Decision Theory for Treatment Choice Problems with Partial Identification*

José Luis Montiel Olea† Chen Qiu‡ Jörg Stoye§

July 2024

Abstract

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

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

†Email: jlo67@cornell.edu

‡Email: cq62@cornell.edu

§Email: stoye@cornell.edu

1 Introduction

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

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

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

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

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

Minimax Regret. A decision rule is minimax-regret (MMR) optimal if it attains the lowest worst-case expected regret, where an action's regret is its welfare loss relative to the action that would be optimal if payoff-relevant parameters were known. In some point-identified treatment choice problems, the MMR rule is both essentially unique and nonrandomized (Canner, 1970; Stoye, 2009a; Tetenov, 2012). This may not not be true under partial identification. To make this point, we specialize our framework to the class of treatment choice problems recently studied by Yata (2021). For cases where the identified set for payoff relevant parameters is large enough, we discover ***infinitely many*** MMR optimal decision rules, ***all of which*** randomize the policy recommendation for some or all data realizations. This presents a challenge for the application of MMR, at least if one hopes for the resulting recommendation to be unique. Moreover, as we explain later by means of an example, different MMR optimal rules can lead to meaningfully different policy choices for the same data.

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

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

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

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

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

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

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

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

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

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

2 Notation and Framework

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

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

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

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

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

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

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

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

contrast is nonnegative and to preserve the status quo otherwise.

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

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

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

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

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

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

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

where $\theta : \mathbb{R}^d \to \mathbb{R}$ is an unknown Lipschitz function with known constant C. For simplicity, in the example we write μ_i for $\theta(x_i)$ henceforth, thus $Y_i \sim N(\mu_i, \sigma_i^2)$. We also assume w.l.o.g. that countries are arranged in nondecreasing order of $\|x_i - x_0\|$. While our analysis extends to $x_1 = x_0$, we focus on the case of $x_1 \neq x_0$, so that the sign of μ_0 is not necessarily identified. Finally, we assume that $(x_1, \ldots, x_n)$ are distinct. Even if this were not the case in raw data, one would presumably want to induce it (by adding fixed effects, whose size can be bounded) because the Lipschitz constraint would otherwise imply that countries with same x exhibit no heterogeneity whatsoever. □

3 Challenges

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

3.1 Everything is Admissible

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

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

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

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

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

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

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

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

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

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

and the corresponding upper and lower bounds as

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

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

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

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

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

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

Its extrema can be written as intersection bounds:

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

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

□

We are now ready to state the first main result.

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

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

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

is (welfare-)admissible.

Proof. See Appendix A.1. □

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

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

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

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

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

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

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

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

3.2 Maximin Welfare is Ultra-Pessimistic

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

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

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

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

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

then the no-data rule $d_{\textit{no-data}}(y) := 0$ 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 $U(\theta) \le 0$, it is optimal to preserve the status quo; substituting in for this response, we find that $\inf_{\theta\in\Theta} \mathbb{E}_{m(\theta)}[W(d(Y), \theta)] \le \inf_{\theta\in\Theta,U(\theta)\le 0} W(0, \theta)$ for any rule $d \in \mathcal{D}_n$. Under condition (3.4), this upper bound is attained by $d_{\mathsf{no\text{-}data}}$.

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

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

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

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

3.3 Minimax Regret Admits Many Solutions

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

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

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

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

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

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

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

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

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

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

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

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

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

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

and

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

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

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

(i) There are infinitely many MMR optimal rules.

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

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

Proof. See Appendix A.3. □

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

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

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

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

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

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

is not MMR optimal.

