Econometric Theory, **0**, 2026, 1–37.
doi:10.1017/S0266466626100486

ROBUST BAYES TREATMENT CHOICE WITH PARTIAL IDENTIFICATION

ANDRÉS ARADILLAS FERNÁNDEZ
Massachusetts Institute of Technology

JOSÉ LUIS MONTIEL OLEA
Cornell University

CHEN QIU
Cornell University

JÖRG STOYE
Cornell University

SERDIL TINDA
Stanford University

We study a class of binary treatment choice problems with partial identification through the lens of robust (multiple prior) Bayesian analysis. We use a convenient set of prior distributions to derive ex-ante and ex-post robust Bayes decision rules, both for decision makers who can randomize and for decision makers who cannot. Our main messages are as follows: First, ex-ante and ex-post robust Bayes decision rules do not agree in general, whether or not randomized rules are allowed. Second, randomized treatment assignment for some data realizations can be optimal in both ex-ante and, perhaps more surprisingly, ex-post problems. Therefore, it is usually with loss of generality to exclude randomized rules from consideration, even when regret is evaluated ex post. We apply our results to a stylized problem where a policy maker uses experimental data to choose whether to implement a new policy in a population of interest, but is concerned about the external validity of the experiment at hand, and to aggregation of data generated by multiple randomized control trials in different sites to make a policy choice in a population for which no experimental data are available.

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

We would like to thank the co-editor and three anonymous referees, Karun Adusumilli, Tim Armstrong, Tim Christensen, Patrik Guggenberger, Kei Hirano, Toru Kitagawa, and seminar audiences at Penn State for helpful feedback, comments, and suggestions. We gratefully acknowledge financial support from NSF grant SES-2315600. Address correspondence to Jörg Stoye, Cornell University, USA, email: stoye@cornell.edu.

© The Author(s), 2026. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

of action. Such *treatment choice problems with partial identification* have recently received growing interest (e.g., see D'Adamo, 2021, Ishihara and Kitagawa, 2021, Yata, 2021, Christensen, Moon, and Schorfheide, 2026, Kido, 2022, or Manski, 2025). Several interesting problems that arise in empirical research can be recast using this framework. See Montiel Olea, Qiu, and Stoye (2026) and the end of this section for references.

This article applies the *robust Bayes* approach—which interpolates between Bayesian and agnostic minimax analyses by evaluating minimax risk over a set of priors—to a class of treatment choice problems with partial identification. The use of the robust Bayes approach has drawn recent attention in problems that feature partial identification (Giacomini and Kitagawa, 2021; Giacomini, Kitagawa, and Read, 2021; Christensen et al., 2026). Indeed, due to partial identification, integrating Bayesian and minimax elements into decision making can be particularly attractive (see Poirier, 1998; Moon and Schorfheide, 2012, and the references therein).

The robust Bayes approach can be applied *ex-ante* or *ex-post*, depending on whether the multiple priors are used to evaluate payoffs before or after seeing the data.[1] These concepts represent two different ways of resolving model ambiguity and sampling uncertainty, and both have been proposed to improve Bayesian robustness in decision problems.

By the well-known *dynamic consistency* of Bayesian decision making, ex-ante and ex-post robust Bayes coincide with each other—and with standard Bayes optimality—if the set of priors is a singleton. Indeed, this equivalence is used to calculate Bayes optimal decisions in practice because ex-post rules are usually computed but ex-ante Bayes optimality is claimed. As pointed out for the present context by Giacomini et al. (2021), they do not in general agree otherwise.[2] However, to what extent this inequivalence affects treatment choice problem with partial identification is far from clear. The class of priors that we consider has a Cartesian product structure resembling *rectangularity*, a condition under which maximin welfare loss is known to be dynamically consistent.[3] While this result does not apply here—the priors are not rectangular in the strict technical sense, and the theoretical results were not established for regret loss—one might wonder if the conclusion holds anyway or else, what qualitative and quantitative relationships exist between ex-ante and ex-post robust Bayes in the presence of partial identification. Using a convenient class of priors allows us to exhaustively answer these questions in a case that we believe holds some interest.

[1] Giacomini et al. (2021) discuss both notions; they refer to the ex-ante and ex-post problems as "Gamma-minimax" and "Conditional Gamma-minimax," respectively. Christensen et al. (2026) focus on the ex-post problem for treatment choice problems with partial identification in a restricted class of decision rules.

[2] For an estimation problem with a quadratic loss, Kitagawa (2012, Appendix B) derives the ex-post Γ-minimax estimator and shows that it is not ex-ante Γ-minimax optimal.

[3] See Epstein and Schneider (2003) and also Wakai (2007), Amarante and Siniscalchi (2019), and the references therein.

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

To this end, we use the same framework as Yata (2021) and Montiel Olea et al. (2026) but impose a simple instance of Giacomini and Kitagawa's (2021) set of priors, namely, a symmetric and uniform two-point prior for reduced-form parameters and no restriction at all on unidentified parameters given reduced-form parameters. Working with regret, we formally define ex-ante and ex-post robust Bayes and, following Berger (1985), label them as "Γ-minimax regret" (Γ-MMR) and "Γ-posterior expected regret" (Γ-PER).[4] We then precisely characterize when these notions coincide and when they disagree. The main qualitative insights are as follows:

- Ex-ante Γ-MMR and ex-post Γ-PER criteria do frequently disagree. This conclusion does not hinge on whether or not randomization rules are allowed. If randomization is permitted, the criteria differ whenever updating the prior for the reduced-form parameter does not (ex-ante almost surely) resolve ambiguity regarding the sign of the optimal action. Even if randomization is not allowed, the criteria will disagree whenever the *identification power* of the model (in a sense we make precise) is sufficiently small compared to the *informativeness of the data*.
- Randomized rules are general solution concepts in both Γ-MMR and Γ-PER problems. That is, optimal rules often can or even must randomize even if regret is evaluated ex-post. For the ex-ante Γ-MMR problem, whenever the identification power of the model is sufficiently small compared to the informativeness of the data, (1) infinitely many optimal rules exist and (2) many of them are randomized. For the ex-post Γ-PER criterion, one might conjecture by analogy to single-prior Bayes inference that the optimal rules do not randomize. However, as long as (1) there exists ambiguity regarding the sign of the optimal action at the location of the two-point prior for the reduced-form parameter and (2) randomized rules are allowed, the optimal rule is unique and is randomized. Hence, exclusion of randomized rules is with loss of generality and cannot be justified by simply evaluating regret ex-post—it must come from other considerations, be they second-order preferences or logistical or convenience concerns.

Our point is not to advocate for either notion of robust Bayes criterion. We also do not aim to solve for robust Bayes criteria for more general sets of priors as this would get much more involved but (we suspect) not much more instructive to illustrate the points discussed above. What we hope to illustrate is when, and how, Γ-MMR and Γ-PER criteria differ. We also relate these results to timing assumptions in a fictitious game between the policy maker and an adversarial Nature.

[4] Among others, see Savage (1951), Manski (2004), Stoye (2012b), and Montiel Olea et al. (2026) for justifications of focusing on regret in treatment choice problems. In particular, while minimax loss can be an attractive alternative to minimax regret, it leads to trivial recommendations in treatment choice settings including our examples.

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

Several auxiliary findings might be of independent interest. First, for Γ-MMR, even if we restrict the set of decision rules to be a class of non-randomized threshold rules based on the "efficient" linear index, the optimal threshold is not always zero. Given the apparent symmetry of the problem, we find this feature rather surprising. Second, whenever the dimension of the signal is larger than 1, there always exist (regardless of the parameter space and the variance of the signals) non-randomized linear-index threshold rules (with a threshold equal to zero) that are Γ-MMR optimal (among all decision rules). This is in stark contrast to Montiel Olea et al. (2026), in which no linear index rule is globally minimax regret optimal if the degree of partial identification is severe. The intuition is that the prior much reduces the state space; the signal space then becomes so rich relative to the state space that even linear threshold rules can effectively mimic randomization.

The literature on treatment choice with partially identified parameters has been growing since Manski (2004) and Dehejia (2005). For partial identification with known distribution of data, Manski (2000, 2005, 2007a) and Stoye (2007) find minimax regret optimal treatment rules. For finite-sample minimax regret results with model ambiguity and sampling uncertainty, see Stoye (2012a, 2012b), Yata (2021), Ishihara and Kitagawa (2021), and Montiel Olea et al. (2026); Kido's (2023) analysis is asymptotic. Bayes and robust Bayes approaches are analyzed by Chamberlain (2011), Giacomini and Kitagawa (2021), Giacomini et al. (2021), and Christensen et al. (2026), among others. Earlier investigations of ex-ante and ex-post Γ-minimax estimators include DasGupta and Studden (1989) and Betrò and Ruggeri (1992). See Vidakovic (2000) for a review. For different settings with point-identified welfare, finite- and large-sample results on optimal treatment choice rules were derived by Canner (1970), Chen and Guggenberger (2025), Hirano and Porter (2009, 2020), Kitagawa, Lee, and Qiu (2026), Schlag (2006), Stoye (2009), and Tetenov (2012b). There is also a large literature on optimal policy learning with covariates containing results with point identified (Bhattacharya and Dupas, 2012; Kitagawa and Tetenov, 2018, 2021; Athey and Wager, 2021; Kitagawa et al., 2021; Mbakop and Tabord-Meehan, 2021; Kitagawa and Wang, 2023; Ida et al., 2026) as well as partially identified (Kallus and Zhou, 2018; Ben-Michael et al., 2025; D'Adamo, 2021; Adjaho and Christensen, 2022; Ben-Michael, Imai, and Jiang, 2022; Christensen et al., 2026; Kido, 2022; Lei, Sahoo, and Wager, 2023) parameters. Guggenberger, Mehta, and Pavlov (2024) and Manski and Tetenov (2023) analyze related problems but focus on quantile, as opposed to expected, loss; Song (2014) considers partial identification but mean squared error regret loss.

The rest of this article is organized as follows. Section 2 sets up the problem, provides examples, and defines both versions of robust Bayes optimality. Section 3 contains complete solutions for all aforementioned scenarios and relates them to timing assumptions in the "Games against Nature" interpretation of minimax theory. Section 4 concludes. Proofs and auxiliary results are collected in the Appendix.

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

2. FRAMEWORK

2.1. Actions, Payoffs, Statistical Model, and Decisions

Our setup follows Montiel Olea et al. (2026), who in turn follow Ferguson (1967) and others. Consider a policy maker who needs to choose an *action* $a \in [0,1]$ interpreted as probability of assigning treatment in the target population.[5] Her payoff when taking action $a \in [0,1]$ is captured by the welfare function

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

where $\theta \in \Theta$ is an unknown state of the world or *parameter* and the functions $W(1,\cdot) : \Theta \to \mathbb{R}$ and $W(0,\cdot) : \Theta \to \mathbb{R}$ are known. Here, we may interpret $W(1,\theta)$ and $W(0,\theta)$ as the welfare of actions $a = 1$ (treating everyone in the population) and action $a = 0$ (treating no one in the population). Therefore, (2.1) implies that welfare is linear in actions, a standard assumption in the literature. Denote by $U(\theta) := W(1,\theta) - W(0,\theta)$ the *welfare contrast* at θ. If $U(\theta)$ were known to the policy maker, her optimal action would simply be

$$\mathbf{1}\{U(\theta) \geq 0\}. \tag{2.2}$$

The policy maker does not know $U(\theta)$ but can learn about θ. Specifically, we assume that she observes a random vector $Y \in \mathbb{R}^n$ with multivariate normal distribution

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

where the function $m(\cdot) : \Theta \to \mathbb{R}^n$ and the positive-definite matrix Σ are known.

Our focus is on the case when the data are not entirely informative about the sign of $U(\theta)$: Even if the policy maker perfectly learned $m(\theta)$, she could not (necessarily) pin down the sign of $U(\theta)$. To formally model such *treatment choice problems with (decision-relevant) partial identification*, let

$$M := \{\mu \in \mathbb{R}^n : m(\theta) = \mu, \theta \in \Theta\} \tag{2.4}$$

collect all the means of Y that can be generated as θ ranges over Θ. We refer to elements $\mu \in M$ as *reduced-form* parameters because they are identified in the statistical model (2.3) without further assumptions. Define the *identified set* for the welfare contrast given μ as

$$I(\mu) := \{u \in \mathbb{R} : U(\theta) = u, m(\theta) = \mu, \theta \in \Theta\} \tag{2.5}$$

and the corresponding upper and lower bounds as

$$\overline{I}(\mu) := \sup I(\mu), \quad \underline{I}(\mu) := \inf I(\mu). \tag{2.6}$$

[5] Randomization could be i.i.d. across future potential treatment recipients, *fractional* in the sense of randomly assigning a certain fraction of the treatment population (in this sense, $a \in [0,1]$ can also be interpreted as the fraction of the population receiving the treatment), or an "all or nothing" randomization for the entire treatment population. While these might not be practically equivalent in all applications, they are in the current decision theoretic framework. See Manski and Tetenov (2007) for an exception in the related literature.

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

Henceforth, when we refer to a treatment choice problem with partial identification, we mean that there exists some nonempty open set in M such that for all μ in that open set, $\underline{I}(\mu) < 0 < \bar{I}(\mu)$. For simplicity, we also assume that the infimum and supremum in (2.5) are attained.

A *decision rule* $d : \mathbb{R}^n \to [0,1]$ is a (measurable) mapping from data Y to the unit interval $[0,1]$. We call d *non-randomized* if it (almost surely, a.s.) maps into $\{0,1\}$; otherwise, we call d *randomized*, including if it randomizes for some but not all data realizations. We use $\mathcal{D}_n$ to denote the set of all decision rules, and we consider decision rules the same if they a.s. agree. As a result, a rule is *unique* only up to a.s. agreement. The *oracle policy* $\mathbf{1}\{U(\theta) \geq 0\}$ is of special interest and for any given θ is contained in $\mathcal{D}_n$, but is not feasible in the statistical sense because $U(\theta)$ is not known. In general, there will not be an unambiguously best feasible decision rule, a problem that gave rise to a large literature on different optimality criteria and their implementation. Before introducing the robust Bayes approach, we give two examples that fit into our general framework.

2.2. Examples

Example 1 (Stoye, 2012a). The one-dimensional version of the general setup has been frequently analyzed before (Manski, 2000; Brock, 2006; Stoye, 2012a; Tetenov, 2012a; Kitagawa, Lee, and Qiu, 2023). One motivation for it is to think of a policy maker who uses experimental data to choose whether to implement a new policy in a population of interest, but is concerned about the external validity of the experiment at hand. The treatment effect of action $a = 1$ is $\mu^* \in \mathbb{R}$, while the effect of action $a = 0$ is normalized to 0; thus, the policy maker's expected payoff equals $W(a, \mu^*) := a \cdot \mu^*$. The policy maker observes a realization of the one-dimensional statistic

$$\hat{\mu} \sim N(\mu, \sigma^2), \tag{2.7}$$

where $\sigma > 0$ is known and $\mu \in \mathbb{R}$ is an *identifiable treatment effect*, that is, it could be perfectly learned from infinite data. In this example, $\theta = (\mu, \mu^*)^\top$, $\Theta \subseteq \mathbb{R}^2$, $m(\theta) = \mu$, and $U(\theta) = \mu^*$.

Since the target population and the population from which data (2.7) are collected can be different, partial identification naturally arises. We assume that the identifiable and true treatment effects are constrained by $|\mu^* - \mu| \leq k$ for some known $k \geq 0$, implying

$$I(\mu) = [\mu - k, \mu + k], \quad \bar{I}(\mu) = \mu + k, \quad \underline{I}(\mu) = \mu - k, \quad \forall \mu \in \mathbb{R}.$$

The planner must choose a statistical decision rule $d \in \mathcal{D}_1 : \mathbb{R} \to [0,1]$ that maps observed data $\hat{\mu}$ to an action $a \in [0,1]$. Taking (2.7) as an approximation, this stylized example could reflect model uncertainty (e.g., a treatment effect is estimated in a possibly somewhat misspecified model), external validity concerns (e.g., a randomized clinical trial was performed on volunteers), or a shift in the environment (e.g., we transfer estimates from study populations to treatment

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

populations with slightly different covariates or are concerned about distributional drift over time).

