

# Robust Bayes Treatment Choice with Partial Identification\*

Andres Aradillas Fernandez<sup>†</sup>    José Luis Montiel Olea<sup>‡</sup>    Chen Qiu<sup>§</sup>  
Jörg Stoye<sup>¶</sup>    Serdil Tinda<sup>||</sup>

July 2024

## Abstract

We study a class of binary treatment choice problems with partial identification, through the lens of robust (multiple prior) Bayesian analysis. We use a convenient set of prior distributions to derive ex-ante and ex-post robust Bayes decision rules, both for decision makers who can randomize and for decision makers who cannot.

Our main messages are as follows: First, ex-ante and ex-post robust Bayes decision rules do not tend to agree in general, whether or not randomized rules are allowed. Second, randomized treatment assignment for some data realizations can be optimal in both ex-ante and, perhaps more surprisingly, ex-post problems. Therefore, it is usually with loss of generality to exclude randomized rules from consideration, even when regret is evaluated ex-post.

We apply our results to a stylized problem where a policy maker uses experimental data to choose whether to implement a new policy in a population of interest, but is concerned about the external validity of the experiment at hand ([Stoye, 2012a](#)); and to the aggregation of data generated by multiple randomized control trials in different sites to make a policy choice in a population for which no experimental data are available ([Manski, 2020](#); [Ishihara and Kitagawa, 2021](#)).

KEYWORDS: treatment choice, partial identification, robust Bayes,  $\Gamma$  minimax, posterior robustness

---

\*We gratefully acknowledge financial support from NSF grant SES-2315600.

<sup>†</sup>Department of Economics, Cornell University, aa2679@cornell.edu

<sup>‡</sup>Department of Economics, Cornell University, jlo67@cornell.edu

<sup>§</sup>Department of Economics, Cornell University, cq62@cornell.edu

<sup>¶</sup>Department of Economics, Cornell University, stoye@cornell.edu

<sup>||</sup>Department of Economics, Stanford University, serdil@stanford.edu

# 1 Introduction

A policy maker must decide between implementing a new policy or preserving the status quo. Her data provide information about the potential benefits of these two options. Unfortunately, these data only *partially identify* payoff-relevant parameters and may therefore not reveal, even in large samples, the correct course of action. Such *treatment choice problems with partial identification* have recently received growing interest; for example, see [D’Adamo \(2021\)](#), [Ishihara and Kitagawa \(2021\)](#), [Yata \(2021\)](#), [Christensen, Moon, and Schorfheide \(2022\)](#), [Kido \(2022\)](#) or [Manski \(2024\)](#). Several interesting problems that arise in empirical research can be recast using this framework. See [Montiel Olea, Qiu, and Stoye \(2023\)](#) and the end of this section for references.

This paper applies the *robust Bayes* approach—which interpolates between Bayesian and agnostic minimax analyses by evaluating minimax risk over a set of priors—to a class of treatment choice problems with partial identification. The use of the robust Bayes approach has drawn recent attention in problems that feature partial identification ([Giacomini and Kitagawa, 2021](#); [Giacomini, Kitagawa, and Read, 2021](#); [Christensen et al., 2022](#)). Indeed, due to partial identification, integrating Bayesian and minimax elements into decision making can be particularly attractive (see [Poirier 1998](#); [Moon and Schorfheide 2012](#) and references therein).

The robust Bayes approach can be applied *ex-ante* or *ex-post*, depending on whether the multiple priors are used to evaluate payoffs before or after seeing the data.<sup>1</sup> 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. This raises the question: What are the qualitative and quantitative relationships between ex-ante and ex-post robust Bayes criteria when applied to treatment choice problems with partial identification? In this paper, we provide an answer to these questions in a general class of treatment choice problems, using a convenient class of priors.

We use the same framework as [Yata \(2021\)](#) and [Montiel Olea et al. \(2023\)](#) but impose a simple instance of [Giacomini and Kitagawa’s \(2021\)](#) set of priors, namely a symmetric and uniform two-

---

<sup>1</sup>[Giacomini et al. \(2021\)](#) discuss both notions; they refer to the ex-ante and ex-post problems as “Gamma-minimax” and “Conditional Gamma-minimax”, respectively. [Christensen et al. \(2022\)](#) focus on the ex-post problem for treatment choice problems with partial identification in a restricted class of decision rules.

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 “ $\Gamma$ -minimax regret” ( $\Gamma$ -MMR) and “ $\Gamma$ -posterior expected regret” ( $\Gamma$ -PER).<sup>2</sup> We then precisely characterize when these notions coincide and when they disagree. The main qualitative insights are as follows:

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

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

Several auxiliary findings might be of independent interest. First, for  $\Gamma$ -MMR, even if we restrict the set of decision rules to be a class of non-randomized threshold rules based on the “efficient” linear

---

<sup>2</sup>Among others, see Savage (1951), Manski (2004), Stoye (2012b), and Montiel Olea et al. (2023) for justifications of focusing on regret in treatment choice problems. In particular, while minimax loss can be an attractive alternative to minimax regret, the aforementioned papers find that it leads to trivial recommendations in treatment choice settings; these findings are easily extended to our example.

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  $\Gamma$ -MMR optimal (among all decision rules). This is in stark contrast to [Montiel Olea et al. \(2023\)](#), in which no linear index rule is globally minimax regret optimal if the degree of partial identification is severe. The intuition is that the prior much reduces the state space; the signal space then becomes so rich relative to the state space that even linear threshold rules can effectively mimic randomization.

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

The rest of this paper is organized as follows. Section 2 sets up the problem, provides examples, and defines both versions of robust Bayes optimality. Section 3 contains complete solutions for all aforementioned scenarios, Section 4 relates them to the “fictitious game” interpretation of minimax theory, and Section 5 concludes. Proofs and auxiliary results are collected in the Appendix.

## 2 Framework

### 2.1 Actions, payoffs, statistical models and decisions

Our setup follows Montiel Olea et al. (2023), who in turn follow Ferguson (1967) and others. Consider a policy maker who needs to choose an *action*  $a \in [0, 1]$  interpreted as probability of assigning treatment in the target population.<sup>3</sup> Her payoff when taking action  $a \in [0, 1]$  is captured by the welfare function

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

where  $\theta \in \Theta$  is an unknown state of the world or *parameter* and the functions  $W(1, \cdot) : \Theta \rightarrow \mathbb{R}$  and  $W(0, \cdot) : \Theta \rightarrow \mathbb{R}$  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  $\theta$ . If  $U(\theta)$  were known to the policy maker, her optimal action would simply be

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

Since  $U(\theta)$  is in fact unknown, the policy maker gathers some empirical evidence to learn about  $\theta$ . We assume that she observes a random vector  $Y \in \mathbb{R}^n$  that follows a multivariate normal distribution:

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

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

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

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

collect all the means of  $Y$  that can be generated as  $\theta$  ranges over  $\Theta$ . We refer to elements  $\mu \in M$  as *reduced-form* parameters because they are identified in the statistical model (2.3) without further

---

<sup>3</sup>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.

assumptions. Define the *identified set* for the welfare contrast given  $\mu$  as

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

and the corresponding upper and lower bounds as

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

Henceforth, when we refer to a treatment choice problem with partial identification, we mean there exists some open set in  $M$  such that for all  $\mu$  in that open set,  $\underline{I}(\mu) < 0 < \bar{I}(\mu)$ .

