

# Decision Theory for Treatment Choice Problems with Partial Identification\*

José Luis Montiel Olea<sup>†</sup> Chen Qiu<sup>‡</sup> Jörg Stoye<sup>§</sup>

January 2026

## Abstract

We apply classical statistical decision theory to a large class of *treatment choice problems with partial identification*. We show that, 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. We uniquely characterize the minimax-regret optimal rule that least frequently randomizes, and show that, in some cases, it can outperform other minimax-regret optimal rules in terms of what we term *profiled regret*. We analyze the implications of our results in the aggregation of experimental estimates for policy adoption, extrapolation of Local Average Treatment Effects, and policy making in the presence of omitted variable bias.

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

---

\*We thank Xavier D'Haultfoeuille (the editor) and three anonymous referees, Karun Adusumilli, Isaiah Andrews, Tim Armstrong, Larry Blume, Tim Christensen, Tommaso Denti, Kei Hirano, Hiro Kaido, Tetsuya Kaji, Charles Manski, Douglas Miller, Francesca Molinari, Alex Tetenov, Kohei Yata, as well as seminar audiences at Aarhus, Berkeley, Binghamton, Bocconi, Bonn, CEMFI (Madrid), Columbia, Cornell, Indiana, LMU Munich, Manchester, Northwestern, Penn, Pittsburgh, and UCL 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, the 2023 Greater NY Metropolitan Area Econometrics Colloquium, CEME (Georgetown), Texas Econometrics Camp 2024, EC<sup>2</sup> 2024 (Amsterdam), the ASSA 2025 Annual Meeting (San Francisco), DAGStat 2025 (Berlin), and Statistische Woche 2025 (Wiesbaden) for feedback. Krishiv Shah, Henry Tan, Maximilian Yap, Finn Ye, Sara Yoo, Naiqi Zhang, and especially Brenda Quesada Prallon provided excellent research assistance. We gratefully acknowledge financial support from the NSF under grant SES-2315600.

<sup>†</sup>Email: jlo67@cornell.edu

<sup>‡</sup>Email: cq62@cornell.edu

<sup>§</sup>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 (2024). 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. We do so in a finite-sample framework characterized by a Gaussian likelihood and partial identification.

**MAIN RESULTS:** Three optimality criteria are routinely used to endorse or discard decision rules: *admissibility*, *maximin welfare*, and *minimax regret*.

*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”. Our first result shows that, whenever problems in our setting exhibit partial identification, *every* decision rule, no matter how exotic, is admissible (Theorem 1). As we discuss later, this result stands in stark contrast with applications of the admissibility criterion to point-identified treatment choice problems, in which admissibility meaningfully refines the class of decision rules (by, for example, discarding rules that randomize the policy recommendation). We also show that our result is not tied to the choice of Gaussian likelihood, but instead to the *bounded completeness* of the statistical model for the available data (Theorem 5).

*Maximin Welfare.* A decision rule is maximin(-welfare) optimal if it attains the highest worst-case expected welfare. Echoing critiques from Savage (1951) to Manski (2004), our second result shows 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). We show this may 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). Our third result shows that for cases where the identified set for payoff relevant parameters is large enough, there are *infinitely many* MMR optimal decision rules, *all of which* randomize the policy recommendation for some or all data realizations (Theorem 3). This presents an important 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.

*Least Randomizing MMR rule:* Finally, 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).

We analyze the regret of the least randomizing MMR rule after profiling out some parameters of the risk function. We specifically analyze *profiled regret*, by considering a parameter of interest and reporting worst-case expected regret at each of these parameter values (in a sense we make precise). We show, in the context of our running example, that our least-randomizing rule can profiled-regret dominate the MMR rule suggested by Stoye (2012a) and recently extended by Yata (2021) for a general class of treatment choice problems with partial identification (Proposition 2). More generally, we show that the rule which uniformly randomizes policy is inadmissible with respect to a profiled regret function that reports worst-case regret for different fixed values of the problem's point-identified parameters (Proposition 3). We finally discuss the extent to which our least-randomizing rule can be obtained by explicitly penalizing randomized policy recommendations in the policy maker's welfare function. We show that, under some conditions, the least-randomizing rule is minimax regret optimal (among a suitably defined class of decision rules) for a penalty function that penalizes all randomized assignments equally (Proposition 4).

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  $i = 1, \dots, n$  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 \(2026\)](#), [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, 2023](#); [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., 2025, 2024](#); [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 main results on admissibility, maximin welfare, minimax regret and the least randomizing MMR rule. The applications are presented in Section 4. Section 5 concludes. Appendix A contains proofs of our main results and Appendix B discusses the notion of profiled regret. Additional proofs and results can be found in the Online Appendix.

## 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 policy maker’s payoff when taking action  $a \in [0, 1]$  is given by the known function  $W(a, \theta)$ , where  $\theta \in \Theta$  is an unknown state of the world or *parameter* (possibly of infinite dimension).

Following most of the literature, we take  $W(\cdot, \theta)$  to be affine, in the sense that<sup>1</sup>

$$W(a, \theta) = aW(1, \theta) + (1 - a)W(0, \theta). \quad (2.1)$$

If the policy maker knew  $\theta$ , her optimal choice of action would be

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

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

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

$$Y \sim N(m(\theta), \Sigma). \quad (2.3)$$

Here, the function  $m(\cdot) : \Theta \rightarrow \mathbb{R}^n$  and the positive definite covariance matrix  $\Sigma$  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 data) may not identify 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. 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*,  $d : \mathbb{R}^n \rightarrow [0, 1]$ , 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

---

<sup>1</sup>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  $W(\cdot, \theta)$ .

implementing a new policy in country  $i = 0$  and observes estimates of the policy's effect for countries  $i = 1, \dots, n$ . Let  $Y = (Y_1, \dots, Y_n)^\top \in \mathbb{R}^n$  denote such experimental estimates and let  $(x_0, \dots, x_n)$  be nonrandom baseline covariates living in a metric (typically Euclidean) space  $\mathcal{X}$  with norm  $\|\cdot\|$ . The policy maker is willing to extrapolate from her data by assuming that  $W(1, \theta) = \theta(x_0)$ ,  $W(0, \theta) = 0$ ,  $U(\theta) = \theta(x_0)$  and that

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

where  $\theta : \mathcal{X} \rightarrow \mathbb{R}$  is an unknown Lipschitz function with known constant  $C$ .<sup>2</sup> For simplicity, in the example we write  $\mu_i$  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  $\mu_0$  is not necessarily identified. We assume that  $(x_1, \dots, x_n)$  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. Finally, we assume  $\|x_1 - x_0\| < \|x_2 - x_0\|$ , i.e., the nearest neighbor of country 0 is unique.<sup>3</sup> □

## 3 Main Results

In statistical decision theory, three criteria are commonly used to recommend decision rules: *admissibility*, *maximin welfare*, and *minimax regret*.<sup>4</sup> In this section, we show that 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:

---

<sup>2</sup>In fairness, while we seemingly have only one free parameter  $C$ , the choice of norm on  $\mathcal{X}$  may not be obvious. This is a common issue in causal inference; see [Imbens and Rubin \(2015, Section 18.5\)](#) for a discussion and suggestions.

<sup>3</sup>If two or more countries are nearest neighbors, their signals can be merged into one more precise signal.

<sup>4</sup>See [Stoye \(2012b\)](#) and references therein for theoretical trade-offs between these criteria.

**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

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

collect all means that can be generated as  $\theta$  ranges over  $\Theta$ . We refer to elements  $\mu \in M$  as *reduced-form* parameters because they are identified in the statistical model (2.3).<sup>5</sup> Define the *identified set* for the welfare contrast as function of  $\mu$  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

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

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).

**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

---

<sup>5</sup>The notation is consistent with our running example, in which the observable moments are  $(\mu_1, \dots, \mu_n)$ .

$\mathcal{S} \subseteq M \subseteq \mathbb{R}^n$  such that<sup>6</sup>

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

RUNNING EXAMPLE—CONTINUED: The identified set for the welfare contrast  $\mu_0$  is

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

Its extrema can be written as intersection bounds:

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

For  $\mu$  sufficiently close to the zero vector, we therefore have nontrivial partial identification.  $\square$

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  $d \in \mathcal{D}_n$  is (welfare-)admissible.*

*Proof.* See Appendix A.1.  $\square$

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 (hence, bounded 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.<sup>7</sup>

While the statement of Theorem 1 makes reference to the Gaussian model in (2.3), the above proof sketch only uses this model's *bounded completeness* (and the mutual absolute continuity of

---

<sup>6</sup>We force  $M$  to have a non-empty interior for simplicity of exposition. This would, for example, exclude equality constraints. For Theorem 3, we can weaken the assumption to allow such cases as long as  $\mathcal{S}$  is rich within  $M$ .

