

# Optimal Discrete Decisions when Payoffs are Partially Identified\*

Timothy Christensen<sup>†</sup> Hyungsik Roger Moon<sup>‡</sup> Frank Schorfheide<sup>§</sup>

April 25, 2022

## Abstract

We derive optimal statistical decision rules for discrete choice problems when the decision maker is unable to discriminate among a set of payoff distributions. In this problem, the decision maker must confront both model uncertainty (about the identity of the true payoff distribution) and statistical uncertainty (the set of payoff distributions must be estimated). We derive *efficient-robust decision rules* which minimize maximum risk or regret over the set of payoff distributions and which use the data to learn efficiently about features of the set of payoff distributions germane to the choice problem. We discuss implementation of these decision rules via the bootstrap and Bayesian methods, for both parametric and semiparametric models. Using a limits of experiments framework, we show that efficient-robust decision rules are optimal and can dominate seemingly natural alternatives. We present applications to treatment assignment using observational data and optimal pricing in environments with rich unobserved heterogeneity.

**Keywords:** Model uncertainty, statistical decision theory, partial identification, treatment assignment, revealed preference

**JEL codes:** C10, C18, C21, C44, D81

---

\*We are grateful to X. Chen, L. Hansen, Y. Kitamura, J. Porter, Q. Vuong, and E. Vytlacil and seminar participants at Chicago Booth, Colorado, Durham, Montreal, Princeton, Queen Mary U. of London, SFU, UCL, Vanderbilt, Virginia, Wisconsin, and Yale for helpful comments and suggestions. The theory and methodology developed in this paper is based on and supersedes the preprint [arXiv:2011.03153](https://arxiv.org/abs/2011.03153) (Christensen, Moon, and Schorfheide, 2020). This material is based upon work supported by the National Science Foundation under Grants No. SES-1919034 (Christensen), SES-1625586 (Moon), and SES-1851634 (Schorfheide).

<sup>†</sup>Department of Economics, New York University. [timothy.christensen@nyu.edu](mailto:timothy.christensen@nyu.edu)

<sup>‡</sup>Department of Economics, University of Southern California, and Yonsei University. [moonr@usc.edu](mailto:moonr@usc.edu)

<sup>§</sup>Department of Economics, University of Pennsylvania, CEPR, NBER, and PIER. [schorf@ssc.upenn.edu](mailto:schorf@ssc.upenn.edu)

# 1 Introduction

Many important policy decisions involve discrete choices. Examples include whether or not to treat an aggregate population or large sub-population, firm or worker decisions at the extensive margin, and pricing policies when, in practice, prices must be expressed in whole currency units. Suppose a decision maker must choose a policy from a discrete choice set. The decision maker has data that may be used to bound, but not point identify, the payoffs associated with some choices. For instance, the decision maker may be deciding whether or not to assign treatment based on observational data which is sufficient to bound, but not point identify, the average treatment effect. How should they proceed?

In this paper, we propose an approach for making optimal discrete statistical (i.e., data-driven) decisions when the payoffs associated with some choices are only partially identified. In the model described in Section 2, the decision maker observes data which may be used to learn about a vector of parameters  $P$ . The decision maker then chooses a policy from among a discrete set. The distribution of payoffs associated with the different policies is indexed by a structural parameter  $\theta$ . A key assumption underlying the analysis in this paper is that  $\theta$  is possibly set-identified, but that the parameters  $P$  may be used to deduce restrictions on  $\theta$ . The decision maker therefore confronts both model uncertainty (the payoff distribution is not point-identified) and statistical uncertainty ( $P$  must be estimated from the data).

We propose a theory of optimal statistical decision making in this setting. We adopt a minimax approach to handle the ambiguity that arises from the partial identification of  $\theta$  conditional on  $P$  and average (or integrated) risk minimization to efficiently estimate  $P$ . Such an asymmetric treatment of parameters was first proposed by Hurwicz (1951). We refer to the resulting optimal decision rules as *efficient-robust* decision rules: they are “robust” in the sense that they minimize maximum risk or regret over the set of payoff distributions indexed by  $\theta$  conditional on  $P$ , and “efficient” in the sense that they use the data to learn efficiently about features of  $P$  germane to the choice problem.

Efficient-robust decision rules take a simple form and can be implemented easily via the bootstrap or Bayesian methods, in both parametric and semiparametric settings. To describe the implementation, fix any choice and consider its maximum risk (or regret) over  $\theta$  conditional on  $P$ . The maximum risk (or regret) is averaged across the bootstrap distribution for an efficient estimator  $\hat{P}$  of  $P$ , a posterior distribution for  $P$  (in parametric models), or a quasi-posterior based on a limited-information criterion for  $P$  (in semiparametric models), conditional on the data. The efficient-robust decision is then to simply to choose whatever choice has smallest average maximum risk (or regret).

We show how to implement efficient-robust decisions in the context of two applications, which we use as running examples. The first considers treatment assignment under partial identification of the average treatment effect (ATE). The second considers optimal pricing

in an environment with rich unobserved heterogeneity, where revealed preference arguments may be used to bound, but not point identify, demand responses under counterfactual prices. In both examples, the maximum risk (or regret) associated with different choices is available in closed form or can otherwise be computed easily by solving a simple optimization problem (e.g. a linear program).

For parametric models, our efficiency criterion extends the asymptotic average risk criterion introduced by [Hirano and Porter \(2009\)](#) for point-identified settings to partially identified settings. Sections 3 and 4 present optimality results for decisions based on parametric and semiparametric models, respectively. Our main results show formally that the proposed bootstrap and Bayesian implementations of efficient-robust decisions are optimal under our asymptotic efficiency criterion. Any decision rule that is asymptotically equivalent to the proposed implementations is optimal as well. Moreover, we show that asymptotic equivalence to the efficient-robust decision is in fact necessary for optimality: any decision whose asymptotic behavior is different from the efficient-robust decision is dominated under our optimality criterion.

An important insight is that “plug-in rules”,<sup>1</sup> which plug an efficient estimator  $\hat{P}$  into the oracle decision rule if  $P$  were known, can be dominated. This finding is in contrast with the point-identified case, where plug-in rules are efficient ([Hirano and Porter, 2009](#)). This finding also provides a formal, large-sample justification for [Manski’s \(2021a\)](#) critique of plug-in rules. For the intuition, consider a treatment assignment problem under partial identification of the ATE. Oracle rules under minimax criteria depend on a robust welfare contrast  $b(P)$  formed from bounds on the ATE. The bounds are non-smooth functions of  $P$  in many empirically relevant settings reviewed in Section 5. This non-smoothness leads to a failure of the  $\delta$ -method that breaks the asymptotic equivalence between plug-in rules, which depend on  $b(\hat{P})$ , and efficient-robust rules, which depend on the average of  $b(\cdot)$  across a bootstrap or posterior distribution.<sup>2</sup>

Within the context of treatment assignment under partial identification, this paper complements prior work by [Manski \(2000, 2007a, 2020, 2021a,b\)](#), [Chamberlain \(2011\)](#), [Stoye \(2012\)](#), [Russell \(2020\)](#), [Ishihara and Kitagawa \(2021\)](#), and [Yata \(2021\)](#). Except for [Chamberlain \(2011\)](#), these works seek decision rules that are optimal under finite-sample minimax regret criteria. We depart from these works in two respects. First, our criteria are local asymptotic criteria rather than finite-sample criteria. A local asymptotic framework enables us to approximate the finite-sample decision problem faced by the decision maker without having to explicitly solve the finite-sample problem, which may be intractable in many important applications. This allows us to relax potentially restrictive parametric assumptions on the

---

<sup>1</sup>[Manski \(2021a,b\)](#) refers to plug-in rules as “as-if” optimization.

<sup>2</sup>If  $b$  depends smoothly on  $P$ , then plug-in rules will be asymptotically equivalent to the efficient-robust decision and therefore optimal. We view this case as the exception rather than the rule.

data-generating process (e.g. Gaussianity) that are typically imposed to derive finite-sample rules and accommodate a much broader class of data-generating processes, including semi-parametric models. As such, our optimality results apply to settings where the decision maker cannot confidently assert that the data are drawn from a given parametric model, or where bounds on the ATE are estimated using a vector of moments or summary statistics (e.g. regression or IV estimates from observational studies) whose exact finite-sample distribution is unknown.<sup>3</sup>

Second, these works use optimality criteria that (in our notation) are minimax over  $(\theta, P)$ , whereas our criterion is minimax over the partially-identified parameter  $\theta$  and averages over the point-identified parameter  $P$ , reflecting the asymmetric parameterization of the problem. This lends a great deal of tractability, allowing us to derive optimal rules for a broad class of empirically relevant models where finite-sample rules are infeasible:<sup>4</sup> one simply needs to be able to compute the bounds as a function of  $P$ .

As we discuss in Section 2, our approach is closely related to the multiple priors framework of [Gilboa and Schmeidler \(1989\)](#) and robust Bayes (or  $\Gamma$ -minimax) decision making ([Robbins, 1951](#); [Berger, 1985](#)) in partially identified models. See [Giacomini, Kitagawa, and Read \(2021\)](#) for a recent review. Our paper makes several contributions in connection with  $\Gamma$ -minimax decision making. First, for parametric models, we provide a large-sample frequentist justification for these types of decision rules, which complements the usual finite-sample robust Bayes justification. Second, we show this optimality carries over to bootstrap-based decisions and quasi-Bayes decisions in semiparametric models. In these latter cases the decisions we derive are explicitly not  $\Gamma$ -minimax.

In Section 5 we discuss implementation in the context of intersection bounds ([Manski, 1990](#)), bounds based on extrapolating IV-like estimands ([Mogstad, Santos, and Torgovitsky, 2018](#)), and nonseparable panel data models ([Honoré and Tamer, 2006](#); [Chernozhukov, Fernández-Val, Hahn, and Newey, 2013](#)). We also illustrate our methods in an empirical application to extrapolation and meta-analysis from [Ishihara and Kitagawa \(2021\)](#), where the decision maker decides whether or not to adopt a job-training program based on RCT evidence from different populations.

Our application to optimal pricing is presented in Section 6. This application builds on prior work on revealed-preference demand theory, including [Blundell, Browning, and Crawford \(2007, 2008\)](#), [Blundell, Kristensen, and Matzkin \(2014, 2017\)](#), [Hoderlein and Stoye \(2015\)](#), [Manski \(2007b, 2014\)](#), and [Kitamura and Stoye \(2018, 2019\)](#). These works are primar-

---

<sup>3</sup>Finite-sample results are sometimes developed for this case assuming Gaussianity of the statistics, by arguing that the studies are sufficiently large that the sampling distribution of the statistics is approximately normal. In these cases, it seems logically consistent to use a large-sample optimality criterion.

<sup>4</sup>Indeed, we are not aware of any work deriving a finite-sample minimax treatment rule under partial identification even in the simplest case of a binary outcome and binary treatment when randomization is not permitted and bounds must be estimated from data (see Section 5.3 of [Manski \(2021a\)](#)).

ily concerned with testing rationality or deriving (sharp) bounds on counterfactual demand responses. We instead focus on using the bounds to solve optimal pricing problem, and using demand data at observed prices efficiently in this context. Finally, Section 7 concludes. Technical assumptions and proofs are relegated to the Appendix.

## 2 Efficient-Robust Decisions and Their Implementation

We begin in Section 2.1 by stating the decision problem considered in this paper. It features a partially-identified decision relevant parameter  $\theta$  and a point-identified parameter  $P$  that determines the identified set for  $\theta$ . Many important examples are covered by our framework. We discuss applications to treatment assignment and optimal pricing in detail.<sup>5</sup> We develop our concept of asymptotically efficient-robust decision rules in Section 2.2. Bayesian and bootstrap implementations are presented in Sections 2.3 and 2.4. In Section 2.5 we discuss the interpretation of our approach in the context of a two-player zero-sum game and relate it to the literature on robust Bayes decision making.