A *decision rule*  $d : \mathbb{R}^n \rightarrow [0, 1]$  is a (measurable) mapping from data  $Y$  to the unit interval  $[0, 1]$ . We call  $d$  *non-randomized* if it (almost surely, a.s.) maps into  $\{0, 1\}$ ; otherwise, we call  $d$  *randomized*, including if it randomizes for some but not all data realizations. We use  $\mathcal{D}_n$  to denote the set of all decision rules, and we consider decision rules the same if they a.s. agree. As a result, a rule is *unique* only up to a.s. agreement. The *oracle policy*  $d_{\text{oracle}}^* := \mathbf{1}\{U(\theta) \geq 0\}$  is of special interest and for any given  $\theta$  is contained in  $\mathcal{D}_n$ , but is not feasible in the statistical sense because  $U(\theta)$  is not known.

In general, there will not be an unambiguously best feasible decision rule, a problem that gave rise to a large literature on different optimality criteria and their implementation. Before introducing the robust Bayes approach, we give two examples that fit into our general framework.

## 2.2 Examples

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

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

where  $\sigma > 0$  is known and where  $\mu \in \mathbb{R}$  is an *observable treatment effect*, in the sense that  $\mu$  can be perfectly pinned down if infinite amount of data is provided. In this example,  $\theta = (\mu, \mu^*)^\top$ ,  $\Theta \subseteq \mathbb{R}^2$ ,

$m(\theta) = \mu$ , and  $U(\theta) = \mu^*$ .

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

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

The goal of the planner is to choose a statistical decision rule  $d \in \mathcal{D}_1 : \mathbb{R} \rightarrow [0, 1]$  that maps observed data (2.7) to an action  $a \in [0, 1]$  to be applied to the target population. Taking the normality assumption as an approximation, this stylized example could reflect model uncertainty (e.g., a treatment effect is estimated in a possibly somewhat misspecified model), external validity concerns (e.g., a randomized clinical trial was performed on volunteers), or a shift in the environment (e.g., we transfer estimates from study populations to treatment populations with slightly different covariates or are concerned about distributional drift over time).  $\square$

**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  $i = 1, \dots, n$ . Let  $Y = (Y_1, \dots, Y_n)^\top \in \mathbb{R}^n$  denote these estimates and let  $(x_0, \dots, x_n)$  be nonrandom,  $d$ -dimensional baseline covariates. The policy maker is willing to extrapolate from her data by assuming that the welfare contrast of interest equals  $U(\theta) = \theta(x_0)$  and that

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

where  $\theta : \mathbb{R}^d \rightarrow \mathbb{R}$  is an unknown Lipschitz function with known constant  $C$ . For notational simplicity, we can further write  $\mu_i$  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  $\beta$ . In this example, the identified set for the welfare contrast  $\mu_0$  is

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

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

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

□

### 2.3 Robust Bayes optimality

Our setting up to here is as in [Montiel Olea et al. \(2023\)](#). We now connect it to the robust Bayes literature by imposing a *set of priors*  $\Gamma$  on  $\theta$ . Specifically, for each prior  $\pi$  on  $\theta$ , let  $\pi_\mu$  be the corresponding prior on the point-identified reduced-form parameter  $\mu \in M$ , obtained as the push-forward measure induced by  $\pi$  and  $m(\cdot)$ , i.e.,

$$\pi_\mu(A) = \pi(m^{-1}(A)), \forall A \in \mathcal{F}_M,$$

where  $m^{-1}(A) := \{\theta \in \Theta : m(\theta) \in A\}$  and  $\mathcal{F}_M$  denotes the  $\sigma$ -algebra of  $M$ . Following [Giacomini and Kitagawa \(2021\)](#), we choose a particular single proper prior  $\pi_\mu$  for  $\mu \in M$  but leave the conditional prior of  $\theta$  given  $\mu$ , denoted as  $\pi_{\theta|\mu}$ , unrestricted except for

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

Then, the class of priors  $\Gamma$  consists of all priors on  $\theta$  induced by the single prior  $\pi_\mu$  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  $\pi_\mu$ ; for tractability, we let  $\pi_\mu$  to be supported on two symmetric points  $\{\bar{\mu}, -\bar{\mu}\}$  with equal probability, where  $\bar{\mu} \in \mathbb{R}^n$  is chosen by the decision maker. Henceforth,  $\Gamma$  is understood to refer to the implied set of priors:

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

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

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

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



as

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

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

Even with the set of priors  $\Gamma$  given and the notion of expected regret chosen as the decision criterion, the robust Bayes literature contains multiple decision criteria that do not in general agree. The difference lies in when (and how) the expectations regarding the unknown parameter  $\theta \in \Theta$  are taken. For a decision rule  $d \in \mathcal{D}_n$ , let

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

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

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

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

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

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

where  $\pi_{\theta|Y}$  is the posterior distribution of  $\theta$  given prior  $\pi$  and data  $Y$ .<sup>4</sup> Then we have the following, alternative optimality criterion (Berger, 1985, Definition 10, p. 205):

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

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

---

<sup>4</sup>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  $\sigma$ -algebra of  $\Theta$ , 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  $\mu$  given  $Y$ .

If the decision maker’s action space is restricted to  $\{0, 1\}$ , i.e., randomization is not allowed, the above definition is revised by replacing  $[0, 1]$  with  $\{0, 1\}$ .

The labelling of  $\Gamma$ -MMR as “ex-ante” versus  $\Gamma$ -PER as “ex-post” can be related to the timing of a fictitious game against an adversarial Nature; see Section 4 for additional discussion. If  $\Gamma$  were a singleton, the criteria would coincide and would also agree with (single-prior) Bayes optimality. They do not in general agree otherwise. The  $\Gamma$ -PER criterion is often easier to solve and is routinely employed to quantify posterior robustness of statistical decisions. However, the term “Gamma minimax” overwhelmingly (and even “robust Bayes” more often than not) refers to  $\Gamma$ -MMR; for example, see Berger (1985, Section 4.7.6).<sup>5</sup>

### 3 Main results

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

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

These conditions are restrictive but turn out to encompass many examples of empirical relevance; see this paper’s introduction and Montiel Olea et al. (2023); Yata (2021) for examples. Under Assumption 1, we let

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

as a normalization.

We say a rule is a *linear-index threshold rule* if it has the form  $\mathbf{1}\{\beta'Y \geq c\}$  for some  $\beta \in \mathbb{R}^n$  and  $c \in \mathbb{R}$ . Linear-index threshold rules are nonrandomized. They are of particular interest because they form a complete class when  $U(\theta)$  is point-identified (Karlin and Rubin, 1956) and have received particular attention in the recent literature (Ishihara and Kitagawa, 2021; Montiel Olea et al., 2023). For both  $\Gamma$ -MMR and  $\Gamma$ -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:

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

---

<sup>5</sup>Giacomini, Kitagawa, and Read (2021) discuss both criteria for general loss functions and refer to Definitions 1 and 2 as the “unconditional  $\Gamma$ -minimax” and “conditional  $\Gamma$ -minimax” problems, respectively. For treatment choice problems with partial identification, Christensen et al. (2022) optimize the  $\Gamma$ -PER criterion, restricting the action space to be  $\{0, 1\}$ .