<sup>7</sup>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 Gaussian and Lebesgue measures in  $\mathbb{R}^n$ ). Indeed, Theorem 5 in Appendix A.5, establishes a stronger version of Theorem 1 that applies to boundedly complete statistical models (in a sense we make precise).<sup>8</sup> It is known that many exponential family distributions (Lehmann and Casella, 1998, Theorem 4.3.1, p.116) as well as some location families (Mattner, 1993, Theorem 2.1) are boundedly complete. However, our proof does not apply to distributions that are not boundedly complete even if they are close to Normal; for example, if  $Y = Z + \epsilon$ , where  $Z \sim N(m(\theta), \Sigma)$  and  $\epsilon$  is independent of  $Z$  and contains i.i.d components with characteristic functions containing zeros such as the uniform distributions in  $[-\eta, \eta]$  for any  $\eta > 0$  (c.f. Mattner 1993, Theorem 2.1 and references therein). While we do not know whether bounded completeness is strictly necessary for this result, it cannot just be dropped. To see this, consider the preceding counterexample with a uniform error term, in which we further impose  $n = 1$ ,  $\Sigma = 0$  and  $\eta = 1$ . Furthermore, suppose  $\bar{I}(\mu) = \mu + 1$  and  $\underline{I}(\mu) = \mu - 1$ . Then the coin-flip rule  $d_{\text{coin-flip}}(Y) = 1/2$  would be dominated by the rule that chooses 0 if  $Y < -2$ , chooses 1 if  $Y > 2$ , and 1/2 otherwise.

*Remark 1.* Theorem 1 shows that the notion of admissibility does not have any refinement power on the class of decision rules considered by the policy maker. We think this result is quite surprising, given that in treatment choice problems with point identification, admissibility does meaningfully refine the class of decision rules. For example, let  $n = 1$ ,  $\Theta = \mathbb{R}$ ,  $m(\theta) = W(1, \theta) = \theta$ , and  $W(0, \theta) = 0$ . In this case, the policy maker observes a noisy signal,  $Y \sim N(\theta, \sigma^2)$ , of the payoff-relevant, point-identified parameter  $\theta \in \mathbb{R}$ . 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 of all threshold rules is *complete* (in the sense that any non-threshold rule is dominated by a threshold rule). Therefore, any decision rule that is not a threshold rule can be dismissed by appealing to the notion of admissibility alone. We think it is quite surprising that introducing partial identification renders *all* decision rules admissible. Indeed, Theorem 1 implies that no rule, regardless how eccentric it may appear, is (welfare-)dominated. This makes it much more challenging to recommend a rule or a class of rules, at least without commitment to a more specific decision-theoretic optimality criterion (such as maximin welfare or minimax regret).

Theorem 1 also admits a more optimistic interpretation. 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

---

<sup>8</sup>We would like to thank Tim Christensen and the Editor for pointing this out.

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.

### 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

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

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<sup>9</sup>), 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,](#)

---

<sup>9</sup>In Manski’s (2004) example,  $W(0, \theta)$  does not depend on  $\theta$ , so that condition (3.4) is trivially satisfied.

2012; Yata, 2021) suggests that, with state spaces that describe real-world decision problems, the concern is more salient for maximin.<sup>10</sup>

### 3.3 Minimax Regret Admits Many Solutions

In view of Theorem 1 and Theorem 2, it seems natural to consider the *minimax regret* (MMR) optimality criterion. It is known that some point-identified treatment choice problems admit nonrandomized, essentially unique MMR optimal rules. In contrast, our results below will show that uniqueness of MMR optimal rules should not be expected to hold generally with partial identification. To make this point, we consider a 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 depend on data through an optimal linear index and ii) under weak additional conditions, any MMR rule that depends on the data through the optimal linear index must be randomized.

Our numerical and analytic findings below (see Figure 2 and Proposition 2) demonstrate that, beyond having the same worst-case regret, two MMR optimal rules might be qualitatively and quantitatively very different. This presents a practical challenge for a policy maker as, given the same data realizations, two MMR optimal rules may recommend different policy actions.

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

$$R(d, \theta) := U(\theta) \{ \mathbf{1}\{U(\theta) \geq 0\} - \mathbb{E}_{m(\theta)}[d(Y)] \}. \quad (3.5)$$

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

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

Solving minimax regret problems is often hard, both analytically and algorithmically. Algorithms exist for certain cases (Yu and Kouvelis, 1995; Chamberlain, 2000; Aradillas Fernández et al., 2025; Guggenberger and Huang, 2025), but it is known that obtaining the minimax solution of a decision

---

<sup>10</sup>Given that we next elaborate on multiplicity of MMR rules, it is of interest to also discuss the potential multiplicity of maximin rules: if  $W(0, \theta)$  is treated as known (e.g., equal to 0), the maximin rule may be unique. However, in settings where both  $W(1, \theta)$  and  $W(0, \theta)$  are unknown and can be arbitrarily “bad”, the typical result is that all decision rules are maximin.

problem—and sometimes even deciding whether a minimax solution exists—is NP hard in general; see [Daskalakis et al. \(2021\)](#). 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)  $\Theta$  is convex, centrosymmetric (i.e.,  $\theta \in \Theta$  implies  $-\theta \in \Theta$ ) and nonempty.  
(ii)  $m(\cdot)$  and  $U(\cdot)$  are linear.

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

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

$$d_{\text{RT}}^*(Y) := d_{\text{RT}}^*((w^*)^\top Y) := \Phi((w^*)^\top Y / \tilde{\sigma}) \quad (3.7)$$

for some uniquely characterized  $\tilde{\sigma} > 0$ .<sup>11</sup> Moreover, it then has two important algebraic properties:

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

and

$$\sup_{\theta \in \Theta} R(d_{\text{RT}}^*, \theta) = \sup_{\theta \in \Theta, m(\theta) = \mathbf{0}} R(d_{\text{RT}}^*, \theta). \quad (3.9)$$

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.<sup>12</sup> 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*

---

<sup>11</sup>For readability, we slightly abuse notation and, when a decision rule  $d$  depends on data  $Y$  only through a simple feature like  $w^\top Y$  or  $Y_1$ , we write  $d(Y)$  as  $d(w^\top Y)$  or  $d(Y_1)$ .

<sup>12</sup>[Yata \(2021\)](#) also provide sufficient conditions under which these conditions hold true. Our proof is constructive, and our results below apply to a large class of models with these properties and shall not be considered as a specific example.

*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  $\bar{I}(\mu) > \bar{I}(\mathbf{0})$  and  $\bar{I}(\mathbf{0})$  is large enough (see (A.2)), 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  $\bar{I}(\mu)$  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.<sup>13</sup> 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  $d_{\text{RT}}^*$  is MMR optimal, then so is the *piecewise linear* rule

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

where  $\rho^* > 0$  is characterized in Appendix A.3, Equation (A.11). 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  $\Sigma$  (or as  $\Sigma$  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  $\bar{I}(\mathbf{0})$  is large enough, then randomization is necessary for MMR optimality; if bounds are furthermore differentiable in reduced-form parameters

---

<sup>13</sup>This condition will always be met if  $\Sigma$  is sufficiently small, holding other parameter values fixed. Heuristically, this means that our multiplicity result is always relevant as sample size becomes large. However, to what extent our result is a useful characterization in a limit experiment is beyond the scope of the current paper and an interesting question that we leave for future research.

at  $\mathbf{0}$ , randomization is necessary for any MMR rule that depends on a linear index of the data.<sup>14</sup> In particular, the threshold rule

$$d_0^*((w^*)^\top Y) = \mathbf{1}\{(w^*)^\top Y \geq 0\} \quad (3.11)$$

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, these are characterized by a weighting of studentized signals that resembles a triangular kernel for small enough  $C$  and turns into a nearest neighbor kernel as  $C$  becomes large.  $\square$

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

(i) *If*

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

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

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

*where  $m_0^* > C \|x_1 - x_0\|$  solves a simple fixed point problem ((C.7) in Online Appendix C.3).*

(ii) *If*

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

*then  $\mathbf{1}\{Y_1 \geq 0\}$  is MMR optimal.*

(iii) *If*

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

---

<sup>14</sup>The class of nonrandomized but otherwise unrestricted rules is not interestingly different from the class of all rules due to the possibility of purifying randomized rules. See Remark 5 for discussion and references. Similarly, the differentiability condition is needed to preclude that some component of  $Y$  can effectively be used as randomization device.

then the rule

$$d_{\text{linear}}^*(Y_1) := \begin{cases} 0, & Y_1 < -\rho^*, \\ \frac{Y_1 + \rho^*}{2\rho^*}, & -\rho^* \leq Y_1 \leq \rho^*, \\ 1, & Y_1 > \rho^*, \end{cases} \quad (3.14)$$

is MMR optimal; here,  $\rho^* < C \|x_1 - x_0\|$  is the unique strictly positive solution to  $\rho^* = C \|x_1 - x_0\| (1 - 2\Phi(\rho^*/\sigma_1))$ . The rule  $d_{RT}^*(Y_1) := \Phi(Y_1/\tilde{\sigma})$ , where

$$\tilde{\sigma} = \sqrt{2C^2 \|x_1 - x_0\|^2 / \pi - \sigma_1^2}, \quad (3.15)$$

is MMR optimal as well, as are all convex combinations of these rules.