### 2.1 Decision Problem and Applications

A decision maker (DM) observes  $X^n \sim F_{n,P}$  taking values in a sample space  $\mathcal{X}^n$ . For instance,  $X^n = (X_1, \dots, X_n)$  may be a random sample of size  $n$  with  $F_{n,P}$  denoting the  $n$ -fold product distribution over independent random variables  $X_i \sim F_P$ . Alternatively,  $X^n$  may be a vector of summary statistics (e.g. sample moments or OLS or IV estimates) reported in a study based on a sample of size  $n$  with  $F_{n,P}$  denoting the sampling distribution of  $X^n$ . The distribution  $F_{n,P}$  from which  $X^n$  is drawn is indexed by an unknown *reduced-form parameter*  $P \in \mathcal{P} \subseteq \mathbb{R}^k$ . To simplify exposition, in this section we will be primarily concerned with parametric models for  $X^n$ . Our approach extends naturally to semiparametric models, as discussed in Section 4.

After observing  $X^n$ , the DM chooses an element  $d$  from a finite set  $\mathcal{D} = \{0, 1, \dots, D\}$ . The DM's *statistical decision rule*  $d_n : \mathcal{X}^n \rightarrow \mathcal{D}$  maps realizations of the data into discrete choices. The DM receives expected utility

$$\mathbb{E}_\theta[u(d, Y, \theta, P)] \quad \text{from choice } d \in \mathcal{D},$$

where the expectation is taken with respect to

$$Y \sim G_\theta$$

for a random vector  $Y$  drawn independently of  $X^n$ . The distribution  $G_\theta$  and expectation  $\mathbb{E}_\theta[\cdot]$  can be conditional on specific covariates, but we suppress this dependence to simplify

---

<sup>5</sup>For applications to forecasting, see our earlier working paper [Christensen et al. \(2020\)](#).

notation. The distribution  $G_\theta$  is indexed by a *structural parameter*  $\theta$ , as described in more detail below. We may interpret the expectation as a forward-looking expectation over a random outcome, or a cross-sectional average across the population in a social planner’s problem. Our notation encompasses welfare regret by setting

$$u(d, y, \theta, P) = W(d, y, \theta, P) - \max_{d' \in \mathcal{D}} \mathbb{E}_\theta[W(d', Y, \theta, P)]$$

for a welfare function  $W(d, y, \theta, P)$ . We therefore do not distinguish between “risk” and “regret” in what follows, except when referring to functional forms for  $u(\cdot)$ .

A key assumption underlying the analysis in this paper is that the structural parameter  $\theta$  is possibly set-identified. That is, distribution  $G_\theta$  from which  $Y$  is drawn is known up to  $\theta \in \Theta_0(P) \subseteq \Theta$ , where  $\Theta_0(\cdot)$  is a known set-valued mapping from the reduced-form parameter space to the structural parameter space. The set of distributions from which  $Y$  could be drawn if  $P$  were the true reduced-form parameter is  $\mathcal{G}_P = \{G_\theta : \theta \in \Theta_0(P)\}$ . Although  $X^n$  can be used to learn about  $P$ , it contains no information about the identity of the true  $\theta$  within  $\Theta_0(P)$ .

To fix ideas, we introduce two applications which are used as running examples. The first application is to treatment assignment under partial identification of the average treatment effect. The second example is to optimal pricing allowing for rich unobserved heterogeneity.

**Application 1: Treatment Assignment under Partial Identification.** Suppose the DM is a social planner whose goal is to choose between introducing a policy intervention (treatment,  $d = 1$ ) or no intervention (no treatment,  $d = 0$ ) in a target population of interest. Let  $Y_{0i}$  and  $Y_{1i}$  denote the untreated and treated potential outcomes associated with the intervention for individual  $i$ . Moreover, let  $Y_i = Y_{1i} - Y_{0i}$  be the treatment effect for individual  $i$ , and  $G_\theta$  be the distribution of  $Y_i$  across the target population. The average treatment effect (ATE) is  $\mathbb{E}_\theta[Y_i]$ . We assume that the observations are independently and identically distributed and drop the  $i$  subscript subsequently.

Following [Manski \(2000, 2004\)](#), it is common to derive treatment rules for a regret criterion based on a utilitarian social welfare function that is linear in the ATE. Failure to treat ( $d = 0$ ) incurs zero regret when the ATE is negative, otherwise the regret is  $-\mathbb{E}_\theta[Y]$ . Similarly, treating ( $d = 1$ ) incurs zero regret when the ATE is positive, otherwise the regret is  $\mathbb{E}_\theta[Y]$ . Formally,

$$\mathbb{E}_\theta[u(d, Y, \theta, P)] = (d - \mathbb{I}[\mathbb{E}_\theta[Y] > 0])\mathbb{E}_\theta[Y]. \quad (1)$$

Suppose the DM observes data or summary statistics  $X^n \sim F_{n,P}$  from observational studies that may be used to bound the ATE for the target population. The reduced-form parameter vector  $P$  typically comprises some population moments. We denote the lower and upper bounds on the ATE by  $b_L(P)$  and  $b_U(P)$ , respectively. The set  $\Theta_0(P)$  consists

of all distributions  $G_\theta$  for which  $b_L(P) \leq \mathbb{E}_\theta[Y] \leq b_U(P)$ . We review several methods for constructing these bounds in a number of different empirical settings in Section 5.  $\square$

**Application 2: Optimal Pricing with Rich Unobserved Heterogeneity.** The next application concerns pricing policies using revealed-preference demand theory, building on [Blundell et al. \(2007, 2008\)](#), [Manski \(2007b, 2014\)](#), and [Kitamura and Stoye \(2018, 2019\)](#), among others. Suppose the DM observes repeated cross sections  $X^n = (X_{b,1}, \dots, X_{b,n})_{b=1}^B$ , which consist of  $n$  observations each. Each  $X_{b,i} \in \mathbb{R}^K$  is the demand of individual  $i$  for  $K$  goods under prices  $q_b$ :

$$X_{b,i} = \arg \max_{x \in \mathcal{B}_b} v_i(x),$$

where  $\mathcal{B}_b = \{x \in \mathbb{R}^K : x'q_b = 1\}$  is the budget set (expenditure is normalized to one). Individuals are heterogeneous in their utility functions. We assume that individuals are identically distributed and drop the  $i$  subscript subsequently.

We assume that the demand system is rationalized by a random utility model with a probability distribution  $\Pi^v$  over utility functions  $v$ . The demand under  $q_b$  of a randomly selected individual may therefore be interpreted as stochastic. The probability mass  $p_b(s)$  of individuals whose demand is in any set  $s \subset \mathcal{B}_b$  at price  $q_b$  is

$$p_b(s) = \int \mathbb{I} \{ \operatorname{argmax}_{x \in \mathcal{B}_b} v(x) \in s \} d\Pi^v(v), \quad s \in \mathcal{B}_b, \quad b = 1, \dots, B \quad (2)$$

(see, e.g., [Kitamura and Stoye \(2018\)](#), henceforth KS18).

Suppose the DM's goal is to choose between the current price vector  $q_1$  and a new price vector  $q_0$ , to maximize welfare or a functional of average demand (e.g. revenue) under  $q_d$ ,  $d \in \{0, 1\}$ . We focus on a binary decision for ease of exposition; multiple choices are discussed in Section 6. In principle, there could be many potential new price vectors, representing a collection of prices rounded to nearest currency units or tax or subsidy rates rounded to the nearest percentage.

For concreteness, consider the following illustration. There are two goods, and the econometrician observes the demand for two price vectors  $q_1$  and  $q_2$ , which generate the status quo budget set  $\mathcal{B}_1$  and a second budget set  $\mathcal{B}_2$  from an earlier period. The DM considers a tax policy ( $d = 0$ ) that changes prices from  $q_1$  to  $q_0$ , which generates a counterfactual budget set  $\mathcal{B}_\gamma$ . An illustration is provided in Figure 1.

Let  $h(X_1)$  denote a policy-relevant functional of demand under prices  $q_1$ , which we take to be the demand for Good 2. Let  $Y = h(X_0)$  be the demand for Good 2 under counterfactual prices  $q_0$ . Suppose that the goal of the policy is to reduce average consumption of Good 2 but that there is a certain benefit  $C$  attached to the status quo. Using  $d = 1$  to denote the current

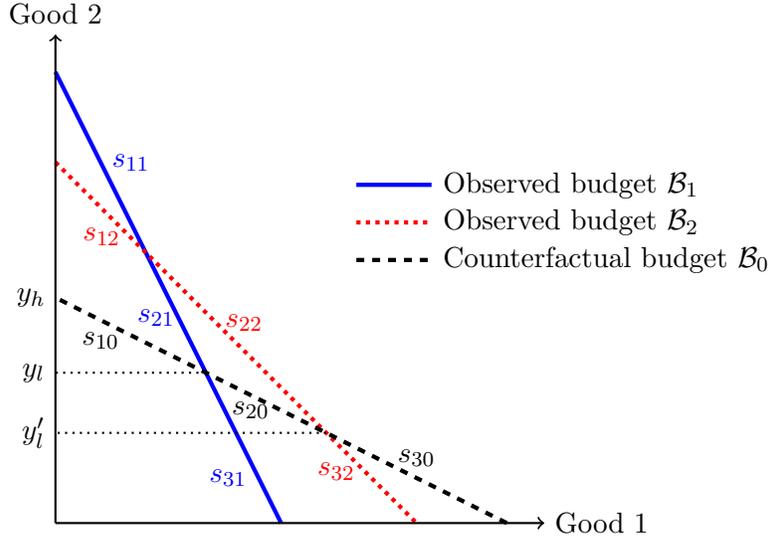


Figure 1: Choices and Budget Lines

policy and  $d = 0$  the alternative policy, we assume that the expected utility of decision  $d$  is

$$\begin{aligned}\mathbb{E}_\theta[u(1, Y, \theta, P)] &= -\mathbb{E}[h(X_1)] + C, \\ \mathbb{E}_\theta[u(0, Y, \theta, P)] &= -\mathbb{E}_\theta[Y].\end{aligned}$$

The observed demand under  $\mathcal{B}_1$  identifies the moment  $\mathbb{E}[h(X_1)]$ , which is a component of  $P$ . Therefore  $\mathbb{E}_\theta[u(1, Y, \theta, P)] = -\mathbb{E}[h(X_1)] + C$  is identified. It is evident in this example, and well known in general, that the counterfactual demand for Good 2 is only partially identified.<sup>6</sup> However, we can construct bounds under the assumption that the counterfactual demands have to satisfy the axiom of revealed stochastic preference. As such, any  $\Pi^\nu$  that is consistent with observed choice data on  $\mathcal{B}_1$  and  $\mathcal{B}_2$  induces a distribution  $G_\theta$  for  $Y$ .

We divide the budget lines in Figure 1 into segments  $s_{jb}$ , called “patches” in KS18. Patch  $s_{11}$  is the segment of  $\mathcal{B}_1$  from the  $y$ -axis to the intersection of  $\mathcal{B}_1$  and  $\mathcal{B}_2$ , patch  $s_{21}$  is the segment of  $\mathcal{B}_1$  from the intersection of  $\mathcal{B}_1$  and  $\mathcal{B}_2$  to the intersection of  $\mathcal{B}_1$  and  $\mathcal{B}_0$ , and so forth.<sup>7</sup>

Consider budget lines  $\mathcal{B}_1$  and  $\mathcal{B}_0$ . First, according to revealed preference, an individual on the  $s_{31}$  patch of the  $\mathcal{B}_1$  budget line (current policy) prefers their bundle over any bundle on the  $s_{11}$  or  $s_{21}$  patches. For any bundle on the  $s_{10}$  patch (counterfactual policy), there exist bundles on the  $s_{11}$  and  $s_{21}$  patches that are strictly preferable because they involve an

<sup>6</sup>See, e.g., Blundell et al. (2007, 2008), Manski (2007b), Blundell et al. (2014, 2017), Hoderlein and Stoye (2015), and Kitamura and Stoye (2019) among many others.