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

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

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

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

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

(i) If

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

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

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

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

(ii) If

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

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

(iii) If

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

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

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

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

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

Proof. See Online Appendix B.2. □

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

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

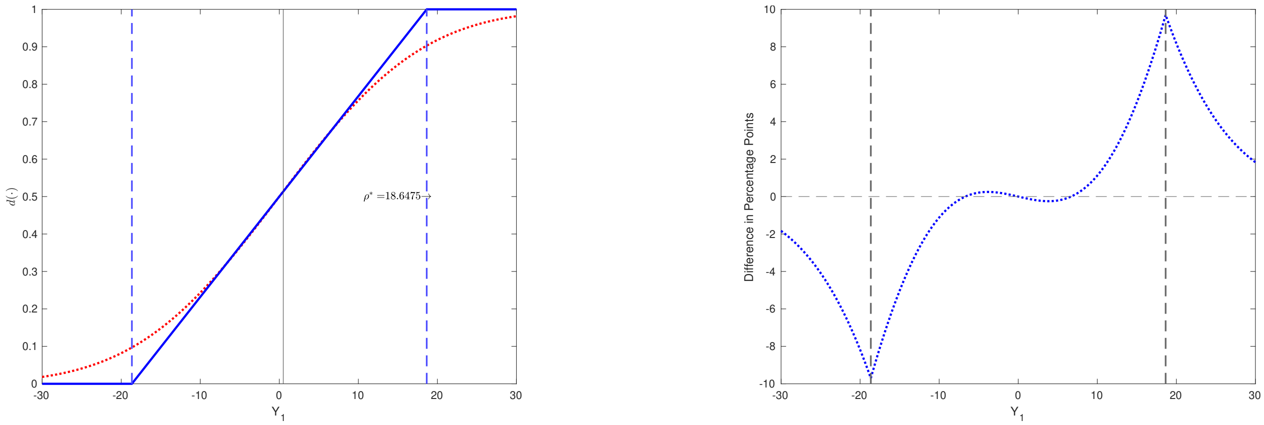

(a) d^*_{RT} (red) and d^*_{linear} (blue)

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

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

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

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

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

4 Opportunities

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

4.1 Profiled Risk

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

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

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

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

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

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

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

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

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

Remark 6 (Not all decision rules are admissible with respect to w-profiled regret)*.* One could say that a decision rule d is w-***profiled regret admissible*** if there is no other rule d' for which

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

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

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

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

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

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

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

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

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

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

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

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

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

expected loss). Consider the loss function

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

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

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

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

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

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

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

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

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

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

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

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

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

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

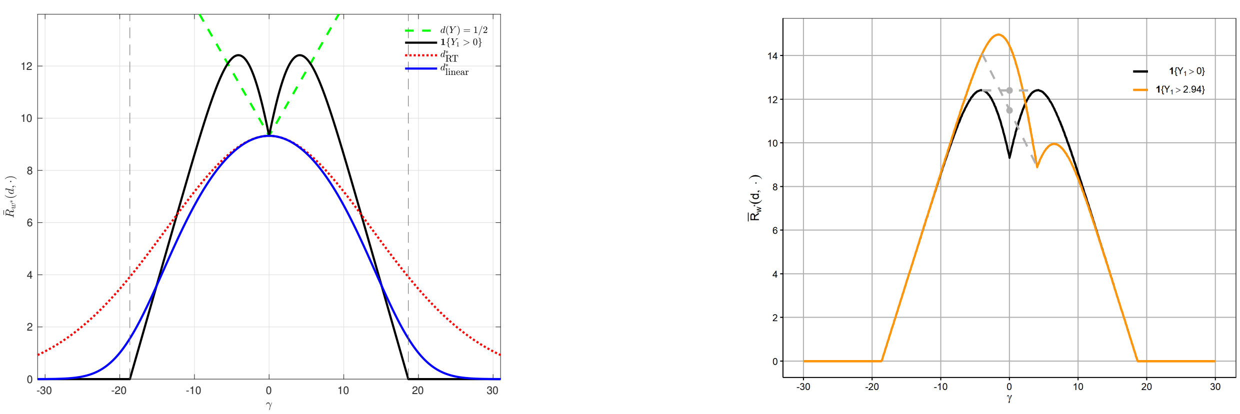

(a) w^*-profiled regret of four decision rules

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

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

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

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

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

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

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

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

and

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

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

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

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

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

4.2 Least Randomizing MMR Optimal Rules

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

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

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

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

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

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

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

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

Proof. See Appendix A.4. □

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

To establish Theorem 4, we first show that for any rule $F \circ w^* \in \tilde{\mathcal{D}}_n$, its expected regret at any parameter θ for which $(w^*)^\top m(\theta) = 0$ equals the MMR value of the problem. If $F \circ w^* \in \tilde{\mathcal{D}}_n$ is MMR optimal, its expected regret must therefore be maximized at $(w^*)^\top m(\theta) = 0$. For any symmetric and unimodal c.d.f. F with $V(d^*_{\mathsf{linear}} \circ w^*) \nsubseteq V(F \circ w^*)$, we can show that a necessary condition for this maximization fails.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