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

### 3.1 Ex-ante Robust Bayes Optimality

**Theorem 1** ( $\Gamma$ -MMR optimal decisions). *Consider a treatment choice problem with welfare function (2.1), statistical model (2.3) and a set of priors (2.9), that satisfies Assumption 1.*

(i) If

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

then  $d_{w,0}^*$  is uniquely  $\Gamma$ -MMR optimal.

(ii) If

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

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

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

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

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

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

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

$$d_{step}^* := \begin{cases} \frac{1}{2} - \beta^*, & w^\top Y < 0, \\ \frac{1}{2} + \beta^*, & w^\top Y \geq 0, \end{cases} \quad (3.8)$$

where

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

$\rho^* > 0$  is unique and solves  $\int_0^1 \Phi\left(\frac{2\rho^*x - \rho^* - \|w\|_{\Sigma}^2}{\|w\|_{\Sigma}}\right) dx = \frac{-\underline{I}(\bar{\mu})}{\bar{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}$ , 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  $\mathbf{1}\{w'Y \geq c\}$ , the optimal thresholds are  $\pm c^*$ , where

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

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

$$d_{w_{t^*}, 0}^*(Y) := \mathbf{1}\{w_{t^*}^{\top} Y \geq 0\},$$

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

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

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  $\Gamma$ -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*  $\mu$  (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  $\Gamma$ -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  $\Gamma$ -MMR optimal rules, all of which satisfy (3.4) and (3.5). Examples include suitably smoothed versions of  $d_{w,0}^*$  like  $d_{\text{RT}}^*$  and  $d_{\text{linear}}^*$  (the functional forms of which showed up in Montiel Olea et al. 2023) as well as  $d_{\text{step}}^*$  (which is a new result).

The contrast between these regimes becomes more intuitive upon inspecting the corresponding least favorable priors. When (3.2) holds, this prior is supported at two points ( $\mu = \bar{\mu}$ ,  $U(\theta) = \bar{I}(\bar{\mu})$ ) versus  $\mu = -\bar{\mu}$ ,  $U(\theta) = \underline{I}(-\bar{\mu})$ ) with equal probability, leading to unique optimal response  $d_{w,0}^*$ . In words, treatment should be assigned to all iff  $\mu = \bar{\mu}$ ; inference toward  $\mu$  and toward  $U(\theta)$  effectively coincide. When (3.3) holds, the least favorable prior is supported on four points ( $\mu = \bar{\mu}$ ,  $U(\theta) = \bar{I}(\bar{\mu})$ ;  $\mu = \bar{\mu}$ ,  $U(\theta) = \underline{I}(\bar{\mu})$ ;  $\mu = -\bar{\mu}$ ,  $U(\theta) = \bar{I}(-\bar{\mu})$  and  $\mu = -\bar{\mu}$ ,  $U(\theta) = \underline{I}(-\bar{\mu})$ ), with a probability profile such that the data are actually uninformative about  $U(\theta)$ . This explains nonuniqueness and randomization in optimal decision rules.<sup>6</sup>

In the case when the model’s identification power is small, we encounter several surprising findings. First, despite the problem’s apparent symmetry,  $d_{w,0}^*$  is *not* optimal even among linear-index threshold rules that use the index  $w^\top Y$ . Instead, Theorem 1(iii) characterizes exactly two optimal thresholds, one positive and one negative. Second, while the particular linear-index threshold rule  $d_{w,0}^*$  is not  $\Gamma$ -MMR optimal, in higher dimensional problems ( $n > 1$ ), there do exist linear-index threshold rules that are. This finding is in stark contrast to Montiel Olea et al. (2023), who find that, in a large class of special cases, *no* linear index threshold rule is MMR optimal. The crucial difference in settings is that the set of priors much constrains the decision theoretic problem’s state space; as a result, the signal space is much richer than the state space, and this can be exploited to mimic randomization without nominally randomizing. Compare Manski’s (2024) abstract observation that, if a policy maker is not allowed to explicitly randomize, sampling uncertainty can be beneficial by providing an implicit randomization device. We note that this phenomenon is reminiscent of classic “purification” results in game theory (Dvoretzky, Wald, and Wolfowitz, 1951; Khan, Rath, and Sun, 2006).<sup>7</sup> 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** ( $\Gamma$ -MMR decisions in Example 1). *In Example 1, the following statements are true:*

---

<sup>6</sup>To be clear, this does not mean that every decision rule becomes optimal. Among other things, rules are still constrained by what one may want to think of as indifference conditions required for mixed strategy equilibrium in a fictitious game; see again Section 4.

<sup>7</sup>We thank Elliot Lipnowski for reminding us of this literature.

(i) If

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

then

$$d_0^*(\cdot) := \mathbf{1}\{\hat{\mu} \geq 0\}$$

is the unique  $\Gamma$ -MMR optimal decision rule.

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

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

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

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

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

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

$$d_{step}^* := \begin{cases} \frac{1}{2} - \frac{\bar{\mu}}{2k(2\Phi(\frac{\bar{\mu}}{\sigma}) - 1)}, & \hat{\mu} < 0, \\ \frac{1}{2} + \frac{\bar{\mu}}{2k(2\Phi(\frac{\bar{\mu}}{\sigma}) - 1)}, & \hat{\mu} \geq 0, \end{cases} \quad (3.15)$$

where

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

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  $\Gamma$ -minimax regret are  $\mathbf{1}\{\hat{\mu} \geq \pm c^*\}$ , where

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

Note that there is no analog to Theorem 1's case (iv); indeed, no linear threshold rule is optimal in case (ii). This is because the scalar nature of the signal  $Y$  shuts down the aforementioned purification

mechanism. See Figure 1 for an illustration of different  $\Gamma$ -MMR optimal rules in Example 1 for selected parameter values of  $k, \sigma$  and  $\bar{\mu}$ .

### 3.2 Ex-post Robust Bayes Optimality

**Theorem 2** ( $\Gamma$ -PER optimal decisions). *Suppose all conditions in Theorem 1 hold true. Then:*

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

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

*Otherwise,  $d_{w,0}^*$  is  $\Gamma$ -PER optimal.*

(ii)  *$d_{w,0}^*$  is always the  $\Gamma$ -PER optimal non-randomized rule.*

The results of Theorem 2 offer some important clarifications regarding  $\Gamma$ -PER in treatment choice problems with partial identification. First, just because regret is evaluated according to the posterior distribution does not mean 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  $\Gamma$ -PER optimal rule is randomized. Therefore, restricting the action space to  $\{0, 1\}$  in such problem is *not* without loss of generality even under the  $\Gamma$ -PER criterion. Comparing results with Theorem 1 also allows for instructive observations on when  $\Gamma$ -PER and  $\Gamma$ -MMR optimal rules agree or disagree; we elaborate these in Corollary 3. Applying Theorem 2 to Example 1, we finally obtain:

**Corollary 2** ( $\Gamma$ -PER rules in Example 1). *In Example 1, the following statements are true:*

(i) *If  $\bar{\mu} < k$ , then  $\Gamma$ -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  $\Gamma$ -PER optimal. Furthermore, it is always the  $\Gamma$ -PER optimal threshold rule.*