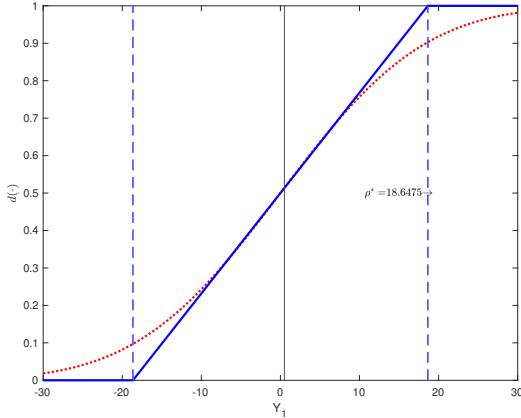
(iv) If Equation (3.13) holds, no linear threshold rule is MMR optimal.

*Proof.* See Online Appendix C.3.  $\square$

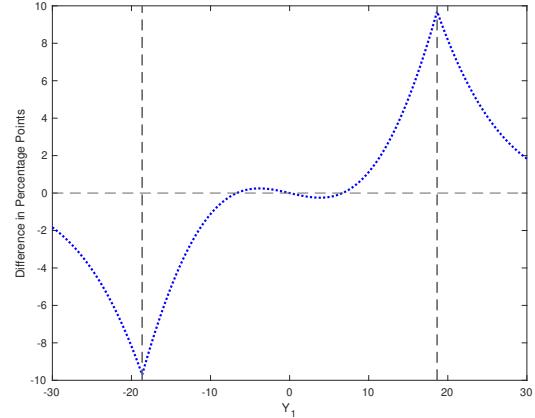
*Remark 2.* All decision rules above converge to the one from case (ii) as  $C\|x_1 - x_0\| \rightarrow \sqrt{\pi/2} \cdot \sigma_1$ .

*Remark 3.* This solution relates to the literature as follows. The problem is within the framework considered by [Yata \(2021\)](#) (who uses results from [Stoye 2012a](#)), and his analysis applies; in particular, our linear index differs from his  $w^*$  only by a more explicit characterization. Although some of our proof steps use algebra from [Stoye \(2012a\)](#), the alternative solutions in (iii), the uniqueness statement, and part (iv) are entirely new. [Ishihara and Kitagawa \(2021\)](#) numerically find a solution within 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. This is because we view their approach primarily as a numerical strategy for finding a minimax solution over a constrained class. Indeed, given a statistical model and a risk function, one can always try to find a minimax decision rule numerically within a class of threshold rules, which may or may not pin down the true minimax rule.

*Remark 4.* If  $\mu$  is exogenous and known, then the decision rule  $d_{\text{known } \mu}^*(\mu) := \max\{\min\{\bar{I}(\mu)/(\bar{I}(\mu) - I(\mu)), 1\}, 0\}$  uniquely attains MMR ([Manski, 2007b](#)). Only our new rule  $d_{\text{linear}}^*$  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.



(a)  $d_{\text{RT}}^*$  (red) and  $d_{\text{linear}}^*$  (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  $x_0 = 0$ ,  $x_1 = -7.5$ ,  $\sigma_1 = 3.9$ ,  $x_2 = 7.9$ ,  $\sigma_2 = 2.4$ ,  $C = 2.5$ ). Panel b): Difference between rules as a function of the nearest neighbor's outcome.

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, our next results below will show that one may plausibly have preferences among the different MMR rules.

**RUNNING EXAMPLE—CONTINUED:** For some parameter values, there are infinitely many MMR rules that depend on the data only through  $Y_1$  and therefore on  $(\mu_1, \dots, \mu_n)^\top$  only through  $\mu_1$ . To better compare and visualize these rules, it appears natural to examine their worst-case performance for different fixed values of  $\mu_1$ . Formally, this leads us to define *profiled regret*

$$\bar{R}(d, \gamma) := \sup_{(\mu_0, \mu_1, \dots, \mu_n)^\top \in \mathbb{R}^{n+1}: \mu_1 = \gamma, \mu_0 \in I(\mu), \mu \in M} \mu_0 (\mathbf{1}\{\mu_0 \geq 0\} - \mathbb{E}_\mu[d(Y)]). \quad (3.16)$$

For the same parameters used in Figure 1, Figure 2 traces out  $\bar{R}(d, \gamma)$  over  $\gamma \in [-30, 30]$  for four different decision rules. The blue (solid) line is  $d_{\text{linear}}^*$  (see Equation (3.10)); the red (dotted) line is  $d_{\text{RT}}^*$  (see Equation (3.7)); the black (bimodal, solid) line is the symmetric threshold rule  $d_0^*(Y) = \mathbf{1}\{Y_1 \geq 0\}$ ; and the green (dashed) line is  $d_{\text{coin-flip}} = 1/2$ .<sup>15</sup>

An immediate use of Figure 2 is to compare decision rules in terms of their worst-case regret.

<sup>15</sup>Appendix C.7.1 presents algebraic and computational details that underlie this figure.

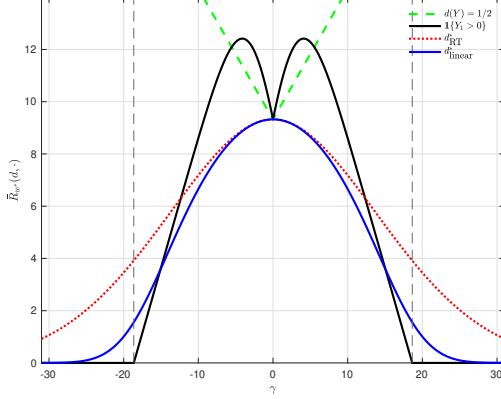


Figure 2: Profiled regret of four decision rules; parameter values are as in Figure 1.

For example, consistent with Theorem 3-(ii),  $d_0^*$  is not MMR optimal (the maximum of the black curve is clearly above those of  $d_{\text{RT}}^*$  and  $d_{\text{linear}}^*$ ). 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. But Figure 2 also reveals interesting differences among MMR optimal rules. In particular,  $d_{\text{linear}}^*$  appears to have smaller profiled regret than  $d_{\text{RT}}^*$  virtually everywhere. Indeed, Proposition 2 in Appendix B shows that whenever condition (iii) of Proposition 1 holds,  $d_{\text{linear}}^*$  profiled-regret dominates  $d_{\text{RT}}^*$  for a large range of values of  $C \|x_1 - x_0\|$  and  $\sigma_1$ , including those used in Figure 2 (though there exists parameter values under which the dominance does not hold for small nonzero  $\gamma$ ). Moreover, the considerable difference between profiled regret functions in Figure 2 continues beyond the figure: The ratio  $\bar{R}(d_{\text{linear}}^*, \gamma)/\bar{R}(d_{\text{RT}}^*, \gamma)$  decays to zero at exponential rate as  $\gamma \rightarrow \pm\infty$ . Finally, Figure 2 illustrates that  $d_{\text{coin-flip}}$  is profiled-regret inadmissible in this example. In fact, profiled-regret inadmissibility of  $d_{\text{coin-flip}}$  holds more generally; see Proposition 3 in Appendix B.

While profiling with respect to  $\mu_1$  is natural in this example, we emphasize that the idea of profiled regret is much more general. In Appendix B, we extend (3.16) to the profiling out of any linear function of  $(\mu_1, \dots, \mu_n)^\top$ .

### 3.4 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. We justify our approach by a decision maker with a lexicographic preference (See Remark 5 below). Alternatively, one may also explicitly incorporate aversion to randomization in the loss function and try to solve the MMR criterion under a modified risk function. See Section C.1 for detailed discussions.