Example 2 (Ishihara and Kitagawa, 2021). This example is taken from Ishihara and Kitagawa's (2021; see also Manski, 2020) "evidence aggregation" framework. A policy maker is interested in implementing a new policy in country $i = 0$ and observes estimates of the policy's effect for countries $i = 1, \ldots, n$. Let $Y = (Y_1, \ldots, Y_n)^\top \in \mathbb{R}^n$ denote these estimates and let $(x_0, \ldots, x_n)$ be non-random, d-dimensional baseline covariates. The policy maker is willing to extrapolate from her data by assuming that the welfare contrast of interest equals $U(\theta) = \theta(x_0)$ and that

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

where $\theta : \mathbb{R}^d \to \mathbb{R}$ is an unknown Lipschitz function with known constant C. For notational simplicity, we can further write μ_i for $\theta(x_i)$. Thus, $Y_i \sim N(\mu_i, \sigma_i^2)$, $Y \sim N(\mu, \Sigma)$, and $U(\theta) = \mu_0$. The policy question is: Given data Y, what proportion of the population in country $i = 0$ should be assigned the new policy?

Let $\|\beta\| := \sqrt{\beta^\top \beta}$ be the Euclidean norm of a vector β. In this example, the identified set for the welfare contrast μ_0 is

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

The lower and upper bounds on $I(\mu)$ are simply intersection bounds:

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

2.3. Robust Bayes Optimality

Our setting up to here is as in Montiel Olea et al. (2026). We now connect it to the robust Bayes literature by imposing a *set of priors* Γ on θ. Following Giacomini and Kitagawa (2021), we choose a particular single proper prior π_μ for $\mu \in M$ but leave the conditional prior of θ given μ, denoted as $\pi_{\theta|\mu}$, unrestricted except for

$$\pi_{\theta|\mu}\{U(\theta) \in I(\mu)\} = 1, \quad \pi_\mu\text{-a.s.} \tag{2.8}$$

Then, the class of priors Γ consists of all priors on θ induced by the single prior π_μ and any conditional prior $\pi_{\theta|\mu}$ that meets (2.8). Intuitively, we choose a single prior on the point-identified parameter and place no new restriction on the partially identified parameter $U(\theta)$. One can pick any proper prior π_μ; for tractability, we let π_μ be supported on two symmetric points $\{\bar{\mu}, -\bar{\mu}\}$ with equal probability, where $\bar{\mu} \in \mathbb{R}^n$ is chosen by the decision maker. Henceforth, Γ is understood to refer to

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

the implied set of priors:

$$\Gamma := \left\{ \pi_\theta = \int \pi_{\theta|\mu} d\pi_\mu : \pi_\mu \sim \text{unif}(\{-\bar{\mu}, \bar{\mu}\}), \pi_{\theta|\mu} \text{ satisfies (2.8)} \right\}. \tag{2.9}$$

While the set of priors Γ broadly puts us into "robust Bayes" territory, it still does not pin down a uniquely best decision rule because the sign of $U(\theta)$ can remain ambiguous. Let

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

be the *regret* of action $a \in [0,1]$. We evaluate decision rules d by their *expected regret*, defined as

$$\begin{aligned} R(d,\theta) &:= \mathbb{E}_{m(\theta)}[L(d(Y),\theta)] \\ &= U(\theta)\left\{\mathbf{1}\{U(\theta) \geq 0\} - \mathbb{E}_{m(\theta)}[d(Y)]\right\}, \end{aligned} \tag{2.10}$$

where for any $x \in \mathbb{R}^n$, $\mathbb{E}_x[\cdot]$ denotes expectation taken over $Y \sim N(x, \Sigma)$.

Even for given set of priors Γ and commitment to expected regret, the Robust Bayes literature contains multiple decision criteria that do not in general agree. The difference lies in when (and how) the expectations regarding the unknown parameter $\theta \in \Theta$ are taken. For decision rule $d \in \mathcal{D}_n$, let

$$r(d,\pi) := \int_{\theta \in \Theta} R(d,\theta) d\pi(\theta)$$

be its Bayes expected regret under a prior $\pi \in \Gamma$. Following Berger (1985, Def. 12, p. 216), we introduce the first robust Bayes optimality notion.

DEFINITION 1 (*Ex-Ante* Γ-Minimax Regret). *A decision rule* $d^* \in \mathcal{D}_n$ *is* Γ*-minimax regret (henceforth* Γ*-MMR) optimal if*

$$\sup_{\pi \in \Gamma} r(d^*,\pi) = \inf_{d \in \mathcal{D}_n} \sup_{\pi \in \Gamma} r(d,\pi). \tag{2.11}$$

An alternative "posterior" robustness notion is also common in Bayesian analysis. For each action $a \in [0,1]$, define *posterior expected regret* under prior π (Berger, 1985, Def. 8, p. 159)

$$\rho(a, \pi_{\theta|Y}) := \int_{\tilde{\theta} \in \Theta} L(a,\tilde{\theta}) d\pi_{\theta|Y}(\tilde{\theta}),$$

where $\pi_{\theta|Y}$ is the posterior distribution of θ given prior π and data Y.[6] Then we have the following alternative optimality criterion (Berger, 1985, Def. 10, p. 205).

[6] In our setting, information from data Y does not revise the conditional prior $\pi_{\theta|\mu}$ (Giacomini and Kitagawa, 2021). For any event A in the σ-algebra of Θ, we therefore have $\pi_{\theta|Y}(A) = \int \pi_{\theta|\mu}(A) d\pi_{\mu|Y}$, where $\pi_{\mu|Y}$ is the posterior distribution of μ given Y.

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

DEFINITION 2 (*Ex-Post* Γ-Posterior Expected Regret). *A decision rule* $d^* \in \mathcal{D}_n$ *is* Γ*-posterior expected regret (henceforth* Γ*-PER) optimal if*

$$\sup_{\pi \in \Gamma} \rho(d^*(Y), \pi_{\theta|Y}) = \inf_{a \in [0,1]} \sup_{\pi \in \Gamma} \rho(a, \pi_{\theta|Y}), \quad \forall Y \in \mathbb{R}^n. \tag{2.12}$$

If the decision maker's action space is restricted to $\{0, 1\}$, that is, randomization is not allowed, the above definitions are revised by replacing $[0, 1]$ with $\{0, 1\}$.

The labeling of Γ-MMR as "ex-ante" versus Γ-PER as "ex-post" can be related to the timing of a fictitious game against an adversarial Nature (see Section 3.3 for additional discussion). If Γ were a singleton, the criteria would coincide and would also agree with (single-prior) Bayes optimality. They do not in general agree otherwise. The term "Gamma minimax" usually (and even "robust Bayes" more often than not) refers to Γ-MMR (e.g., see Berger, 1985, Sect. 4.7.6).[7]

While it is not our agenda to advocate for either criterion, some possible considerations are as follows. The ex-ante approach may be perceived as more suitable if the planner has commitment power and has also been justified axiomatically (Hayashi, 2008; Stoye, 2011). We also find some numerical and theoretical evidence that the Γ-MMR rule may have desirable *frequentist properties*, for example, in states of the world off the prior's support (see discussions at the end of Section 3.2). Regarding computational feasibility, the ex-post approach is often easier due to the applicability of backward induction and is routinely employed to quantify posterior robustness of statistical decisions. On the other hand, recent work on numerical discovery of minimax rules (Aradillas Fernández et al., 2025; Guggenberger and Huang, 2025) may render the ex-ante approach more scalable. See additional discussions in Appendix E.

The two-point structure of π_μ in (2.9) appears in several related treatment choice problems with minimax regret optimality criteria. For example, the least favorable prior (LFP) takes such a form in a completely unconstrained minimax regret problem with point identified (Stoye, 2009) or partially identified (Stoye, 2012a; Yata, 2021; Montiel Olea et al., 2026) welfare contrast. In analogously constrained minimax regret problems in which we restrict $\mu \in [-|\bar{\mu}|, |\bar{\mu}|]$, extending the analyses in the preceding literature would also imply a similar two-point symmetric structure for the LFP. Given these precedents, we think our choice of π_μ represents a general feature of this class of minimax regret decisions, in addition to offering computational tractability.

[7] Giacomini et al. (2021) discuss both criteria for general loss functions and refer to Definitions 1 and 2 as the "unconditional Γ-minimax" and "conditional Γ-minimax" problems, respectively. For treatment choice problems with partial identification, Christensen et al. (2026) optimize the Γ-PER criterion, restricting the action space to be $\{0, 1\}$.

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

3. MAIN RESULTS

In this section, we solve for the two versions of robust Bayes optimality under the set of priors (2.9). Following Yata (2021), we invoke the following assumption.

Assumption 1. (i) Θ is convex, centrosymmetric (i.e., $\theta \in \Theta$ implies $-\theta \in \Theta$) and nonempty.
(ii) $m(\cdot)$ and $U(\cdot)$ are linear.

These conditions are restrictive but encompass many examples of empirical relevance (see this article's introduction, Montiel Olea et al., 2026, and Yata, 2021). Exploiting symmetry of the setting, we also set

$$\bar{I}(\bar{\mu}) + \underline{I}(\bar{\mu}) > 0.$$

This is a normalization because, by Lemma D.1, Assumption 1 implies $\bar{I}(-\bar{\mu}) + \underline{I}(-\bar{\mu}) = -(\underline{I}(\bar{\mu}) + \bar{I}(\bar{\mu}))$ and we could always replace $m(\cdot) \mapsto -m(\cdot)$; furthermore, the case of $\bar{I}(\bar{\mu}) + \underline{I}(\bar{\mu}) = 0$ gives rise to trivial solutions.[8]

We say that a rule is a *linear-index threshold rule* if it has the form $\mathbf{1}\{\beta^\top Y \geq c\}$ for some $\beta \in \mathbb{R}^n$ and $c \in \mathbb{R}$. Linear-index threshold rules are non-randomized. They are of particular interest because they form a complete class when $U(\theta)$ is point-identified (Karlin and Rubin, 1956) and have received particular attention in the recent literature (Ishihara and Kitagawa, 2021; Montiel Olea et al., 2026). For both Γ-MMR and Γ-PER, we will clarify when linear-index threshold rules are optimal and when they are not. For reasons that will become obvious, the following linear-index threshold rule is of particular interest:

$$d^*_{w,0} := d^*_{w,0}(Y) := \mathbf{1}\{w^\top Y \geq 0\}, \; w := \Sigma^{-1}\bar{\mu}. \tag{3.1}$$

For each vector $\beta \in \mathbb{R}^n$, let $\|\beta\|_\Sigma := \sqrt{\beta^\top \Sigma \beta}$. Thus, $\|w\|_\Sigma = \sqrt{w^\top \Sigma w} = \sqrt{\bar{\mu}^\top \Sigma^{-1} \bar{\mu}}$. Denote by $\Phi(\cdot)$ the standard normal cdf and by $\Phi^{-1}(\cdot)$ its inverse, that is, the corresponding quantile function.

3.1. Ex-Ante Robust Bayes Optimality

THEOREM 1 (Γ-MMR Optimal Decisions). *Consider a treatment choice problem with welfare function (2.1), statistical model (2.3), and set of priors (2.9), that satisfy Assumption 1.*

(i) *If*

$$\frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu}) - \underline{I}(\bar{\mu})} \geq \Phi(\|w\|_\Sigma), \tag{3.2}$$

then $d^*_{w,0}$ is uniquely Γ-MMR optimal.

[8] In this case, the identified set for $U(\theta)$ does not change with true value of μ, and under all decision criteria considered here, the solution will be a no-data rule that simply flips a coin.

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

(ii) *If*

$$\frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})} < \Phi(\|w\|_\Sigma), \tag{3.3}$$

then a rule $d \in \mathcal{D}_n$ attains Γ-MMR if, and only if, it implies

$$\mathbb{E}_{-\bar{\mu}}[d^*(Y)] = \frac{-\underline{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})}, \tag{3.4}$$

$$\mathbb{E}_{\bar{\mu}}[d^*(Y)] = \frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})}. \tag{3.5}$$

In particular, any rule of the form $\mathbf{1}\{w^\top Y \geq c\}$ for some $c \in \mathbb{R}$ is not Γ-MMR optimal. The following rules, and any convex combination of them, are Γ-MMR optimal:

$$d^*_{RT} := \Phi\left(\frac{w^\top Y}{\tilde{\sigma}}\right), \tag{3.6}$$

$$d^*_{linear} := \begin{cases} 0, & w^\top Y < -\rho^*, \\ \frac{w^\top Y + \rho^*}{2\rho^*}, & -\rho^* \leq w^\top Y \leq \rho^*, \\ 1, & w^\top Y > \rho^*, \end{cases} \tag{3.7}$$

$$d^*_{step} := \begin{cases} \frac{1}{2} - \beta^*, & w^\top Y < 0, \\ \frac{1}{2} + \beta^*, & w^\top Y \geq 0, \end{cases} \tag{3.8}$$

where

$$\tilde{\sigma} = \sqrt{\left[\frac{\|w\|^2_\Sigma}{\Phi^{-1}\left(\frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})}\right)}\right]^2 - \|w\|^2_\Sigma},$$

$\rho^* > 0$ is unique and solves $\int_0^1 \Phi\left(\frac{2\rho^* x - \rho^* - \|w\|^2_\Sigma}{\|w\|_\Sigma}\right) dx = \frac{-\underline{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})}$, and

$$\beta^* = \frac{\frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})} - \frac{1}{2}}{2\Phi(\|w\|_\Sigma) - 1} \in \left(0, \frac{1}{2}\right).$$

(iii) *In case (3.3), among linear-index threshold rules of the form $\mathbf{1}\{w^\top Y \geq c\}$, the optimal thresholds are $\pm c^*$, where*

$$c^* := \|w\|^2_\Sigma - \|w\|_\Sigma \Phi^{-1}\left(\frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})}\right) > 0.$$

(iv) *In case (3.3) and if $n > 1$, the following linear-index threshold rule is also Γ-MMR optimal:*

$$d^*_{w_{t^*},0}(Y) := \mathbf{1}\{w^\top_{t^*} Y \geq 0\},$$

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

where $w_{t^*} = \Sigma^{-1}(t^*\bar{\mu} + (1-t^*)\dot{\mu})$, $\dot{\mu} \neq 0$ *is such that* $\dot{\mu}^\top \Sigma^{-1}\bar{\mu} = 0$,

$$t^* := \frac{1}{1 \pm \sqrt{\frac{(1-s^*)}{s^*}\frac{\|w\|_\Sigma^2}{\|\Sigma^{-1}\dot{\mu}\|_\Sigma^2}}}, \tag{3.9}$$

and

$$s^* := \frac{\left[\Phi^{-1}\left(\frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}\right)\right]^2}{\|w\|_\Sigma^2} \in (0,1).$$

Thus, $d^*_{w_{t^*},0}$ *is also optimal among all linear threshold rules.*

Theorem 1 reveals that the Γ-MMR rules qualitatively change depending on whether condition (3.2) is met or not. This condition admits an intuitive interpretation: Up to clamping to the unit interval (i.e., values outside this interval are mapped to its edges), the left-hand side of (3.2) equals the unique minimax regret optimal rule for *known* μ (Manski, 2007b); therefore, it arguably measures the model's identification strength. The right-hand side of (3.2) can be interpreted as the informativeness of the data about which of $\{-\bar{\mu}, \bar{\mu}\}$ obtained. Therefore, Theorem 1 says that if the model's identification power is sufficiently large compared to the informativeness of the data, then, the unique Γ-MMR optimal rule is $d^*_{w,0}$, a non-randomized linear index rule with threshold 0 that effectively ignores the partial identification issue. In contrast, if the model's identification power is small compared to the informativeness of data, there are infinitely many Γ-MMR optimal rules, all of which satisfy (3.4) and (3.5). Examples include suitably smoothed versions of $d^*_{w,0}$ like d^*_{RT} and d^*_{linear} (the functional forms of which showed up in Montiel Olea et al., 2026) as well as d^*_{step} (which is a new result).