In Figure 2, we depict  $\Gamma$ -PER optimal rules for Example 1 with the same parameter values considered in Figure 1. We see clearly that  $\Gamma$ -MMR and -PER optimal rules coincide (if randomization is allowed) only in the special case when  $k \leq \bar{\mu}$ , an observation we generalize in Corollary

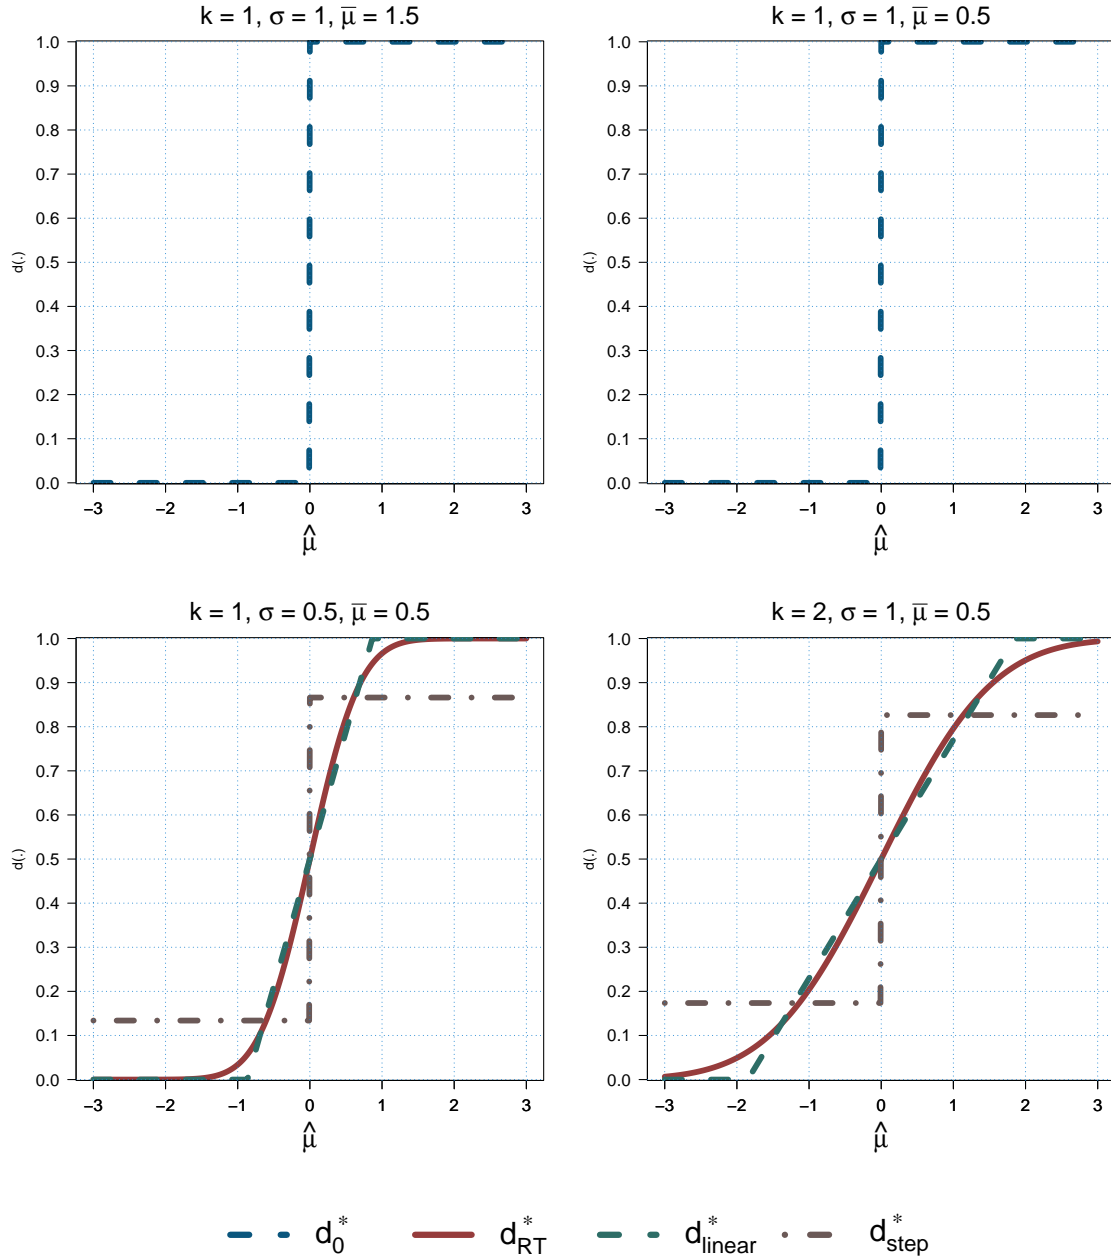


Figure 1:  $\Gamma$ -MMR optimal rules in Example 1

*Notes:* This figure reports the  $\Gamma$ -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  $\Gamma$ -MMR optimal rule is  $d_0^*$ . In the bottom two panels, (3.10) fails. As a result,  $d_0^*$  is no longer  $\Gamma$ -MMR optimal. Instead,  $d_{RT}^*$ ,  $d_{linear}^*$  and  $d_{step}^*$  are all  $\Gamma$ -MMR optimal.



3. More specifically, in the top left panel of Figure 2, as  $k \leq \bar{\mu}$ ,  $\Gamma$ -MMR and  $\Gamma$ -PER coincide and are the non-randomized threshold rule  $d_0^*$ . In the top right panel,  $k > \bar{\mu}$  and the  $\Gamma$ -PER optimal rule becomes  $d_{\text{PER}}^*$ . However, since (3.10) still holds,  $d_0^*$  is still  $\Gamma$ -MMR optimal. For the bottom two panels, as it still holds  $k > \bar{\mu}$ ,  $d_{\text{PER}}^*$  is still  $\Gamma$ -PER optimal. However, the associated parameter values imply (3.10) fails. As a result,  $d_0^*$  is no longer  $\Gamma$ -MMR optimal and many  $\Gamma$ -MMR rules exist. But even in these cases,  $\Gamma$ -MMR and  $\Gamma$ -PER rules differ, as among the class of step function rules (which contain  $d_{\text{PER}}^*$ ), only  $d_{\text{step}}^*$  is  $\Gamma$ -MMR optimal, still different from  $d_{\text{PER}}^*$ .

### 3.3 When Do Ex-ante and Ex-post Agree?

Theorems 1 and 2 clarify to what extent  $\Gamma$ -MMR and  $\Gamma$ -PER criteria yield the same or different optimal rules. We spell this out in Corollary 3, which considers both cases when randomization is allowed and not allowed. The bottom line is that the assessment is sensitive to how the problem is set up. If underlying parameters lead to sufficiently small identification power, the criteria disagree.

