi\left(\frac{\|w\|_\Sigma^2+c}{\|w\|_\Sigma}\right)\right\} > 0.$$

Having shown that $\frac{\partial g(c)}{\partial c} < 0$ when $c \in [0, c^*)$ and $\frac{\partial g(c)}{\partial c} > 0$ when $c \in (c^*, \infty)$, and since $g(c)$ is continuous at $c = c^*$, we conclude that g is decreasing in $[0, c^*]$ and increasing $[c^*, \infty)$. □

Lemma B.6. *In case (3.3) and when $n > 1$, let*

$$d_{w_t,0}(Y) = \mathbf{1}\{w_t^\top Y \geq 0\}, \text{ where } w_t = \Sigma^{-1}(t\bar{\mu} + (1-t)\dot{\mu}), \dot{\mu} \neq \mathbf{0},$$

and $\dot{\mu}^\top\Sigma^{-1}\bar{\mu} = 0$. Then, $\mathbb{E}_{\bar{\mu}}[d_{w_{t^*},0}(Y)] = \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})}$, where t^* is defined in (3.9).

Proof. Write $f(t) := \mathbb{E}_{\bar{\mu}}[d_{w_t,c}(Y)] = \Phi\left(\frac{w_t^\top\bar{\mu}}{\sqrt{w_t^\top\Sigma w_t}}\right)$ and $k^* := \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})}$. Then, it suffices to show that $f(t^*) = k^*$, which is equivalent to showing

$$\frac{\left(w_{t^*}^\top\bar{\mu}\right)^2}{w_{t^*}^\top\Sigma w_{t^*}} = \left(\Phi^{-1}\left(k^*\right)\right)^2.$$

Furthermore, note

$$w_{t^*}^\top \bar{\mu} = (t^* \bar{\mu}^\top \Sigma^{-1} + (1-t^*)\dot{\mu}^\top \Sigma^{-1})\bar{\mu} = t^* \bar{\mu}^\top \Sigma^{-1} \bar{\mu} + (1-t^*) \underbrace{\dot{\mu}^\top \Sigma^{-1} \bar{\mu}}_{=0} = t^* \|w\|_\Sigma^2$$

and

$$\begin{aligned} w_{t^*}^\top \Sigma w_{t^*} &= (t^* \bar{\mu}^\top \Sigma^{-1} + (1-t^*)\dot{\mu}^\top \Sigma^{-1}) \Sigma (t^* \Sigma^{-1} \bar{\mu} + (1-t^*)\Sigma^{-1}\dot{\mu}) \\ &= (t^*)^2 \bar{\mu}^\top \Sigma^{-1} \bar{\mu} + 2t^*(1-t^*) \underbrace{\bar{\mu}^\top \Sigma^{-1} \dot{\mu}}_{=0} + (1-t^*)^2 \underbrace{\dot{\mu}^\top \Sigma^{-1} \dot{\mu}}_{:=\|\Sigma^{-1}\dot{\mu}\|_\Sigma^2} = (t^*)^2 \|w\|_\Sigma^2 + (1-t^*)^2 \|\Sigma^{-1}\dot{\mu}\|_\Sigma^2. \end{aligned}$$

Therefore, substituting in these expressions, t^* should be such that

$$\frac{\left(\Phi^{-1}(k^*)\right)^2}{\|w\|_\Sigma^2} = \frac{(t^*)^2 \|w\|_\Sigma^2}{(t^*)^2 \|w\|_\Sigma^2 + (1-t^*)^2 \|\Sigma^{-1}\dot{\mu}\|_\Sigma^2}.$$

As $s^* := \frac{\left(\Phi^{-1}(k^*)\right)^2}{\|w\|_\Sigma^2} \in (0,1)$ due to (3.3), we can solve the above equation for t^*, and after lengthy algebra, find that

$$t^* = \frac{1}{1 \pm \sqrt{\frac{(1-s^*)}{s^*} \cdot \frac{\|w\|_\Sigma^2}{\|\Sigma^{-1}\dot{\mu}\|_\Sigma^2}}}.$$