To understand how these two "regimes" arise, it is helpful to think of a zero-sum game against an adversarial Nature in which an MMR decision rule and a distribution over parameter values (the LFP) form a Nash equilibrium.[9] If identification is strong in the sense of (3.2), this game has an *informative* equilibrium: the LFP evenly randomizes over two points $(\mu, U(\theta)) \in \{(\bar{\mu}, \bar{I}(\bar{\mu})), (-\bar{\mu}, \underline{I}(-\bar{\mu}))\}$, to which the unique Bayes response (and, therefore, uniquely optimal MMR rule) is $d^*_{w,0}$. In all other cases, the equilibrium is *uninformative*: the LFP is supported on four points $(\mu, U(\theta)) \in \{(\bar{\mu}, \bar{I}(\bar{\mu})), (\bar{\mu}, \underline{I}(\bar{\mu})), (-\bar{\mu}, \bar{I}(-\bar{\mu})), (-\bar{\mu}, \underline{I}(-\bar{\mu}))\}$ with a probability profile such that the posterior expectation of $U(\theta)$ always equals 0. While any decision rule best responds to this prior, not every decision rule is MMR because the LFP must also best respond to the decision rule. In this very structured setting, the latter is guaranteed by conditions (3.4) and (3.5),

[9] Analyzing this game is also how results are formally proved. See Wald (1945) for what may be the first clear statement of this and Lehmann and Casella (1998, Thm. 1.4, Chap. 5) for a formalization. Within the literature on decisions under partial identification, the proof technique was first explicitly used in Stoye (2007) and the equilibrium structure with a noninformative and an informative regime was first encountered in Stoye (2012a).

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

establishing claim (ii) and nonuniqueness of optimal decision rules. Indeed, even the non-randomized threshold rules from Theorem 1(iv) do not reflect any updating along the game's equilibrium path; they rather use uninformative features of the data as randomization device.

In the uninformative equilibrium, we encounter several additional findings. First, despite the problem's apparent symmetry, $d^*_{w,0}$ is *not* optimal even among linear-index threshold rules that use the index $w^\top Y$. Instead, Theorem 1(iii) characterizes exactly two optimal thresholds, one positive and one negative. Second, while the particular linear-index threshold rule $d^*_{w,0}$ is not Γ-MMR optimal, in higher-dimensional problems ($n > 1$), there do exist linear-index threshold rules that are. This finding is in stark contrast to Montiel Olea et al. (2026), who find that, in a large class of special cases, *no* linear index threshold rule is MMR optimal. The crucial difference in settings is that the set of priors much constrains the decision theoretic problem's state space; as a result, the signal space is much richer than the state space, and this can be exploited to mimic randomization without nominally randomizing. Compare Manski's (2025) abstract observation that, if a policy maker is not allowed to explicitly randomize, sampling uncertainty can be beneficial by providing an implicit randomization device. We note that this phenomenon is reminiscent of classic "purification" results in game theory (Dvoretzky, Wald, and Wolfowitz, 1951; Khan, Rath, and Sun, 2006).[10] In contrast, it is not deeply related to the general intuition that "Bayesians don't randomize."

We next apply Theorem 1 to Example 1 and immediately get the following results.

Corollary 1 (Γ-MMR Decisions in Example 1). *In Example 1, the following statements are true:*

(i) *If*

$$\frac{\bar{\mu}+k}{2k} \geq \Phi(\bar{\mu}/\sigma), \tag{3.10}$$

then

$$d^*_0(\cdot) := \mathbf{1}\{\hat{\mu} \geq 0\}$$

is the unique Γ-MMR optimal decision rule.

(ii) *If (3.10) fails, then a rule $d \in \mathcal{D}_1$ attains Γ-MMR if, and only if, it implies*

$$\mathbb{E}[d(\hat{\mu}) \mid \mu = -\bar{\mu}] = \frac{-\bar{\mu}+k}{2k}, \tag{3.11}$$

$$\mathbb{E}[d(\hat{\mu}) \mid \mu = \bar{\mu}] = \frac{\bar{\mu}+k}{2k}. \tag{3.12}$$

[10] We thank Elliot Lipnowski for reminding us of this literature.

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

In particular, no linear threshold rule is Γ*-MMR optimal. The following rules, and any convex combination of them, are all* Γ*-MMR optimal:*

$$d^*_{RT} := \Phi\left(\frac{\hat{\mu}}{\tilde{\sigma}}\right), \tag{3.13}$$

$$d^*_{linear} := \begin{cases} 0, & \hat{\mu} < -\frac{\sigma^2\rho^*}{\bar{\mu}}, \\ \frac{\bar{\mu}\hat{\mu}+\sigma^2\rho^*}{2\sigma^2\rho^*}, & -\frac{\sigma^2\rho^*}{\bar{\mu}} \le \hat{\mu} \le \frac{\sigma^2\rho^*}{\bar{\mu}}, \\ 1, & \hat{\mu} > \frac{\sigma^2\rho^*}{\bar{\mu}}, \end{cases} \tag{3.14}$$

$$d^*_{step} := \begin{cases} \frac{1}{2} - \frac{\bar{\mu}}{2k\left(2\Phi\left(\frac{\bar{\mu}}{\sigma}\right)-1\right)}, & \hat{\mu} < 0, \\ \frac{1}{2} + \frac{\bar{\mu}}{2k\left(2\Phi\left(\frac{\bar{\mu}}{\sigma}\right)-1\right)}, & \hat{\mu} \ge 0, \end{cases} \tag{3.15}$$

where

$$\tilde{\sigma} = \sigma\sqrt{\left[\frac{\bar{\mu}}{\sigma\,\Phi^{-1}\left(\frac{\bar{\mu}+k}{2k}\right)}\right]^2 - 1},$$

and $\rho^* > 0$ *is unique and solves* $\int_0^1 \Phi\left(\frac{2\rho^* x - \rho^* - (\bar{\mu}/\sigma)^2}{\bar{\mu}/\sigma}\right) dx = \frac{-\bar{\mu}+k}{2k}$.

(iii) *In case (ii), the best linear threshold rules in terms of* Γ*-minimax regret are* $\mathbf{1}\{\hat{\mu} \ge \pm c^*\}$*, where*

$$c^* = \bar{\mu} - \sigma\,\Phi^{-1}\left(\frac{\bar{\mu}+k}{2k}\right).$$

Note that there is no analog to Theorem 1's case (iv); indeed, no linear threshold rule is optimal in case (ii). This is because the scalar nature of the signal Y shuts down the aforementioned purification mechanism. See Figure 1 for an illustration of different Γ-MMR optimal rules in Example 1 for selected parameter values of k, σ, and $\bar{\mu}$.

3.2. Ex-Post Robust Bayes Optimality

THEOREM 2 (Γ-PER Optimal Decisions). *Suppose all conditions in Theorem 1 hold true. Then:*

(i) *If* $\underline{I}(\bar{\mu}) < 0 < \overline{I}(\bar{\mu})$*, the unique* Γ*-PER optimal rule is*

$$d^*_{\mathrm{PER}}(Y) = \begin{cases} \frac{-\underline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})}, & \text{if } w^\top Y < 0, \\ \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})}, & \text{if } w^\top Y \ge 0. \end{cases}$$

Otherwise, $d^*_{w,0}$ *is* Γ*-PER optimal.*

(ii) $d^*_{w,0}$ *is always the* Γ*-PER optimal non-randomized rule.*

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

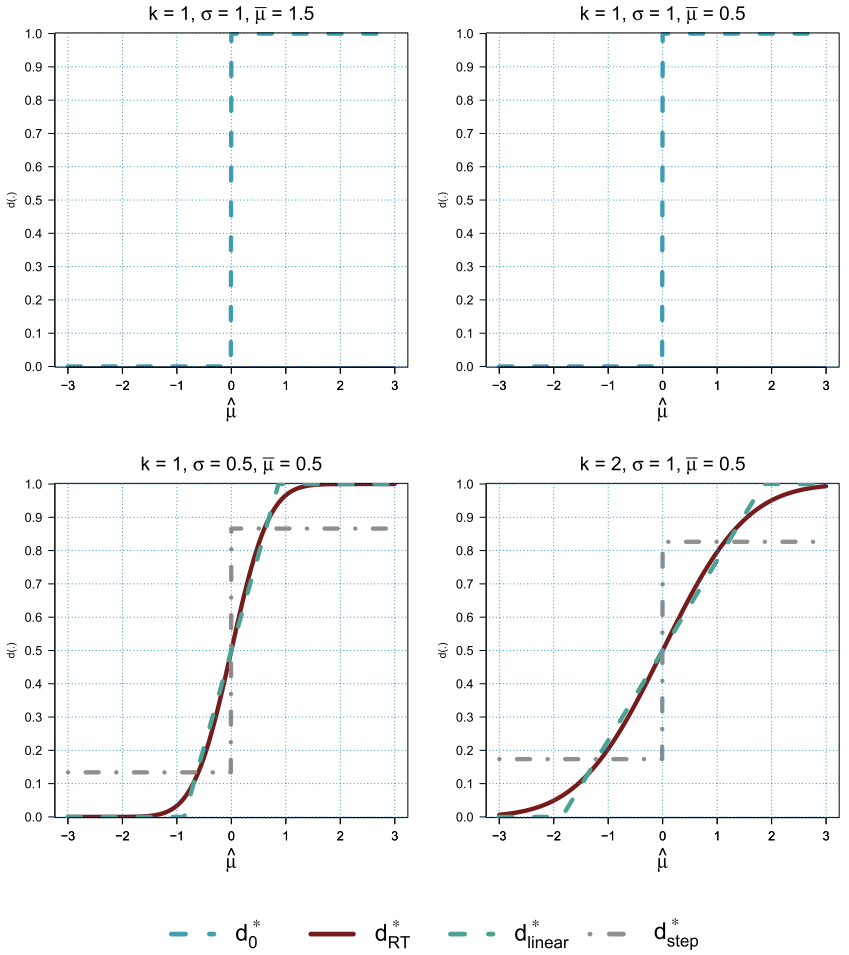


FIGURE 1. Γ-MMR optimal rules in Example 1.

Note: This figure reports the Γ-MMR optimal rules in Example 1 for various combinations of parameter values. In the top two panels, the combinations of parameter values satisfy (3.10). Therefore, the unique Γ-MMR optimal rule is d_0^*. In the bottom two panels, (3.10) fails. As a result, d_0^* is no longer Γ-MMR optimal. Instead, d_{RT}^*, d_{linear}^*, and d_{step}^* are all Γ-MMR optimal.

The results of Theorem 2 offer some important clarifications regarding Γ-PER in treatment choice problems with partial identification. First, even if regret is evaluated according to the posterior distribution, it is not necessarily true that optimal rules are non-randomized. In fact, whenever there is model ambiguity regarding the sign of $U(\theta)$ (i.e., $\underline{I}(\bar{\mu}) < 0 < \bar{I}(\bar{\mu})$), the unique Γ-PER optimal rule is randomized. Therefore, restricting the action space to $\{0, 1\}$ in such problems is

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

not without loss of generality even under the Γ-PER criterion. Comparing results with Theorem 1 also allows for instructive observations on when Γ-PER and Γ-MMR optimal rules agree or disagree; we elaborate these in Corollary 3. Applying Theorem 2 to Example 1, we finally obtain the following corollary.

COROLLARY 2 (Γ-PER Rules in Example 1). *In Example 1, the following statements are true:*

(i) *If* $\bar{\mu} < k$, *then* Γ*-PER is uniquely minimized by*

$$d^*_{\text{PER}}(\hat{\mu}) = \begin{cases} \frac{k+\bar{\mu}}{2k} & \text{if } \hat{\mu} \geq 0 \\ \frac{k-\bar{\mu}}{2k} & \text{if } \hat{\mu} < 0. \end{cases}$$

(ii) *If* $\bar{\mu} \geq k$, *then* d^*_0 *is* Γ*-PER optimal. Furthermore, it is always the* Γ*-PER optimal threshold rule.*

In Figure 2, we depict Γ-PER optimal rules for Example 1 with the same parameter values considered in Figure 1. We see clearly that Γ-MMR and -PER optimal rules coincide (if randomization is allowed) only in the special case when $k \leq \bar{\mu}$, an observation we generalize in Corollary 3. More specifically, in the top left panel of Figure 2, as $k \leq \bar{\mu}$, Γ-MMR and -PER coincide and are the non-randomized threshold rule d^*_0. In the top right panel, $k > \bar{\mu}$ and the Γ-PER optimal rule becomes d^*_{PER}. However, since (3.10) still holds, d^*_0 is still Γ-MMR optimal. For the bottom two panels, as it still holds $k > \bar{\mu}$, d^*_{PER} is still Γ-PER optimal. However, the associated parameter values imply that (3.10) fails. As a result, d^*_0 is no longer Γ-MMR optimal and many Γ-MMR rules exist. But even in these cases, Γ-MMR and -PER rules differ, as among the class of step function rules (which contain d^*_{PER}), only d^*_{step} is Γ-MMR optimal, still different from d^*_{PER}.

Applying Montiel Olea et al. (2026, Thm. 1) to the current setting, we can conclude that all rules, including both Γ-MMR and Γ-PER optimal rules, are at least admissible. Moreover, by definition, the Γ-MMR rule will have the lower value of the (ex ante) game. Therefore, it might be more useful to compare the (frequentist) profiled regrets (Montiel Olea et al., 2026) of ex-post Γ-PER and other rules as a function of the true but unknown mean μ of data $\hat{\mu}$. In the context of Example 1, we can report $\bar{R}(d, \mu) := \sup_{\mu^* \in [\mu-k, \mu+k]} R(d, \mu, \mu^*)$ as μ varies for each rule d, where $R(d, \mu, \mu^*)$ is the expected regret of rule d as a function of μ and μ^*. We report several findings (we emphasize that these are not obvious: we took an ex-ante perspective but did not restrict ourselves to those values of μ that the prior allows). First, in both panels of Figure 3, the profiled expected regret of the Γ-PER rule exceeds that of the Γ-MMR rule. Therefore, in this particular example, Γ-MMR arguably outperforms Γ-PER from a broader frequentist point of view. Second, we in fact prove that, whenever $\bar{\mu}$ is sufficiently small and k is sufficiently large, the Γ-PER rule is profiled-regret dominated, that is, there exists a rule $d \neq d^*_{\text{PER}}$ such that

$$\bar{R}(d, \mu) \leq \bar{R}(d^*_{\text{PER}}, \mu) \qquad \textbf{(3.16)}$$

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

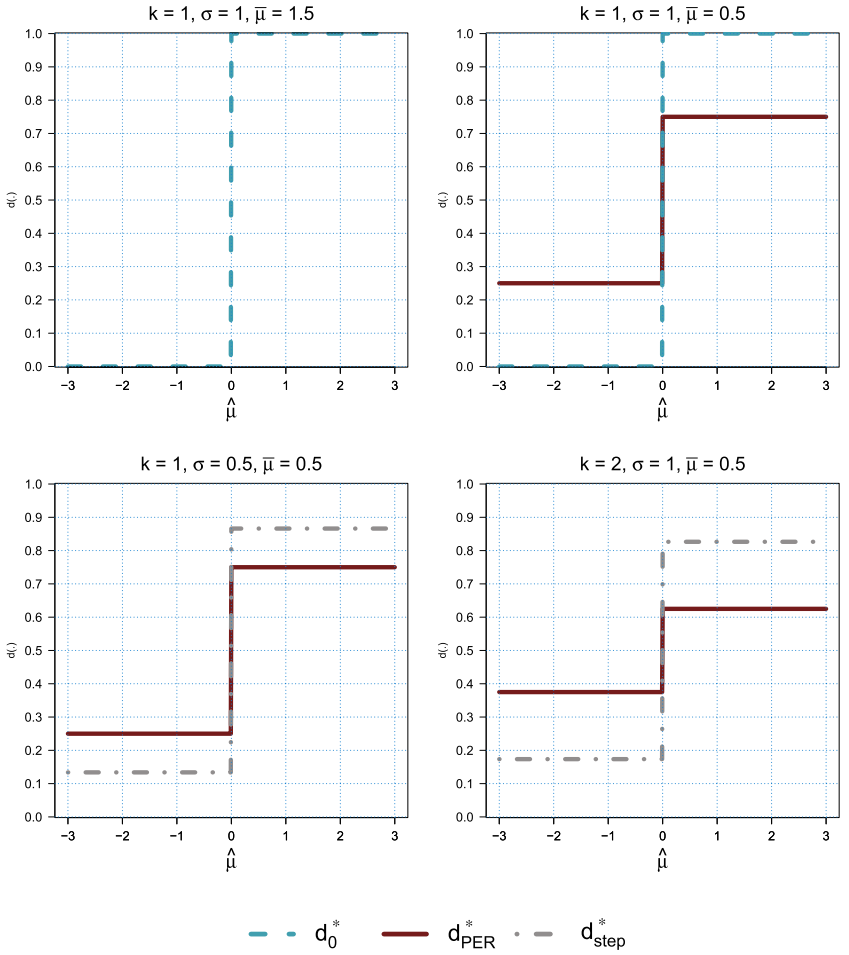


FIGURE 2. Γ-PER optimal rules in Example 1.