5 Further Applications

5.1 Extrapolating Local Average Treatment Effects

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

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

The policy maker observes

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

where

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

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

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

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

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

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

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

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

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

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

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

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

5.2 Decision-theoretic Breakdown Analysis

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

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

where Y is observed outcome, D is the binary treatment, and e is a projection residual. The infeasible optimal treatment policy is $\mathbf{1}\{\beta_{\mathsf{long}} \geq 0\}$. However, W is unobserved, so that the policy maker can only estimate the "medium" regression

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

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

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

where

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

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

Diegert et al. (2022) use this result to ask: How strong does omitted variables bias have to be to potentially overturn findings based on β_{med}? At population level, the answer is that this can happen if $|k| > |\beta_{\mathsf{med}}|$, 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 $\hat{\beta}_{\mathsf{med}} \sim N(\beta_{\mathsf{med}}, \sigma^2)$. Then our results apply upon letting $\theta = (\beta_{\mathsf{long}}, \beta_{\mathsf{med}})^\top \in \mathbb{R}^2$, $U(\theta) = \beta_{\mathsf{long}}$ and $m(\theta) = \beta_{\mathsf{med}}$. 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

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

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

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

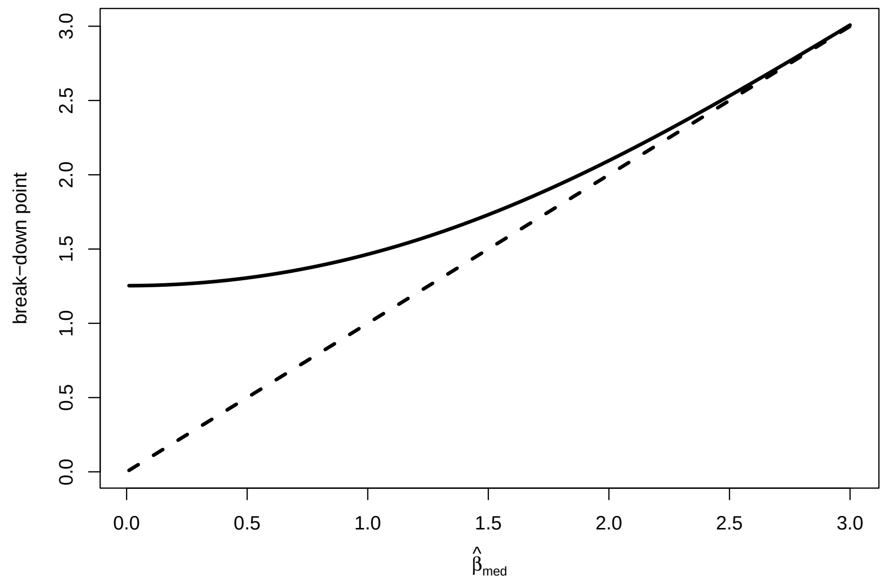

Figure 3: $\bar{k}$ (solid) and $\tilde{k}$ (dashed) as a function of $\hat{\beta}_{\mathsf{med}}$

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

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

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

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

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

6 Conclusion

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

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

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

References

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Yu, G. and P. Kouvelis (1995): "On min-max optimization of a collection of classical discrete optimization problems," in *Minimax and applications*, Springer, 157–171.

A Proofs of Main Results

A.1 Proof of Theorem 1

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

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

for all $\theta \in \Theta$.

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

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

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

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

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

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

contradicting (A.1). We conclude that

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

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

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

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

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

Step 3: Consider now the family of distributions

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

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

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

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

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

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

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

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

A.2 Proof of Theorem 2

First, we can bound

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

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

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

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

A.3 Proof of Theorem 3

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

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

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

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

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

where

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

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

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

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

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

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

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

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

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

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

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

$$\begin{aligned}
&\sup_{\theta\in\Theta} \left(U(\theta)\left(\mathbf{1}\{U(\theta)\geq 0\} - \mathbb{E}_\gamma[d((w^*)^\top Y)]\right)\right) && \text{(A.8)}\\
&= \sup_{\gamma\in\Gamma_{w^*}} \left(\sup_{\theta\in\Theta,\,(w^*)^\top m(\theta)=\gamma} U(\theta)\left(\mathbf{1}\{U(\theta)\geq 0\} - \mathbb{E}_\gamma[d((w^*)^\top Y)]\right)\right)\\
&= \sup_{\gamma\in\Gamma_{w^*}} \left(\sup_{U^*\in ISU_{w^*}(\gamma)} U^*\left(\mathbf{1}\{U^*\geq 0\} - \mathbb{E}_\gamma[d((w^*)^\top Y)]\right)\right), && \text{(A.9)}
\end{aligned}$$