<sup>7</sup>Like Kitamura and Stoye (2018), we suppose for simplicity that the distribution of demand is continuous so we can disregard the “intersection patches” formed at the intersections of budget lines.

increase in consumption of at least one good. We deduce that an individual consuming on the  $s_{31}$  patch will choose their consumption from the  $s_{20}$  or  $s_{30}$  patch if the budget line shifts from  $\mathcal{B}_1$  to  $\mathcal{B}_7$ . This implies that their counterfactual consumption of Good 2 does not exceed  $y_l$ .

Second, the axiom of revealed preference does not restrict the  $\mathcal{B}_0$  consumption of individuals who under the current policy consume on the  $s_{11}$  or  $s_{21}$  patches. Thus, under the counterfactual budget  $\mathcal{B}_0$  their consumption of Good 2 lies between 0 and  $y_h$ .

Using the revealed preference argument, we obtain the (sharp) bounds for  $\mathbb{E}_\theta[u(0, Y, \theta, P)] = -\mathbb{E}_\theta[Y]$ :

$$-y_h(1 - p_1(s_{31})) - y_l p_1(s_{31}) \leq \mathbb{E}_\theta[u(0, Y, \theta, P)] \leq 0. \quad (3)$$

The choice probability  $p_1(s_{31})$  is identified from observed demand under  $\mathcal{B}_1$ . Now consider budget lines  $\mathcal{B}_2$  and  $\mathcal{B}_0$ . Repeating the revealed preference argument, we obtain a second set of bounds:

$$-y_h(1 - p_2(s_{32})) - y'_l p_2(s_{32}) \leq \mathbb{E}_\theta[u(0, Y, \theta, P)] \leq 0, \quad (4)$$

where  $p_2(s_{32})$  is identified from observed demand under  $\mathcal{B}_2$ . Combining (3) and (4), leads to the intersection bound

$$-y_h + \max\{(y_h - y_l)p_1(s_{31}), (y_h - y'_l)p_2(s_{32})\} \leq \mathbb{E}_\theta[u(0, Y, \theta, P)] \leq 0. \quad (5)$$

Whichever of (3) and (4) is larger depends on the relative size of  $y_l$  and  $y'_l$  and the probabilities  $p_1(s_{31})$  and  $p_2(s_{32})$ .<sup>8</sup>

In this example the vector of reduced-form parameters is  $P = (\mathbb{E}[h(X_1)], p_1(s_{31}), p_2(s_{32}))$ . The set  $\Theta_0(P)$  consists of all distributions  $G_\theta$  for  $Y$  for which

$$0 \leq \mathbb{E}_\theta[Y] \leq y_h - \max\{(y_h - y_l)p_1(s_{31}), (y_h - y'_l)p_2(s_{32})\}.$$

Linear Programming techniques developed in [Kitamura and Stoye \(2019\)](#) can be used to generalize these calculations to more complicated settings with a larger number of goods and a larger number of observed budget sets—see Section 6 below.  $\square$

## 2.2 Optimality Criterion and Efficient-Robust Decision Rules

We now define our optimality criterion and our notion of efficient-robust decision rules. We adopt the terminology of statistical decision theory from [Wald \(1950\)](#) and interpret negative utility as “loss” and negative expected utility as “risk”. Our setup features two types of

---

<sup>8</sup>Revealed preference restrictions also imply that the patch probabilities associated with the observed budget sets  $\mathcal{B}_1$  and  $\mathcal{B}_2$  need to satisfy  $p_2(s_{22}) + p_2(s_{32}) \geq p_1(s_{21}) + p_1(s_{31})$ .

parameters: a point-identified parameter  $P$  and a set-identified parameter  $\theta \in \Theta_0(P)$  that affects the risk associated with the decision. We adopt a minimax approach to handle the ambiguity that arises from the partial identification of  $\theta$  conditional on  $P$  and average (or integrated) risk minimization to efficiently estimate  $P$ . Such an asymmetric treatment of parameters was first proposed by [Hurwicz \(1951\)](#). We will elaborate on the connection in more detail in [Section 2.3](#). We will first consider the case in which  $P$  is known, and then extend the analysis to the case of unknown  $P$ , which will lead us to the definition of asymptotic efficient-robustness.

**Known  $P$ .** First suppose that  $P$  is known to the DM. To handle the ambiguity about  $\theta \in \Theta_0(P)$ , the DM evaluates choices  $d \in \mathcal{D}$  by their maximum risk

$$R(d, P) := \sup_{\theta \in \Theta_0(P)} \mathbb{E}_\theta[-u(d, Y, \theta, P)]. \quad (6)$$

The minimax decision for the model  $\mathcal{G}_P$  minimizes maximum risk:

$$d^o(P) = \arg \min_{d \in \mathcal{D}} R(d, P)$$

if the argmin is unique, otherwise  $d^o(P)$  is chosen randomly from  $\arg \min_{d \in \mathcal{D}} R(d, P)$ . The choice can be interpreted as the equilibrium in a two-player zero-sum game in which an adversarial “nature” chooses  $\theta$  in response to the DM’s choice of  $d$  to maximize  $\mathbb{E}[u(d, Y, \theta, P)]$ . We refer to  $d^o(P)$  as the *oracle* decision, as it represents the DM’s optimal choice if  $P$  was known. The oracle decision is an infeasible first-best, as in any application  $P$  is unknown and will need to be estimated from data. Nevertheless, it serves as a useful benchmark, as its maximum risk, namely  $\min_{d \in \mathcal{D}} R(d, P)$ , provides a lower bound on the minimax risk attainable when  $P$  is unknown.

**Unknown  $P$ .** When the DM does not know  $P$ , the DM’s decision will be data-dependent. For a decision rule  $d_n : \mathcal{X}^n \rightarrow \mathcal{D}$ ,<sup>9</sup> the DM incurs *frequentist (or average) maximum risk*

$$\mathbb{E}_P[R(d_n(X^n), P)], \quad (7)$$

where  $\mathbb{E}_P[\cdot]$  denotes expectation with respect to  $X^n \sim F_{n,P}$ . Our goal is to construct decision rules that use the data efficiently, in the sense that their average maximum risk is as close to that of the oracle over a range of data-generating processes.

Following [Hirano and Porter \(2009\)](#), we use a local asymptotic framework based on perturbations to the data-generating process of the same order as sampling variation. The idea of this asymptotic framework is that it should better approximate the finite-sample problem faced in practice, where one does not know  $P$  with certainty.

---

<sup>9</sup>We use the  $n$  subscript to denote that a decision rule may also depend on sample size  $n$ .

Formally, fix any limiting  $P_0 \in \mathcal{P}$  and let  $P_{n,h} = P_0 + h/\sqrt{n}$  denote its local perturbation for  $h \in \mathbb{R}^k$  in a sample of size  $n$ . For the asymptotic analysis it is convenient to express the frequentist maximum risk in excess of the oracle risk (the infeasible first-best) when  $P_{n,h}$  is the true reduced-form parameter. We scale the excess risk by  $\sqrt{n}$  to ensure the large-sample limit is not degenerate.<sup>10</sup> In the limit we obtain:

$$\mathcal{R}(\{d_n\}; P_0, h) := \lim_{n \rightarrow \infty} \sqrt{n} \left( \mathbb{E}_{P_{n,h}} [R(d_n(X^n), P_{n,h})] - \min_{d \in \mathcal{D}} R(d, P_{n,h}) \right), \quad (8)$$

which is non-negative, and where the limit exists for large class  $\mathbb{D}$  of decision rules whose asymptotic behavior is well-defined—see Section 3. This *asymptotic frequentist maximum risk* depends on the perturbation parameter  $h$ . We now integrate over  $h$  which leads to the following *asymptotic average maximum risk* criterion:

$$\mathcal{R}(\{d_n\}; P_0) := \int \lim_{n \rightarrow \infty} \sqrt{n} \left( \mathbb{E}_{P_{n,h}} [R(d_n(X^n), P_{n,h})] - \min_{d \in \mathcal{D}} R(d, P_{n,h}) \right) dh. \quad (9)$$

Averaging over  $h$  allows us to solve for the optimal decision in the treatment effect and optimal pricing applications and, more generally, any discrete choice application.<sup>11</sup> We define efficient-robustness as follows:

**Definition 1** *A sequence of decision rules  $\{d_n\}$  is asymptotically efficient-robust if*

$$\mathcal{R}(\{d_n\}; P_0) = \inf_{\{d'_n\} \in \mathbb{D}} \mathcal{R}(\{d'_n\}; P_0) \quad \text{for all } P_0 \in \mathcal{P}.$$

**Example.** To illustrate the scaling and centering of the asymptotic maximum risk, consider the following stylized example of Application 1. Let  $P \in \mathcal{P} = \mathbb{R}$  and suppose that the ATE  $\mathbb{E}_\theta[Y]$  is bounded by  $b_L(P) = (P)_-$  and  $b_U(P) = (1+c)(P)_+$ , where  $c > -1$  is an asymmetry parameter,  $(a)_+ = \max\{a, 0\}$ , and  $(a)_- = \min\{a, 0\}$ . For instance, for  $c = 1$  and  $P = -1$  ( $P = 1$ ) we obtain the interval  $[-1, 0]$  ( $[0, 2]$ ). Consider the regret criterion (1). The maximum risk of  $d \in \{0, 1\}$  is  $R(0, P) = (1+c)(P)_+$  and  $R(1, P) = -(P)_-$ . Therefore,

$$d^o(P) = \mathbb{I}[P \geq 0] \quad \text{and} \quad \min_{d \in \mathcal{D}} R(d, P) = \min \{(1+c)(P)_+, -(P)_-\} = 0. \quad (10)$$

In any finite sample, the DM cannot determine with certainty whether  $P$  is positive. Let  $P_0 = 0$ , so that the local asymptotic experiment mimics the finite-sample counterpart. This

---

<sup>10</sup>The scaling depends on the structure of the problem. Here we implicitly assume that there exists a  $\sqrt{n}$  consistent estimator of  $P$ .

<sup>11</sup>An alternative to (9) would be to take the maximum over  $h$ , which has the downside that the optimal decisions become more difficult to compute. Because we are considering non-standard decision problems in this paper, different optimality concepts generally lead to different decisions.

leads to  $P_{n,h} = h/\sqrt{n}$  for  $h \in \mathbb{R}$ . Thus,

$$\mathbb{E}_{P_{n,h}} [R(d_n(X^n), P_{n,h})] = \frac{1}{\sqrt{n}} \left( (1 - \mathbb{E}_{P_{n,h}} [d_n(X^n)])(1+c)(h)_+ - \mathbb{E}_{P_{n,h}} [d_n(X^n)](h)_- \right). \quad (11)$$

The asymptotic average maximum risk is obtained by plugging the minimum in (10) and the expression in (11) into (9) to obtain

$$\begin{aligned} \mathcal{R}(\{d_n\}; P_0) &= (1+c) \int_0^\infty h \left[ 1 - \left( \lim_{n \rightarrow \infty} \mathbb{E}_{P_{n,h}} [d_n(X^n)] \right) \right] dh - \int_{-\infty}^0 h \left( \lim_{n \rightarrow \infty} \mathbb{E}_{P_{n,h}} [d_n(X^n)] \right) dh, \end{aligned}$$

where the limit is a well-defined function of  $h$  under the regularity conditions below.  $\square$

In the remainder of this section we present a Bayesian implementation and a bootstrap implementation and revisit the two applications.