Note: This figure reports the Γ-PER optimal rules for the same parameter values considered in Figure 1. In the top left panel, Γ-PER and Γ-MMR optimal rules coincide and are both d_0^*. For the rest of the panels, the unique Γ-PER optimal rules are all d_{PER}^* and are different from any Γ-MMR optimal rules.

for all $\mu \in \mathbb{R}$ with the inequality strict for some μ (see Lemma D.2 in Appendix D for an exact statement). Intuitively, the Γ-PER rule mixes between a coin flip rule and the naive threshold rule (d_0^*). When $\bar{\mu}$ is small, Γ-PER rule is more analogous to the coin flip rule, which Montiel Olea et al. (2026) show to be dominated in terms of profiled regret; also, its profiled regret fails to vanish as $\mu \to \pm\infty$ even though the optimal treatment is known ex ante in this case. Furthermore, Figure 3 reveals that the Γ-MMR rule d_{linear}^*, though not globally MMR optimal, in some cases has essentially the same profiled risk function as the least randomizing global

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

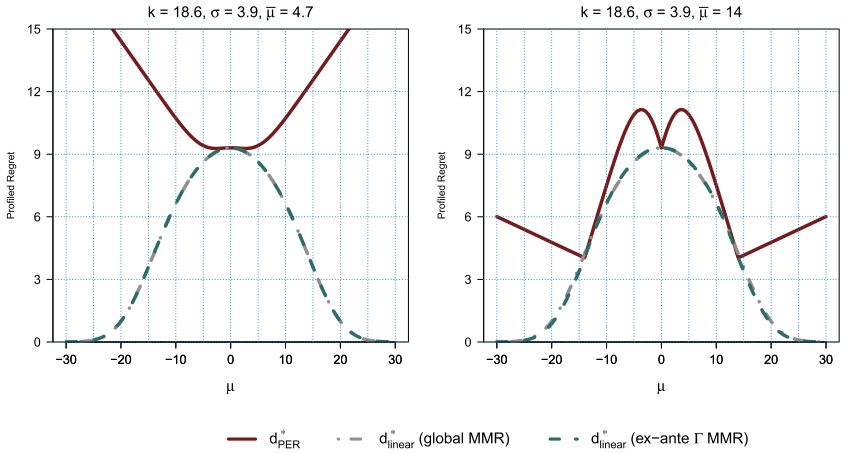


FIGURE 3. Profiled regret of Γ-PER and other rules in Example 1.
Note: This figure reports the (frequentist) profiled regrets, as a function of the true mean μ of data $\hat{\mu}$, of Γ-PER rule, Γ-MMR linear rule, and the least randomizing global MMR optimal rule (Montiel Olea et al., 2026) for the parameter values considered in Montiel Olea et al. (2026, Fig. 2). In the left plot, we see that Γ-PER rule is visually dominated. In both plots, the profiled regret of Γ-MMR linear rule and the least randomizing global MMR optimal rule look very similar and are essentially overlapping with each other.

MMR optimal rule (Montiel Olea et al., 2026). This may lend a Robust Bayes interpretation to the latter.

3.3. Ex-Ante and Ex-Post in the Game against Nature

The different approaches analyzed here can all be expressed as different timing assumptions in the statistical game. The distinction is quite obvious for ex-ante versus ex-post: In the former (and original Waldian) perspective, this game is simultaneous move; in particular, Nature moves before data Y are realized. In the ex-post perspective, Nature sees Y before choosing a prior. It is immediately clear that this may be easier to solve because it allows for backward induction. There is also an immediate sense that solutions might not agree, as we indeed found.

But whether the decision maker is allowed to (or at least wants to) randomize or not can equally be thought of as changing the game's timing, providing another way to think about the action space for the ex-ante and ex-post approaches. The more standard perspective is that the decision maker may randomize over decision rules and Nature must move before learning the outcome of this randomization. By the nature of zero-sum games, this setup will frequently yield randomized solutions. In consequence, it is essential to define the decision maker's action space as $[0, 1]$. In contrast, if Nature is allowed to move after the decision maker's randomization is realized, then any incentive to randomize is gone and we may as well restrict the action space to $\{0, 1\}$.

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

Theorems 1 and 2 clarify that these distinctions actually matter in an interesting example. We spell this out in Corollary 3, which considers both cases when randomization is allowed and not allowed. The bottom line is that the assessment is quite sensitive to how the problem is set up. If underlying parameters lead to sufficiently small identification power, the criteria disagree.

COROLLARY 3. *Consider a treatment choice problem with welfare function (2.1) and statistical model (2.3) that satisfies Assumption 1.*

(i) *Suppose randomization is allowed. Then the* Γ*-MMR and* Γ*-PER optimal rules coincide if, and only if,* $\underline{I}(\bar{\mu}) \geq 0$.
(ii) *Suppose randomization is not allowed. Then the* Γ*-MMR and* Γ*-PER optimal rules coincide if, and only if,*

$$\frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu}) - \underline{I}(\bar{\mu})} \geq \Phi(\|w\|_\Sigma). \tag{3.17}$$

Thus, if randomization is allowed, Γ-MMR and Γ-PER optimal rules coincide only in the somewhat trivial case in which we a priori know that $U(\theta)$ and μ have the same sign, so that optimal treatment choice reduces to Bayesian inference on the point identified μ. They disagree in all other cases, including in settings where there are infinitely many Γ-MMR optimal rules.

One might expect more agreement once randomization is excluded; after all, this leads to a much simpler action space. Part (ii) shows that there is some truth to this: The condition for agreement changes from $\underline{I}(\bar{\mu}) \geq 0$ to the strictly weaker (3.17). However, the criteria continue to disagree in many cases. It may be instructive to think of these cases in terms of "comparative statics." For example, consider holding all parameters of the problem fixed but scaling the signal variance Σ by a positive scalar, say to reflect a change in sample size. Then (3.17) will hold if, and only if, Σ is *large* enough; hence, as long as $\underline{I}(\bar{\mu}) \geq 0$, increasing sample size will eventually cause disagreement between Γ-MMR and -PER even if randomization is excluded. Similarly, for fixed Σ, Γ-MMR and -PER rules will always disagree if the model's identification power is *small* enough.

4. CONCLUSION

We studied treatment choice problems that display partial identification through the lens of the robust Bayes criteria. To do so, we take the general framework of Yata (2021) and others and embed in it a simple example of the set of priors advocated by Giacomini and Kitagawa (2021). We describe and contrast (ex-ante) Γ-minimax regret and (ex-post) Γ-posterior expected regret and analytically derive optimal solutions with and without randomization.

Our results contain two key messages that we think are valuable to the literature. First, with partial identification and multiple priors, ex-ante and ex-post assessments do *not* agree in general, whether or not randomized rules are allowed.

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

This may at first seem expected due to dynamic inconsistency of multiple prior Bayes criteria, but was not obvious in view of the set of prior's specific structure. Second, randomization can be optimal in both ex-ante and ex-post problems—it is with loss of generality to exclude them even when regret is evaluated ex-post. The contrast between the results also illustrates a need to better understand the comparative advantages—whether from a theoretical or practical perspective—of using one criterion over the other.

An obvious limitation lies in our use of a convenient but restrictive prior on μ. As we discovered in Section 3.2, the Γ-MMR rule also performs well across other values of μ; however, this may be related to the fact that the globally LFP in this setting has a two-point structure on μ as well and would obviously not generalize to arbitrary uses of restrictive priors. We provide some insight on more general priors in Appendix E. In short, the tractability advantage of Γ-PER may become pronounced in such settings, although we hope that recent computational developments (Aradillas Fernández et al., 2025; Guggenberger and Huang, 2025) will attenuate this concern.

A. PROOFS OF MAIN RESULTS

Proofs will frequently claim and verify equilibria of the fictitious game against Nature. Recall that, from basic facts about zero-sum games, if a decision rule *uniquely* best responds to some least favorable prior, it must be the *unique* equilibrium rule.

A.1. Proof of Theorem 1

A.1.1. *Statement (i).* The expected regret of decision rule d is

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

Recall that, by the definition of Γ, we have $\pi_\mu \sim \text{unif}(\{-\bar{\mu}, \bar{\mu}\})$ and, given $\mu = \pm\bar{\mu}$, $\pi_{\theta|\mu}(U(\theta) \in I(\mu)) = 1$. Therefore, the Bayes expected regret of d under prior $\pi \in \Gamma$ equals

$$\begin{aligned} r(d,\pi) = \frac{1}{2} \cdot \left[\int_{\tilde{\theta}\in\Theta} U(\tilde{\theta}) \left(\mathbf{1}\{U(\tilde{\theta}) \geq 0\} - \mathbb{E}_{\bar{\mu}}[d]\right) d\pi_{\theta|\bar{\mu}}(\tilde{\theta}) \right] \\ + \frac{1}{2} \cdot \left[\int_{\tilde{\theta}\in\Theta} U(\tilde{\theta}) \left(\mathbf{1}\{U(\tilde{\theta}) \geq 0\} - \mathbb{E}_{-\bar{\mu}}[d]\right) d\pi_{\theta|-\bar{\mu}}(\tilde{\theta}) \right], \end{aligned}$$

where $\mathbb{E}_{\bar{\mu}}[d] := \mathbb{E}_{\bar{\mu}}[d(Y)]$ and $\mathbb{E}_{-\bar{\mu}}[d] := \mathbb{E}_{-\bar{\mu}}[d(Y)]$. One can easily solve for

$$\begin{aligned} \sup_{\pi\in\Gamma} r(d,\pi) = \frac{1}{2} \max\left\{ \bar{I}(\bar{\mu})(1 - \mathbb{E}_{\bar{\mu}}[d]), -\underline{I}(\bar{\mu})\mathbb{E}_{\bar{\mu}}[d] \right\} \\ + \frac{1}{2} \max\left\{ \bar{I}(-\bar{\mu})(1 - \mathbb{E}_{-\bar{\mu}}[d]), -\underline{I}(-\bar{\mu})\mathbb{E}_{-\bar{\mu}}[d] \right\}. \end{aligned}$$

The least favorable prior π^* equals

$$\pi^* = \left\{ \pi^*_\mu \sim \text{unif}(\{-\bar{\mu}, \bar{\mu}\}), \pi^*_{\theta|\bar{\mu}}(U(\theta) = \bar{I}(\bar{\mu})) = 1, \pi^*_{\theta|-\bar{\mu}}(U(\theta) = \underline{I}(-\bar{\mu})) = 1 \right\}. \tag{A.1}$$

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

Lemma B.1 shows that the unique Bayes rule against π^* is

$$d^*_{w,0} = \mathbf{1}\{w^\top Y \geq 0\}, \quad \text{where } w = \Sigma^{-1}\bar{\mu}.$$

Lemma B.2 establishes that $\sup_{\pi\in\Gamma} r(d^*_{w,0},\pi) = r(d^*_{w,0},\pi^*)$ as long as (3.2) holds true. This establishes the claim.

A.1.2. *Statement (ii).*

Step 1: We show that when (3.3) holds, any rule $d \in \mathcal{D}_n$ is Γ-MMR optimal if

$$\mathbb{E}_{\bar{\mu}}[d] = \frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}, \quad \mathbb{E}_{-\bar{\mu}}[d] = \frac{-\underline{I}(\bar{\mu})}{\bar{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}. \tag{A.2}$$

The least favorable prior π^* is such that $\mu \sim \text{unif}(\{-\bar{\mu},\bar{\mu}\})$, and when $\mu = \bar{\mu}$,

$$U(\theta) = \begin{cases} \bar{I}(\bar{\mu}), & \text{with probability (w.p.) } p_1, \\ \underline{I}(\bar{\mu}), & \text{w.p. } 1-p_1, \end{cases}$$

where $p_1 > 0$ is such that $p_1\bar{I}(\bar{\mu}) + (1-p_1)\underline{I}(\bar{\mu}) = 0$, and when $\mu = -\bar{\mu}$,

$$U(\theta) = \begin{cases} \bar{I}(-\bar{\mu}), & \text{w.p. } p_2, \\ \underline{I}(-\bar{\mu}), & \text{w.p. } 1-p_2, \end{cases}$$

where p_2 is such that $p_2\bar{I}(-\bar{\mu}) + (1-p_2)\underline{I}(-\bar{\mu}) = 0$. Lemma B.3 establishes that any decision rule is Bayes against this prior (intuitively because the data are uninformative), and Lemma B.4 further shows that, for any rule d that satisfies (A.2), $\sup_{\pi\in\Gamma} r(d,\pi) = r(d,\pi^*)$ obtains. This establishes the claim.

Step 2: We next verify the "only if" statement. Recall that the least favorable prior π^* must best respond to any optimal decision rule, that is, for any MMR optimal rule d, π^* must solve

$$\begin{aligned}\sup_{\pi\in\Gamma} r(d,\pi) = {} & \frac{1}{2}\max\left\{\bar{I}(\bar{\mu})(1-\mathbb{E}_{\bar{\mu}}[d]), -\underline{I}(\bar{\mu})\mathbb{E}_{\bar{\mu}}[d]\right\} \\ & + \frac{1}{2}\max\left\{\bar{I}(-\bar{\mu})(1-\mathbb{E}_{-\bar{\mu}}[d]), -\underline{I}(-\bar{\mu})\mathbb{E}_{-\bar{\mu}}[d]\right\}.\end{aligned}$$

This, however, requires that Nature is indifferent between $\bar{I}(\bar{\mu})$ and $\underline{I}(\bar{\mu})$ when $\mu = \bar{\mu}$ and similarly between $\bar{I}(-\bar{\mu})$ and $\underline{I}(-\bar{\mu})$ when $\mu = -\bar{\mu}$. That is, we must have

$$\bar{I}(\bar{\mu})(1-\mathbb{E}_{\bar{\mu}}[d]) = -\underline{I}(\bar{\mu})\mathbb{E}_{\bar{\mu}}[d], \quad \bar{I}(-\bar{\mu})(1-\mathbb{E}_{-\bar{\mu}}[d]) = -\underline{I}(-\bar{\mu})\mathbb{E}_{-\bar{\mu}}[d],$$

which is equivalent to (A.2).