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

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

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

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

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

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

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

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

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

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

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

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

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

It follows that

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

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

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

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

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

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

where

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

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

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

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

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

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

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

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

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

This is because Step 1 implies that

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

and Step 2 implies that

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

where the last equality follows from Step 3. Consequently,

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

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

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

and

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

$$\sup_{\theta\in\Theta} U(\theta)\left(\mathbf{1}\{U(\theta)\geq 0\} - \mathbb{E}_{m(\theta)}[d_m(Y)]\right) \tag{A.22}$$

$$\geq \sup_{\theta\in\Theta:m(\theta)=\mathbf{0}} U(\theta)\left(\mathbf{1}\{U(\theta)\geq 0\} - \mathbb{E}_{m(\theta)}[d_m(Y)]\right)$$

$$= \max\left\{-\underline{I}(\mathbf{0})\mathbb{E}_0[d_m(Y)], \overline{I}(\mathbf{0})(1-\mathbb{E}_0[d_m(Y)])\right\}$$

$$\geq \overline{I}(\mathbf{0})/2, \tag{A.23}$$

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

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

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

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

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

$$= \sup_{\mu\in M, \overline{I}(\mu)>0} \overline{I}(\mu)\left(1 - \mathbb{E}_{\mu}[\mathbf{1}\{w^\top Y\geq 0\}]\right)$$

$$= \sup_{\mu\in M, \overline{I}(\mu)>0} \overline{I}(\mu)\Phi\left(-\frac{w^\top\mu}{\sqrt{w^\top\Sigma w}}\right)$$

$$:= \sup_{\mu\in M, \overline{I}(\mu)>0} \mathrm{g}_w(\mu).$$

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

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

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

A.4 Proof of Theorem 4

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

B Online Appendix

B.1 Lemmas for Theorem 3

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

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

where

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

□

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

where

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

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

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

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

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

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

and this expression equals

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

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

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

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

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

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

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

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

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

Furthermore, by definition of the minimax problem,

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

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

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

and

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

Hence, we conclude that

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

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

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

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

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

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

Proof. Note

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

is equivalent to

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

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

To see (i), further write

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

Note $\Phi(x) < \frac{1}{2}$ and $\phi'(x) > 0$ for all $x < 0$. Thus, $\Phi\left(-\frac{\rho}{\sigma}\right) - \phi'\left(-\frac{\rho}{\sigma}\right) < \frac{1}{2}$ for all $\rho > 0$. It follows that $\frac{\partial}{\partial \rho}\mathbf{g} > 0$, i.e. $\mathbf{g}$ is strictly increasing in ρ for all $\rho > 0$. Furthermore, note

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

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

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

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

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

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

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

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

and

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

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

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

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

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

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

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

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

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

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

where the second equality for $g^{(2)}_{\mathsf{linear}}(\gamma)$ uses that $\phi'(x) = -x\phi(x)$ for all $x \in \mathbb{R}$, the third one follows from integration by change-of-variable, and the fourth equality follows from change-of-variable again and $\phi(x) = \phi(-x)$ for all $x \in \mathbb{R}$. As $A > 0$, the sign of $g^{(2)}_{\mathsf{linear}}(\gamma)$ is determined by

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

Furthermore, we can write

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

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

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

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

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

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

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

is not MMR optimal.