### 2.3 Bayesian Interpretation and Implementation

Let  $\pi$  denote a strictly positive, smooth density on  $\mathcal{P}$ . Integrating the frequentist maximum risk criterion (7) across  $P$  using  $\pi$  yields the *integrated maximum risk* criterion

$$\int \mathbb{E}_P [R(d_n(X^n), P)] d\pi(P). \quad (12)$$

After observing the data  $X^n$ , the DM can form a posterior  $\pi_n(P|X^n)$  for  $P$ . Standard arguments (e.g. Wald, 1950, Chapter 5.1) imply that the integrated maximum risk can be minimized by minimizing the posterior maximum risk

$$\bar{R}_n(d|X^n) = \int R(d, P) d\pi_n(P|X^n) \quad (13)$$

with respect to  $d$  for almost every realization of the data. This leads to the Bayes decision rule

$$d_n^*(X^n; \pi) = \arg \min_{d \in \mathcal{D}} \bar{R}_n(d|X^n) \quad (14)$$

if the argmin is unique, otherwise  $d_n^*(X^n; \pi)$  is chosen randomly from  $\arg \min_{d \in \mathcal{D}} \bar{R}_n(d|X^n)$ . We include  $\pi$  as an argument in the decision rule  $d_n^*(\cdot)$  to indicate that the decision depends on the prior in finite samples.

We formally prove in Section 3 that Bayes decisions  $d_n^*(X^n; \pi)$  are optimal under the asymptotic efficiency criterion (9). We therefore refer to  $d_n^*(X^n; \pi)$  as an *efficient-robust* decision rule in what follows. In regard to the prior, we only require that its density is strictly positive and continuous at  $P_0$ . To see the connection between (12) and (9), use a change-of-

variables to express the prior for  $h$  as  $n^{-1/2}\pi(P_0 + h/\sqrt{n})$ . This prior becomes uniform as  $n \rightarrow \infty$ , which is the weight function underlying (9). Minimizing integrated maximum risk is a convenient device that leads to the tractable decision rule (14) with desirable large-sample (frequentist) efficiency properties.

**Application 1: Treatment Assignment under Partial Identification (continued).**

Recall that the set  $\Theta_0(P)$  consists of all probability distributions of  $Y = Y_1 - Y_0$  for which  $b_L(P) \leq \mathbb{E}_\theta[Y] \leq b_U(P)$ . The maximum risk of  $d \in \{0, 1\}$  is

$$\begin{aligned} R(0, P) &= \sup_{\theta \in \Theta_0(P)} (\mathbb{E}_\theta[Y])_+ = (b_U(P))_+, \\ R(1, P) &= \sup_{\theta \in \Theta_0(P)} ((\mathbb{E}_\theta[Y])_+ - \mathbb{E}_\theta[Y]) = -(b_L(P))_-. \end{aligned}$$

Because  $R(1, P) \leq R(0, P)$  if and only if  $-(b_L(P))_- \leq (b_U(P))_+$ , we see that

$$d^o(P) = \mathbb{I}[(b_U(P))_+ + (b_L(P))_- \geq 0] \tag{15}$$

is an oracle decision.<sup>12</sup>

Averaging the maximum risks  $R(1, P)$  and  $R(0, P)$  across the posterior for  $P$  yields

$$\bar{R}_n(0|X^n) = \int (b_U(P))_+ d\pi_n(P|X^n), \quad \bar{R}_n(1|X^n) = \int -(b_L(P))_- d\pi_n(P|X^n).$$

It follows that

$$d_n^*(X^n; \pi) = \mathbb{I} \left[ \int (b_U(P))_+ + (b_L(P))_- d\pi_n(P|X^n) \geq 0 \right] \tag{16}$$

is an efficient-robust decision. That is, the efficient-robust decision is to treat if the posterior mean of the “robust” welfare contrast  $b(P) := (b_U(P))_+ + (b_L(P))_-$  is positive, don’t treat if it is negative, and choose (possibly randomly) between treatment and no treatment when it is zero. The above rule deterministically chooses treatment in the latter case, though any (possibly randomized) tie-breaking rule leads to an optimal decision under our asymptotic efficiency criterion. We discuss different bounds and how to form posteriors in a number of different empirical settings in Section 5.  $\square$

**Application 2: Optimal Pricing with Rich Unobserved Heterogeneity (continued).**

In the example given above we saw  $\mathbb{E}_\theta[u(1, Y, \theta, P)] = -\mathbb{E}[h(X_1)] + C$  was point-identified and

---

<sup>12</sup>The oracle decision is not unique because both  $d = 0$  and  $d = 1$  are optimal when  $(b_U(P))_+ + (b_L(P))_- = 0$ . Any decision that sets  $d = 1$  (respectively, 0) when  $(b_U(P))_+ + (b_L(P))_- > 0$  (respectively,  $< 0$ ) and randomizes between  $d = 0$  and  $d = 1$  when  $(b_U(P))_+ + (b_L(P))_- = 0$  is an oracle decision and has the same maximum risk  $\min_{d \in \{0, 1\}} R(d, P)$  as  $d^o(P)$ .

we derived the sharp lower bound for  $\mathbb{E}_\theta[u(0, Y, \theta, P)]$  in (5). The maximum risk of  $d \in \{0, 1\}$  is

$$\begin{aligned} R(0, P) &= y_h - \max \{ (y_h - y_l)P_1(s_{31}), (y_h - y'_l)P_2(s_{32}) \}, \\ R(1, P) &= \mathbb{E}[h(X_1)] - C, \end{aligned}$$

where  $P = (\mathbb{E}[h(X_1)], P_1(s_{31}), P_2(s_{32}))$ . Define

$$b(P) = y_h - \max \{ (y_h - y_l)P_1(s_{31}), (y_h - y'_l)P_2(s_{32}) \} - \mathbb{E}[h(X_1)] + C.$$

An oracle decision is

$$d^o(P) = \mathbb{I}[b(P) \geq 0].$$

Averaging the contrast  $b(P)$  across the posterior for  $P$  yields

$$\bar{R}_n(0|X^n) - \bar{R}_n(1|X^n) = \int b(P) d\pi_n(P|X^n).$$

It follows that an efficient-robust decision is to adopt the tax policy if the posterior mean of  $b(P)$  is positive

$$d^*(X^n; \pi) = \mathbb{I} \left[ \int b(P) d\pi_n(P|X^n) \geq 0 \right].$$

The posterior expectation can be constructed using a Bayesian Generalized Method of Moments (GMM) approach. Because there are no over-identifying moment conditions here, the method proposed in [Chamberlain and Imbens \(2003\)](#) would be suitable.  $\square$

To implement the efficient-robust decision, the DM only needs to be able to compute maximum risk  $R(d, P)$  as a function of  $P$  for each choice  $d$ . Therefore, our framework is only restrictive insofar as the optimization problem (6) must be simple enough to solve across repeated draws from the posterior for  $P$ . In many cases, such as the treatment choice or pricing policies examples,  $R(d, P)$  may be computed in closed form or, when analytical solutions aren't available, by linear programming; see Sections 5 and 6.

## 2.4 Bootstrap Implementation

We now present an alternative implementation of efficient-robust decision rules using the bootstrap. Let  $\hat{P}$  denote an efficient estimator of  $P$ . Define the bootstrap average maximum risk

$$R_n^*(d) = \mathbb{E}_n^* \left[ R(d, \hat{P}^*) \right],$$

where  $\mathbb{E}_n^*$  denotes expectation with respect to the bootstrap version  $\hat{P}^*$  of  $\hat{P}$  conditional on the data  $X^n$ . The bootstrap efficient-robust decision rule is

$$d_n^{**}(X^n) = \arg \min_{d \in \mathcal{D}} R_n^*(d)$$

if the argmin is unique (with a random selection from the argmin otherwise). This rule is asymptotically equivalent to  $d_n^*$  under mild conditions ensuring asymptotic equivalence of the bootstrap distribution for  $\hat{P}^*$  and posterior distribution for  $P$ . As such,  $d_n^{**}$  will inherit the asymptotic efficiency properties of the efficient-robust decision rule  $d_n^*$ .

**Applications (continued).** Averaging the maximum risks  $R(1, P)$  and  $R(0, P)$  across the bootstrap distribution of  $\hat{P}^*$  for the treatment assignment under partial identification yields

$$R_n^*(1|X^n) = \mathbb{E}_n^* \left[ -(b_L(\hat{P}^*))_- \right], \quad R_n^*(0|X^n) = \mathbb{E}_n^* \left[ (b_U(\hat{P}^*))_+ \right].$$

Therefore,

$$d^{**}(X^n) = \mathbb{I} \left[ \mathbb{E}_n^* \left[ (b_U(\hat{P}^*))_+ + (b_L(\hat{P}^*))_- \right] \geq 0 \right]$$

minimizes bootstrap average maximum risk. A similar argument for the optimal pricing example yields  $d^{**}(X^n) = \mathbb{I} [\mathbb{E}_n^* [b(P)] > 0]$ .  $\square$

## 2.5 Further Discussion

In the remainder of this section we discuss the interpretation of our approach in the context of a two-player zero-sum game and relate it to the literature on robust Bayes decision making.

**Two-Player Game.** Minimax decisions can be interpreted as optimal strategies in a zero-sum game played between the DM and an adversary (“nature”). Recall the case of “Known  $P$ ” discussed in Section 2.2. In a zero-sum game the supremum over  $\theta \in \Theta_0(P)$  in the maximum risk definition (6) can be viewed as nature’s best response to the DM choosing  $d$ .

The interpretation of the “Unknown  $P$ ” case is more delicate. Consider the Bayesian construction of the decision rule. Substitute the definition of  $R(d, P)$  from (6) into the integrated maximum risk criterion (12) and manipulate the resulting expression as follows:

$$\begin{aligned} & \min_{d: \mathcal{X}^n \rightarrow \mathcal{D}} \int_{\mathcal{P}} \mathbb{E}_P \left[ \sup_{\theta \in \Theta_0(P)} \mathbb{E}_\theta [-u(d(X^n), Y, \theta, P)] \right] d\pi(P) \\ &= \int_{\mathcal{X}^n} \min_{d(X^n) \in \mathcal{D}} \left[ \int_{\mathcal{P}} \left( \sup_{\theta \in \Theta_0(P)} \mathbb{E}_\theta [-u(d(X^n), Y, \theta, P)] \right) d\pi_n(P|X^n) \right] d\pi_n(X^n). \end{aligned}$$

On the right-hand side we factorize the joint distribution of  $(X^n, P)$  into the posterior  $\pi_n(P|X^n)$  and the marginal distribution of the data  $\pi_n(X^n)$ . The exchange of minimization

with respect to  $d(X^n)$  and integration over  $X^n$  is satisfied under weak regularity conditions because the DM makes a decision conditional on  $X^n$ . We used this calculation previously to obtain the Bayes decision  $d_n^*(X^n; \pi)$  in (14).

From a minimax perspective, the non-standard feature of our setup is that the supremum over  $\theta$  is taken inside the posterior expectation. This can be justified by assuming that in a zero-sum game nature is allowed to choose  $\theta$  conditional on knowing  $X^n$  and  $P$ , after the prior  $\pi(\cdot)$  has been set and  $(P, X^n)$  have been sampled from their joint distribution.

**Relationship to Robust Bayesian Decision Making.** Our setup is closely related to robust Bayes (or  $\Gamma$ -minimax) decision making (Robbins, 1951; Berger, 1985) and the multiple priors framework of Gilboa and Schmeidler (1989). A key difference is that we distinguish two groups of parameters,  $\theta$  and  $P$ , and apply the minimax reasoning only to  $\theta$  because only it is partially identified. As mentioned previously, the approach of treating groups of parameters differently dates back to Hurwicz (1951). He argued that in econometrics and other applications researchers might consider multiple priors rather than a single prior for some of the parameters and referred to it as generalized Bayes-minimax principle. He provided a two-parameter example in which the marginal prior distribution for one of the parameters,  $P$  in our notation, is fixed, whereas a large family of priors is considered for the conditional distribution of the second parameter, which would be  $\theta$ , given  $P$ .