□

C Technical Lemmas Supporting Ex-post Analysis

Lemma C.1. *Suppose all the conditions of Theorem 1 hold. If $\underline{I}(\bar{\mu}) < 0 < \overline{I}(\bar{\mu})$, then the unique Γ-PER optimal rule is*

$$d^*_{\mathrm{PER}}(Y) = \begin{cases} \frac{-\underline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}, & \text{if } w^\top Y < 0, \\ \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}, & \text{if } w^\top Y \geq 0. \end{cases}$$

Otherwise, $d^*_{w,0}$ is Γ-PER optimal.

Proof. Let $f(Y \mid \mu)$ be the likelihood of Y. Then, for each action $a \in [0,1]$, we have

$$
\begin{aligned}
&\sup_{\pi\in\Gamma}\int_{\tilde{\theta}\in\Theta} L(a,\tilde{\theta})d\pi_{\theta|Y}(\tilde{\theta}) \\
\propto&\frac{1}{2}\left(\sup_{\pi_{\theta|\bar{\mu}}}\int_{\tilde{\theta}\in\Theta} U(\tilde{\theta})(\mathbf{1}\{U(\tilde{\theta})\geq 0\}-a)d\pi_{\theta|\bar{\mu}}(\tilde{\theta})\right) f(Y\mid\bar{\mu}) \\
+&\frac{1}{2}\left(\sup_{\pi_{\theta|-\bar{\mu}}}\int_{\tilde{\theta}\in\Theta} U(\tilde{\theta})(\mathbf{1}\{U(\tilde{\theta})\geq 0\}-a)d\pi_{\theta|-\bar{\mu}}(\tilde{\theta})\right) f(Y\mid-\bar{\mu}) \\
=&\frac{1}{2}f(Y\mid\bar{\mu})\max\left\{\overline{I}(\bar{\mu})(1-a),-\underline{I}(\bar{\mu})a\right\} \\
+&\frac{1}{2}f(Y\mid-\bar{\mu})\max\left\{\overline{I}(-\bar{\mu})(1-a),-\underline{I}(-\bar{\mu})a\right\}.
\end{aligned}
$$

Therefore, when randomization is allowed, the Γ-PER optimal rule can be found by solving

$$\inf_{a\in[0,1]} V_\Gamma(a), \tag{C.1}$$

where

$$V_\Gamma(a) := f(Y\mid\bar{\mu})\max\left\{\overline{I}(\bar{\mu})(1-a),-\underline{I}(\bar{\mu})a\right\}+f(Y\mid-\bar{\mu})\max\left\{\overline{I}(\bar{\mu})a,-\underline{I}(\bar{\mu})(1-a)\right\},$$

since $\overline{I}(-\bar{\mu})=-\underline{I}(\bar{\mu})$ and $-\underline{I}(-\bar{\mu})=\overline{I}(\bar{\mu})$ by Lemma D.1. Note that, for each action $a\in[0,1]$,

$$
\begin{aligned}
\overline{I}(\bar{\mu})(1-a)\geq-\underline{I}(\bar{\mu})a\Leftrightarrow a\leq\frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})}, \\
\overline{I}(\bar{\mu})a\geq-\underline{I}(\bar{\mu})(1-a)\Leftrightarrow a\geq\frac{-\underline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})},
\end{aligned}
$$

where $\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})\geq 0$ and $\overline{I}(\bar{\mu})>0$ due to our normalization $\overline{I}(\bar{\mu})+\underline{I}(\bar{\mu})>0$. The conclusion follows from discussing the following two cases.