To formalize this, observe that  $d_{\text{RT}}^*$  and  $d_{\text{linear}}^*$  can be considered smoothed versions of  $d_0^*$  in a sense that we now make precise. Let  $F : \mathbb{R} \rightarrow [0, 1]$  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 restrict attention to c.d.f.'s whose associated distributions are *symmetric* ( $F(-x) = 1 - F(x)$ ) and *unimodal* ( $F(\cdot)$  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)

$$V(F \circ w^*) := (\sup\{x \in \mathbb{R} : F(x) = 0\}, \inf\{x \in \mathbb{R} : F(x) = 1\}). \quad (3.17)$$

All MMR decision rules considered in this paper are in  $\tilde{\mathcal{D}}_n$ . We next show that  $d_{\text{linear}}^*$  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  $V(d_{\text{linear}}^* \circ w^*) \subseteq V(F \circ w^*)$ , with equality if and only if  $F = d_{\text{linear}}^*$ .*

*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  $d_{\text{linear}}^*$ , 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  $\theta$  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 enough,  $d_{\text{linear}}^*$  can be expressed as (3.14). Of note, an identified set for  $\mu_0$  given the true mean of  $Y_1$  (i.e.,  $\mu_1$ ) is

$$[ \mu_1 - C \|x_1 - x_0\|, \mu_1 + C \|x_1 - x_0\| ],$$

which can be estimated naturally by

$$[ Y_1 - C \|x_1 - x_0\|, Y_1 + C \|x_1 - x_0\| ]. \quad (3.18)$$

As  $d_{\text{linear}}^*$  only randomizes when  $|Y_1| < \rho^*$ , we see that interval (3.18) contains 0 whenever  $d_{\text{linear}}^*$  randomizes. Equivalently,  $d_{\text{linear}}^*$  always (never) implements the new policy when (3.18) is to the right (left) of zero. In this sense, the estimated identified set is explicitly used for decision making.<sup>16</sup> In contrast,  $d_{\text{RT}}^*$  always randomizes the policy recommendation, although for large  $Y_1$  the fraction of population assigned to treatment will be large.  $\square$

*Remark 5.* Consider a decision maker who wishes to pick an optimal rule according to MMR but is hesitant to implement randomized rules due to additional inconvenience cost from implementation or concerns over *ex-post* fairness (i.e., units in the same population may receive different treatments). In this case, it is indeed possible to define the MMR problem among nonrandomized rules. However, if we do not place any restriction to the class of nonrandomized rules, game-theoretic purification arguments (Dvoretzky et al., 1951; Khan et al., 2006) suggest existence of nonrandomized solutions that approximate the risk profile of the randomized MMR optimal rule. Moreover, these solutions would be unattractive and unnatural; for example, they cannot be monotone due to Theorem 3(iii). In light of these observations, we think there is little value in pursuing solutions over nonrandomized rules in our setup, unless one focuses on a specific class of benign or reasonable nonrandomized rules, e.g. as in Ishihara and Kitagawa (2021). Furthermore, within the class of rules  $\tilde{\mathcal{D}}_n$ , our least randomizing MMR rule can be justified by a decision maker who possesses a lexicographic preference with a priority given to minimizing MMR criterion over minimizing the inconvenience cost of randomization (in our case, the cost is proxied by the Lebesgue measure of the realization of  $(w^*)^\top Y$  taking a value in  $(0, 1)$ ). Such lexicographic ordering of statistical decisions has ample

<sup>16</sup>One may consider a scenario in which, to aid interpretability of decision rules, the decision maker does not wish to randomize whenever the estimated identified set for  $\mu_0$  (3.18) does not contain zero, i.e. the estimate suggests that the sign of  $\mu_0$  is identified. This restricts decision rules to the following set:

$$\{d \in \tilde{\mathcal{D}}_n : F \in \mathcal{F}, F(x) = 0 \text{ for all } x \leq -C \|x_1 - x_0\| \text{ and } F(x) = 1 \text{ for all } x \geq C \|x_1 - x_0\|\} \quad (3.19)$$

As (3.19) is a subset of  $\tilde{\mathcal{D}}_n$  containing  $d_{\text{linear}}^*$ , Theorem 4 implies that  $d_{\text{linear}}^*$  is also least-randomizing MMR in (3.19). Moreover,  $d_{\text{linear}}^*$  is currently the only known rule in set (3.19) that is also globally MMR optimal.

precedents in the literature: for statistical decision rules, we first remove dominated rules and then choose among undominated ones according to further optimality criteria (Wald, 1950; Manski, 2021); in the Neyman-Pearson paradigm of selecting hypothesis tests, one first controls size and then maximizes power; in choosing estimators, it is customary to focus on unbiased estimators, among which the one that minimizes variances is regarded as optimal.

## 4 Further Applications

### 4.1 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),<sup>17</sup> 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  $\Theta$  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

$$\begin{pmatrix} \hat{m}_1 \\ \hat{m}_2 \end{pmatrix} \sim N \left( \begin{pmatrix} m_1(\theta) \\ m_2(\theta) \end{pmatrix}, \Sigma \right), \quad (4.1)$$

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  $\Sigma$  is positive definite. Suppose a new instrument expanded the complier subpopulation through shifting  $p(1)$  to  $p(1) + \alpha$  for some  $\alpha > 0$ . Mogstad et al. (2018) show that one parameter of interest is the ‘‘policy-relevant treatment

---

<sup>17</sup>See, for example, Assumption I and Equation (2) in Mogstad and Torgovitsky (2018). Also see Imbens and Angrist (1994) for additional references.

effect” (Heckman and Vytlacil, 2005)

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

Suppose the welfare contrast  $U(\theta)$  equals  $\text{PRTE}(\alpha)$ <sup>18</sup>, which can be written as

$$U(\theta) := \text{PRTE}(\alpha) = \frac{m_1(\theta)}{\alpha + m_2(\theta)} + \frac{1}{\alpha + m_2(\theta)} \int_{p(1)}^{p(1)+\alpha} \text{MTE}(v) dv. \quad (4.2)$$

Hence, the decision maker wants to find an optimal treatment policy given partial identification of parameter (4.2) in model (4.1). In Online Appendix C.7.3, we verify that Theorem 1 applies. Therefore, any decision rule is admissible in this example. For example, implementing a policy for large values (and more surprisingly, for very negative values) of the IV estimator would be admissible, as would be the approach of Christensen et al. (2022), who discuss the same application.

## 4.2 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.<sup>19</sup> In general, if there exists selection on unobservables (that is,  $D$  is correlated with  $W$ ), then  $\beta_{\text{long}}$  is only partially identified. Specifically, Diegert et al. (2022,

---

<sup>18</sup>This could be motivated by a hypothetical policy question in which  $W(0, \theta) = 0$  and  $W(1, \theta) = \text{PRTE}(\alpha)$ . Note there exist other policy questions of interest and other versions of  $U(\theta)$ , different from  $\text{PRTE}(\alpha)$  but still partially identified, may arise.

<sup>19</sup>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.

Theorem 4) show that the identified set of  $\beta_{\text{long}}$  given  $\beta_{\text{med}}$  is

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

where

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

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

[Diegert et al. \(2022\)](#) use this result to ask: How strong does omitted variables bias have to be to potentially overturn findings based on  $\beta_{\text{med}}$ ? 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  $\hat{\beta}_{\text{med}} \sim N(\beta_{\text{med}}, \sigma^2)$ . The results of [Diegert et al. \(2022\)](#) imply an estimated breakdown point  $\tilde{k}(\hat{\beta}_{\text{med}}) := \hat{\beta}_{\text{med}}$  for positive  $\hat{\beta}_{\text{med}}$ . 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

$$d_{\text{linear}}^*(\hat{\beta}_{\text{med}}) := \begin{cases} 0, & \hat{\beta}_{\text{med}} < -\rho^*, \\ \frac{\hat{\beta}_{\text{med}} + \rho^*}{2\rho^*}, & -\rho^* \leq \hat{\beta}_{\text{med}} \leq \rho^*, \\ 1, & \hat{\beta}_{\text{med}} > \rho^*, \end{cases} \quad (4.3)$$

where  $\rho^* < k$  is the unique strictly positive solution to  $\rho^* = k(1 - 2\Phi(-\rho^*/\sigma))$ . 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?<sup>20</sup> Due to its least randomizing property,  $d_{\text{linear}}^*$  implies the tightest possible answer to this question. Specifically, MMR supports

---

<sup>20</sup>Informal exploration of this question goes back at least to [Stoye \(2009b, see Table 3\)](#).