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

(i) *Suppose randomization is allowed. Then the  $\Gamma$ -MMR and  $\Gamma$ -PER optimal rules coincide if, and only if,  $\underline{I}(\bar{\mu}) \geq 0$ .*

(ii) *Suppose randomization is not allowed. Then the  $\Gamma$ -MMR and  $\Gamma$ -PER optimal rules coincide if, and only if,*

$$\frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu}) - \underline{I}(\bar{\mu})} \geq \Phi(\|w\|_{\Sigma}). \quad (3.16)$$

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

One might conjecture that the two criteria agree more once randomization is excluded; after all, this leads to a much simpler action space. Part (ii) shows that there is some truth to this: The condition for agreement changes from  $\underline{I}(\bar{\mu}) \geq 0$  to the strictly weaker (3.16). However, the criteria continue to disagree in many cases. It may be instructive to think of these cases in terms of “comparative statics.” For example, consider holding all parameters of the problem fixed but scaling the signal variance  $\Sigma$  by a positive scalar, for example to reflect a change in sample size. Then (3.16) will hold if, and only if,  $\Sigma$  is *large* enough; hence, as long as  $\underline{I}(\bar{\mu}) \geq 0$ , increasing

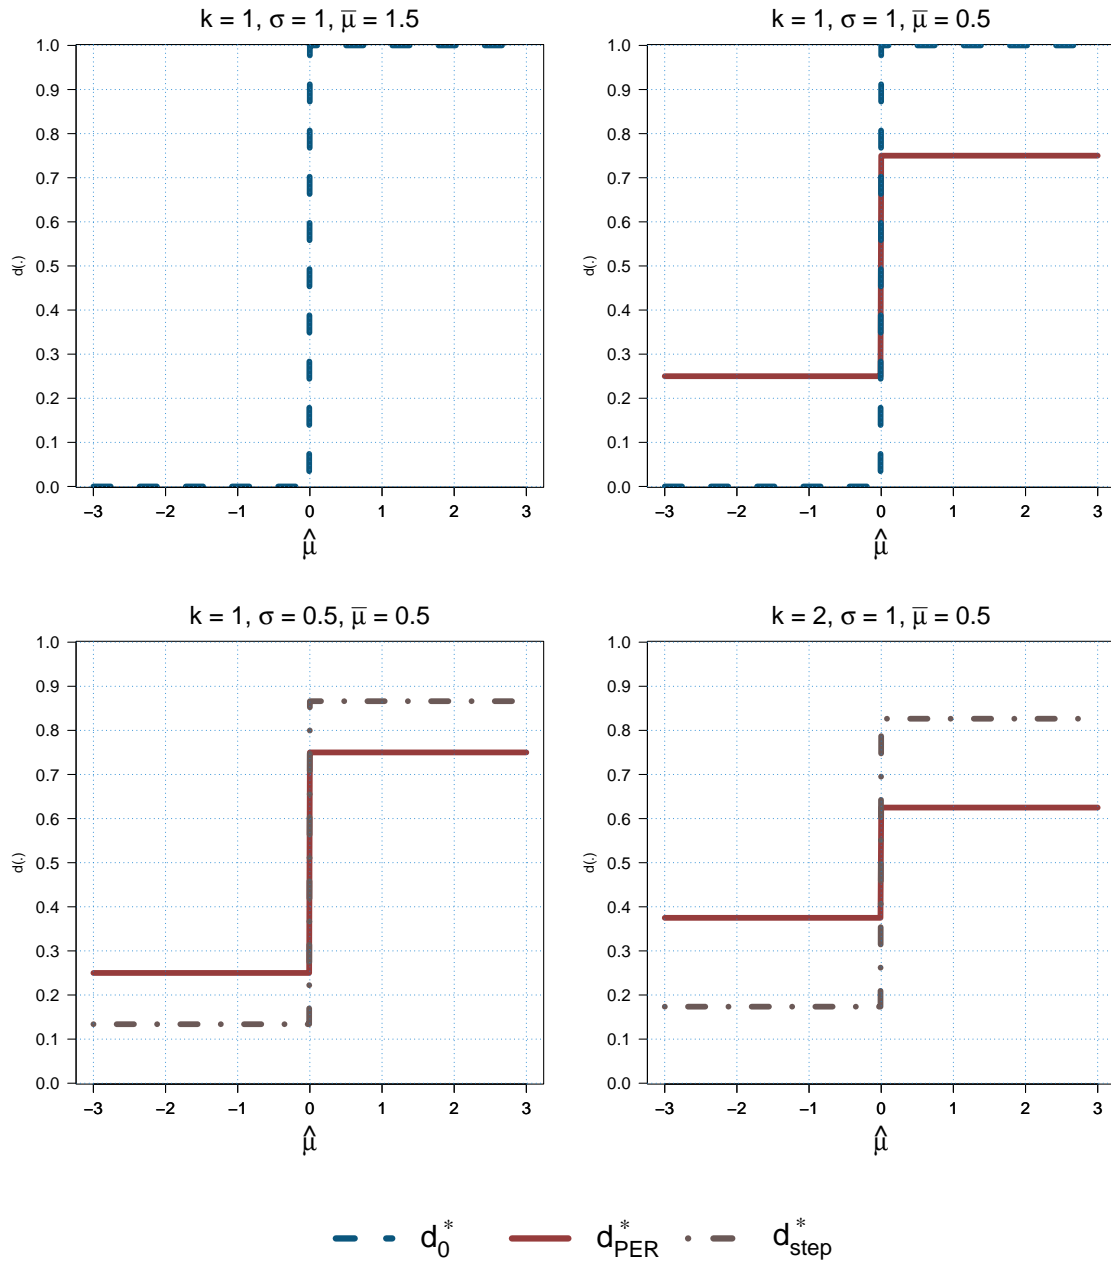


Figure 2:  $\Gamma$ -PER optimal rules in Example 1

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

sample size will eventually cause disagreement between  $\Gamma$ -MMR and  $\Sigma$ -PER even if randomization is excluded. Similarly, for fixed  $\Sigma$ ,  $\Gamma$ -MMR and  $\Sigma$ -PER rules will always disagree if the model’s identification power is *small* enough.

## 4 Game-theoretic Interpretation

As Wald (1940) was first to discover, it can be instructive to think of minimax decision rules as equilibrium rules in a fictitious, zero-sum game against Nature. In particular, any equilibrium of the game in which the decision maker picks a decision rule and an adversarial Nature picks a prior characterizes a minimax (regret) decision rule and corresponding prior; furthermore, known features of zero-sum games may allow one to claim uniqueness of the solution and the like.<sup>8</sup> While usually stated for “agnostic” minimax rules, it is easy to see that these facts extend to our setting. They are liberally used in this paper’s proofs; we here use them to briefly discuss some intuitions and contrast ex-ante versus ex-post robust Bayes, but also randomization versus no randomization, in terms of the fictitious game’s timing.

First, this perspective gives further intuition for why MMR rules tend to randomize. Indeed, nontrivial simultaneous move zero-sum games frequently fail to have pure strategy equilibria, and that is precisely what happens here. Furthermore, in those cases where optimal decision rules randomize, the equilibria are “pooling” equilibria in which data are noninformative. This explains why these rules also tend to be nonunique; what disciplines them is Nature’s best-response condition, not learning from data in equilibrium. This is also true for some nonrandomized threshold rules from Theorem 1(iv); these avoid randomization because uninformative features of the data can be used as randomization device, not because any updating occurs in equilibrium.