Case 1: $\underline{I}(\bar{\mu})\geq 0$**.** In this case, $\frac{-\underline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})}\leq a\leq\frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})}$ for all $a\in[0,1]$. Thus, (C.1) becomes

$$\inf_{a\in[0,1]}\underbrace{\overline{I}(\bar{\mu})}_{>0}\left\{f(Y\mid\bar{\mu})+\left[f(Y\mid-\bar{\mu})-f(Y\mid\bar{\mu})\right]a\right\}.$$

The $\Gamma-$PER optimal rule is $\mathbf{1}\left\{f(Y\mid-\bar{\mu})-f(Y\mid\bar{\mu})\leq 0\right\}$, which is $d^*_{w,0}$ after further algebra.

Case 2: $\underline{I}(\bar{\mu}) < 0 < \overline{I}(\bar{\mu})$**.** In this case, we have (again, note $\overline{I}(\bar{\mu}) + \underline{I}(\bar{\mu}) > 0$):

$$0 < \frac{-\underline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})} < \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})} < 1.$$

Therefore:

- When $a \in \left[0, \frac{-\underline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}\right)$,

$$\begin{aligned} V_\Gamma(a) =& f(Y \mid \bar{\mu})\overline{I}(\bar{\mu})(1-a) - f(Y \mid -\bar{\mu})\underline{I}(\bar{\mu})(1-a) \\ =& \underbrace{\left\{ f(Y \mid \bar{\mu})\overline{I}(\bar{\mu}) + f(Y \mid -\bar{\mu})\left(-\underline{I}(\bar{\mu})\right)\right\}}_{>0}(1-a). \end{aligned}$$

  Thus, $V_\Gamma(\cdot)$ is strictly decreasing in $a \in \left[0, \frac{-\underline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}\right)$.

- When $a \in \left[\frac{-\underline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}, \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}\right)$,

$$\begin{aligned} V_\Gamma(a) =& f(Y \mid \bar{\mu})\overline{I}(\bar{\mu})(1-a) + f(Y \mid -\bar{\mu})\overline{I}(\bar{\mu})a \\ =& f(Y \mid \bar{\mu})\overline{I}(\bar{\mu}) + \underbrace{\overline{I}(\bar{\mu})}_{>0}\left(f(Y \mid -\bar{\mu}) - f(Y \mid \bar{\mu})\right)a. \end{aligned}$$

  Thus, if $f(Y \mid -\bar{\mu}) > f(Y \mid \bar{\mu})$, $V_\Gamma(\cdot)$ is strictly increasing in a, and if $f(Y \mid -\bar{\mu}) < f(Y \mid \bar{\mu})$, $V_\Gamma(\cdot)$ is strictly decreasing in a. When $f(Y \mid -\bar{\mu}) = f(Y \mid \bar{\mu})$, $V_\Gamma(\cdot)$ is constant in a.

- When $a \in \left[\frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}, 1\right]$,

$$V_\Gamma(a) = \underbrace{\left(f(Y \mid -\bar{\mu})\overline{I}(\bar{\mu}) - f(Y \mid \bar{\mu})\underline{I}(\bar{\mu})\right)}_{>0} a,$$

  implying that $V_\Gamma(\cdot)$ is strictly increasing in a.

In sum, if $f(Y \mid -\bar{\mu}) > f(Y \mid \bar{\mu})$, $V_\Gamma(\cdot)$ is first strictly decreasing and then increasing, with the minimum achieved at $a = \frac{-\underline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}$; if $f(Y \mid -\bar{\mu}) < f(Y \mid \bar{\mu})$, $V_\Gamma(\cdot)$ is first strictly decreasing and increasing with the minimum achieved at $a = \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}$; If $f(Y \mid -\bar{\mu}) = f(Y \mid \bar{\mu})$ (which only happens in a null set), $V_\Gamma(\cdot)$ achieves minimum at any point between $[\frac{-\underline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}, \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}]$. Therefore,

the unique $\Gamma-$PER optimal rule is

$$d(Y) = \begin{cases} \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})}, & \text{if } f(Y \mid -\bar{\mu}) \leq f(Y \mid \bar{\mu}) \Longleftrightarrow \bar{\mu}^\top \Sigma^{-1} Y \geq 0, \\ \frac{-\underline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})} & \text{if } f(Y \mid -\bar{\mu}) > f(Y \mid \bar{\mu}) \Longleftrightarrow \bar{\mu}^\top \Sigma^{-1} Y < 0. \end{cases}$$