Suppose one combines the unique prior  $\pi(\cdot)$  with a family  $\Lambda$  of conditional priors  $\lambda(\theta|P)$  for  $\theta$  given  $P$  with support  $\Theta_\theta(P)$  for each  $P \in \mathcal{P}$ . Then one can define  $\Gamma = \{\lambda(\theta|P) \cdot \pi(P) : \lambda \in \Lambda\}$ . The  $\Gamma$ -minimax decision rule minimizes the maximum Bayes risk over  $\gamma = \lambda \cdot \pi$  on each  $X^n$  trajectory

$$\begin{aligned} & \min_{d: \mathcal{X}^n \rightarrow \mathcal{D}} \sup_{\gamma \in \Gamma} \int_{\mathcal{X}^n} \left( \int_{\mathcal{P}} \left[ \int_{\Theta} \mathbb{E}_\theta[-u(d(X^n), Y, \theta, P)] d\lambda(\theta|P) \right] d\pi_n(P|X^n) \right) d\pi_n(X^n) \\ &= \int_{\mathcal{X}^n} \min_{d(X^n) \in \mathcal{D}} \sup_{\gamma \in \Gamma} \left( \int_{\mathcal{P}} \left[ \int_{\Theta} \mathbb{E}_\theta[-u(d(X^n), Y, \theta, P)] d\lambda(\theta|P) \right] d\pi_n(P|X^n) \right) d\pi_n(X^n). \end{aligned}$$

The key insight is that the marginal distribution of the data,  $\pi_n(X^n)$ , does not depend on nature's choice of  $\gamma$  because conditional on  $P$  the distribution of  $X^n$  does not depend on  $\theta$ . This allows us to move the supremum over  $\gamma \in \Gamma$  inside of the integral. Moreover, it implies that the posterior of  $\theta|(P, X^n)$  is equal to the prior distribution of  $\theta|P$ .

Let  $\theta_*(P) \in \arg \sup_{\theta \in \Theta_\theta(P)} \mathbb{E}_\theta[-u(d, Y, \theta, P)]$ , assuming that the arg sup is non-empty, and let  $\lambda_*(\theta|P)$  be a pointmass at  $\theta_*(P)$ . Then, by construction

$$\int_{\mathcal{P}} \left[ \int_{\Theta} \mathbb{E}_\theta[-u(d, Y, \theta, P)] d\lambda_*(\theta|P) \right] d\pi_n(P|X^n) = \int_{\mathcal{P}} \left[ \sup_{\theta \in \Theta_\theta(P)} \mathbb{E}_\theta[-u(d, Y, \theta, P)] \right] d\pi_n(P|X^n). \quad (17)$$

Moreover, for every  $P$

$$\sup_{\theta \in \Theta_0(P)} \mathbb{E}_\theta[-u(d, Y, \theta, P)] \geq \int \mathbb{E}_\theta[-u(d, Y, \theta, P)] d\lambda(\theta|P). \quad (18)$$

Assuming that  $\lambda_*(\theta|P) \cdot \pi(P) \in \Gamma$ , we can deduce from (17) and (18) that

$$\int_{\mathcal{P}} \left[ \sup_{\theta \in \Theta_0(P)} \mathbb{E}_\theta[-u(d, Y, \theta, P)] \right] d\pi_n(P|X^n) = \sup_{\gamma \in \Gamma} \int_{\mathcal{P}} \left[ \int_{\Theta} \mathbb{E}_\theta[-u(d, Y, \theta, P)] \lambda(\theta|P) d\theta \right] d\pi_n(P|X^n).$$

Thus, the efficient-robust decision rule may be viewed as a  $\Gamma$ -minimax decision.

### 3 Optimality Results

We now present the main optimality results for the parametric case. Section 3.1 contains two main results. First, Theorem 1 shows that the Bayes decision rules  $d_n^*(\cdot; \pi)$ , defined in (14), are optimal under our asymptotic efficiency criterion (9). Moreover, any decision that behaves asymptotically like  $d_n^*(\cdot; \pi)$  is also optimal. Optimality of the bootstrap-based decision  $d_n^{**}$  then follows from this fact, under suitable regularity conditions. Second, Theorem 2 shows that any decision whose asymptotic behavior is different from  $d_n^*(\cdot; \pi)$  is dominated by  $d_n^*(\cdot; \pi)$  under our optimality criterion. Specifically, we show in Section 3.2 that plug-in decisions are dominated, in applications in which the maximum risk  $R(d, P)$  defined in (6) is a non-smooth function of  $P$ . As we discuss further in Section 5, this finding has important implications for treatment assignment under partial identification, where in many empirically relevant applications the maximum risk is non-smooth in  $P$ .

#### 3.1 Optimality of Bayes and Bootstrap Decisions

Appendix A presents Assumptions 1, 2, and 3, which are the main regularity conditions. Assumption 1 imposes continuity and directional differentiability conditions on the maximum risk. Assumption 2 states that the model for  $X^n$  is locally asymptotically normal at each  $P_0 \in \mathcal{P}$ . Assumption 3 contains high-level consistency and asymptotic normality assumptions for functionals of the posterior distribution. These conditions implicitly restrict the priors  $\pi$  to be positive and continuous at every  $P_0 \in \mathcal{P}$ . Let

$$\mathbb{D} = \left\{ \{d_n\} : d_n(X^n) \overset{P_{n,h}}{\rightsquigarrow} Q_{P_0,h} \text{ for all } h \in \mathbb{R}^k, P_0 \in \mathcal{P} \right\},$$

where  $\overset{P_{n,h}}{\rightsquigarrow}$  denotes convergence in distribution along a sequence  $\{F_{n,P_{n,h}}\}$  with  $X^n \sim F_{n,P_{n,h}}$  for each  $n$ , and  $Q_{P_0,h}$  is the (possibly degenerate) limiting probability measure on  $\mathcal{D}$ . The set  $\mathbb{D}$  represents the set of all “well behaved” sequences of decision rules that converge in

distribution under each sequence  $\{P_{n,h}\}$  for each fixed perturbation direction  $h$ . In general,  $Q_{P_0,h}$  will depend on the sequence  $\{d_n\}$ . We say that two sequences of decisions  $\{d_n\}$  and  $\{d'_n\}$  are *asymptotically equivalent* if  $d_n(X^n)$  and  $d'_n(X^n)$  have the same asymptotic distribution along  $\{P_{n,h}\}$  for all  $P_0 \in \mathcal{P}$  and  $h \in \mathbb{R}^k$ .

**Theorem 1** *Suppose that Assumptions 1 and 2 are satisfied. Let  $\Pi$  be the class of priors such that Assumption 3 holds.*

- (i) *For any  $\pi \in \Pi$ , the Bayes decision  $\{d_n^*(\cdot; \pi)\}$  is asymptotically efficient-robust.*
- (ii) *For any  $\pi, \pi' \in \Pi$ , the Bayes decisions  $\{d_n^*(\cdot; \pi)\}$  and  $\{d_n^*(\cdot; \pi')\}$  are asymptotically equivalent.*
- (iii) *If  $\{d_n\}$  is asymptotically equivalent to  $\{d_n^*(\cdot; \pi)\}$  for a  $\pi \in \Pi$ , then  $\{d_n\}$  is asymptotically efficient-robust.*

According to parts (i) and (ii) of Theorem 1, all Bayes decisions based on priors that satisfy some mild regularity conditions are asymptotically equivalent and also efficient-robust. Part (iii) implies that a decision that is not a Bayes decision but is asymptotically equivalent to a Bayes decision under a prior  $\pi \in \Pi$  is asymptotically efficient-robust. Asymptotic optimality of the bootstrap implementation discussed in Section 2.4 follows from this result.

Asymptotic equivalence to a Bayes decision is in fact a *necessary* condition for optimality under a side condition ruling out the absence of ties. Say  $\{d_n\}$  and  $\{d'_n\}$  fail to be asymptotically equivalent at  $P_0$  if  $d_n(X^n)$  and  $d'_n(X^n)$  have different asymptotic distributions along  $\{P_{n,h_0}\}$  for some  $h_0 \in \mathbb{R}^k$ . Let  $\rho_{d,P_0}[\cdot]$  denote the directional derivative of  $R(d, P)$  with respect to  $P$  at  $P_0$  and let  $\mathbb{E}^*$  denote expectation with respect to  $Z^* \sim N(0, I_0^{-1})$  where  $I_0 = I_0(P_0)$  is the (asymptotic) information matrix at  $P_0$ —see Appendix A for definitions.

**Theorem 2** *Suppose that Assumptions 1 and 2 are satisfied. Let  $\Pi$  be the class of priors such that Assumption 3 holds. Suppose that  $\{d_n\} \in \mathbb{D}$  is not asymptotically equivalent to  $\{d_n^*(\cdot; \pi)\}$  for  $\pi \in \Pi$ . Then at any  $P_0$  at which asymptotic equivalence fails,*

$$\mathcal{R}(\{d_n\}; P_0) > \mathcal{R}(\{d_n^*(\cdot; \pi)\}; P_0) \quad \text{for all } \pi \in \Pi$$

*provided either of the following hold:*

- (i)  *$\arg \min_{d \in \mathcal{D}} R(d, P_0)$  is a singleton;*
- (ii)  *$\arg \min_{d \in \mathcal{D}} R(d, P_0)$  is not a singleton, and  $\arg \min_{d \in \mathcal{D}} \mathbb{E}^*[\rho_{d,P_0}[Z^* + z]]$  is a singleton for almost every  $z$ .*

We conclude by giving a slightly stronger optimality statement than that presented in

Theorem 1. Define

$$\mathcal{R}_n(d_n; P_0) = \int_{\{h: P_{n,h} \in \mathcal{P}\}} \sqrt{n} \left( \mathbb{E}_{P_{n,h}} [R(d_n(X^n), P_{n,h})] - \min_{d \in \mathcal{D}} R(d, P_{n,h}) \right) \pi(P_{n,h}) dh,$$

where  $\pi$  is the density used to construct  $d_n^*(X^n; \pi)$ . Under mild conditions permitting the exchange of limits and integration (see condition (A.5) in the Appendix), we have

$$\lim_{n \rightarrow \infty} \mathcal{R}_n(d_n; P_0) = \pi(P_0) \cdot \mathcal{R}(\{d_n\}; P_0).$$

As the scale factor  $\pi(P_0)$  is independent of the sequence  $\{d_n\}$ , the criterions  $\lim_{n \rightarrow \infty} \mathcal{R}_n(d_n; P_0)$  and  $\mathcal{R}(\{d_n\}; P_0)$  induce the same rankings over sequences of decision rules.

**Corollary 1** *Suppose that Assumptions 1 and 2 are satisfied. Let  $\Pi$  be the class of priors such that Assumption 3 holds. Let  $\{d_n\} \in \mathbb{D}$  be asymptotically equivalent to  $\{d_n^*(\cdot; \pi)\}$  for a  $\pi \in \Pi$  and satisfy condition (A.5). Then for all  $P_0 \in \mathcal{P}$ ,*

$$\lim_{n \rightarrow \infty} \mathcal{R}_n(d_n; P_0) = \liminf_{n \rightarrow \infty} \inf_{d'_n} \mathcal{R}_n(d'_n; P_0),$$

where the infimum is over all measurable  $d'_n : \mathcal{X}^n \rightarrow \mathcal{D}$ .