Step 3: We show a rule of form $\mathbf{1}\left\{w^\top Y \geq c\right\}$ for some $c \in \mathbb{R}$ cannot be Γ-MMR optimal when (3.3) holds. Note $w = \Sigma^{-1}\bar{\mu}$. Thus,

$$\mathbb{E}_\mu[\mathbf{1}\{w^\top Y \geq c\}] = 1 - \Phi\left(\frac{c - w^\top\mu}{\sqrt{w^\top\Sigma w}}\right).$$

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

Suppose by contradiction that a rule $\mathbf{1}\left\{w^\top Y \geq c\right\}$ is optimal, then, by statement (ii), we have

$$1-\Phi\left(\frac{c-w^\top\bar{\mu}}{\sqrt{w^\top\Sigma w}}\right)=\frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})} \tag{A.3}$$

$$1-\Phi\left(\frac{c+w^\top\bar{\mu}}{\sqrt{w^\top\Sigma w}}\right)=\frac{-\underline{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})}. \tag{A.4}$$

By symmetry, (A.3) and (A.4) can both hold only if $c=0$. But (3.3) then implies that

$$\Phi\left(\frac{w^\top\bar{\mu}}{\sqrt{w^\top\Sigma w}}\right)=\Phi(\|w\|_\Sigma)>\frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})},$$

so that (A.3) cannot in fact hold when $c=0$, a contradiction.

Step 4: We verify that d^*_{RT}, d^*_{linear}, and d^*_{step} are all Γ-MMR optimal. Due to symmetry, it suffices to show that

$$\mathbb{E}_{\bar{\mu}}[d^*_{\mathrm{RT}}]=\mathbb{E}_{\bar{\mu}}[d^*_{\mathrm{linear}}]=\mathbb{E}_{\bar{\mu}}[d^*_{\mathrm{step}}]=\frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})}.$$

To see $\mathbb{E}_{\bar{\mu}}[d^*_{\mathrm{RT}}]=\frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})}$, consider the random threshold rule $\mathbf{1}\left\{w^\top Y\geq\xi\right\}$, where $\xi\sim N(0,\tilde{\sigma}^2)$ is independent of Y. As $w^\top Y-\xi\sim N\left(\|w\|_\Sigma^2,\|w\|_\Sigma^2+\tilde{\sigma}^2\right)$, algebra shows

$$\mathbb{E}_{\bar{\mu}}\left[\mathbf{1}\left\{w^\top Y\geq\xi\right\}\right]=\frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})}$$

as required. To see $\mathbb{E}_{\bar{\mu}}[d^*_{\mathrm{linear}}]=\frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})}$, let $\rho>0$ and

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

Applying Lemma C.7 in Montiel Olea et al. (2026) and integration by parts yield

$$f(\rho):=\mathbb{E}_{\bar{\mu}}[d_{\mathrm{linear},\rho}]=1-\int_0^1\Phi\left(\frac{2\rho x-\rho-\|w\|_\Sigma^2}{\|w\|_\Sigma}\right)dx=1-\frac{\|w\|_\Sigma}{2\rho}\int_{\frac{-\rho-\|w\|_\Sigma^2}{\|w\|_\Sigma}}^{\frac{\rho-\|w\|_\Sigma^2}{\|w\|_\Sigma}}\Phi(t)\,dt.$$

Note that $\lim_{\rho\downarrow0}f(\rho)=1-\Phi(-\|w\|_\Sigma)=\Phi(\|w\|_\Sigma)>\frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})}$, while L'Hopital's rule implies

$$\lim_{\rho\to\infty}f(\rho)=1-\frac{1}{2}\lim_{\rho\to\infty}\left\{\Phi\left(\frac{\rho-\|w\|_\Sigma^2}{\|w\|_\Sigma}\right)-\Phi\left(\frac{-\rho-\|w\|_\Sigma^2}{\|w\|_\Sigma}\right)\right\}=\frac{1}{2}<\frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})},$$

where the last inequality follows from $\bar{I}(\bar{\mu})+\underline{I}(\bar{\mu})>0$ and $1>\Phi(\|w\|_\Sigma)>\frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})}$. Furthermore, by applying the chain rule, $\frac{\partial f(\rho)}{\partial\rho}<0$. Therefore, $f(\cdot)$ is strictly decreasing in

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

$(0,\infty)$. We conclude that there must exist some unique $\rho^* > 0$ such that $f(\rho^*) = \frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})}$, or equivalently, $\int_0^1 \Phi\left(\frac{2\rho^* x - \rho^* - \|w\|_\Sigma^2}{\|w\|_\Sigma}\right) dx = \frac{-\underline{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})}$, which implies $\mathbb{E}_{\bar{\mu}}[d^*_{\text{linear}}] = \frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})}$.

Finally, we verify that $\mathbb{E}_{\bar{\mu}}[d^*_{\text{step}}] = \frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})}$. For any $\beta \in \left(0, \frac{1}{2}\right)$, consider the following step function rule:

$$d_{\text{step},\beta} := \left(\frac{1}{2} - \beta\right)\mathbf{1}\left\{w^\top Y < 0\right\} + \left(\frac{1}{2} + \beta\right)\mathbf{1}\left\{w^\top Y \geq 0\right\}.$$

One then has

$$\begin{aligned}\mathbb{E}_{\bar{\mu}}[d_{\text{step},\beta}] &= \left(\frac{1}{2} - \beta\right)\Phi\left(-\|w\|_\Sigma\right) + \left(\frac{1}{2} + \beta\right)\left(1 - \Phi\left(-\|w\|_\Sigma\right)\right) \\ &= \frac{1}{2} + \beta\left(2\Phi\left(\|w\|_\Sigma\right) - 1\right).\end{aligned}$$

Setting $\mathbb{E}_{\bar{\mu}}[d_{\text{step},\beta^*}] = \frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})}$ yields $\beta^* = \frac{\frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})} - \frac{1}{2}}{2\Phi(\|w\|_\Sigma)-1}$. As $\bar{I}(\bar{\mu}) + \underline{I}(\bar{\mu}) > 0$ and (3.3) holds, $\frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})} > \frac{1}{2}$ and therefore $\beta^* > 0$. Furthermore, $\beta^* < \frac{1}{2}$ holds due to (3.3) as well. Since $d^*_{\text{step}} = d_{\text{step},\beta^*}$, we conclude that $\mathbb{E}_{\bar{\mu}}[d^*_{\text{step}}] = \frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})}$.

A.1.3. *Statement (iii).* For any rule of form $d_{w,c}(Y) := \mathbf{1}\{w^\top Y \geq c\}$, where $w = \Sigma^{-1}\bar{\mu}$ and $c \in \mathbb{R}$, we may calculate

$$\mathbb{E}_{\bar{\mu}}[d_{w,c}(Y)] = 1 - \Phi\left(\frac{c - w^\top \bar{\mu}}{\sqrt{w^\top \Sigma w}}\right)$$

and (due to Lemma D.1, recalling $\|w\|_\Sigma^2 = w^\top \Sigma w = \bar{\mu}^\top \Sigma^{-1}\bar{\mu}$)

$$\begin{aligned}g(c) := \sup_{\pi \in \Gamma} r(d_{w,c}, \pi) = \frac{1}{2}\max\left\{\bar{I}(\bar{\mu})\Phi\left(\frac{-\|w\|_\Sigma^2 + c}{\|w\|_\Sigma}\right), -\underline{I}(\bar{\mu})\Phi\left(\frac{\|w\|_\Sigma^2 - c}{\|w\|_\Sigma}\right)\right\} \\ + \frac{1}{2}\max\left\{-\underline{I}(\bar{\mu})\Phi\left(\frac{\|w\|_\Sigma^2 + c}{\|w\|_\Sigma}\right), \bar{I}(\bar{\mu})\Phi\left(\frac{-\|w\|_\Sigma^2 - c}{\|w\|_\Sigma}\right)\right\}.\end{aligned}$$

Lemma B.5 shows that g is decreasing on $[0, c^*]$ and increasing on $[c^*, \infty)$, implying that the optimal threshold rule is d_{w,c^*} when $c \in [0,\infty)$. By symmetry, $d_{w,-c^*}$ is optimal when $c \in (-\infty, 0]$, and $d_{w,-c^*}$ and d_{w,c^*} share the same worst-case expected regret.

A.1.4. *Statement (iv).* In case (3.3) and when $n > 1$, there exists $\dot{\mu} \neq \mathbf{0}$ such that $\dot{\mu}^\top \Sigma^{-1}\bar{\mu} = 0$, that is, $\dot{\mu}$ is orthogonal to $\Sigma^{-1}\bar{\mu}$. For any $t \in \mathbb{R}$, let

$$d_{w_t,0}(Y) = \mathbf{1}\{w_t^\top Y \geq 0\}, \text{ where } w_t = \Sigma^{-1}(t\bar{\mu} + (1-t)\dot{\mu}). \tag{A.5}$$

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

Lemma B.6 shows that when $t = t^*$, $\mathbb{E}_{\tilde{\mu}}[d_{w_{t^*},0}(Y)] = \frac{\bar{I}(\tilde{\mu})}{\bar{I}(\tilde{\mu})-\underline{I}(\tilde{\mu})}$. By symmetry, one then also has $\mathbb{E}_{-\tilde{\mu}}[d_{w_{t^*},c}(Y)] = \frac{-\underline{I}(\tilde{\mu})}{\bar{I}(\tilde{\mu})-\underline{I}(\tilde{\mu})}$. Applying statement (ii) yields that $d_{w_{t^*},0}$ is MMR optimal.

A.2. Proof of Corollary 1

In Example 1, $\bar{I}(\tilde{\mu}) = \tilde{\mu} + k$, $\underline{I}(\tilde{\mu}) = \tilde{\mu} - k$, $\Sigma = \sigma^2$. The results of the corollary follow directly from Theorem 1(i)–(iii).

A.3. Proof of Theorem 2

Recall that the Γ-PER optimal rule solves

$$\inf_{a\in[0,1]} \sup_{\pi\in\Gamma} \int_{\tilde{\theta}\in\Theta} L(a,\tilde{\theta}) d\pi_{\theta|Y}(\tilde{\theta}), \quad \forall Y \in \mathbb{R}^n,$$

where $L(a,\theta) = U(\theta)(\mathbf{1}\{U(\theta) \geq 0\} - a)$ and $\pi_{\theta|Y}$ is the posterior distribution of θ given Y. If randomization is not allowed, the Γ-PER optimal rule solves

$$\inf_{a\in\{0,1\}} \sup_{\pi\in\Gamma} \int_{\tilde{\theta}\in\Theta} L(a,\tilde{\theta}) d\pi_{\theta|Y}(\tilde{\theta}), \quad \forall Y \in \mathbb{R}^n.$$

Statement (i) then follows from Lemma C.1; statement (ii) follows from Lemma C.2.

A.4. Proof of Corollary 2

Directly follows from Theorem 2.

A.5. Proof of Corollary 3

(i) "If": If $\underline{I}(\tilde{\mu}) \geq 0$, then Theorems 1(i) and 2(i) apply and establish that $d^*_{w,0}$ is both the unique Γ-MMR and the unique Γ-PER optimal rule.

"Only if": If $\underline{I}(\tilde{\mu}) < 0$, then d^*_{PER} is uniquely Γ-PER optimal by Theorem 2(i). If condition (3.2) holds as well, then Theorem 1(i) implies that $d^*_{w,0}$ is uniquely Γ-MMR optimal; hence, Γ-MMR and Γ-PER optimal rules disagree. If condition (3.3) applies, then, by Theorem 1(ii), a Γ-MMR optimal rule must satisfy (3.4) and (3.5). But d^*_{PER} can be written as

$$d^*_{\text{PER}} = d_{\text{step},\beta_{\text{PER}}} = \begin{cases} \frac{1}{2} - \beta_{\text{PER}}, & \text{if } w^\top Y < 0, \\ \frac{1}{2} + \beta_{\text{PER}}, & \text{if } w^\top Y \geq 0, \end{cases}$$

where $\beta_{\text{PER}} = \frac{\bar{I}(\tilde{\mu})}{\bar{I}(\tilde{\mu})-\underline{I}(\tilde{\mu})} - \frac{1}{2}$. Step 4 in the proof of Theorem 1(ii) shows that the only Γ-MMR optimal rule of form $d_{\text{step},\beta}$ is $d^*_{\text{step}} \neq d^*_{\text{PER}}$.

(ii) "If": When randomization is not allowed, Theorem 2(ii) shows that $d^*_{w,0}$ is always Γ-PER optimal. Under condition (3.17), Theorem 1(i) establishes that $d^*_{w,0}$ is Γ-MMR optimal as well.

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

"Only if": Again, $d^*_{w,0}$ is always Γ-PER optimal when randomization is not allowed. If (3.17) fails (i.e., (3.3) holds) and $n > 1$, it follows by Theorem 1(iv) that the randomized rule $d^*_{w_{I^*},0}$ is Γ-MMR optimal. If $n = 1$, Theorem 1(iii) implies that $d^*_{w,0}$ is not optimal even among linear threshold rules. Therefore, when (3.3) holds, Γ-MMR and Γ-PER optimal rules disagree with and without randomization.

B. TECHNICAL LEMMAS SUPPORTING EX-ANTE ANALYSIS

LEMMA B.1. *The Bayes rule supported by prior (A.1) is* $d^*_{w,0}$.

Proof. Note $\bar{I}(\bar{\mu}) > 0$ due to $\bar{I}(\bar{\mu}) + \underline{I}(\bar{\mu}) > 0$ and $\underline{I}(-\bar{\mu}) = -\bar{I}(\bar{\mu})$ by Lemma D.1. Given π^*, the Bayes optimal rule must solve the posterior problem

$$\min_{a\in[0,1]} \int_{\tilde{\theta}\in\Theta} L(a,\tilde{\theta}) d\pi^*_{\theta|Y}(\tilde{\theta}),$$

$$\int_{\tilde{\theta}\in\Theta} L(a,\tilde{\theta}) d\pi^*_{\theta|Y}(\tilde{\theta}) \propto \bar{I}(\bar{\mu})(1-a)\cdot\frac{1}{2}\cdot f(Y|\bar{\mu}) + \underline{I}(-\bar{\mu})(-a)\cdot\frac{1}{2}\cdot f(Y|-\bar{\mu}),$$

where $f(Y|\bar{\mu})$ and $f(Y|-\bar{\mu})$ are the likelihood of Y at $\bar{\mu}$ and $-\bar{\mu}$. This problem is equivalent to

$$\min_{a\in[0,1]} \bar{I}(\bar{\mu}) f(Y|\bar{\mu}) + a\underbrace{\bar{I}(\bar{\mu})}_{>0}(f(Y|-\bar{\mu}) - f(Y|\bar{\mu})).$$

Since $\bar{I}(\bar{\mu}) > 0$, the unique Bayes optimal rule is $\mathbf{1}\{f(Y|-\bar{\mu}) - f(Y|\bar{\mu}) \le 0\}$, which is equivalent to $d^*_{w,0}$ after further algebra. □

LEMMA B.2. *Consider decision rule* $d^*_{w,0}$ *and prior* π^* *defined in (A.1). If (3.2) holds, then*

$$\sup_{\pi\in\Gamma} r(d^*_{w,0},\pi) = r(d^*_{w,0},\pi^*). \tag{B.1}$$

Proof. As $w^\top Y \sim N(w^\top\mu, w^\top\Sigma w)$, algebra shows

$$\mathbb{E}_\mu[d^*_{w,0}] = \Phi\left(\frac{w^\top\mu}{\sqrt{w^\top\Sigma w}}\right)$$

for all $\mu \in M$. In particular,

$$\mathbb{E}_{\bar{\mu}}[d^*_{w,0}] = \Phi\left(\frac{w^\top\bar{\mu}}{\sqrt{w^\top\Sigma w}}\right), \quad \mathbb{E}_{-\bar{\mu}}[d^*_{w,0}] = \Phi\left(-\frac{w^\top\bar{\mu}}{\sqrt{w^\top\Sigma w}}\right).$$