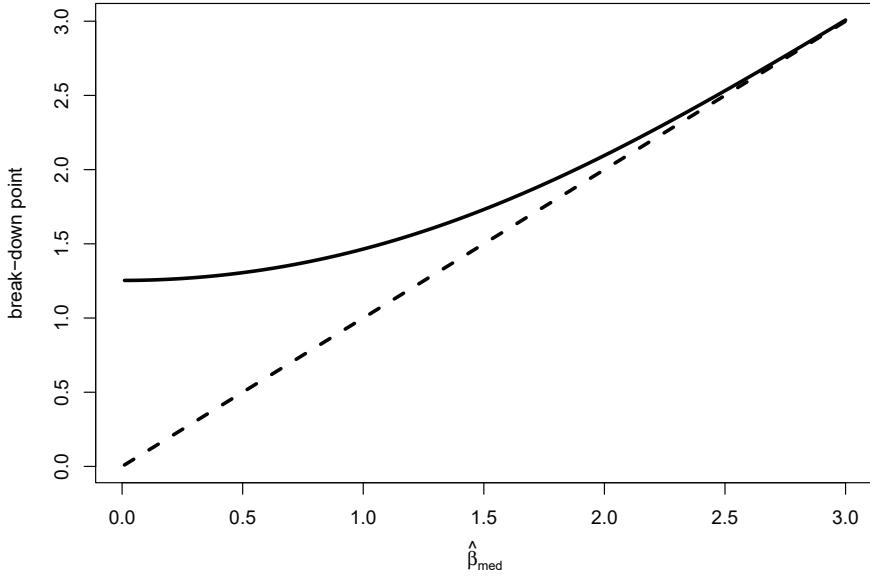


Figure 3:  $\bar{k}$  (solid) and  $\tilde{k}$  (dashed) as a function of  $\hat{\beta}_{\text{med}}$

non-randomized policy assignment up to the “decision theoretic breakdown point”

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

In contrast, the implied breakdown point of  $d_{\text{RT}}^*$  is a constant across all values of  $\hat{\beta}_{\text{med}}$ , as  $d_{\text{RT}}^*$  always randomizes regardless of data realizations when  $k > \sqrt{\frac{\pi}{2}}\sigma$ . Figure 3 displays both  $\tilde{k}(\hat{\beta}_{\text{med}})$  and  $\bar{k}(\hat{\beta}_{\text{med}})$  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.

## 5 Conclusion

In this paper, we used statistical decision theory to study treatment choice problems with partial identification. 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 provide a decision rule that is *least randomizing* in a large class of MMR optimal rules including all known ones. We show, in the context of our running example, that our least-randomizing rule can profiled-regret dominate other MMR rules.

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.

## A Proofs of Main Results

### A.1 Proof of Theorem 1

It suffices to verify conditions A1-A2 in Theorem 5 for the class of distributions  $N(m(\theta), \Sigma), \theta \in \Theta$ .

A1: Consider the class of distributions  $N(m(\theta), \Sigma), \theta \in m^{-1}(\mathcal{S})$ . As  $\theta$  ranges over  $m^{-1}(\mathcal{S})$ ,  $m(\theta)$  ranges over  $\mathcal{S}$ . Thus, it suffices to show that

$$Y \sim N(\mu, \Sigma), \quad \mu \in \mathcal{S}, \tag{A.1}$$

is *complete* (Casella and Berger, 2002, Definition 6.2.21) and therefore also bounded complete. Define the vector  $\tilde{\mu} := \Sigma^{-1}\mu$  ranging over the set

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

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

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

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

A2: Consider any set  $E$  such that  $P(E) = 0$  for all  $P \sim N(m(\theta), \Sigma)$  as  $\theta$  ranges over  $m^{-1}(\mathcal{S})$ . By Skorohod (2012) (Theorem 2, p. 83), the Gaussian measures in  $\mathbb{R}^n$  with the same positive definite covariance matrix  $\Sigma$  are equivalent (mutually absolute continuous with each other), implying

$P(E) = 0$  for all  $P \sim N(m(\theta), \Sigma)$  such that  $\theta \in \Theta \setminus m^{-1}(\mathcal{S})$  as well.

## A.2 Proof of Theorem 2

First, we can bound

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

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

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

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

Recall that we impose Assumption 1, the existence of a MMR optimal rule  $d^*$  that depends on the data only through  $(w^*)^\top Y$  and that satisfies (3.8) and (3.9), that

$$\bar{I}(\mathbf{0}) > \sqrt{\pi/2} \cdot \sqrt{(w^*)^\top \Sigma w^*} \cdot s_{w^*}(0), \quad (\text{A.2})$$

and finally that  $\bar{I}(\mu) > \bar{I}(\mathbf{0})$  for some  $\mu \in M$ .

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

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

$$\mathbf{R} := \inf_{d \in \mathcal{D}_n} \sup_{\theta \in \Theta} \{U(\theta) (\mathbf{1}\{U(\theta) \geq 0\} - \mathbb{E}_{m(\theta)}[d(Y)])\}. \quad (\text{A.3})$$

STEP 1 (Minimax Regret Value): We first show that

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

where

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

Lemma C.1 shows that 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  $\Theta$  imply

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

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

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

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  $\Gamma_{w^*} := \{\gamma \in \mathbb{R} : (w^*)^\top m(\theta) = \gamma, \theta \in \Theta\}$  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

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

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

$$\sup_{\theta \in \Theta} (U(\theta) (\mathbf{1}\{U(\theta) \geq 0\} - \mathbb{E}_\gamma[d((w^*)^\top Y)])) \tag{A.4}$$

$$\begin{aligned} &= \sup_{\gamma \in \Gamma_{w^*}} \left( \sup_{\theta \in \Theta, (w^*)^\top m(\theta) = \gamma} U(\theta) (\mathbf{1}\{U(\theta) \geq 0\} - \mathbb{E}_\gamma[d((w^*)^\top Y)]) \right) \\ &= \sup_{\gamma \in \Gamma_{w^*}} \left( \sup_{U^* \in ISU_{w^*}(\gamma)} U^* (\mathbf{1}\{U^* \geq 0\} - \mathbb{E}_\gamma[d((w^*)^\top Y)]) \right), \end{aligned} \tag{A.5}$$

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

$$\begin{aligned}\underline{k}_{w^*}(\gamma) &:= \inf ISU_{w^*}(\gamma) = \inf\{U(\theta) : (w^*)^\top m(\theta) = \gamma, \theta \in \Theta\} \\ \bar{k}_{w^*}(\gamma) &:= \sup ISU_{w^*}(\gamma) = \sup\{U(\theta) : (w^*)^\top m(\theta) = \gamma, \theta \in \Theta\}.\end{aligned}\tag{A.6}$$

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

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

and therefore that  $ISU_{w^*}(\gamma) \subseteq [-\bar{k}_{w^*}(-\gamma), \bar{k}_{w^*}(\gamma)]$ ,  $\forall \gamma \in \Gamma_{w^*}$ . Equation (A.5) then implies that (A.4) is bounded above by

$$\sup_{\gamma \in \Gamma_{w^*}} \left( \sup_{U^* \in [-\bar{k}_{w^*}(-\gamma), \bar{k}_{w^*}(\gamma)]} U^* (\mathbf{1}\{U^* \geq 0\} - \mathbb{E}_\gamma[d((w^*)^\top Y)]) \right). \tag{A.7}$$

Lemma C.2 shows that, under Assumption 1,  $\bar{k}_{w^*}(\gamma)$  is concave (and therefore  $-\bar{k}_{w^*}(-\gamma)$  is convex). Furthermore, Lemma C.3 shows that, under Assumption 1, the superdifferential (Rockafellar, 1997, p. 214-215) of the function  $\bar{k}_{w^*}(\cdot)$  at  $\gamma = 0$ ,

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

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 C.3 established that  $\Gamma_{w^*}$  is symmetric around 0. The definition of superdifferential then gives

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

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.7) then implies that (A.4) is further bounded above by