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

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

Write

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

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

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

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

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

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

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

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

B.2 Proof of Proposition 1

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

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

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

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

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

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

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

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

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

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

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

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

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

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

where we used centrosymmetry of Θ and where

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

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

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

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

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

$$\sup_{\mu \in M, \mu_0 \in I(\mu), \mu_0 \geq 0} \mu_0 \Phi\left(-\frac{w_{m_0}^\top \mu}{\sqrt{w_{m_0}^\top \Sigma w_{m_0}}}\right) = \sup_{\mu_0 \geq 0} g(\mu_0, m_0),$$

$$g(\mu_0, m_0) = \mu_0 \Phi\left(-\frac{\sum_{j=1}^n \frac{\max\{m_0 - C\|x_j - x_0\|, 0\}(\mu_0 - C\|x_j - x_0\|)}{\sigma_j^2}}{\sqrt{\sum_{j=1}^n \frac{\max^2\{m_0 - C\|x_j - x_0\|, 0\}}{\sigma_j^2}}}\right).$$

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

B.3 Additional Results

B.3.1 Profiled Regret in the Running Example

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

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

Specifically, we look at three rules:

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

defined in (3.7)-(3.11).

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

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

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

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

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

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

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

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

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

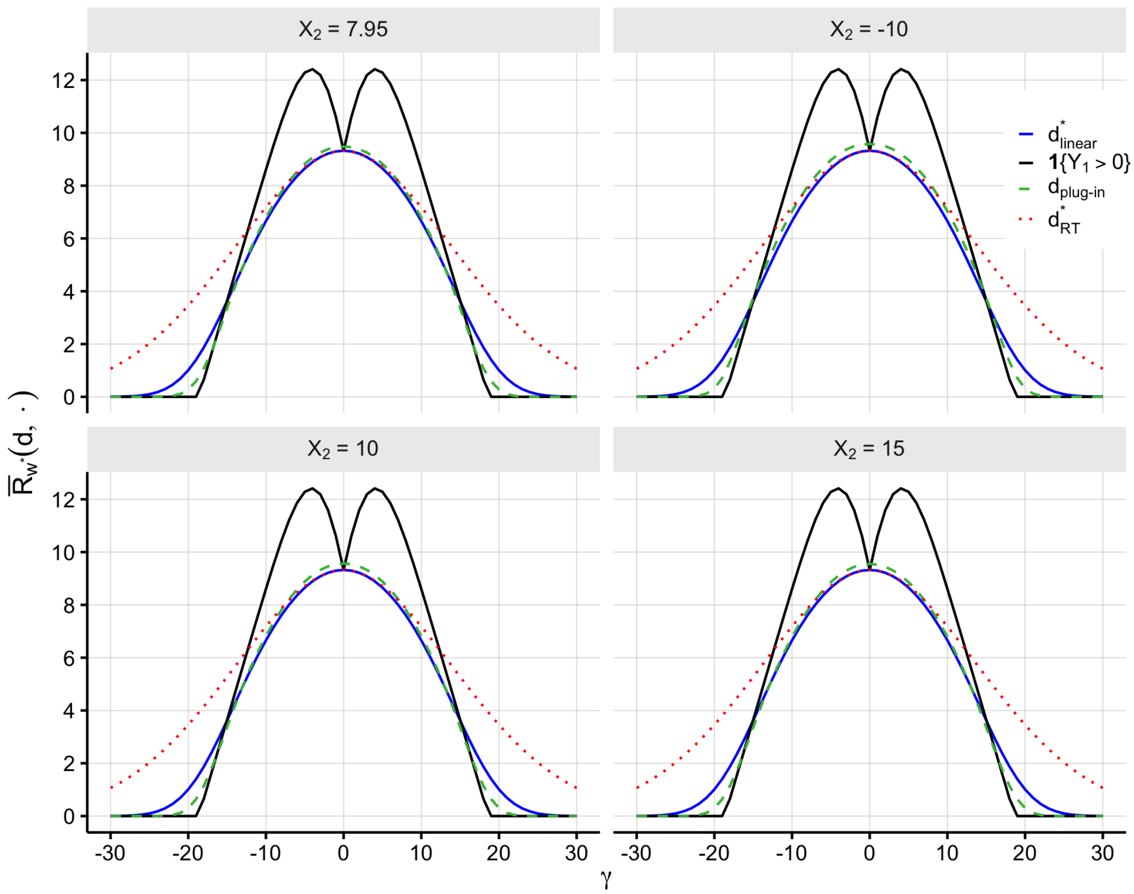

Figure 4: w^*-profiled regrets of $d_{\text{plug-in}}$ in (4.14) and other rules. The top left plot uses parameters of the running examples from Figure 1; the rest of the plots change the value of x_2 while keeping other parameters the same.

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

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

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

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

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

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

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

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

with extrema

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

and similarly

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

Let

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

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

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

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