It follows that

$$\sup_{\pi\in\Gamma} r(d^*_{w,0},\pi) = \frac{1}{2}\max\left\{\bar{I}(\bar{\mu})\Phi\left(-\frac{w^\top\bar{\mu}}{\sqrt{w^\top\Sigma w}}\right), -\underline{I}(\bar{\mu})\Phi\left(\frac{w^\top\bar{\mu}}{\sqrt{w^\top\Sigma w}}\right)\right\}$$
$$+\frac{1}{2}\max\left\{\bar{I}(-\bar{\mu})\Phi\left(\frac{w^\top\bar{\mu}}{\sqrt{w^\top\Sigma w}}\right), -\underline{I}(-\bar{\mu})\Phi\left(-\frac{w^\top\bar{\mu}}{\sqrt{w^\top\Sigma w}}\right)\right\},$$

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

and

$$r(d^*_{w,0},\pi^*) = \frac{1}{2}\bar{I}(\bar{\mu})\Phi\left(-\frac{w^\top\bar{\mu}}{\sqrt{w^\top\Sigma w}}\right) - \frac{1}{2}\underline{I}(-\bar{\mu})\Phi\left(-\frac{w^\top\bar{\mu}}{\sqrt{w^\top\Sigma w}}\right).$$

So, (B.1) holds as long as

(i) $\bar{I}(\bar{\mu})\Phi\left(-\frac{w^\top\bar{\mu}}{\sqrt{w^\top\Sigma w}}\right) \geq -\underline{I}(\bar{\mu})\Phi\left(\frac{w^\top\bar{\mu}}{\sqrt{w^\top\Sigma w}}\right)$,
(ii) $\bar{I}(-\bar{\mu})\Phi\left(\frac{w^\top\bar{\mu}}{\sqrt{w^\top\Sigma w}}\right) \leq -\underline{I}(-\bar{\mu})\Phi\left(-\frac{w^\top\bar{\mu}}{\sqrt{w^\top\Sigma w}}\right)$,

both of which are the same as (3.2) after further algebra, recalling that $\bar{I}(-\bar{\mu}) = -\underline{I}(\bar{\mu})$ and $-\underline{I}(-\bar{\mu}) = \bar{I}(\bar{\mu})$ by Lemma D.1, $w = \Sigma^{-1}\bar{\mu}$, and $\|w\|_\Sigma = \sqrt{\bar{\mu}^\top\Sigma^{-1}\bar{\mu}}$. □

LEMMA B.3. *Consider the prior π^* such that $\mu \sim \text{unif}(\{-\bar{\mu},\bar{\mu}\})$, and when $\mu = \bar{\mu}$,*

$$U(\theta) = \begin{cases} \bar{I}(\bar{\mu}), & \text{w.p. } p_1, \\ \underline{I}(\bar{\mu}), & \text{w.p. } 1-p_1, \end{cases}$$

where p_1 is such that $p_1\bar{I}(\bar{\mu}) + (1-p_1)\underline{I}(\bar{\mu}) = 0$, and when $\mu = -\bar{\mu}$,

$$U(\theta) = \begin{cases} \bar{I}(-\bar{\mu}), & \text{w.p. } p_2, \\ \underline{I}(-\bar{\mu}), & \text{w.p. } 1-p_2, \end{cases}$$

where p_2 is such that $p_2\bar{I}(-\bar{\mu}) + (1-p_2)\underline{I}(-\bar{\mu}) = 0$. Given this prior, any decision rule is Bayes optimal under (3.3).

Proof. As (3.3) holds, we have $\underline{I}(\bar{\mu}) < 0 < \bar{I}(\bar{\mu})$. Analogous to Lemma B.1, a Bayes rule must solve

$$\min_{a\in[0,1]} \quad \frac{1}{2}\cdot f(Y \mid \bar{\mu})\left[p_1\bar{I}(\bar{\mu})(1-a) + (1-p_1)(-\underline{I}(\bar{\mu}))a\right]$$
$$+\frac{1}{2}\cdot f(Y \mid -\bar{\mu})\left[p_2\bar{I}(-\bar{\mu})(1-a) + (1-p_2)(-\underline{I}(-\bar{\mu}))a\right].$$

Since $p_1\bar{I}(\bar{\mu}) + (1-p_1)\underline{I}(\bar{\mu}) = 0$ and $p_2\bar{I}(-\bar{\mu}) + (1-p_2)\underline{I}(-\bar{\mu}) = 0$, the objective is constant in a, hence the claim. □

LEMMA B.4. *Consider the prior π^* such that $\mu \sim \text{unif}(\{-\bar{\mu},\bar{\mu}\})$, when $\mu = \bar{\mu}$,*

$$U(\theta) = \begin{cases} \bar{I}(\bar{\mu}), & \text{w.p. } p_1, \\ \underline{I}(\bar{\mu}), & \text{w.p. } 1-p_1, \end{cases}$$

where p_1 is such that $p_1\bar{I}(\bar{\mu}) + (1-p_1)\underline{I}(\bar{\mu}) = 0$, and when $\mu = -\bar{\mu}$,

$$U(\theta) = \begin{cases} \bar{I}(-\bar{\mu}), & \text{w.p. } p_2, \\ \underline{I}(-\bar{\mu}), & \text{w.p. } 1-p_2, \end{cases}$$

where p_2 is such that $p_2\bar{I}(-\bar{\mu}) + (1-p_2)\underline{I}(-\bar{\mu}) = 0$. Then, when (3.3) holds, we have

$$\sup_{\pi\in\Gamma} r(d,\pi) = r(d,\pi^*)$$

for any decision rule $d \in \mathcal{D}_n$ such that (A.2) is true.

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

Proof. Note again that (3.3) implies $\underline{I}(\bar{\mu}) < 0$ and $\bar{I}(\bar{\mu}) > 0$. For any decision rule $d \in \mathcal{D}_n$ such that (A.2) is true, algebra shows

$$\bar{I}(\bar{\mu})(1 - \mathbb{E}_{\bar{\mu}}[d]) = -\underline{I}(\bar{\mu})\mathbb{E}_{\bar{\mu}}[d] = -\frac{\bar{I}(\bar{\mu})\underline{I}(\bar{\mu})}{\bar{I}(\bar{\mu}) - \underline{I}(\bar{\mu})},$$

$$\bar{I}(-\bar{\mu})(1 - \mathbb{E}_{-\bar{\mu}}[d]) = -\underline{I}(-\bar{\mu})\mathbb{E}_{-\bar{\mu}}[d] = -\frac{\bar{I}(\bar{\mu})\underline{I}(\bar{\mu})}{\bar{I}(\bar{\mu}) - \underline{I}(\bar{\mu})},$$

implying $\sup_{\pi \in \Gamma} r(d, \pi) = -\frac{\bar{I}(\bar{\mu})\underline{I}(\bar{\mu})}{\bar{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}$. Meanwhile, for any $d \in \mathcal{D}_n$ such that (A.2) is true, we may also calculate the following that yields the desired conclusion:

$$r(d, \pi^*) = \frac{1}{2}p_1\bar{I}(\bar{\mu}) + \frac{1}{2}p_2\bar{I}(-\bar{\mu}) = -\frac{\bar{I}(\bar{\mu})\underline{I}(\bar{\mu})}{\bar{I}(\bar{\mu}) - \underline{I}(\bar{\mu})} = \sup_{\pi \in \Gamma} r(d, \pi).$$

□

LEMMA B.5. *In case (3.3), the function*

$$g(c) = \frac{1}{2}\max\left\{\bar{I}(\bar{\mu})\Phi\left(\frac{-\|w\|_\Sigma^2 + c}{\|w\|_\Sigma}\right), -\underline{I}(\bar{\mu})\Phi\left(\frac{\|w\|_\Sigma^2 - c}{\|w\|_\Sigma}\right)\right\}$$
$$+ \frac{1}{2}\max\left\{-\underline{I}(\bar{\mu})\Phi\left(\frac{\|w\|_\Sigma^2 + c}{\|w\|_\Sigma}\right), \bar{I}(\bar{\mu})\Phi\left(\frac{-\|w\|_\Sigma^2 - c}{\|w\|_\Sigma}\right)\right\}$$

is decreasing in $[0, c^*]$ *and increasing in* $[c^*, \infty)$.

Proof. Note when $c \in [0, c^*)$ and (3.3) holds, we have $-\underline{I}(\bar{\mu}) > 0$,

$$-\underline{I}(\bar{\mu})\Phi\left(\frac{\|w\|_\Sigma^2 - c}{\|w\|_\Sigma}\right) > \bar{I}(\bar{\mu})\Phi\left(\frac{-\|w\|_\Sigma^2 + c}{\|w\|_\Sigma}\right),$$
$$-\underline{I}(\bar{\mu})\Phi\left(\frac{\|w\|_\Sigma^2 + c}{\|w\|_\Sigma}\right) > \bar{I}(\bar{\mu})\Phi\left(\frac{-\|w\|_\Sigma^2 - c}{\|w\|_\Sigma}\right),$$

and therefore,

$$g(c) = -\frac{\underline{I}(\bar{\mu})}{2}\left[\Phi\left(\frac{\|w\|_\Sigma^2 - c}{\|w\|_\Sigma}\right) + \Phi\left(\frac{\|w\|_\Sigma^2 + c}{\|w\|_\Sigma}\right)\right],$$

and

$$\frac{\partial g(c)}{\partial c} = -\frac{\underline{I}(\bar{\mu})}{2\|w\|_\Sigma}\left[\phi\left(\frac{\|w\|_\Sigma^2 + c}{\|w\|_\Sigma}\right) - \phi\left(\frac{\|w\|_\Sigma^2 - c}{\|w\|_\Sigma}\right)\right] < 0.$$

When $c = c^*$, note $\bar{I}(\bar{\mu})\Phi\left(\frac{-\|w\|_\Sigma^2 + c^*}{\|w\|_\Sigma}\right) = -\underline{I}(\bar{\mu})\Phi\left(\frac{\|w\|_\Sigma^2 - c^*}{\|w\|_\Sigma}\right)$ and

$$g(c^*) = -\frac{\underline{I}(\bar{\mu})}{2}\left[\Phi\left(\frac{\|w\|_\Sigma^2 - c^*}{\|w\|_\Sigma}\right) + \Phi\left(\frac{\|w\|_\Sigma^2 + c^*}{\|w\|_\Sigma}\right)\right].$$

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

When $c \in (c^*, \infty)$, note

$$-\underline{I}(\bar{\mu})\Phi\left(\frac{\|w\|_\Sigma^2 - c}{\|w\|_\Sigma}\right) < \bar{I}(\bar{\mu})\Phi\left(\frac{-\|w\|_\Sigma^2 + c}{\|w\|_\Sigma}\right),$$
$$-\underline{I}(\bar{\mu})\Phi\left(\frac{\|w\|_\Sigma^2 + c}{\|w\|_\Sigma}\right) > \bar{I}(\bar{\mu})\Phi\left(\frac{-\|w\|_\Sigma^2 - c}{\|w\|_\Sigma}\right),$$

so that

$$g(c) = \frac{1}{2}\left\{\bar{I}(\bar{\mu})\Phi\left(\frac{-\|w\|_\Sigma^2 + c}{\|w\|_\Sigma}\right) - \underline{I}(\bar{\mu})\Phi\left(\frac{\|w\|_\Sigma^2 + c}{\|w\|_\Sigma}\right)\right\},$$

and

$$\frac{\partial g(c)}{\partial c} = \frac{1}{2\|w\|_\Sigma}\left\{\bar{I}(\bar{\mu})\phi\left(\frac{-\|w\|_\Sigma^2 + c}{\|w\|_\Sigma}\right) - \underline{I}(\bar{\mu})\phi\left(\frac{\|w\|_\Sigma^2 + c}{\|w\|_\Sigma}\right)\right\} > 0.$$

Having shown that $\frac{\partial g(c)}{\partial c} < 0$ when $c \in [0, c^*)$ and $\frac{\partial g(c)}{\partial c} > 0$ when $c \in (c^*, \infty)$, and since $g(c)$ is continuous at $c = c^*$, we conclude that g is decreasing in $[0, c^*]$ and increasing in $[c^*, \infty)$. □

LEMMA B.6. *In case (3.3) and when $n > 1$, let*

$$d_{w_t,0}(Y) = \mathbf{1}\{w_t^\top Y \geq 0\}, \text{ where } w_t = \Sigma^{-1}(t\bar{\mu} + (1-t)\dot{\mu}), \dot{\mu} \neq \mathbf{0},$$

and $\dot{\mu}^\top \Sigma^{-1}\bar{\mu} = 0$. Then, $\mathbb{E}_{\bar{\mu}}[d_{w_{t^*},0}(Y)] = \frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}$, where t^* is defined in (3.9).

Proof. Write $f(t) := \mathbb{E}_{\bar{\mu}}[d_{w_t,c}(Y)] = \Phi\left(\frac{w_t^\top \bar{\mu}}{\sqrt{w_t^\top \Sigma w_t}}\right)$ and $k^* := \frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}$. Then, it suffices to show that $f(t^*) = k^*$, which is equivalent to showing

$$\frac{\left(w_{t^*}^\top \bar{\mu}\right)^2}{w_{t^*}^\top \Sigma w_{t^*}} = \left(\Phi^{-1}\left(k^*\right)\right)^2.$$

Furthermore, note

$$w_{t^*}^\top \bar{\mu} = (t^*\bar{\mu}^\top \Sigma^{-1} + (1-t^*)\dot{\mu}^\top \Sigma^{-1})\bar{\mu} = t^*\bar{\mu}^\top \Sigma^{-1}\bar{\mu} + (1-t^*)\underbrace{\dot{\mu}^\top \Sigma^{-1}\bar{\mu}}_{=0} = t^*\|w\|_\Sigma^2$$

and

$$\begin{aligned} w_{t^*}^\top \Sigma w_{t^*} &= (t^*\bar{\mu}^\top \Sigma^{-1} + (1-t^*)\dot{\mu}^\top \Sigma^{-1})\Sigma(t^*\Sigma^{-1}\bar{\mu} + (1-t^*)\Sigma^{-1}\dot{\mu}) \\ &= (t^*)^2\bar{\mu}^\top \Sigma^{-1}\bar{\mu} + 2t^*(1-t^*)\underbrace{\bar{\mu}^\top \Sigma^{-1}\dot{\mu}}_{=0} + (1-t^*)^2\underbrace{\dot{\mu}^\top \Sigma^{-1}\dot{\mu}}_{:=\|\Sigma^{-1}\dot{\mu}\|_\Sigma^2} \\ &= (t^*)^2\|w\|_\Sigma^2 + (1-t^*)^2\|\Sigma^{-1}\dot{\mu}\|_\Sigma^2. \end{aligned}$$

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

Therefore, substituting in these expressions, t^* should be such that

$$\frac{\left(\Phi^{-1}(k^*)\right)^2}{\|w\|_\Sigma^2} = \frac{(t^*)^2\|w\|_\Sigma^2}{(t^*)^2\|w\|_\Sigma^2+(1-t^*)^2\|\Sigma^{-1}\dot{\mu}\|_\Sigma^2}.$$

As $s^* := \frac{\left(\Phi^{-1}(k^*)\right)^2}{\|w\|_\Sigma^2} \in (0,1)$ due to (3.3), we can solve the above equation for t^*, and after lengthy algebra, find that

$$t^* = \frac{1}{1 \pm \sqrt{\frac{(1-s^*)}{s^*} \cdot \frac{\|w\|_\Sigma^2}{\|\Sigma^{-1}\dot{\mu}\|_\Sigma^2}}}.$$

□

C. TECHNICAL LEMMAS SUPPORTING EX-POST ANALYSIS

LEMMA C.1. *Suppose all the conditions of Theorem 1 hold. If $\underline{I}(\bar{\mu}) < 0 < \bar{I}(\bar{\mu})$, then the unique Γ-PER optimal rule is*

$$d^*_{\mathrm{PER}}(Y) = \begin{cases} \frac{-\underline{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})}, & \text{if } w^\top Y < 0, \\ \frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})}, & \text{if } w^\top Y \geq 0. \end{cases}$$

Otherwise, $d^*_{w,0}$ is Γ-PER optimal.