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

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

STEP 3 (“Linear Embedding” Minimax Regret Problem): The previous step and the fact that  $\Gamma_{w^*} \subseteq \mathbb{R}$  imply that (A.4) is bounded above by

$$\inf_{d \in \mathcal{D}} \sup_{\gamma \in \mathbb{R}} \left( \sup_{U^* \in [s_{w^*}(0)\gamma - \bar{k}_{w^*}(0), s_{w^*}(0)\gamma + \bar{k}_{w^*}(0)]} U^* (\mathbf{1}\{U^* \geq 0\} - \mathbb{E}_\gamma[d(\hat{\gamma})]) \right), \quad (\text{A.8})$$

where

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

Lemma C.4 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.8) 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:

$$d_{\text{RT}}^*(\hat{\gamma}) := \Phi \left( \hat{\gamma} / \sqrt{\frac{2 \cdot \bar{k}_{w^*}(0)^2}{\pi \cdot s_{w^*}(0)^2} - (w^*)^\top \Sigma w^*} \right) \quad (\text{A.9})$$

$$d_{\text{linear}}^*(\hat{\gamma}) := \begin{cases} 0, & \hat{\gamma} < -\rho^*, \\ \frac{\hat{\gamma} + \rho^*}{2\rho^*}, & -\rho^* \leq \hat{\gamma} \leq \rho^*, \\ 1, & \hat{\gamma} > \rho^*, \end{cases} \quad (\text{A.10})$$

where  $\rho^* < \frac{\bar{k}_{w^*}(0)}{s_{w^*}(0)}$  is the unique strictly positive solution to

$$\left( \frac{s_{w^*}(0)}{2 \cdot \bar{k}_{w^*}(0)} \right) \rho^* - \frac{1}{2} + \Phi \left( -\frac{\rho^*}{\sqrt{(w^*)^\top \Sigma w^*}} \right) = 0. \quad (\text{A.11})$$

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  $d^\star \in \mathcal{D}$  solves the linear embedding problem, then

$$d^\star \circ w^*(Y) := d^\star((w^*)^\top Y) \in \mathcal{D}_n$$

is MMR optimal in the original decision problem (A.3).

This is because Step 1 implies that

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

and Step 2 implies that

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

where the last equality follows from Step 3. Consequently,

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

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

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

and

$$\bar{k}_{w^*}(0) > \sqrt{\pi/2} \cdot \sqrt{(w^*)^\top \Sigma w^*} \cdot s_{w^*}(0). \quad (\text{A.12})$$

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

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

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

$$\bar{I}(\mu) \leq \bar{I}(\mathbf{0}) + s_{w^*}(0)(w^*)^\top \mu, \text{ for all } \mu \in M. \quad (\text{A.13})$$

If  $s_{w^*}(0) = 0$ , Equation (A.13) then implies  $\bar{I}(\mu) \leq \bar{I}(\mathbf{0})$  for all  $\mu \in M$ , contradicting the assumption that there exists  $\mu \in M$  such that  $\bar{I}(\mu) > \bar{I}(\mathbf{0})$ .

Second, as Step 1 showed that  $\bar{I}(\mathbf{0}) = \bar{k}_{w^*}(0)$ , (A.12) is directly implied by (A.2).

CONCLUSION: Steps 1-5 imply there are infinitely many rules that solve the problem (A.3).

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

Consider any decision rule  $d_m(\cdot)$  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  $d_m(\cdot) \in \{0, 1\}$  for all data realizations. Then we must have  $d_m(w^\top Y) = \mathbf{1}\{w^\top Y \geq c\}$  for some  $-\infty \leq c \leq \infty$ . The worst-case expected regret of such a rule satisfies

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

where we used that  $\underline{I}(\mathbf{0}) = -\bar{I}(\mathbf{0})$  by centrosymmetry. The last inequality is strict unless  $c = 0$ . As  $\bar{I}(\mathbf{0})/2$  is the MMR value of the problem,  $d_m(\cdot)$  cannot be MMR optimal if  $c \neq 0$ . For  $w = w^*$ , substantial additional algebra that we relegate to Lemma C.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:

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

Note that  $g_w(\mathbf{0}) = \bar{I}(\mathbf{0})/2$  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  $g_w(\mu)$ . Because  $M$  contains an open set including  $\mathbf{0}$  by Definition 2 and Assumption 1 and since we assumed differentiability of  $\bar{I}(\mu)$  at  $\mathbf{0}$ , this requires a first-order condition

$$\frac{\partial g_w(\mu)}{\partial \mu_j} \Big|_{\mu=\mathbf{0}} = \frac{1}{2} \frac{\partial \bar{I}(\mathbf{0})}{\partial \mu_j} - \frac{w_j}{\sqrt{w^\top \Sigma w}} \bar{I}(\mathbf{0}) \phi(0) = 0, \quad j = 1 \dots n. \quad (\text{A.14})$$

To simplify expressions, change co-ordinates (if necessary) so that  $w^* = (1, 0, \dots, 0)^\top$ . Because  $w^*$  must fulfil (A.14), we have that  $\frac{\partial \bar{I}(\mathbf{0})}{\partial \mu_j} = 0$  for  $j = 2, \dots, n$ . But this, in turn, means that (A.14) requires  $w_2 = \dots = w_n = 0$ . 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  $w^* = (1, 0, \dots, 0)$  solves (A.14) 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

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

where  $F \in \mathcal{F}$  is a symmetric and unimodal c.d.f (thus weakly increasing as well),  $\hat{\gamma} := (w^*)^\top Y \sim N(\gamma, \sigma^2)$ , 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

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

where  $\bar{k}_w^*$  is defined in (A.6). Letting  $g_F(\gamma) := \bar{k}_{w^*}(\gamma) (1 - \mathbb{E}_\gamma[F(\hat{\gamma})])$ . Since  $\hat{\gamma} \sim N(\gamma, \sigma^2)$ , we may further calculate (using integration by parts)

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

Therefore,  $g_F(\gamma) = \bar{k}_{w^*}(\gamma) \int \Phi\left(\frac{s-\gamma}{\sigma}\right) dF(s)$ . Note that  $F(-x) = 1 - F(x)$  for all  $x \in \mathbb{R}$ ; hence,  $\int \Phi\left(\frac{s}{\sigma}\right) dF(s) = \frac{1}{2}$  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 C.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.<sup>21</sup> Then  $0 \in \partial g_F(0)$  is necessary for optimality. To show that it fails, compute the generalized gradient as<sup>22</sup>

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

---

<sup>21</sup>The generalized gradient of  $g : \mathbb{R} \rightarrow \mathbb{R}$  at  $x$  equals  $\partial g(x) := \{\xi \in \mathbb{R} : \limsup_{y \rightarrow x, t \downarrow 0} \frac{f(y+tv) - f(y)}{t} \geq \xi v, \forall v \in \mathbb{R}\}$ . See Clarke (1990, p. 27).

<sup>22</sup>We can verify, following the same steps in the proof for Lemma C.8, that conditions of Proposition 2.3.13 in Clarke (1990) are satisfied, so that the chain rule can be applied.

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.8) (Lemma C.4). In particular, evaluating  $g_{\text{linear}}^{(1)}(\gamma)$  at  $\gamma = 0$ , where  $g_{\text{linear}}(\gamma)$  is defined in Lemma C.7 with  $k = \frac{\bar{k}_{w^*}(0)}{\tilde{s}_{w^*}(0)}$  and  $\sigma^2 = (w^*)^\top \Sigma w^*$ , one finds

$$\int \phi\left(\frac{s}{\sigma}\right) d(d_{\text{linear}}^*(s)) = \int_{-\rho^*}^{\rho^*} \phi\left(\frac{s}{\sigma}\right) \frac{1}{2\rho^*} ds = \frac{s_{w^*}(0)\sigma}{2\bar{k}_{w^*}(0)}, \quad (\text{A.15})$$

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

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

where the first 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  $d_{\text{linear}}^*$  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.15), 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.

## A.5 Proof of Theorem 5: General Theorem on Admissibility.

In this section, we prove a general theorem regarding the admissibility of treatment choice rules in partially identified models. Let  $\theta \in \Theta$ ,  $U(\cdot)$ , and  $m(\cdot)$  be defined as in Section 2 and the set  $\mathcal{S}$  as in Definition 2. We observe a random vector  $Y \in \mathcal{Y} \subseteq \mathbb{R}^n$  whose distribution  $P_{m(\theta)}$  is in

$$\mathcal{P} := \{P_{m(\theta)} : \theta \in \Theta\} \quad (\text{A.16})$$

defined over the common sample space  $\mathcal{Y}$  endowed with a  $\sigma$ -algebra  $\mathcal{A}$ .

**Theorem 5.** *Consider a treatment choice problem with payoff function (2.1) and a class of statistical models (A.16) that exhibit nontrivial partial identification in the sense of Definition 2. In addition, suppose the following two conditions hold true:*

*A1 The family of distributions  $\mathcal{P}_{\mathcal{S}} := \{P_{m(\theta)} : \theta \in m^{-1}(\mathcal{S})\}$  is bounded complete, i.e., for all bounded functions  $f : \mathbb{R}^n \rightarrow \mathbb{R}$ ,  $\mathbb{E}_{m(\theta)}[f(Y)] = 0$  for all  $\theta \in m^{-1}(\mathcal{S})$  implies  $f(y) = 0$ , a.e.  $\mathcal{P}_{\mathcal{S}}$ .*

*A2 For any set  $E \in \mathcal{A}$ ,  $P(E) = 0$  for all  $P \in \mathcal{P}_{\mathcal{S}}$  implies  $P(E) = 0$  for all  $P \in \mathcal{P} \setminus \mathcal{P}_{\mathcal{S}}$ .*

Then, every decision rule  $d \in \mathcal{D}_n$  is (welfare-)admissible.