□

Lemma C.2. *Suppose all conditions of Theorem 1 hold true. Then, $d^*_{w,0}$ is always the Γ-PER optimal non-randomized rule.*

Proof. When randomization is not allowed, we need to solve (by a similar derivation to Lemma C.1) $\inf_{a\in\{0,1\}} V_\Gamma(a)$, where $V_\Gamma(a)$ is defined in (C.1). As

$$\begin{aligned} V_\Gamma(1) &= -f(Y \mid \bar{\mu})\underline{I}(\bar{\mu}) + f(Y \mid -\bar{\mu})\overline{I}(\bar{\mu}), \\ V_\Gamma(0) &= f(Y \mid \bar{\mu})\overline{I}(\bar{\mu}) - f(Y \mid -\bar{\mu})\underline{I}(\bar{\mu}), \end{aligned}$$

the optimal action would be $a = 1$ if and only if $V_\Gamma(1) \leq V_\Gamma(0)$, which is equivalent to $d^*_{w,0}$ after further algebra. □

D Additional Results

Lemma D.1. *Consider a treatment choice problem with welfare function (2.1), statistical model (2.3) and a set of priors (2.9), that satisfies Assumption 1. Then, the following statements hold:*

(i) $\overline{I}(-\bar{\mu}) = -\underline{I}(\bar{\mu})$, $\underline{I}(-\bar{\mu}) = -\overline{I}(\bar{\mu})$;

(ii) $\overline{I}(-\bar{\mu}) + \underline{I}(-\bar{\mu}) = -(\underline{I}(\bar{\mu}) + \overline{I}(\bar{\mu}))$.

Proof. **Statement (i).** Here, we only show that $\overline{I}(-\bar{\mu}) = -\underline{I}(\bar{\mu})$. An analogous argument proves $\underline{I}(-\bar{\mu}) = -\overline{I}(\bar{\mu})$. By definition,

$$\overline{I}(-\bar{\mu}) = \sup_{\{\theta\in\Theta : m(\theta) = -\bar{\mu}\}} U(\theta).$$

Since Θ is centrosymmetric and both $U(\cdot)$ and $m(\cdot)$ are linear, we have

$$\begin{aligned}\sup_{\{\theta\in\Theta:m(\theta)=-\bar{\mu}\}} U(\theta) &= \sup_{\{-\theta\in\Theta:-m(-\theta)=-\bar{\mu}\}} -U(-\theta)\\ &= \sup_{\{\tilde{\theta}\in\Theta:-m(\tilde{\theta})=-\bar{\mu}\}} -U(\tilde{\theta})\\ &= -\left\{\inf_{\{\tilde{\theta}\in\Theta:m(\tilde{\theta})=\bar{\mu}\}} U(\tilde{\theta})\right\}\\ &= -\underline{I}(\bar{\mu})\end{aligned}$$

as required.

Statement (ii). Directly follows from summing both equalities in statement (i). □

Lemma D.2. *In Example 1, suppose $k > \sqrt{\frac{\pi}{2}}\sigma$. Then, if $\bar{\mu} > 0$ is sufficiently small, the corresponding $\Gamma-$PER optimal rule is dominated in terms of profiled regret.*

Proof. Denote by $d^*_{MMR,\text{linear}}$ the least randomizing global MMR optimal rule derived in Montiel Olea et al. (2023). We aim to show that when $k > \sqrt{\frac{\pi}{2}}\sigma$ and $\bar{\mu} > 0$ is sufficiently small, the associated Γ-PER optimal rule d^*_{PER} is such that

$$\bar{R}(d^*_{\text{PER}}, \mu) \geq \bar{R}(d^*_{MMR,\text{linear}}, \mu), \text{ for all } \mu \geq 0, \tag{D.1}$$

with the inequality strict for all $\mu > 0$. A symmetry argument then immediately implies that d^*_{PER} is dominated in terms of profiled regret.