### 3.2 Plug-in Rules are Dominated

A seemingly natural alternative to  $d_n^*(X^n; \pi)$  is to plug an efficient estimator  $\hat{P}$  of  $P$  into the oracle decision rule, which yields the plug-in rule  $d^o(\hat{P})$ . Manski (2021a,b) refers to this approach as “as-if optimization”: the estimated parameters  $\hat{P}$  are treated “as if” they are the true parameters, with  $d^o(\hat{P})$  minimizing maximum risk (or regret) over  $\Theta_0(\hat{P})$ . This approach also has connections with anticipated utility (Kreps, 1998; Cogley and Sargent, 2008).

Consider the treatment assignment example. Hirano and Porter (2009) show that plug-in rules  $d^o(\hat{P}) = \mathbb{I}[g(\hat{P}) \geq 0]$  are optimal under asymptotic efficiency criteria similar to ours when the ATE is point identified and a welfare contrast  $g$  depends smoothly on a regularly estimable parameter  $P$ . For the intuition, by the Bernstein–von Mises theorem and differentiability of  $g$ , the posterior mean  $\bar{g}_n$  of  $g(P)$  is (first-order) asymptotically equivalent to  $g(\hat{P})$  under suitable regularity conditions:  $\bar{g}_n = g(\hat{P}) + o_p(n^{-1/2})$ . As such,  $d^o(\hat{P})$  is asymptotically equivalent to  $d_n^*$  and inherits its optimality.

Under partial identification, the oracle and efficient-robust decisions depend on  $P$  through the robust welfare contrast  $b(P) = (b_U(P))_+ + (b_L(P))_-$ . Here  $b$  may fail to be differentiable, either because of the outer  $(\cdot)_+$  and  $(\cdot)_-$  operations or because the bounds  $b_L(P)$  and  $b_U(P)$  can often depend non-smoothly on  $P$ . As we discuss further in Section 5, non-smoothness of the bounds in  $P$  is a fairly common feature across a wide range of empirical settings—see also

the optimal pricing application in Sections 2.1 and 6. When  $b$  is non-smooth in  $P$ , the plug-in term  $b(\hat{P})$  and the posterior mean  $\bar{b}_n$  of  $b(P)$  may fail to be asymptotically equivalent. In consequence, the oracle plug-in rule  $d^o(\hat{P}) = \mathbb{I}[b(\hat{P}) \geq 0]$  and efficient-robust decision  $d_n^*(X^n) = \mathbb{I}[\bar{b}_n \geq 0]$  may fail to be asymptotically equivalent, in which case  $d_n^*$  will dominate  $d^o(\hat{P})$  under our asymptotic efficiency criterion.

**Example (continued).** Recall the example from Section 2.2, in which  $b_L(P) = (P)_-$  and  $b_U(P) = (1+c)(P)_+$ , where  $c > -1$ . The robust welfare contrast is  $b(P) = P(1 + c\mathbb{I}[P \geq 0])$ . Moreover, recall from (10) that irrespective of the value of  $c$ , the oracle decision is given by  $d^o(P) = \mathbb{I}[P \geq 0]$ . The oracle plug-in rule is therefore  $d^o(\hat{P}) = \mathbb{I}[\sqrt{n}\hat{P} \geq 0]$ .

For the efficient-robust decision, suppose  $\sqrt{n}(\hat{P} - P_0) \overset{P_{n,h}}{\rightsquigarrow} N(h, 1)$  and the posterior for  $P$  is  $N(\hat{P}, n^{-1})$ . The posterior mean  $\bar{b}_n$  of  $b(P)$  is

$$\bar{b}_n = \sqrt{n}\hat{P}(1 + c\Phi(\sqrt{n}\hat{P})) + c\phi(\sqrt{n}\hat{P}),$$

where  $\phi$  and  $\Phi$  are the standard normal pdf and cdf. The efficient-robust decision is therefore

$$d_n^*(X^n; \pi) = \mathbb{I}\left[\sqrt{n}\hat{P} \geq -\frac{c\phi(\sqrt{n}\hat{P})}{1 + c\Phi(\sqrt{n}\hat{P})}\right].$$

Inspecting the formulas for  $d_n^*$  and  $d^o(\hat{P})$ , we see that the two decisions are equivalent in the symmetric case ( $c = 0$ ). However, when  $c > 0$ , the efficient-robust decision is to treat if  $\sqrt{n}\hat{P}$  exceeds a negative threshold. Intuitively, there is relatively more upside from treating than downside from not treating when  $\hat{P}$  is close to zero and  $c > 0$ . The efficient-robust decision recognizes this and recommends treatment more aggressively than the plug-in rule.

We again center the local asymptotic experiment around  $P_0 = 0$  to mimic the finite-sample counterpart where the DM does not know the sign of  $P$  with certainty. In this case,  $\sqrt{n}\hat{P}$  converges in distribution to a  $N(h, 1)$  random variable along each  $\{P_{n,h}\}$ , so this difference between the decisions persists asymptotically. As such, Theorem 2 implies that the efficient-robust decision dominates the plug-in rule under our asymptotic efficiency criterion.<sup>13</sup>

Figure 2 plots the asymptotic frequentist maximum risk (8) for  $d_n^*$  and  $d^o(\hat{P})$  as a function of the perturbation direction  $h$  when  $P_0 = 0$ . The area under the curves is our asymptotic efficiency criterion (9). Figure 2a plots the symmetric case ( $c = 0$ ), where the two decisions are equivalent. Figures 2b and 2c plot asymmetric cases for  $c = 1, 2$ . Here the area under the blue curve (corresponding to  $d_n^*$ ) is smaller than under the black dotted curve (corresponding to  $d^o(\hat{P})$ ), as predicted by Theorem 2. Similar results are obtained for other values of  $c$ .

---

<sup>13</sup>If  $P_0 \neq 0$  then  $\sqrt{n}\hat{P}$  diverges and both rules choose the optimal treatment with probability approaching one.

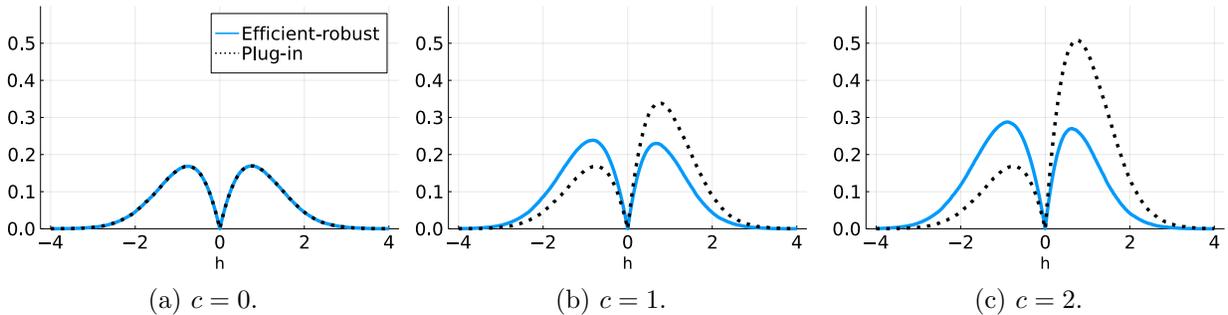


Figure 2: Asymptotic frequentist maximum risk (8) for  $d_n^*$  (efficient-robust) and  $d^o(\hat{P})$  (oracle plug-in) a function of  $h$  for different values of  $c$ .

Inspecting Figure 2, we see that the efficient-robust decision also dominates the plug-in rule under an asymptotic criterion that takes the maximum, rather than the average, of (8) over  $h$ , namely

$$\mathcal{M}(\{d_n\}; P_0) := \sup_h \lim_{n \rightarrow \infty} \sqrt{n} \left( \mathbb{E}_{P_{n,h}} [R(d_n(X^n), P_{n,h})] - \min_{d \in \mathcal{D}} R(d, P_{n,h}) \right).$$

Moreover, the relative efficiency of  $d_n^*$  is increasing in the level of asymmetry  $c$  for both this criterion and our asymptotic efficiency criterion (9).

The failure of asymptotic equivalence of the efficient-robust and plug-in rules when  $P_0 = 0$  is explained by a failure of differentiability of the robust welfare contrast  $b(P)$  at  $P = 0$ . Note that

$$\dot{b}_0[h] := \lim_{t \downarrow 0} \frac{b(th) - b(0)}{t} = \begin{cases} (1+c)h & \text{if } h \geq 0, \\ h & \text{if } h < 0. \end{cases}$$

Therefore,  $b(P)$  is differentiable at  $P = 0$  when  $c = 0$  and directionally differentiable otherwise. When  $c = 0$ , the plug-in contrast  $b(\hat{P})$  and the posterior mean  $\bar{b}_n$  of  $b(P)$  are equal, so the plug-in and efficient-robust decisions agree. When  $c \neq 0$ ,  $\sqrt{n}b(\hat{P}) = \sqrt{n}\hat{P}(1 + c\mathbb{I}[\sqrt{n}\hat{P} \geq 0])$  while  $\sqrt{n}\bar{b}_n = \sqrt{n}\hat{P}(1 + c\Phi(\sqrt{n}\hat{P})) + c\phi(\sqrt{n}\hat{P})$ . As we see, the lack of differentiability of  $b(P)$  at  $P = 0$  leads to a failure of the  $\delta$ -method, whereby the asymptotic distribution of  $\sqrt{n}b(\hat{P})$  and  $\sqrt{n}\bar{b}_n$  will be different and, consequently, the efficient-robust and plug-in rules are not asymptotically equivalent.

Section 5 presents a number of different empirical settings in which the bounds  $b_L(P)$  and  $b_U(P)$ , and therefore the robust welfare contrast  $b(P)$ , can fail to be differentiable in  $P$ . Similar reasoning applies in these settings: a lack of differentiability means the efficient-robust and plug-in rules may fail to be asymptotically equivalent, in which case the efficient-robust rules will dominate the plug-in rules under our asymptotic efficiency criterion.  $\square$

## 4 Semiparametric Models

This section extends our approach to semiparametric models. This extension is relevant for handling a number of cases that may arise in practice, including the treatment assignment and optimal pricing applications presented in Sections 5 and 6. For instance, the DM may not have grounds for asserting that the data are drawn from a given parametric model. Alternatively,  $X^n$  may be a vector of moments or summary statistics whose exact finite-sample distribution is unknown. We first state the model in Section 4.1 and then generalize our notion of asymptotically efficient-robust decision rules to semiparametric models in Section 4.2. Section 4.3 describes a quasi-Bayesian implementation of efficient-robust decisions in this setting. Section 4.4 presents the main optimality results. In particular, Theorem 3 shows that quasi-Bayes decision rules  $d_n^*(\cdot; \pi)$  are asymptotically efficient-robust. Moreover, Theorem 4 shows that any decision whose asymptotic behavior is different from  $d_n^*(\cdot; \pi)$  is dominated by  $d_n^*(\cdot; \pi)$ .

### 4.1 Model

Let  $X^n \sim F_{n,\beta}$  with  $\beta = (P, \eta)$  for  $P \in \mathcal{P} \subseteq \mathbb{R}^k$  and  $\eta \in \mathcal{H}$ , an infinite-dimensional space. For instance, in a GMM model the parameter  $\eta$  is the marginal distribution of each observation  $X_i$  and  $\mathcal{H} = \{\eta : \int g(x, P) d\eta(x) = 0 \text{ for some } P \in \mathcal{P}\}$  for a vector of moment functions  $g$ .

We again assume that given  $P$ , the structural parameter takes values in a set  $\Theta_0(P)$ . Therefore, the set of payoff distributions is indexed only by the parametric component  $P$ . The nonparametric component  $\eta$  is a nuisance parameter. For instance,  $P$  may be a vector of population moments used to construct bounds on  $\theta$ . The nuisance parameter  $\eta$  represents other features of the distribution of  $X^n$  that are irrelevant for the DM's decision problem.