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

$$U(\theta)\mathbb{E}_{m(\theta)}[d'(Y)] \geq U(\theta)\mathbb{E}_{m(\theta)}[d(Y)] \quad (\text{A.17})$$

for all  $\theta \in \Theta$ , with a strict inequality for some  $\theta$ .

STEP 1: We first show that (A.17) 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

$$U(\theta^*)\mathbb{E}_{m(\theta^*)}[d'(Y)] > U(\theta^*)\mathbb{E}_{m(\theta^*)}[d(Y)],$$

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

$$\mathbb{E}_{m(\theta^*)}[d'(Y)] > \mathbb{E}_{m(\theta^*)}[d(Y)]. \quad (\text{A.18})$$

Define  $\mu^* := m(\theta^*)$ . Since  $\mu^* \in \mathcal{S}$ , there must exist some  $\theta_{\mu^*} \in \Theta$ , such that  $m(\theta_{\mu^*}) = \mu^* = m(\theta^*)$

and  $U(\theta_{\mu^*}) < 0$ . Therefore, (A.18) implies

$$U(\theta_{\mu^*})\mathbb{E}_{m(\theta_{\mu^*})}[d'(Y)] < U(\theta_{\mu^*})\mathbb{E}_{m(\theta_{\mu^*})}[d(Y)],$$

contradicting (A.17). We conclude that

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

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

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

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

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

STEP 3: The conclusion from step 2 implies that  $\mathbb{E}_{m(\theta)}[d'(Y) - d(Y)] = 0$  for any  $\theta \in m^{-1}(\mathcal{S})$ . Consider now the family of distributions

$$Y \sim \mathcal{P}_{\mathcal{S}} = \{P_{m(\theta)} : \theta \in m^{-1}(\mathcal{S})\},$$

which is bounded complete under A1. Let  $E := \{y \in \mathcal{Y} : d'(y) - d(y) \neq 0\}$ . The definition of bounded completeness implies that  $P(E) = 0$  for all  $P \in \mathcal{P}_{\mathcal{S}}$ . It follows by A2 that  $P(E) = 0$  for all  $P \in \mathcal{P} \setminus \mathcal{P}_{\mathcal{S}}$ . Therefore,  $P(E) = 0$  for all  $P \in \mathcal{P}$ .

CONCLUSION: We find that  $\mathbb{E}_{m(\theta)}[d'(Y)] = \mathbb{E}_{m(\theta)}[d(Y)]$  for all  $\theta \in \Theta$ . This means that  $d'$  cannot dominate  $d$ . A contradiction.  $\square$

## B Profiled Risk

Given the findings in Section 3, in this section, we show that profiling out some parameters of expected welfare or regret may yield interesting insights. We illustrate this idea by exploring a *profiled expected regret* criterion. Beyond allowing to better visualize risk profiles of decision rules, there are several additional values the profiling approach can provide in our problem. First, as we have shown in Figure 2, two MMR optimal rules may display drastically different regret profiles. In fact, as we show in Proposition 2, one dominates the other in terms of profiled regret for many

parameter values. Therefore, the notion of profiled risk may help decision makers better distinguish two rules that cannot otherwise be ranked according to the MMR criterion alone. Second, as we will show below, the profiled approach may render some rules, e.g., the no-data rule  $d_{\text{coin-flip}}$ , inadmissible (in a redefined sense; see Proposition 3).<sup>23</sup> While we could profile out any known function  $h(m(\theta))$ , except for Proposition 3 we simplify exposition by restricting attention to linear indices. Thus, for a vector  $w \in \mathbb{R}^n \setminus \{\mathbf{0}\}$ , let

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

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

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

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

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

$$\bar{R}_w(d, \gamma) := \sup_{\theta \in \Theta : w^\top m(\theta) = \gamma} U(\theta) (\mathbf{1}\{U(\theta) \geq 0\} - \mathbb{E}_{m(\theta)}[d(Y)]), \quad (\text{B.1})$$

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. For example, when extrapolating local average treatment effects in Section 4.1, Theorem 3 does not apply and there are multiple reasonable alternatives for parameters to be profiled out. One may 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

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

---

<sup>23</sup>Although  $d_{\text{coin-flip}}$  may also be ruled out by the MMR criterion, being inadmissible is a stronger statement than not being MMR optimal. For example, it would be unclear whether  $d_{\text{coin-flip}}$  can be optimal for a decision maker with a robust Bayes criterion for a class of priors advocated by [Giacomini and Kitagawa \(2021\)](#). As a Bayes rule with respect to our profiled risk may be viewed as a robust Bayes criterion under some conditions,  $d_{\text{coin-flip}}$  being inadmissible under the profiled risk criterion implies that it also cannot be Robust Bayes optimal for a large class of priors in [Giacomini and Kitagawa \(2021\)](#)'s framework (see Remark 7 below for related discussions).

can be viewed as the population [Anderson and Rubin \(1949\)](#) statistic for the null hypothesis of the local average treat effect equal to  $\beta_0$ . Thus, the profiled regret function reports the worst-case regret as one keeps constant the population analog of that statistic. One may also profile out the intention-to-treat effect  $m_1(\theta)$ , the local average treatment effect  $m_1(\theta)/m_2(\theta)$ , or even the entire reduced-form parameter vector  $(m_1(\theta), m_2(\theta))^\top$ . In this paper, we do not argue for one specific criterion for profiling out, but believe researchers should have the freedom to choose the device judiciously depending on the specific contexts of their problems.

*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

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

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.<sup>24</sup> Moreover, the following proposition shows that, in the running example,  $d_{\text{linear}}^*$  indeed  $w^*$ -profiled regret dominates  $d_{\text{RT}}^*$  for a large range of parameter values, including those used in [Figure 2](#). Let  $\phi(\cdot)$  denote the p.d.f. of a standard normal random variable, and let  $0 < \rho^* \leq C\|x_1 - x_0\|$  be the threshold after which  $d_{\text{linear}}^*$  implements the policy with probability one. Let  $\gamma$  be the effect of the policy of interest in the nearest neighbor to Country 0 (which, by assumption, we have set to be Country 1). Let  $\bar{R}_{w^*}(d, \gamma)$  be the profiled-regret associated to the nearest-neighbor weights; that is, the largest value of expected regret for decision rule  $d$  for a fixed value of  $\gamma$ .

*Proposition 2.* Suppose  $C\|x_1 - x_0\| > \sqrt{\frac{\pi}{2}}\sigma_1$ .

- (i) If  $\phi\left(\frac{\rho^*}{\sigma_1}\right) \left(\frac{C\|x_1 - x_0\|}{\sqrt{\frac{\pi}{2}}\sigma_1}\right)^3 \leq \phi(0)$ , then  $\bar{R}_{w^*}(d_{\text{linear}}^*, \gamma) \leq \bar{R}_{w^*}(d_{\text{RT}}^*, \gamma)$  for all  $\gamma$  with the inequality strict for  $\gamma \neq 0$ .
- (ii) If  $\phi\left(\frac{\rho^*}{\sigma_1}\right) \left(\frac{C\|x_1 - x_0\|}{\sqrt{\frac{\pi}{2}}\sigma_1}\right)^3 > \phi(0)$ , then  $\bar{R}_{w^*}(d_{\text{linear}}^*, \gamma) < \bar{R}_{w^*}(d_{\text{RT}}^*, \gamma)$  for all  $|\gamma| > \underline{\gamma}$ , where  $\underline{\gamma} > 0$  is unique in  $(0, \infty)$  such that  $\Phi\left(-\frac{\underline{\gamma}}{C\|x_1 - x_0\|}\right) = \int_0^1 \Phi\left(\frac{2\rho^*x - \rho^* - \underline{\gamma}}{\sigma}\right) dx$ .

*Proof.* See [Online Appendix C.4](#). □

---

<sup>24</sup>Simple algebra shows that the profiled regret of  $d_{\text{coin-flip}}$  is *minimized* at  $\gamma = 0$ , where it coincides with the *maximized* regret of any MMR optimal rule. This can also be verified if  $C\|x_1 - x_0\| \leq \sqrt{\pi/2} \cdot \sigma_1$ . See [Figure 2](#).