Proof. Let $f(Y \mid \mu)$ be the likelihood of Y. Then, for each action $a \in [0,1]$, we have

$$\begin{aligned}
&\sup_{\pi\in\Gamma}\int_{\tilde{\theta}\in\Theta} L(a,\tilde{\theta})d\pi_{\theta|Y}(\tilde{\theta}) \\
&\quad\propto \frac{1}{2}\left(\sup_{\pi_{\theta|\bar{\mu}}}\int_{\tilde{\theta}\in\Theta} U(\tilde{\theta})(\mathbf{1}\{U(\tilde{\theta})\geq 0\}-a)d\pi_{\theta|\bar{\mu}}(\tilde{\theta})\right)f(Y\mid\bar{\mu}) \\
&\qquad+\frac{1}{2}\left(\sup_{\pi_{\theta|-\bar{\mu}}}\int_{\tilde{\theta}\in\Theta} U(\tilde{\theta})(\mathbf{1}\{U(\tilde{\theta})\geq 0\}-a)d\pi_{\theta|-\bar{\mu}}(\tilde{\theta})\right)f(Y\mid-\bar{\mu}) \\
&\quad= \frac{1}{2}f(Y\mid\bar{\mu})\max\left\{\bar{I}(\bar{\mu})(1-a), -\underline{I}(\bar{\mu})a\right\} \\
&\qquad+\frac{1}{2}f(Y\mid-\bar{\mu})\max\left\{\bar{I}(-\bar{\mu})(1-a), -\underline{I}(-\bar{\mu})a\right\}.
\end{aligned}$$

Therefore, when randomization is allowed, the Γ-PER optimal rule can be found by solving

$$\inf_{a\in[0,1]} V_\Gamma(a), \tag{C.1}$$

where

$$V_\Gamma(a) := f(Y\mid\bar{\mu})\max\left\{\bar{I}(\bar{\mu})(1-a), -\underline{I}(\bar{\mu})a\right\} + f(Y\mid-\bar{\mu})\max\left\{\bar{I}(\bar{\mu})a, -\underline{I}(\bar{\mu})(1-a)\right\},$$

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

since $\overline{I}(-\bar{\mu}) = -\underline{I}(\bar{\mu})$ and $-\underline{I}(-\bar{\mu}) = \overline{I}(\bar{\mu})$ by Lemma D.1. Note that, for each action $a \in [0,1]$,

$$\overline{I}(\bar{\mu})(1-a) \geq -\underline{I}(\bar{\mu})a \Leftrightarrow a \leq \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})},$$
$$\overline{I}(\bar{\mu})a \geq -\underline{I}(\bar{\mu})(1-a) \Leftrightarrow a \geq \frac{-\underline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})},$$

where $\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu}) \geq 0$ and $\overline{I}(\bar{\mu}) > 0$ due to our normalization $\overline{I}(\bar{\mu}) + \underline{I}(\bar{\mu}) > 0$. The conclusion follows from discussing the following two cases.

Case 1: $\underline{I}(\bar{\mu}) \geq 0$.

In this case, $\frac{-\underline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})} \leq a \leq \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})}$ for all $a \in [0,1]$. Thus, (C.1) becomes

$$\inf_{a\in[0,1]} \underbrace{\overline{I}(\bar{\mu})}_{>0} \{f(Y \mid \bar{\mu}) + [f(Y \mid -\bar{\mu}) - f(Y \mid \bar{\mu})]a\}.$$

The Γ-PER optimal rule is $\mathbf{1}\{f(Y \mid -\bar{\mu}) - f(Y \mid \bar{\mu}) \leq 0\}$, which is $d^*_{w,0}$ after further algebra.

Case 2: $\underline{I}(\bar{\mu}) < 0 < \overline{I}(\bar{\mu})$.

In this case, we have (again, note $\overline{I}(\bar{\mu}) + \underline{I}(\bar{\mu}) > 0$):

$$0 < \frac{-\underline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})} < \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu}) - \underline{I}(\bar{\mu})} < 1.$$

Therefore:

- When $a \in \left[0, \frac{-\underline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})}\right)$,

$$\begin{aligned} V_\Gamma(a) &= f(Y \mid \bar{\mu})\overline{I}(\bar{\mu})(1-a) - f(Y \mid -\bar{\mu})\underline{I}(\bar{\mu})(1-a) \\ &= \underbrace{\left\{f(Y \mid \bar{\mu})\overline{I}(\bar{\mu}) + f(Y \mid -\bar{\mu})\left(-\underline{I}(\bar{\mu})\right)\right\}}_{>0}(1-a). \end{aligned}$$

  Thus, $V_\Gamma(\cdot)$ is strictly decreasing in $a \in \left[0, \frac{-\underline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})}\right)$.

- When $a \in \left[\frac{-\underline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})}, \frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})}\right)$,

$$\begin{aligned} V_\Gamma(a) &= f(Y \mid \bar{\mu})\overline{I}(\bar{\mu})(1-a) + f(Y \mid -\bar{\mu})\overline{I}(\bar{\mu})a \\ &= f(Y \mid \bar{\mu})\overline{I}(\bar{\mu}) + \underbrace{\overline{I}(\bar{\mu})}_{>0}(f(Y \mid -\bar{\mu}) - f(Y \mid \bar{\mu}))\,a. \end{aligned}$$

  Thus, if $f(Y \mid -\bar{\mu}) > f(Y \mid \bar{\mu})$, $V_\Gamma(\cdot)$ is strictly increasing in a, and if $f(Y \mid -\bar{\mu}) < f(Y \mid \bar{\mu})$, $V_\Gamma(\cdot)$ is strictly decreasing in a. When $f(Y \mid -\bar{\mu}) = f(Y \mid \bar{\mu})$, $V_\Gamma(\cdot)$ is constant in a.

- When $a \in \left[\frac{\overline{I}(\bar{\mu})}{\overline{I}(\bar{\mu})-\underline{I}(\bar{\mu})}, 1\right]$,

$$V_\Gamma(a) = \underbrace{(f(Y \mid -\bar{\mu})\overline{I}(\bar{\mu}) - f(Y \mid \bar{\mu})\underline{I}(\bar{\mu}))}_{>0}\,a,$$

  implying that $V_\Gamma(\cdot)$ is strictly increasing in a.

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

In sum, if $f(Y \mid -\bar{\mu}) > f(Y \mid \bar{\mu})$, $V_\Gamma(\cdot)$ is first strictly decreasing and then increasing, with the minimum achieved at $a = \frac{-\underline{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})}$; if $f(Y \mid -\bar{\mu}) < f(Y \mid \bar{\mu})$, $V_\Gamma(\cdot)$ is first strictly decreasing and increasing with the minimum achieved at $a = \frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})}$; if $f(Y \mid -\bar{\mu}) = f(Y \mid \bar{\mu})$ (which only happens in a null set), $V_\Gamma(\cdot)$ achieves minimum at any point between $[\frac{-\underline{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})}, \frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})}]$. Therefore, the unique Γ-PER optimal rule is

$$d(Y) = \begin{cases} \frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})}, & \text{if } f(Y \mid -\bar{\mu}) \le f(Y \mid \bar{\mu}) \Longleftrightarrow \bar{\mu}^\top \Sigma^{-1} Y \ge 0, \\ \frac{-\underline{I}(\bar{\mu})}{\bar{I}(\bar{\mu})-\underline{I}(\bar{\mu})} & \text{if } f(Y \mid -\bar{\mu}) > f(Y \mid \bar{\mu}) \Longleftrightarrow \bar{\mu}^\top \Sigma^{-1} Y < 0. \end{cases}$$

□

LEMMA C.2. *Suppose all conditions of Theorem 1 hold true. Then, $d^*_{w,0}$ is always the Γ-PER optimal non-randomized rule.*

Proof. When randomization is not allowed, we need to solve (by a similar derivation to Lemma C.1) $\inf_{a\in\{0,1\}} V_\Gamma(a)$, where $V_\Gamma(a)$ is defined in (C.1). As

$$V_\Gamma(1) = -f(Y \mid \bar{\mu})\underline{I}(\bar{\mu}) + f(Y \mid -\bar{\mu})\bar{I}(\bar{\mu}),$$
$$V_\Gamma(0) = f(Y \mid \bar{\mu})\bar{I}(\bar{\mu}) - f(Y \mid -\bar{\mu})\underline{I}(\bar{\mu}),$$

the optimal action would be $a = 1$ if and only if $V_\Gamma(1) \le V_\Gamma(0)$, which is equivalent to $d^*_{w,0}$ after further algebra. □

D. ADDITIONAL RESULTS

LEMMA D.1. *Consider a treatment choice problem with welfare function (2.1), statistical model (2.3), and a set of priors (2.9), which satisfy Assumption 1. Then, the following statements hold:*

(i) $\bar{I}(-\bar{\mu}) = -\underline{I}(\bar{\mu})$, $\underline{I}(-\bar{\mu}) = -\bar{I}(\bar{\mu})$;
(ii) $\bar{I}(-\bar{\mu}) + \underline{I}(-\bar{\mu}) = -(\underline{I}(\bar{\mu}) + \bar{I}(\bar{\mu}))$.

Proof. Statement (i). Here, we only show that $\bar{I}(-\bar{\mu}) = -\underline{I}(\bar{\mu})$. An analogous argument proves $\underline{I}(-\bar{\mu}) = -\bar{I}(\bar{\mu})$. By definition,

$$\bar{I}(-\bar{\mu}) = \sup_{\{\theta\in\Theta: m(\theta) = -\bar{\mu}\}} U(\theta).$$

Since Θ is centrosymmetric and both $U(\cdot)$ and $m(\cdot)$ are linear, we have

$$\begin{aligned} \sup_{\{\theta\in\Theta: m(\theta) = -\bar{\mu}\}} U(\theta) &= \sup_{\{-\theta\in\Theta: -m(-\theta) = -\bar{\mu}\}} -U(-\theta) \\ &= \sup_{\{\tilde{\theta}\in\Theta: -m(\tilde{\theta}) = -\bar{\mu}\}} -U(\tilde{\theta}) \\ &= -\left\{ \inf_{\{\tilde{\theta}\in\Theta: m(\tilde{\theta}) = \bar{\mu}\}} U(\tilde{\theta}) \right\} \\ &= -\underline{I}(\bar{\mu}) \end{aligned}$$

as required.

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

Statement (ii). Directly follows from summing both equalities in statement (i). □

LEMMA D.2. *In Example 1, suppose* $k > \sqrt{\frac{\pi}{2}}\sigma$. *Then, if* $\bar{\mu} > 0$ *is sufficiently small, the corresponding* Γ*-PER optimal rule is dominated in terms of profiled regret.*

Proof. Denote by $d^*_{MMR,\text{linear}}$ the least randomizing global MMR optimal rule derived in Montiel Olea et al. (2026). We aim to show that when $k > \sqrt{\frac{\pi}{2}}\sigma$ and $\bar{\mu} > 0$ is sufficiently small, the associated Γ-PER optimal rule d^*_{PER} is such that

$$\bar{R}(d^*_{\text{PER}}, \mu) \geq \bar{R}(d^*_{MMR,\text{linear}}, \mu), \text{ for all } \mu \geq 0, \tag{D.1}$$

with the inequality strict for all $\mu > 0$. A symmetry argument then immediately implies that d^*_{PER} is dominated in terms of profiled regret.

Step 1: Pick any $0 < \bar{\mu} < \sqrt{\frac{\pi}{2}}\sigma$. We show that $\bar{R}(d^*_{\text{PER}}, \mu) = (\mu + k)\left(1 - \mathbb{E}_\mu\left[d^*_{\text{PER}}(\hat{\mu})\right]\right)$ for all $\mu \geq 0$. By results in Montiel Olea et al. (2026, Appendix B.3.1), we may write the profiled regret of a rule d as

$$\bar{R}(d, \mu) = \begin{cases} (-\mu + k)\mathbb{E}_\mu\left[d(\hat{\mu})\right], & \text{if } \mu < -k, \\ \max\left\{(\mu + k)\left(1 - \mathbb{E}_\mu\left[d(\hat{\mu})\right]\right), (-\mu + k)\mathbb{E}_\mu\left[d(\hat{\mu})\right]\right\}, & \text{if } -k \leq \mu \leq k, \\ (\mu + k)(1 - \mathbb{E}_\mu\left[d(\hat{\mu})\right]), & \text{if } \mu > k. \end{cases}$$

Thus, it suffices to show

$$\max\left\{(\mu + k)\left(1 - \mathbb{E}_\mu\left[d^*_{\text{PER}}(\hat{\mu})\right]\right), (-\mu + k)\mathbb{E}_\mu\left[d^*_{\text{PER}}(\hat{\mu})\right]\right\} = (\mu + k)\left(1 - \mathbb{E}_\mu\left[d^*_{\text{PER}}(\hat{\mu})\right]\right)$$

for any $0 \leq \mu \leq k$. As $\sqrt{\frac{\pi}{2}}\sigma < k$, Corollary 2 implies that the corresponding Γ-PER rule for any $0 < \bar{\mu} < \sqrt{\frac{\pi}{2}}\sigma$ is

$$\begin{aligned} d^*_{\text{PER}}(\hat{\mu}) &= \frac{k + \overline{\mu}}{2k}\mathbf{1}\{\hat{\mu} \geq 0\} + \frac{k - \overline{\mu}}{2k}\mathbf{1}\{\hat{\mu} < 0\} \\ &= \frac{1}{2} + \frac{\overline{\mu}}{2k}\left(\mathbf{1}\{\hat{\mu} \geq 0\} - \mathbf{1}\{\hat{\mu} < 0\}\right), \end{aligned}$$

and further algebra shows $\mathbb{E}_\mu\left[d^*_{\text{PER}}(\hat{\mu})\right] = \frac{1}{2} + \frac{\overline{\mu}}{2k}\left(1 - 2\Phi\left(-\frac{\mu}{\sigma}\right)\right)$. Then,

$$(\mu + k)\left(1 - \mathbb{E}_\mu\left[d^*_{\text{PER}}(\hat{\mu})\right]\right) \geq (-\mu + k)\mathbb{E}_\mu\left[d^*_{\text{PER}}(\hat{\mu})\right]$$

if and only if

$$(\mu + k)\left(\frac{1}{2} - \frac{\overline{\mu}}{2k}\left(1 - 2\Phi\left(-\frac{\mu}{\sigma}\right)\right)\right) \geq (-\mu + k)\left(\frac{1}{2} + \frac{\overline{\mu}}{2k}\left(1 - 2\Phi\left(-\frac{\mu}{\sigma}\right)\right)\right),$$

which is equivalent to

$$\frac{\mu}{\left(1 - 2\Phi\left(-\frac{\mu}{\sigma}\right)\right)} \geq \overline{\mu}. \tag{D.2}$$

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

Note the left-hand side of (D.2) is increasing in μ, and $\lim_{\mu \downarrow 0}\left(\frac{\mu}{(1-2\Phi(-\frac{\mu}{\sigma}))}\right) = \frac{\sigma}{2\phi(0)} = \sigma\sqrt{\frac{\pi}{2}} > \overline{\mu}$. As $0 < \bar{\mu} < \sqrt{\frac{\pi}{2}}\sigma$, we conclude that (D.2) indeed holds, and it follows

$$\bar{R}(d^*_{\mathrm{PER}}, \mu) = (\mu + k)\left(1 - \mathbb{E}_\mu\left[d^*_{\mathrm{PER}}(\hat{\mu})\right]\right), \tag{D.3}$$

for all $\mu \geq 0$. As a result, we conclude that, for any $0 < \bar{\mu} < \sqrt{\frac{\pi}{2}}\sigma$, we have

$$\bar{R}(d^*_{\mathrm{PER}}, \mu) = (\mu + k)\left(\frac{1}{2} - \frac{\overline{\mu}}{2k}\left(1 - 2\Phi\left(-\frac{\mu}{\sigma}\right)\right)\right), \text{ for all } \mu \geq 0.$$

Step 2: We show that $\bar{R}(d^*_{\mathrm{PER}}, \mu)$ is strictly increasing in μ for any $0 < \bar{\mu} < \sqrt{\frac{\pi}{2}}\sigma$ small enough. To see this, note

$$\begin{aligned}\frac{\partial \bar{R}(d^*_{\mathrm{PER}}, \mu)}{\partial \mu} &= \frac{1}{2} - \frac{\overline{\mu}}{2k}\left(1 - 2\Phi\left(-\frac{\mu}{\sigma}\right)\right) \\ &\quad - (\mu + k)\frac{\overline{\mu}}{k}\frac{1}{\sigma}\phi\left(\frac{\mu}{\sigma}\right) \\ &= \frac{1}{2} - \frac{\overline{\mu}}{2k}\left(1 - 2\Phi\left(-\frac{\mu}{\sigma}\right) + 2(\mu + k)\frac{1}{\sigma}\phi\left(\frac{\mu}{\sigma}\right)\right),\end{aligned}$$

in which note (1): $0 < 1 - 2\Phi\left(-\frac{\mu}{\sigma}\right) \leq 1$ for all $\mu \geq 0$ and (2): $2(\mu + k)\frac{1}{\sigma}\phi\left(\frac{\mu}{\sigma}\right) > 0$ is first increasing and then decreasing with a unique and finite maximum $C_k > 0$. Therefore, we conclude that for any $0 < \overline{\mu} < \min\left(\frac{k}{(1+C_k)}, \sqrt{\frac{\pi}{2}}\sigma\right)$, we have $\frac{\partial \bar{R}(d^*_{\mathrm{PER}}, \mu)}{\partial \mu} > 0$ for all $\mu \geq 0$. That is, $\bar{R}(d^*_{\mathrm{PER}}, \mu)$ is strictly increasing in $\mu \in [0, \infty]$.