### 4.2 Optimality Criterion and Efficient-Robust Decision Rules

In the parametric case, our asymptotic efficiency criterion (9) integrates the excess maximum risk incurred by  $d_n(X^n)$  relative to the oracle under  $P_{n,h}$  using Lebesgue measure on the local perturbations  $h \in \mathbb{R}^k$  of  $P_0$ . This approach does not extend easily to perturbations of  $(P_0, \eta_0)$  in the semiparametric case due to measure-theoretic complications that arise in infinite-dimensional spaces.

We therefore form our asymptotic efficiency criterion by averaging across local perturbations of  $P_0$  within an *approximately least-favorable submodel*, in which  $X^n$  carries the least information about  $P$  of all parametric submodels. The problem of parameter estimation in the least-favorable submodel is asymptotically equivalent to the problem of estimating  $P$  in full the semiparametric model. A formal definition of the least-favorable submodel is presented in Appendix A. For now, we let  $t \mapsto \beta_{(P_0, \eta_0)}(t)$  denote the mapping from an open

neighborhood  $\mathcal{P}(P_0, \eta_0)$  of  $P_0$  into  $\mathcal{P} \times \mathcal{H}$  under the least-favorable submodel at  $(P_0, \eta_0)$ . Our asymptotic efficiency criterion is analogous to criterion (9), namely

$$\mathcal{R}(\{d_n\}; (P_0, \eta_0)) := \int \lim_{n \rightarrow \infty} \sqrt{n} \left( \mathbb{E}_{\beta_{(P_0, \eta_0)}(P_{n,h})} [R(d_n(X^n), P_{n,h})] - \min_{d \in \mathcal{D}} R(d, P_{n,h}) \right) dh, \quad (19)$$

where  $\mathbb{E}_{\beta_{(P_0, \eta_0)}(P_{n,h})}$  denotes expectation with respect to  $X^n \sim F_{n, \beta_{(P_0, \eta_0)}(P_{n,h})}$ .

In this context, we say a sequence of decision rules  $\{d_n\}$  is asymptotically efficient-robust if it minimizes  $\mathcal{R}(\cdot; (P_0, \eta_0))$  over the class  $\mathbb{D}$  of decision rules for which the limit in (19) is well defined, for all  $(P_0, \eta_0) \in \mathcal{P} \times \mathcal{H}$ .

### 4.3 Quasi-Bayesian Implementation

Efficient-robust decisions are formed similarly to the parametric case, but we now replace the (parametric) posterior with a *quasi-posterior* formed from a limited-information criterion for  $P$  and a prior  $\pi$  on  $\mathcal{P}$ .

Suppose first that the DM may compute an efficient estimator  $\hat{P}$  of  $P$  from  $X^n$  such that for each  $(P_0, \eta_0) \in \mathcal{P} \times \mathcal{H}$ ,  $\sqrt{n}(\hat{P} - P_0) \rightsquigarrow N(0, I_0^{-1})$  with  $I_0 = I_0(P_0, \eta_0)$  the semiparametric information bound. Also suppose that the DM may compute a consistent estimator  $\hat{I}$  of  $I_0$ . Following [Doksum and Lo \(1990\)](#) [Kim \(2002\)](#), the DM could use a limited information  $N(\hat{P}, (n\hat{I})^{-1})$  quasi-likelihood for  $P$  which leads to the *quasi-posterior*

$$\pi_n(P|X^n) \propto e^{-\frac{1}{2}(P-\hat{P})'(n\hat{I})(P-\hat{P})} \pi(P),$$

for some strictly positive, continuously differentiable density  $\pi$ . Other limited-information criteria  $Q_n(P)$  (e.g. GMM, minimum distance, or simulated method of moments) could also be used following [Chernozhukov and Hong \(2003\)](#) (see also [Chen, Christensen, and Tamer \(2018\)](#)), in which case

$$\pi_n(P|X^n) \propto e^{Q_n(P)} \pi(P).$$

While we are deliberately vague about  $Q_n$ , we require that the quasi-posterior for  $\sqrt{n}(P - \hat{P})$  should be well approximated by a  $N(0, I_0^{-1})$  distribution in large samples.<sup>14</sup> We therefore implicitly require that the criterion  $Q_n$  is “optimal”, in the sense that maximizing  $Q_n(P)$  leads to a semiparametrically efficient estimator  $\hat{P}$  of  $P$ .

The *quasi-posterior maximum risk*  $\bar{R}(d|X^n)$  is calculated by averaging  $R(d, P)$  across the quasi-posterior, analogously to display (13). The *quasi-Bayes decision*  $d_n^*(X^n; \pi)$  is chosen to minimize the quasi-posterior maximum risk  $\bar{R}(d|X^n)$ , analogously to (14). We again include

---

<sup>14</sup>This is an implication of Assumption 5(ii).

$\pi$  as an argument of  $d_n^*(\cdot)$  to indicate that the decision depends on the prior in finite samples. As we will show formally, the decisions  $d_n^*(X^n; \pi)$  are optimal under the criterion (19) for any prior  $\pi$  in the class of priors  $\Pi$  for which our regularity conditions hold. We therefore refer to  $d_n^*(X^n; \pi)$  as an *efficient-robust* decision rule in what follows.

Unlike the parametric case, here the efficient-robust decision cannot be justified on the basis of robust Bayes analysis or maxmin expected utility because  $Q_n(P)$  is not the true log-likelihood. A formal Bayesian approach would require specifying a prior on  $\mathcal{P} \times \mathcal{H}$  then forming a marginal posterior for  $P$ .<sup>15</sup> Our limited-information approach is computationally simple and avoids the delicate issue of specifying priors in infinite-dimensional parameter spaces.

#### 4.4 Optimality of Quasi-Bayes Decisions

Appendix A presents Assumptions 1, 4, and 5, which are the main regularity conditions. Assumption 4 defines the approximately least-favorable model  $\{F_{n,\beta(P_0,\eta_0)}(P_{n,h})\}$  for  $X^n$  and presents a notion of local asymptotic normality for it. Assumption 5 contains consistency and asymptotic normality assumptions for functionals of the quasi-posterior.

Let  $\mathbb{D}$  denote all sequences  $\{d_n\}$  that converge in distribution along  $\{F_{n,\beta(P_0,\eta_0)}(P_{n,h})\}$  for each  $(P_0, \eta_0) \in \mathcal{P} \times \mathcal{H}$  and each  $h \in \mathbb{R}^k$ . Say  $\{d_n\}$  and  $\{d'_n\}$  are *asymptotically equivalent* if  $d_n(X^n)$  and  $d'_n(X^n)$  have the same asymptotic distribution along  $\{F_{n,\beta(P_0,\eta_0)}(P_{n,h})\}$  for all  $P_0 \in \mathcal{P}$ ,  $\eta_0 \in \mathcal{H}$ , and  $h \in \mathbb{R}^k$ .

**Theorem 3** *Suppose that Assumptions 1 and 4 are satisfied. Let  $\Pi$  be the class of priors such that Assumption 5 holds.*

- (i) *For any  $\pi \in \Pi$ , the quasi-Bayes decision  $\{d_n^*(\cdot; \pi)\}$  is asymptotically efficient-robust.*
- (ii) *For any  $\pi, \pi' \in \Pi$ , the quasi-Bayes decisions  $\{d_n^*(\cdot; \pi)\}$  and  $\{d_n^*(\cdot; \pi')\}$  are asymptotically equivalent.*
- (iii) *If  $\{d_n\}$  is asymptotically equivalent to  $\{d_n^*(\cdot; \pi)\}$  for a  $\pi \in \Pi$ , then  $\{d_n\}$  is asymptotically efficient-robust.*

As in the parametric case, asymptotic equivalence to  $d_n^*$  is necessary for optimality under a side condition ruling out the absence of ties. Say asymptotic equivalence of  $\{d_n\}$  and  $\{d'_n\}$  fails at  $(P_0, \eta_0)$  if  $d_n(X^n)$  and  $d'_n(X^n)$  have different asymptotic distributions along  $\{F_{n,\beta(P_0,\eta_0)}(P_{n,h})\}$  for some  $h_0 \in \mathbb{R}^k$ . Let  $\mathbb{E}^*$  denote expectation with respect to  $Z^* \sim N(0, I_0^{-1})$  where  $I_0 = I_0(P_0, \eta_0)$  is the semiparametric information at  $(P_0, \eta_0)$ , and let  $\dot{\rho}_{d,P_0}[\cdot]$  denote the directional derivative of  $R(d, P)$  with respect to  $P$  at  $P_0$ —see Appendix A for definitions.

---

<sup>15</sup>See, for instance, the Bayesian exponentially tilted empirical likelihood approach of Schennach (2005) or the Bayesian GMM approach of Shin (2015).

**Theorem 4** *Suppose that Assumptions 1 and 4 are satisfied. Let  $\Pi$  be the class of priors such that Assumption 5 holds. Suppose that  $\{d_n\} \in \mathbb{D}$  is not asymptotically equivalent to  $\{d_n^*(\cdot; \pi)\}$  for  $\pi \in \Pi$ . Then at any  $(P_0, \eta_0)$  at which asymptotic equivalence fails,*

$$\mathcal{R}(\{d_n\}; (P_0, \eta_0)) > \mathcal{R}(\{d_n^*(\cdot; \pi)\}; (P_0, \eta_0)) \quad \text{for all } \pi \in \Pi$$

*provided either of the following hold:*

- (i)  $\arg \min_{d \in \mathcal{D}} R(d, P_0)$  is a singleton;
- (ii)  $\arg \min_{d \in \mathcal{D}} R(d, P_0)$  is not a singleton, and  $\arg \min_{d \in \mathcal{D}} \mathbb{E}^*[\dot{\rho}_{d, P_0}[Z^* + z]]$  is a singleton for almost every  $z$ .

## 5 Treatment Assignment under Partial Identification

This section expands on our running example of treatment assignment under partial identification. Section 5.1 presents an empirical application to extrapolation and meta-analyses. Section 5.2 reviews important additional examples of lower and upper bounds  $b_L(P)$  and  $b_U(P)$  on the ATE, shows how non-differentiability of the bounds in  $P$  arises naturally, and describes how to implement our methods.

### 5.1 Empirical Application: Extrapolation and Meta-Analyses

Following Ishihara and Kitagawa (2021) (see also Manski (2020)), we consider the following problem. Suppose the ATE  $\tau_0$  in the target population is unknown, but the researcher has estimates  $\hat{P}_k$  of the ATE  $P_k$  in populations  $k = 1, \dots, K$  from a meta-analysis. These ATEs may be extrapolated to deduce bounds on  $\tau_0$ . How should the DM use the estimates  $(\hat{P}_k)_{k=1}^K$  to inform whether or not to treat the target population?

We first describe how to deduce bounds on  $\tau_0$ , then how to implement efficient-robust decisions in this context. We then present an empirical application which revisits the empirical example of Ishihara and Kitagawa (2021).

**Extrapolation Bounds.** Suppose it is known that  $(\tau_0, P_1, \dots, P_K) \in \mathcal{T}$ , where the set  $\mathcal{T} \subset \mathbb{R}^{K+1}$  is informed by functional-form or shape restrictions on the ATEs across the  $K+1$  populations. This leads to the variational bounds

$$b_L(P) = \inf_{(\tau_0, P) \in \mathcal{T}} \tau_0, \quad b_U(P) = \sup_{(\tau_0, P) \in \mathcal{T}} \tau_0$$

on  $\tau_0$  as functions of  $P = (P_k)_{k=1}^K$ .

Consider Example 2 of [Ishihara and Kitagawa \(2021\)](#), in which

$$\mathcal{T} = \{\tau \in \mathbb{R}^{K+1} : |\tau_i - \tau_j| \leq C\|x_i - x_j\| \text{ for all } i, j = 0, \dots, K\}, \quad (20)$$

where  $C$  is a pre-specified constant and  $x_k$  is a vector of characteristics that describes population  $k$ . This choice of  $\mathcal{T}$  leads to the intersection bounds<sup>16</sup>

$$b_L(P) = \max_{1 \leq k \leq K} (P_k - C\|x_0 - x_k\|), \quad b_U(P) = \min_{1 \leq k \leq K} (P_k + C\|x_0 - x_k\|).$$

Note that the presence of the maxima and minima mean the bounds are only directionally differentiable in  $P$ .