Second, ex-ante and ex-post rules can be distinguished in terms of the game being played. In the ex-ante (and original Waldian) perspective, this game is simultaneous move; in particular, Nature moves before data  $Y$  are realized. In the ex-post perspective, Nature sees  $Y$  before choosing a prior. It is immediately clear that this may be easier to solve because it allows for backward induction. There is also an immediate sense that solutions might not agree, as we indeed found.

Finally, whether the decision maker is allowed to (or at least wants to) randomize or not can also be thought of as changing the timing of the fictitious game. Specifically, the more standard perspective is that the decision maker may randomize over decision rules and Nature must move before learning the outcome of this randomization. This setup will often yield randomized solutions.

---

<sup>8</sup>See Lehmann and Casella (1998, Theorem 1.4, Chap. 5) for a formal statement. The first explicit use of this technique in the literature on decisions under partial identification is in Stoye (2007).

In contrast, if Nature is allowed to move after the decision maker’s randomization realized, then any incentive to randomize is gone and we may as well restrict the action space to  $\{0, 1\}$ . Again, there is no reason to expect that solutions agree.

## 5 Conclusions

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)  $\Gamma$ -minimax regret and (ex-post)  $\Gamma$ -posterior expected regret and analytically derive optimal solutions with and without randomization.

Our results contain two key messages that we think are valuable to the literature. First, with partial identification and multiple priors, ex-ante and ex-post assessments do *not* agree in general, whether or not randomized rules are allowed. This may at first seem expected due to dynamic inconsistency of multiple prior Bayes criteria, but is actually not obvious in view of the set of prior’s specific structure. Second, randomization can be optimal in both ex-ante and ex-post problems—it is with loss of generality to exclude them even when regret is evaluated ex-post. We conjecture, but leave to future research, that our findings qualitatively extend to more general sets of priors. The contrast between the results also illustrates a need to better understand the comparative advantages—whether from a theoretical or practical perspective—of using one criterion over the other.

## References

- ADJAHO, C. AND T. CHRISTENSEN (2022): “Externally Valid Treatment Choice,” *arXiv preprint arXiv:2205.05561*.
- ATHEY, S. AND S. WAGER (2021): “Efficient policy learning with observational data,” *Econometrica*, 89, 133–161.
- BEN-MICHAEL, E., D. J. GREINER, K. IMAI, AND Z. JIANG (2021): “Safe policy learning through extrapolation: Application to pre-trial risk assessment,” *arXiv preprint arXiv:2109.11679*.
- BEN-MICHAEL, E., K. IMAI, AND Z. JIANG (2022): “Policy learning with asymmetric utilities,” *arXiv preprint arXiv:2206.10479*.
- BERGER, J. (1985): *Statistical decision theory and Bayesian analysis*, Springer.

- BHATTACHARYA, D. AND P. DUPAS (2012): “Inferring welfare maximizing treatment assignment under budget constraints,” *Journal of Econometrics*, 167, 168–196.
- BROCK, W. A. (2006): “Profiling problems with partially identified structure,” *The Economic Journal*, 116, F427–F440.
- CANNER, P. L. (1970): “Selecting One of Two Treatments when the Responses are Dichotomous,” *Journal of the American Statistical Association*, 65, 293–306.
- CHAMBERLAIN, G. (2011): “Bayesian Aspects of Treatment Choice,” in *The Oxford Handbook of Bayesian Econometrics*, Oxford University Press.
- CHEN, H. AND P. GUGGENBERGER (2024): “A note on minimax regret rules with multiple treatments in finite samples,” Tech. rep., The Pennsylvania State University.
- CHRISTENSEN, T., H. R. MOON, AND F. SCHORFHEIDE (2022): “Optimal Discrete Decisions when Payoffs are Partially Identified,” *arXiv preprint arXiv:2204.11748*.
- D’ADAMO, R. (2021): “Policy Learning Under Ambiguity,” *arXiv preprint arXiv:2111.10904*.
- DEHEJIA (2005): “Program evaluation as a decision problem,” *Journal of Econometrics*, 125, 141–173.
- DVORETZKY, A., A. WALD, AND J. WOLFOWITZ (1951): “Elimination of Randomization in Certain Statistical Decision Procedures and Zero-Sum Two-Person Games,” *The Annals of Mathematical Statistics*, 22, 1 – 21.
- FERGUSON, T. (1967): *Mathematical Statistics: A Decision Theoretic Approach*, vol. 7, Academic Press New York.
- GIACOMINI, R. AND T. KITAGAWA (2021): “Robust Bayesian Inference for Set-Identified Models,” *Econometrica*, 89, 1519–1556.
- GIACOMINI, R., T. KITAGAWA, AND M. READ (2021): “Robust Bayesian Analysis for Econometrics,” *CEPR Discussion Paper No. 16488*.
- GUGGENBERGER, P., N. MEHTA, AND N. PAVLOV (2024): “Minimax regret treatment rules with finite samples when a quantile is the object of interest,” Tech. rep., The Pennsylvania State University.
- HIRANO, K. AND J. R. PORTER (2009): “Asymptotics for statistical treatment rules,” *Econometrica*, 77, 1683–1701.
- (2020): “Asymptotic analysis of statistical decision rules in econometrics,” in *Handbook of Econometrics, Volume 7A*, ed. by S. N. Durlauf, L. P. Hansen, J. J. Heckman, and R. L. Matzkin, Elsevier, vol. 7 of *Handbook of Econometrics*, 283–354.
- IDA, T., T. ISHIHARA, K. ITO, D. KIDO, T. KITAGAWA, S. SAKAGUCHI, AND S. SASAKI (2022): “Choosing Who Chooses: Selection-Driven Targeting in Energy Rebate Programs,” Tech. rep.,

- National Bureau of Economic Research.
- ISHIHARA, T. AND T. KITAGAWA (2021): “Evidence Aggregation for Treatment Choice,” ArXiv:2108.06473 [econ.EM], <https://doi.org/10.48550/arXiv.2108.06473>.
- KALLUS, N. AND A. ZHOU (2018): “Confounding-robust policy improvement,” *Advances in neural information processing systems*, 31.
- KARLIN, S. AND H. RUBIN (1956): “The theory of decision procedures for distributions with monotone likelihood ratio,” *The Annals of Mathematical Statistics*, 272–299.
- KHAN, M. A., K. P. RATH, AND Y. SUN (2006): “The Dvoretzky-Wald-Wolfowitz Theorem and Purification in Atomless Finite-action Games,” *International Journal of Game Theory*, 34, 91–104.
- KIDO, D. (2022): “Distributionally Robust Policy Learning with Wasserstein Distance,” *arXiv preprint arXiv:2205.04637*.
- (2023): “Locally Asymptotically Minimax Statistical Treatment Rules Under Partial Identification,” *arXiv preprint arXiv:2311.08958*.
- KITAGAWA, T., S. LEE, AND C. QIU (2022): “Treatment Choice with Nonlinear Regret,” *arXiv preprint arXiv:2205.08586*.
- (2023): “Treatment choice, mean square regret and partial identification,” *The Japanese Economic Review*, 74, 573–602.
- KITAGAWA, T., S. SAKAGUCHI, AND A. TETENOV (2021): “Constrained classification and policy learning,” *arXiv preprint*.
- KITAGAWA, T. AND A. TETENOV (2018): “Who should be treated? Empirical welfare maximization methods for treatment choice,” *Econometrica*, 86, 591–616.
- (2021): “Equality-Minded Treatment Choice,” *Journal of Business Economics and Statistics*, 39, 561–574.
- KITAGAWA, T. AND G. WANG (2023): “Who should get vaccinated? Individualized allocation of vaccines over SIR network,” *Journal of Econometrics*, 232, 109–131.
- LEHMANN, E. L. AND G. CASELLA (1998): *Theory of point estimation*, Springer Science & Business Media, second ed.
- LEI, L., R. SAHOO, AND S. WAGER (2023): “Policy Learning under Biased Sample Selection,” *arXiv preprint arXiv:2304.11735*.
- MANSKI, C. F. (2000a): “Identification problems and decisions under ambiguity: empirical analysis of treatment response and normative analysis of treatment choice,” *Journal of Econometrics*, 95, 415–442.
- (2000b): “Identification problems and decisions under ambiguity: Empirical analysis of