Step 1 Pick any $0 < \bar{\mu} < \sqrt{\frac{\pi}{2}}\sigma$. We show that $\bar{R}(d^*_{\text{PER}}, \mu) = (\mu + k)\left(1 - \mathbb{E}_\mu\left[d^*_{\text{PER}}(\hat{\mu})\right]\right)$ for all $\mu \geq 0$. By results in Montiel Olea et al. (Appendix B.3.1, 2023), we may write the profiled regret of a rule d as

$$\begin{aligned}&\bar{R}(d,\mu)\\ &= \begin{cases}(-\mu + k)\mathbb{E}_\mu\left[d(\hat{\mu})\right], & \text{if } \mu < -k,\\ \max\left\{(\mu+k)\left(1-\mathbb{E}_\mu\left[d(\hat{\mu})\right]\right), (-\mu+k)\mathbb{E}_\mu\left[d(\hat{\mu})\right]\right\}, & \text{if } -k \leq \mu \leq k,\\ (\mu+k)(1-\mathbb{E}_\mu\left[d(\hat{\mu})\right]), & \text{if } \mu > k.\end{cases}\end{aligned}$$

Thus, it suffices to show

$$\begin{aligned}&\max\left\{(\mu+k)\left(1-\mathbb{E}_\mu\left[d^*_{\mathrm{PER}}(\hat{\mu})\right]\right),(-\mu+k)\mathbb{E}_\mu\left[d^*_{\mathrm{PER}}(\hat{\mu})\right]\right\}\\=&(\mu+k)\left(1-\mathbb{E}_\mu\left[d^*_{\mathrm{PER}}(\hat{\mu})\right]\right)\end{aligned}$$

for any $0 \le \mu \le k$. As $\sqrt{\frac{\pi}{2}}\sigma < k$, Corollary 2 implies that the corresponding $\Gamma-$PER rule for any $0 < \bar{\mu} < \sqrt{\frac{\pi}{2}}\sigma$ is

$$\begin{aligned}d^*_{\mathrm{PER}}(\hat{\mu}) &= \frac{k+\overline{\mu}}{2k}\mathbf{1}\left\{\hat{\mu}\ge 0\right\}+\frac{k-\overline{\mu}}{2k}\mathbf{1}\left\{\hat{\mu}<0\right\}\\&=\frac{1}{2}+\frac{\overline{\mu}}{2k}\left(\mathbf{1}\left\{\hat{\mu}\ge 0\right\}-\mathbf{1}\left\{\hat{\mu}<0\right\}\right),\end{aligned}$$

and further algebra shows $\mathbb{E}_\mu\left[d^*_{\mathrm{PER}}(\hat{\mu})\right] = \frac{1}{2}+\frac{\overline{\mu}}{2k}\left(1-2\Phi\left(-\frac{\mu}{\sigma}\right)\right)$. Then,

$$(\mu+k)\left(1-\mathbb{E}_\mu\left[d^*_{\mathrm{PER}}(\hat{\mu})\right]\right) \ge (-\mu+k)\mathbb{E}_\mu\left[d^*_{\mathrm{PER}}(\hat{\mu})\right]$$

if and only if

$$(\mu+k)\left(\frac{1}{2}-\frac{\overline{\mu}}{2k}\left(1-2\Phi\left(-\frac{\mu}{\sigma}\right)\right)\right) \ge (-\mu+k)\left(\frac{1}{2}+\frac{\overline{\mu}}{2k}\left(1-2\Phi\left(-\frac{\mu}{\sigma}\right)\right)\right),$$

which is equivalent to

$$\frac{\mu}{\left(1-2\Phi\left(-\frac{\mu}{\sigma}\right)\right)} \ge \overline{\mu}. \tag{D.2}$$