Part i) of Proposition 2 gives conditions on the Lipschitz constant ( $C$ ), the variance of the estimated effect in Country 1 ( $\sigma_1^2$ ), and the difference between the covariates between Countries 1 and 0 ( $\|x_1 - x_0\|$ ) under which  $d_{\text{linear}}^*$   $w^*$ -profiled regret dominates  $d_{\text{RT}}^*$ . Part ii) of Proposition 2 shows that for those parameter configurations for which we cannot establish  $w^*$ -profiled regret dominance (which we later show can only happen when  $C \|x_1 - x_0\| \downarrow \sigma_1$ ), the decision rule  $d_{\text{linear}}^*$  still has smaller  $w^*$ -profiled regrets than  $d_{\text{RT}}^*$  for all  $|\gamma|$  large enough. Therefore, our new result shows there is a strong sense that  $d_{\text{linear}}^*$  performs better than  $d_{\text{RT}}^*$  under the  $w^*$ -profiled regret criterion, despite the fact that both of them are equivalent under the standard minimax regret criterion. Finally, we also argue that Proposition 2 will continue to hold (with suitably redefined parameter values) in a model as general as those studied in Theorem 3, as long as the upper and lower bounds of the identified set of  $U(\theta)$  given  $\gamma \in \Gamma_{w^*} = \{\gamma \in \mathbb{R} | (w^*)^\top m(\theta) = \gamma, \theta \in \Theta\}$  are affine with the same positive gradients. Under some conditions, a sufficient condition for this to happen is for  $\gamma$  to be unbounded.

Below, we also present a general result on the profiled-regret inadmissibility of  $d_{\text{coin-flip}}$  in a large class of general models when we profile out with respect to the reduced-form parameter  $m(\theta)$ .<sup>25</sup> Just as before, let  $\bar{I}(\mu)$  and  $\underline{I}(\mu)$  be the upper and lower bounds of the identified set for the welfare contrast given the reduced-form parameter  $\mu$ .

*Proposition 3.* *In the setup of Section 3.3, let the profiled regret (with respect to the reduced-form parameter  $\mu \in M$ ) of a rule  $d$  be*

$$\bar{R}_r(d, \mu) := \sup_{\theta \in \Theta, m(\theta) = \mu} R(d, \theta).$$

*Suppose all conditions of Theorem 3 hold true, and  $\bar{I}(\mathbf{0}) \leq \max \{\bar{I}(\mu), -\underline{I}(\mu)\}$  for all  $\mu \in M$  with the inequality strict for some  $\mu \in M$ . Then,  $d_{\text{coin-flip}}$  is dominated in terms of  $\bar{R}_r(d, \mu)$ .*

*Proof.* See Online Appendix C.5. □

One important remark about Proposition 3 is that, while it is possible to rule out  $d_{\text{coin-flip}}$  using the minimax regret criterion, such a rule could still be optimal under some other criterion such as Robust Bayes (for some class of priors). An implication of Proposition 3 is that  $d_{\text{coin-flip}}$  could never be Robust Bayes optimal for a class of priors where i) the prior over the reduced-form parameter is fixed, and ii) the priors of the partially identified parameters are arbitrary.

---

<sup>25</sup>We thank a referee for suggesting this result.

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

$$\inf_{d \in \mathcal{D}_n} \int_{\Gamma_w} \bar{R}_w(d, \gamma) d\pi^*(\gamma), \quad (\text{B.2})$$

the weighted average of  $w$ -profiled regret for some prior  $\pi^*$  over  $\Gamma_w$ . If  $\pi^*$  has full support and profiled regret is continuous in  $\gamma$  (for any  $d$ ), then any solution to (B.2) 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 (B.2) can be interpreted as  $\Pi^*$ -minimax decision rules in the sense of [Berger \(1985, Definition 13, p. 216\)](#), where  $\Pi^*$  is a class of priors over  $\Theta$ .<sup>26</sup> Let  $\mathcal{D}_w$  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  $\Pi^*$ -minimax problem

$$\inf_{d \in \mathcal{D}_w} \left( \sup_{\pi \in \Pi^*} \int_{\Theta} R(d, \theta) d\pi(\theta) \right). \quad (\text{B.3})$$

Let  $\Pi^*$  collect all priors over  $\Theta$  for which  $w^\top m(\theta) \sim \pi^*$ , where  $\pi^*$  is the prior over  $\Gamma_w$  used in (B.2). This class of priors has recently been advocated by [Giacomini and Kitagawa \(2021\)](#). Their Theorem 2 establishes that, for any  $d \in \mathcal{D}_w$ ,

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

Thus, the problem in (B.3) is equivalent to minimizing average  $w$ -profiled regret over decision rules that depend on the data only through  $w^\top Y$ .

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

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

and

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

In principle, these are the value functions of two infinite-dimensional, nonlinear optimization

---

<sup>26</sup>To avoid notational conflict, we write  $\Pi^*$ -minimax instead of the more common  $\Gamma$ -minimax.

problems. However, they can be recast as finite dimensional. For example,  $\bar{k}_w^+(\gamma)$  equals

$$\bar{I}_w^+(\gamma) := \sup_{\mu \in M} \bar{I}(\mu)(1 - \mathbb{E}_\mu[d(Y)]) \quad \text{s.t. } w^\top \mu = \gamma, \quad \bar{I}(\mu) \geq 0. \quad (\text{B.4})$$

This problem has a scalar choice variable, one linear equality constraint, and one potentially nonlinear inequality constraint. The bottleneck is evaluation of  $\bar{I}(\mu)$ . The running example admits a closed-form solution for  $\bar{I}(\mu)$ , so that evaluating (B.4) is easy. More generally, the computational cost of evaluating (B.4) can be reduced by imposing more structure on the parameter space  $\Theta$ . For instance, when  $\Theta$  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  $\bar{I}(\mu)$  can be shown to be concave. This means that under Assumption 1, the optimization problem in (B.4) is convex.

## Data Availability Statement

The code underlying this research is available on Zenodo at <https://dx.doi.org/10.5281/zenodo.17711688>.

## 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. BLANCHET, J. L. MONTIEL OLEA, C. QIU, J. STOYE, AND L. TAN (2025): “ $\epsilon$ -Minimax Solutions of Statistical Decision Problems,” *arXiv preprint arXiv:2509.08107*.

ARADILLAS FERNÁNDEZ, A., J. L. MONTIEL OLEA, C. QIU, J. STOYE, AND S. TINDA (2024): “Robust Bayes Treatment Choice with Partial Identification,” *arXiv preprint arXiv:2408.11621*.

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 (2025): “Safe policy learning through

extrapolation: Application to pre-trial risk assessment,” *Journal of the American Statistical Association*, in press.

BEN-MICHAEL, E., K. IMAI, AND Z. JIANG (2024): “Policy learning with asymmetric utilities,” *Journal of the American Statistical Association*, 119, 3045–3058.

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 (2026): “A note on minimax regret rules with multiple treatments in finite samples,” *Econometric Theory*, forthcoming.

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): “Orthogonal Policy Learning Under Ambiguity,” *arXiv preprint arXiv:2111.10904*.

DASKALAKIS, C., S. SKOULAKIS, AND M. ZAMPETAKIS (2021): “The complexity of constrained min-max optimization,” in *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, 1466–1478.

DEHEJIA, R. (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*.

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.

GUGGENBERGER, P. AND J. HUANG (2025): “On the numerical approximation of minimax regret rules via fictitious play,” *arXiv preprint arXiv:2503.10932*.

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.

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.

IMBENS, G. W. AND D. B. RUBIN (2015): *Causal inference in statistics, social, and biomedical sciences*, Cambridge university press.

ISHIHARA, T. AND T. KITAGAWA (2021): “Evidence Aggregation for Treatment Choice,” *arXiv preprint 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., S. LEE, AND C. QIU (2022): “Treatment Choice with Nonlinear Regret,” *arXiv preprint arXiv:2205.08586*.

KITAGAWA, T., S. SAKAGUCHI, AND A. TETENOV (2021): “Constrained classification and policy learning,” *arXiv preprint arXiv:2106.12886*.

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.

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*.

MANSKI, C. F. (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*, 31, 345–352.

——— (2021): “Econometrics for decision making: Building foundations sketched by Haavelmo and Wald,” *Econometrica*, 89, 2827–2853.

——— (2024): “Identification and Statistical Decision Theory,” *Econometric Theory*, 41, 1–17.

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.

MATTNER, L. (1993): “Some incomplete but boundedly complete location families,” *The Annals of Statistics*, 21, 2158–2162.

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/EC02006-2.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 preprint*

*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.