- treatment response and normative analysis of treatment choice,” *Journal of Econometrics*, 95, 415–442.
- (2004): “Statistical treatment rules for heterogeneous populations,” *Econometrica*, 72, 1221–1246.
- (2005): *Social choice with partial knowledge of treatment response*, Princeton University Press.
- (2007a): *Identification for prediction and decision*, Harvard University Press.
- (2007b): “Minimax-regret treatment choice with missing outcome data,” *Journal of Econometrics*, 139, 105–115.
- (2020): “Towards credible patient-centered meta-analysis,” *Epidemiology*, 31, 345–352.
- (2024): “Identification and Statistical Decision Theory,” *Econometric Theory*, in press.
- MANSKI, C. F. AND A. TETENOV (2007): “Admissible treatment rules for a risk-averse planner with experimental data on an innovation,” *Journal of Statistical Planning and Inference*, 137, 1998–2010.
- (2023): “Statistical decision theory respecting stochastic dominance,” *The Japanese Economic Review*, 74, 447–469.
- MBAKOP, E. AND M. TABORD-MEEHAN (2021): “Model selection for treatment choice: Penalized welfare maximization,” *Econometrica*, 89, 825–848.
- MONTIEL OLEA, J. L., C. QIU, AND J. STOYE (2023): “Decision Theory for Treatment Choice Problems with Partial Identification,” ArXiv: 2312.17623.
- MOON, H. R. AND F. SCHORFHEIDE (2012): “Bayesian and Frequentist Inference in Partially Identified Models,” *Econometrica*, 80, 755–782.
- POIRIER, D. J. (1998): “Revising Beliefs in Nonidentified Models,” *Econometric Theory*, 14, 483–509.
- SAVAGE, L. (1951): “The theory of statistical decision,” *Journal of the American Statistical Association*, 46, 55–67.
- SCHLAG, K. H. (2006): “ELEVEN - Tests needed for a Recommendation,” Tech. rep., European University Institute Working Paper, ECO No. 2006/2, <https://cadmus.eui.eu/bitstream/handle/1814/3937/ECO2006-2.pdf>.
- SONG, K. (2014): “Point Decisions for Interval-Identified Parameters,” *Econometric Theory*, 30, 334–356.
- STOYE, J. (2007): “Minimax Regret Treatment Choice With Incomplete Data and many Treatments,” *Econometric Theory*, 23, 190–199.
- (2009): “Minimax regret treatment choice with finite samples,” *Journal of Econometrics*,

151, 70–81.

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

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

TETENOV, A. (2012a): “Measuring precision of statistical inference on partially identified parameters,” *Discuss. Pap., Coll. Carlo Alberto, Torino*.

——— (2012b): “Statistical treatment choice based on asymmetric minimax regret criteria,” *Journal of Econometrics*, 166, 157–165.

WALD, A. (1940): “The Fitting of Straight Lines if Both Variables are Subject to Error,” *The Annals of Mathematical Statistics*, 11, 284–300.

YATA, K. (2021): “Optimal Decision Rules Under Partial Identification,” ArXiv:2111.04926 [econ.EM], <https://doi.org/10.48550/arXiv.2111.04926>.

## A Proofs of main results

We liberally invoke the game theoretic characterization alluded to in Section 4. Thus, proofs will frequently claim and verify equilibria of the fictitious game. Recall that, from basic facts about zero-sum games, if a decision rule *uniquely* best responds to some least favorable prior, it must be the *unique* equilibrium rule.

### A.1 Proof of Theorem 1

#### A.1.1 Statement (i)

The expected regret of decision rule  $d$  is

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

Recall that, by the definition of  $\Gamma$ , 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

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



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

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

The least favorable prior  $\pi^*$  equals

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

Lemma B.1 shows that the unique Bayes rule against  $\pi^*$  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  $\Gamma$ -MMR optimal if

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

The least favorable prior  $\pi^*$  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  $\pi^*$  must best respond to any optimal decision rule, i.e., for any MMR optimal rule  $d$ ,  $\pi^*$  must solve

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

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  $\mathbf{1}\{w^\top Y \geq c\}$  for some  $c \in \mathbb{R}$  cannot be  $\Gamma$ -MMR optimal when (3.3) holds. Note  $w = \Sigma^{-1}\bar{\mu}$ . Thus,

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

Suppose by contradiction that a rule  $\mathbf{1}\{w^\top Y \geq c\}$  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  $d_{\text{RT}}^*$ ,  $d_{\text{linear}}^*$  and  $d_{\text{step}}^*$  are all  $\Gamma$ -MMR optimal. Due to symmetry, it suffices to show that

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

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

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

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

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

Applying Lemma B.8 in [Montiel Olea et al. \(2023\)](#) and integration by parts yield

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

Note that  $\lim_{\rho \downarrow 0} f(\rho) = 1 - \Phi(-\|w\|_\Sigma) = \Phi(\|w\|_\Sigma) > \frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}$ , while L'Hopital's rule implies

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

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  $(0, \infty)$ .

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,  $\int_0^1 \Phi\left(\frac{2\rho^* x - \rho^* - \|w\|_\Sigma^2}{\|w\|_\Sigma}\right) dx = \frac{-\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}$ , 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  $\beta \in (0, \frac{1}{2})$ , consider the following step function rule:

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

One then has

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

Setting  $\mathbb{E}_{\bar{\mu}}[d_{\text{step},\beta^*}] = \frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}$  yields  $\beta^* = \frac{\bar{I}(\bar{\mu})}{2\Phi(\|w\|_{\Sigma}) - 1} - \frac{1}{2}$ . 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  $\dot{\mu} \neq \mathbf{0}$  is such that  $\dot{\mu}^{\top}\Sigma^{-1}\bar{\mu} = 0$ , i.e.,  $\dot{\mu}$  is orthogonal to  $\Sigma^{-1}\bar{\mu}$ . For any  $t \in \mathbb{R}$ , let

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

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

## A.2 Proof of Corollary 1

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

## A.3 Proof of Theorem 2

Recall that the  $\Gamma$ -PER optimal rule solves

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

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

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

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

## A.4 Proof of Corollary 2

Directly follows from Theorem 2.