Note the left hand side of (D.2) is increasing in μ, and $\lim_{\mu\downarrow 0}\left(\frac{\mu}{\left(1-2\Phi\left(-\frac{\mu}{\sigma}\right)\right)}\right) = \frac{\sigma}{2\phi(0)} = \sigma\sqrt{\frac{\pi}{2}} > \overline{\mu}$. As $0 < \bar{\mu} < \sqrt{\frac{\pi}{2}}\sigma$, we conclude that (D.2) indeed holds, and it follows

$$\bar{R}(d^*_{\mathrm{PER}},\mu) = (\mu+k)\left(1-\mathbb{E}_\mu\left[d^*_{\mathrm{PER}}(\hat{\mu})\right]\right), \tag{D.3}$$

for all $\mu \ge 0$. As a result, we conclude that, for any $0 < \bar{\mu} < \sqrt{\frac{\pi}{2}}\sigma$, we have

$$\bar{R}(d^*_{\mathrm{PER}},\mu) = (\mu+k)\left(\frac{1}{2}-\frac{\overline{\mu}}{2k}\left(1-2\Phi\left(-\frac{\mu}{\sigma}\right)\right)\right), \text{ for all } \mu \ge 0.$$

Step 2 We show that $\bar{R}(d^*_{\text{PER}}, \mu)$ is strictly increasing in μ for any $0 < \bar{\mu} < \sqrt{\frac{\pi}{2}}\sigma$ small enough. To see this, note

$$\begin{aligned}\frac{\partial \bar{R}(d^*_{\text{PER}}, \mu)}{\partial \mu} &= \frac{1}{2} - \frac{\overline{\mu}}{2k}\left(1 - 2\Phi\left(-\frac{\mu}{\sigma}\right)\right) \\ &\quad - (\mu + k)\frac{\overline{\mu}}{k}\frac{1}{\sigma}\phi\left(\frac{\mu}{\sigma}\right) \\ &= \frac{1}{2} - \frac{\overline{\mu}}{2k}\left(1 - 2\Phi\left(-\frac{\mu}{\sigma}\right) + 2(\mu + k)\frac{1}{\sigma}\phi\left(\frac{\mu}{\sigma}\right)\right),\end{aligned}$$

in which note (1): $0 < 1 - 2\Phi\left(-\frac{\mu}{\sigma}\right) \leq 1$ for all $\mu \geq 0$, and (2): $2(\mu + k)\frac{1}{\sigma}\phi\left(\frac{\mu}{\sigma}\right) > 0$ is first increasing and then decreasing with a unique and finite maximum $C_k > 0$. Therefore, we conclude that for any $0 < \overline{\mu} < \min\left(\frac{k}{(1+C_k)}, \sqrt{\frac{\pi}{2}}\sigma\right)$, we have $\frac{\partial \bar{R}(d^*_{\text{PER}}, \mu)}{\partial \mu} > 0$ for all $\mu \geq 0$. That is, $\bar{R}(d^*_{\text{PER}}, \mu)$ is strictly increasing in $\mu \in [0, \infty]$.

Step 3 Montiel Olea et al. (2023) have already shown that

$$\sup_{\mu \geq 0} \bar{R}(d^*_{MMR,\text{linear}}, \mu) = \bar{R}(d^*_{MMR,\text{linear}}, 0) = \frac{k}{2}.$$

Then, by the conclusion from step 3 and noting $\bar{R}(d^*_{\text{PER}}, 0) = \frac{k}{2}$, we see that (D.1) indeed holds with the inequality strict for all $\mu > 0$, completing the proof. □