Step 3: Montiel Olea et al. (2026) have already shown that

$$\sup_{\mu \geq 0} \bar{R}(d^*_{MMR,\,\mathrm{linear}}, \mu) = \bar{R}(d^*_{MMR,\,\mathrm{linear}}, 0) = \frac{k}{2}.$$

Then, by the conclusion from step 3 and noting $\bar{R}(d^*_{\mathrm{PER}}, 0) = \frac{k}{2}$, we see that (D.1) indeed holds with the inequality strict for all $\mu > 0$, completing the proof. □

E. GENERAL NATURE OF OUR MAIN RESULTS

In the main text, we derived finite-sample Γ-MMR and -PER optimal rules for a class of priors such that $\pi_\mu \sim \text{unif}(\{-\overline{\mu}, \overline{\mu}\})$. In this section, we discuss the implications of our results when π_μ has a more general structure (e.g., with a continuous support).

It is relatively straightforward to extend our Γ-PER results (Theorem 2) to other forms of π_μ. Let

$$V_\Gamma(a) := V_\Gamma(a, Y) := \sup_{\pi \in \Gamma} \int_{\tilde{\theta} \in \Theta} L(a, \tilde{\theta}) d_{\theta|Y}(\tilde{\theta}).$$

The ex-post Γ-PER criterion aims to solve $\min_{a \in [0,1]} V_\Gamma(a)$. Given a general π_μ and the normal likelihood of Y, we can derive the posterior distribution of the reduced-form parameter μ given Y, written as $\pi_{\mu|Y}$, using the standard Bayes rule. In light of the structure

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

of the class of priors Γ, we see

$$V_\Gamma(a) = \int \sup_{\mu^* \in I(x)} \left\{\mu^* \left[\mathbf{1}\left\{\mu^* \geq 0\right\} - a\right]\right\} d\pi_{\mu|Y}(x)$$
$$= \int \max\left\{\bar{I}(x)(1-a), -\underline{I}(x)a\right\} d\pi_{\mu|Y}(x).$$

Note $\bar{I}(x)(1-a) \geq -\underline{I}(x)a$ if and only if $\bar{I}(x) \geq a\left(\bar{I}(x) - \underline{I}(x)\right)$, equivalent to (assuming $\bar{I}(x) - \underline{I}(x) > 0$ for all x in the support of $\pi_{\mu|Y}$) $a \leq \frac{\bar{I}(x)}{\bar{I}(x)-\underline{I}(x)}$. Thus, for each $a \in [0,1]$, we have

$$V_\Gamma(a) = \int \left[\mathbf{1}\left\{a \leq \frac{\bar{I}(x)}{\bar{I}(x)-\underline{I}(x)}\right\}\bar{I}(x)(1-a) - \underline{I}(x)a\mathbf{1}\left\{a > \frac{\bar{I}(x)}{\bar{I}(x)-\underline{I}(x)}\right\}\right] d\pi_{\mu|Y}(x)$$
$$= (1-a)\int_{\left\{x\in\mathbb{R}^n : \frac{\bar{I}(x)}{\bar{I}(x)-\underline{I}(x)} \geq a\right\}} \bar{I}(x) d\pi_{\mu|Y}(x) - a\int_{\left\{x\in\mathbb{R}^n : \frac{\bar{I}(x)}{\bar{I}(x)-\underline{I}(x)} < a\right\}} \underline{I}(x) d\pi_{\mu|Y}(x).$$

The key observation is that the value of $V_\Gamma(a)$ for each $a \in [0,1]$ can be numerically solved, although it may not have a closed-form solution. Therefore, in general, we can still find $d^*_{\text{PER}}(Y)$ by numerically solving $\min_{a\in[0,1]} V_\Gamma(a)$. Note as the solution of $\min_{a\in[0,1]} V_\Gamma(a)$ may or may not be at the corners $\{0,1\}$, it is in general *not* the case that the associated Γ-PER rule always randomizes.

In contrast, the Γ-MMR optimal rule is more difficult to find once we allow general forms of π_μ. Similar to the unconstrained MMR optimality problem, one often needs to resort to the "guess-and-verify" strategy by forming a least favorable prior—unlike the unconstrained scenario, for the Γ-MMR problem, the LFP has to come from Γ, which restricts Nature's strategy significantly. The particular form of $\pi_\mu \sim \text{unif}(\{-\overline{\mu}, \overline{\mu}\})$ makes our guess of the LFP much easier, while for other cases, it may be more difficult. That said, although the exact analytic form of Γ-MMR optimal rule is difficult to acquire for general π_μ, one may utilize recent computational techniques (Aradillas Fernández et al., 2025; Guggenberger and Huang, 2025) to numerically find an approximately Γ-MMR optimal rule. find an ε-Γ-MMR optimal rule with near optimal convergence property. Moreover, we think that the main qualitative insights regarding Γ-MMR optimal rules may extend to other forms of π_μ. For example, we conjecture that whenever the identification power of the model is sufficiently large compared to the informativeness of the data (in a sense to be suitably defined), Γ-MMR optimal rule will take a form of a non-randomized threshold rule; otherwise, the optimal rule may be randomized. We leave the verification of these conjectures for future research.

FUNDING STATEMENT

This study was funded by the NSF (Grant No. SES-2315600).

COMPETING INTEREST STATEMENT

The authors declare that no competing interests exist.

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

REFERENCES

Adjaho, C., & Christensen, T. (2022). *Externally valid treatment choice*. Preprint, arXiv:2205.05561.

Amarante, M., & Siniscalchi, M. (2019). Recursive maxmin preferences and rectangular priors: A simple proof. *Economic Theory Bulletin*, 7, 125–129.

Aradillas Fernández, A., Blanchet, J., Olea, J. L. M., Qiu, C., Stoye, J., & Tan, L. (2025). *ϵ-Minimax solutions of statistical decision problems*. Preprint, arXiv:2509.08107.

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

Ben-Michael, E., Greiner, D. J., Imai, K., & Jiang, Z. (2025). Safe policy learning through extrapolation: Application to pre-trial risk assessment. *Journal of the American Statistical Association*, 120(551), 1386–1399. https://doi.org/10.1080/01621459.2025.2489135.

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

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

Betrò, B., & Ruggeri, F. (1992). Conditional¡-minimax actions under convex losses. *Communications in Statistics-Theory and Methods*, 21, 1051–1066.

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

Brock, W. A. (2006). Profiling problems with partially identified structure. *The Economic Journal*, 116, F427–F440.

Canner, P. L. (1970). Selecting one of two treatments when the responses are dichotomous. *Journal of the American Statistical Association*, 65, 293–306.

Chamberlain, G. (2011). Bayesian aspects of treatment choice. In J. Geweke, G. Koop, & H. Van Dijk (Eds.), *The Oxford handbook of Bayesian econometrics* (online ed., Oxford Academic, November 21,. 2012). Oxford University Press.

Chen, H., & Guggenberger, P. (2025). A note on minimax regret rules with multiple treatments in finite samples. *Econometric Theory*. Forthcoming.

Christensen, T., Moon, H. R., & Schorfheide, F. (2026). Optimal discrete decisions when payoffs are partially identified. *Review of Economic Studies*. Preprint, arXiv:2204.11748.

D'Adamo, R. (2021). *Policy learning under ambiguity*. Preprint, arXiv:2111.10904.

DasGupta, A., & Studden, W. J. (1989). Frequentist behavior of robust Bayes estimates of normal means. *Statistics & Risk Modeling*, 7, 333–362.

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

Dvoretzky, A., Wald, A., & Wolfowitz, J. (1951). Elimination of randomization in certain statistical decision procedures and zero-sum two-person games. *The Annals of Mathematical Statistics*, 22, 1–21.

Epstein, L. G., & Schneider, M. (2003). Recursive multiple-priors. *Journal of Economic Theory*, 113, 1–31.

Ferguson, T. (1967). *Mathematical statistics: A decision theoretic approach* (Vol. 7). Academic Press.

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

Giacomini, R., Kitagawa, T., & Read, M. (2021). Robust Bayesian analysis for econometrics. CEPR Discussion Paper No. 16488.

Guggenberger, P., & Huang, J. (2025). On the numerical approximation of minimax regret rules via fictitious play. Preprint, arXiv:2503.10932.

Guggenberger, P., Mehta, N., & Pavlov, N. (2024). Minimax regret treatment rules with finite samples when a quantile is the object of interest. Tech. rep., The Pennsylvania State University.

Hayashi, T. (2008). Regret aversion and opportunity dependence. *Journal of Economic Theory*, 139, 242–268.

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

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

Hirano, K., & Porter, J. R. (2020). Asymptotic analysis of statistical decision rules in econometrics. In S. N. Durlauf, L. P. Hansen, J. J. Heckman, & R. L. Matzkin (Eds.), *Handbook of econometrics, volume 7A* (volume 7 of Handbook of Econometrics, pp. 283–354). Elsevier.

Ida, T., Ishihara, T., Ito, K., Kido, D., Kitagawa, T., Sakaguchi, S., & Sasaki, S. (2026). Choosing who chooses: Selection-driven targeting in energy rebate programs. *Econometrica*, 94, 225–247. https://doi.org/10.3982/ECTA21180.

Ishihara, T., & Kitagawa, T. (2021). *Evidence aggregation for treatment choice*. Preprint, arXiv:2108.06473.

Kallus, N., & Zhou, A. (2018). Confounding-robust policy improvement. In *Advances in neural information processing systems* (Vol. 31), 9289–9299.

Karlin, S., & Rubin, H. (1956). The theory of decision procedures for distributions with monotone likelihood ratio. *The Annals of Mathematical Statistics*, 27(2), 272–299.

Khan, M. A., Rath, K. P., & Sun, Y. (2006). The Dvoretzky-Wald-Wolfowitz theorem and purification in atomless finite-action games. *International Journal of Game Theory*, 34, 91–104.

Kido, D. (2022). *Distributionally robust policy learning with Wasserstein distance*. Preprint, arXiv:2205.04637.

Kido, D. (2023). *Locally asymptotically minimax statistical treatment rules under partial identification*. Preprint, arXiv:2311.08958.

Kitagawa, T. (2012). *Estimation and inference for set-identified parameters using posterior lower probability*. Manuscript, UCL.

Kitagawa, T., Lee, S., & Qiu, C. (2023). Treatment choice, mean square regret and partial identification. *The Japanese Economic Review*, 74, 573–602.

Kitagawa, T., Lee, S., & Qiu, C. (2026). Treatment choice with nonlinear regret. *Biometrika*, 113(2), asag008, https://doi.org/10.1093/biomet/asag008.

Kitagawa, T., Sakaguchi, S., & Tetenov, A. (2021). *Constrained classification and policy learning*. Preprint, https://arxiv.org/abs/2106.12886.

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

Kitagawa, T., & Tetenov, A. (2021). Equality-minded treatment choice. *Journal of Business Economics and Statistics*, 39, 561–574.

Kitagawa, T., & Wang, G. (2023). Who should get vaccinated? Individualized allocation of vaccines over SIR network. *Journal of Econometrics*, 232, 109–131.

Lehmann, E. L., & Casella, G. (1998). *Theory of point estimation*. (2nd ed.) Springer Science & Business Media.

Lei, L., Sahoo, R., & Wager, S. (2023). *Policy learning under biased sample selection*. Preprint, arXiv:2304.11735.

Manski, C. F. (2000). Identification problems and decisions under ambiguity: Empirical analysis of treatment response and normative analysis of treatment choice. *Journal of Econometrics*, 95, 415–442.

Manski, C. F. (2004). Statistical treatment rules for heterogeneous populations. *Econometrica*, 72, 1221–1246.

Manski, C. F. (2005). *Social choice with partial knowledge of treatment response*. Princeton University Press.

Manski, C. F. (2007a). *Identification for prediction and decision*. Harvard University Press.

Manski, C. F. (2007b). Minimax-regret treatment choice with missing outcome data. *Journal of Econometrics*, 139, 105–115.

Manski, C. F. (2020). Towards credible patient-centered meta-analysis. *Epidemiology*, 31, 345–352.

Manski, C. F. (2025). Identification and statistical decision theory. *Econometric Theory*, 41(4), 977–993, doi:10.1017/S0266466624000197.

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

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press

Manski, C. F., & Tetenov, A. (2023). Statistical decision theory respecting stochastic dominance. *The Japanese Economic Review*, 74, 447–469.

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

Montiel Olea, J. L., Qiu, C., & Stoye, J. (2026). Decision theory for treatment choice problems with partial identification. *The Review of Economic Studies*, rdag015.

Moon, H. R., & Schorfheide, F. (2012). Bayesian and frequentist inference in partially identified models. *Econometrica*, 80, 755–782.

Poirier, D. J. (1998). Revising beliefs in nonidentified models. *Econometric Theory*, 14, 483–509.

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

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

Song, K. (2014). Point decisions for interval-identified parameters. *Econometric Theory*, 30, 334–356.

Stoye, J. (2007). Minimax regret treatment choice with incomplete data and many treatments. *Econometric Theory*, 23, 190–199.

Stoye, J. (2009). Minimax regret treatment choice with finite samples. *Journal of Econometrics*, 151, 70–81.

Stoye, J. (2011). Axioms for minimax regret choice correspondences. *Journal of Economic Theory*, 146, 2226–2251.

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

Stoye, J. (2012b). New perspectives on statistical decisions under ambiguity. *Annual Review of Economics*, 4, 257–282.

Tetenov, A. (2012a). Measuring precision of statistical inference on partially identified parameters. Discussion Paper, Collegio Carlo Alberto, Torino.

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

Vidakovic, B. (2000). Γ-minimax: A paradigm for conservative robust Bayesians. In D. R. Insua & F. Ruggeri (Eds.), *Robust Bayesian analysis* (pp. 241–259). Springer.

Wakai, K. (2007). A note on recursive multiple-priors. *Journal of Economic Theory*, 135, 567–571.

Wald, A. (1945). Statistical decision functions which minimize the maximum risk. *Annals of Mathematics*, 46, 265–280.

Yata, K. (2021). *Optimal decision rules under partial identification*. Preprint, arXiv:2111.04926.

https://doi.org/10.1017/S0266466626100486 Published online by Cambridge University Press