## A.5 Proof of Corollary 3

(i) “If”: If  $\underline{I}(\bar{\mu}) \geq 0$ , then Theorem 1(i) and Theorem 2(i) apply and establish that  $d_{w,0}^*$  is both the unique  $\Gamma$ -MMR and the unique  $\Gamma$ -PER optimal rule.

“Only if”: If  $\underline{I}(\bar{\mu}) < 0$ , then  $d_{\text{PER}}^*$  is uniquely  $\Gamma$ -PER optimal by Theorem 2(i). If condition (3.2) holds as well, then Theorem 1(i) implies that  $d_{w,0}^*$  is uniquely  $\Gamma$ -MMR optimal; hence,  $\Gamma$ -MMR and  $\Gamma$ -PER optimal rules disagree. If condition (3.3) applies, then, by Theorem 1(ii), a  $\Gamma$ -MMR optimal rule must satisfy (3.4)-(3.5). But  $d_{\text{PER}}^*$  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  $\beta_{\text{PER}} = \frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu}) - \underline{I}(\bar{\mu})} - \frac{1}{2}$ . Step 4 for the proof of Theorem 1(ii) shows that the only  $\Gamma$ -MMR optimal rule of form  $d_{\text{step}, \beta}$  is  $d_{\text{step}, \beta}^* \neq d_{\text{PER}}^*$ .

(ii) “If”: When randomization is not allowed, Theorem 2(ii) shows that  $d_{w,0}^*$  is always  $\Gamma$ -PER optimal. Under condition (3.16), Theorem 1(i) establishes that  $d_{w,0}^*$  is  $\Gamma$ -MMR optimal as well.

“Only if”: Again,  $d_{w,0}^*$  is always  $\Gamma$ -PER optimal when randomization is not allowed. If (3.16) fails (i.e., (3.3) holds) and  $n > 1$ , it follows by Theorem 1(iv) that the randomized rule  $d_{w_{t^*},0}^*$  is  $\Gamma$ -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,  $\Gamma$ -MMR and  $\Gamma$ -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  $\pi^*$ , the Bayes optimal rule must solve the posterior problem

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

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

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  $\mathbf{1}\{f(Y|-\bar{\mu}) - f(Y|\bar{\mu}) \leq 0\}$ , which is equivalent to  $d_{w,0}^*$  after further algebra.  $\square$

**Lemma B.2.** *Consider decision rule  $d_{w,0}^*$  and prior  $\pi^*$  defined in (A.1). If (3.2) holds, then*

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

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

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

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

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

It follows that

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

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

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

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

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  $\|w\|_\Sigma = \sqrt{\bar{\mu}^\top \Sigma^{-1} \bar{\mu}}$ .  $\square$

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

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

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

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

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

$$\begin{aligned} \min_{a \in [0,1]} & \frac{1}{2} \cdot f(Y | \bar{\mu}) [p_1 \bar{I}(\bar{\mu})(1-a) + (1-p_1)(-\underline{I}(\bar{\mu}))a] \\ & + \frac{1}{2} \cdot f(Y | -\bar{\mu}) [p_2 \bar{I}(-\bar{\mu})(1-a) + (1-p_2)(-\underline{I}(-\bar{\mu}))a]. \end{aligned}$$

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.  $\square$

**Lemma B.4.** Consider the prior  $\pi^*$  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.

*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  $[c^*, \infty)$ .

*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

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

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  $[c^*, \infty)$ .  $\square$

**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

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

Furthermore, note

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

and

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

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

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

As  $s^* := \frac{(\Phi^{-1}(k^*))^2}{\|w\|_\Sigma^2} \in (0, 1)$  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  $\Gamma$ -PER optimal rule is*

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

Otherwise,  $d_{w,0}^*$  is  $\Gamma$ -PER optimal.

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

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

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

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

where

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

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

$$\begin{aligned}\bar{I}(\bar{\mu})(1-a) \geq -\underline{I}(\bar{\mu})a &\Leftrightarrow a \leq \frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}, \\ \bar{I}(\bar{\mu})a \geq -\underline{I}(\bar{\mu})(1-a) &\Leftrightarrow a \geq \frac{-\underline{I}(\bar{\mu})}{\bar{I}(\bar{\mu}) - \underline{I}(\bar{\mu})},\end{aligned}$$

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

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

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

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

**Case 2:**  $\underline{I}(\bar{\mu}) < 0 < \bar{I}(\bar{\mu})$ . In this case, we have (again, note  $\bar{I}(\bar{\mu}) + \underline{I}(\bar{\mu}) > 0$ ):

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

Therefore:

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

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

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

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

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

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

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

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

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

In sum, if  $f(Y | -\bar{\mu}) > f(Y | \bar{\mu})$ ,  $V_\Gamma(\cdot)$  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  $f(Y | -\bar{\mu}) < f(Y | \bar{\mu})$ ,  $V_\Gamma(\cdot)$  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  $f(Y | -\bar{\mu}) = f(Y | \bar{\mu})$  (which only happens in a null set),  $V_\Gamma(\cdot)$  achieves minimum at any point between  $\left[ \frac{-\underline{I}(\bar{\mu})}{\bar{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}, \frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu}) - \underline{I}(\bar{\mu})} \right]$ . Therefore, the unique  $\Gamma$ -PER optimal rule is

$$d(Y) = \begin{cases} \frac{\bar{I}(\bar{\mu})}{\bar{I}(\bar{\mu}) - \underline{I}(\bar{\mu})}, & \text{if } f(Y | -\bar{\mu}) \leq f(Y | \bar{\mu}) \iff \bar{\mu}^\top \Sigma^{-1} Y \geq 0, \\ \frac{-\underline{I}(\bar{\mu})}{\bar{I}(\bar{\mu}) - \underline{I}(\bar{\mu})} & \text{if } f(Y | -\bar{\mu}) > f(Y | \bar{\mu}) \iff \bar{\mu}^\top \Sigma^{-1} Y < 0. \end{cases}$$

□

**Lemma C.2.** *Suppose all conditions of Theorem 1 hold true. Then,  $d_{w,0}^*$  is always the  $\Gamma$ -PER optimal non-randomized rule.*

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

$$\begin{aligned} V_\Gamma(1) &= -f(Y | \bar{\mu})\underline{I}(\bar{\mu}) + f(Y | -\bar{\mu})\bar{I}(\bar{\mu}), \\ V_\Gamma(0) &= f(Y | \bar{\mu})\bar{I}(\bar{\mu}) - f(Y | -\bar{\mu})\underline{I}(\bar{\mu}), \end{aligned}$$

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

## D Additional results

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

$$(i) \quad \bar{I}(-\bar{\mu}) = -\underline{I}(\bar{\mu}), \quad \underline{I}(-\bar{\mu}) = -\bar{I}(\bar{\mu});$$

$$(ii) \quad \bar{I}(-\bar{\mu}) + \underline{I}(-\bar{\mu}) = -(\underline{I}(\bar{\mu}) + \bar{I}(\bar{\mu})).$$

*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  $\Theta$  is centrosymmetric and both  $U(\cdot)$  and  $m(\cdot)$  is linear, we have

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

as required.

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