**Implementation.** Suppose the DM observes  $X^n = (\hat{P}_k, s_k)_{k=1}^K$  where  $\hat{P}_k$  is an estimate of  $P_k$  and  $s_k$  is its standard error. This is a setting in which our semiparametric theory applies as the finite-sample distributions of the  $\hat{P}_k$  are typically unknown in practice.

The efficient-robust decision may be implemented using a Gaussian quasi-likelihood that treats each  $\hat{P}_k$  as an independent  $N(P_k, s_k^2)$  random variable. This approach is well justified in an asymptotic framework in which the  $K$  studies are independent, each  $s_k^{-1}(\hat{P}_k - P_k)$  is asymptotically  $N(0, 1)$ , and all samples from which the  $\hat{P}_k$  are estimated become large at the same rate. A flat prior on  $\mathcal{P} = \mathbb{R}^K$  yields a quasi-posterior under which the  $P_k$  are independent  $N(\hat{P}_k, s_k^2)$ . To implement the efficient-robust decision, draw  $P_k^* \sim N(\hat{P}_k, s_k^2)$  independently for  $k = 1, \dots, K$  and compute the robust welfare contrast  $b(P^*)$  with  $P^* = (P_k^*)_{k=1}^K$ . This computation is repeated across a large number of independent draws of  $P^*$ . The decision is then to treat if the mean  $\bar{b}_n$  of  $b(P^*)$  across draws is positive, otherwise don't treat.

**Application.** As an empirical illustration, we revisit [Ishihara and Kitagawa \(2021\)](#). They consider a subset of  $K = 14$  studies from the database of [Card, Kluve, and Weber \(2017\)](#), each of which is an RCT looking at the impact of job training programs on employment. Studies are implemented in a number of different countries and in groups that differ by characteristics  $x_k$ . In their application,  $x_k$  consists of gender (males, females, or both), age (youths, adults, or both), OECD membership status, GDP growth (standardized) and unemployment (standardized). For each study we have a treatment effect estimate  $\hat{P}_k$  and a standard error  $s_k$  of  $\hat{P}_k$ .<sup>17</sup>

We consider the hypothetical question of whether to roll out a job-training program in two populations: German male youths in 2010 and German female youths in 2010 (with GDP

---

<sup>16</sup>The identified set  $\Theta_0(P)$  depends on characteristics of the target population and the populations underlying the meta study. We suppress this dependence to simplify notation.

<sup>17</sup>[Ishihara and Kitagawa \(2021\)](#) compute minimax (with respect to  $\psi = (\theta, P)$ ) regret treatment rules among a class of linear aggregation rules under the assumption that the  $\hat{P}_k$  are independent  $N(P_k, s_k^2)$ . Our semiparametric approach treats the normality assumption as valid only in an asymptotic sense.

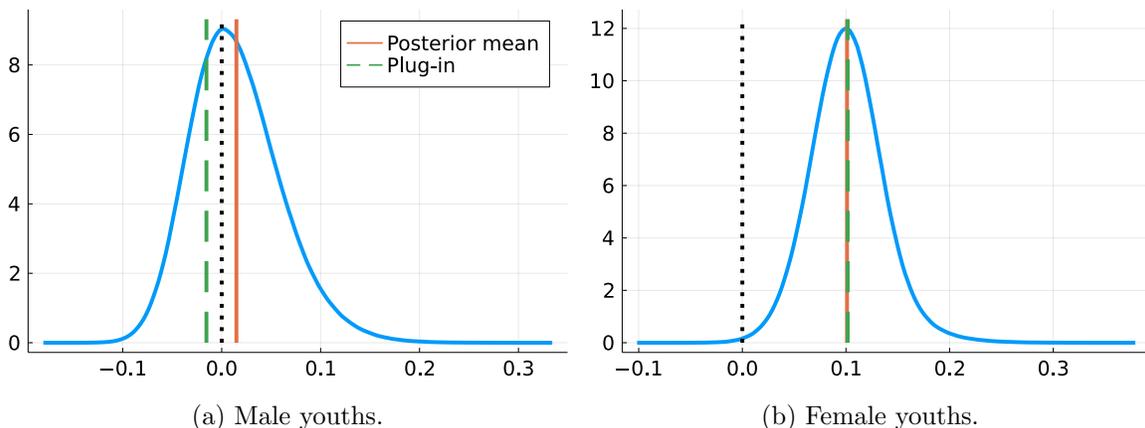


Figure 3: Quasi-posterior distribution of the robust welfare contrast  $b(P)$ , (quasi-)posterior mean of  $b(P)$ , and plug-in value  $b(\hat{P})$ .

growth 3.48% and unemployment rate 9.45%). We compute the efficient-robust decision as described above, increasing the Lipschitz constant from  $C = 0.025$  to  $C = 0.25$  so that the identified sets  $\Theta_0(P)$  are not empty. Figures 3a and 3b plot the distribution of the robust welfare contrast  $b(P)$  for males and females, respectively, across 5 million draws from the quasi-posterior for  $P$ . Both figures also display the quasi-posterior mean  $\bar{b}_n$  of  $b(P)$  whose sign determines the efficient-robust decision (16), and the plug-in value  $b(\hat{P})$  whose sign determines the oracle plug-in rule  $d^o(\hat{P})$  using the oracle decision (15).

In Figure 3a (male youths) we see that the distribution of  $b(P)$  has a mode located at zero and a pronounced right-skew. Recall that  $b_L(P) = \max_{1 \leq k \leq K} (P_k - C\|x_0 - x_k\|)$ . When evaluated at  $\hat{P}$ , the largest value of  $\hat{P}_k - C\|x_0 - x_k\|$  is  $-0.3190$  (corresponding to the US study) and the second-largest is  $-0.3298$  (corresponding to the Brazilian study). As the maximum is not well separated relative to the dispersion of the distribution from which  $P$  is sampled (the average  $s_k$  across the  $K$  studies is 0.034), the distribution of  $b_L(P)$  is right-skewed because it behaves like a maximum of Gaussians rather than a Gaussian. The term  $b_U(P)$  behaves similarly, but with a less pronounced distortion because the corresponding minimum is better separated. As a consequence of this asymmetry, the quasi-posterior mean  $\bar{b}_n$  is positive and the efficient-robust decision is to treat. By contrast, the plug-in value  $b(\hat{P})$  is negative so the oracle plug-in decision is to not treat.

The situation is different for female youths (Figure 3b). Here the minima and maxima that characterize  $b_L(P)$  and  $b_U(P)$  are relatively more well separated, so the distribution of  $b_L(P)$  and  $b_U(P)$  across draws is close to Gaussian. In consequence, the quasi-posterior mean  $\bar{b}_n$  of  $b(P)$  and the plug-in value  $b(\hat{P})$  are almost identical, and the efficient-robust and oracle plug-in decisions are both to treat.

## 5.2 Additional Examples

We now review some additional empirically relevant approaches for constructing bounds on the ATE. Each of these constructions leads to bounds  $b_L(P)$  and  $b_U(P)$  that will in general be only directionally differentiable in  $P$ . We then discuss implementation of the efficient-robust decision rules.

**Intersection Bounds.** Suppose  $X^n$  consists of data from  $K$  observational studies. In each study  $k$  we can consistently estimate lower and upper bounds  $b_{L,k}$  and  $b_{U,k}$  on the ATE. Standard approaches, e.g., following [Manski \(1990\)](#), yield bounds  $b_{L,k}$  and  $b_{U,k}$  as a function of population moments  $P_k$ . Assuming each study is estimating the same treatment effect across homogeneous sub-populations, we obtain the intersection bounds

$$b_L(P) = \max_{1 \leq k \leq K} b_{L,k}, \quad b_U(P) = \min_{1 \leq k \leq K} b_{U,k},$$

with  $P = (P_k)_{k=1}^K$ . While the bounds  $b_{L,k}$  and  $b_{U,k}$  may themselves be smooth in  $P_k$ , the presence of the min and max operations makes the intersection bounds  $b_L(P)$  and  $b_U(P)$  only directionally differentiable in  $P$ .

**Bounds via IV-like Estimands.** Building on [Heckman and Vytlacil \(1999, 2005\)](#), [Mogstad et al. \(2018\)](#) present an approach for bounding the ATE and other causal effects using IV-like estimates from observational studies. Suppose that treatment is determined by  $D = \mathbb{I}[U \leq \nu(Z)]$  where  $U \sim \text{Uniform}(0, 1)$  and  $Z = (X, Z_0)$  collects control variables  $X$  and instrumental variables  $Z_0$ . According to [Heckman and Vytlacil \(1999, 2005\)](#), the ATE may be expressed as a functional  $\Gamma_0(m)$ , where  $m = (m_0, m_1)$  are the marginal treatment response (MTR) functions

$$m_d(u, x) = \mathbb{E}[Y_d | U = u, X = x], \quad d \in \{0, 1\},$$

and

$$\Gamma_0(m) = \mathbb{E} \left[ \int_0^1 m_1(u, X) du - \int_0^1 m_0(u, X) du \right].$$

[Mogstad et al. \(2018\)](#) show the MTR functions, and hence the identified set for the ATE, can be disciplined if we know the value of certain IV-like estimands. For ease of exposition, suppose there is a single IV estimand

$$\beta_{IV} = \frac{\text{Cov}(Y, Z_0)}{\text{Cov}(D, Z_0)},$$

resulting from using  $Z_0$  as an instrument for treatment status dummy  $D$  in the observational study. The IV estimand may be expressed as  $\beta_{IV} = \Gamma_\beta(m)$  where

$$\Gamma_\beta(m) = \mathbb{E} \left[ \int_0^1 m_0(u, X) s(0, Z_0) \mathbb{I}[u > p(z_0)] du + \int_0^1 m_1(u, X) s(1, Z_0) \mathbb{I}[u \leq p(z_0)] du \right]$$

with  $s(d, z) = \frac{z_0 - \mathbb{E}[Z_0]}{\text{Cov}(Z_0, D)}$  and where  $p(z_0) = \mathbb{E}[D|Z_0 = z_0]$  is the propensity score. In view of the above, [Mogstad et al. \(2018\)](#) derive (more general versions of) the bounds

$$b_L(P) = \inf_{m \in \mathcal{S}: \Gamma_\beta(m) = \beta_{IV}} \Gamma_0(m), \quad b_U(P) = \sup_{m \in \mathcal{S}: \Gamma_\beta(m) = \beta_{IV}} \Gamma_0(m),$$

where  $\mathcal{S}$  is a class of functions and  $P = (\text{Cov}(Y, Z_0), \text{Cov}(D, Z_0), \mathbb{E}[Z_0], p)$ . The propensity score  $p$ , and therefore  $P$ , is finite-dimensional if  $Z_0$  has finite support (e.g. binary  $Z_0$ ), which is often the case in applications—see, e.g., [Brinch, Mogstad, and Wiswall \(2017\)](#).

[Mogstad et al. \(2018\)](#) show that  $b_L(P)$  and  $b_U(P)$  may be expressed as the optimal values of linear programs parameterized by  $P$ . It is known from [Milgrom and Segal \(2002\)](#) (see also [Mills \(1956\)](#) and [Williams \(1963\)](#) for linear programs) that the value of optimization problems may be only directionally differentiable in parameters.

**Non-separable Panel Data Models.** Suppose the outcome for individual  $i$  at date  $t$  is of the form  $Y_{it} = g(X_{it}, \alpha_i, \varepsilon_{it})$  where  $X_{it}$  is a vector of covariates,  $\alpha_i$  is a latent individual effect, and  $\varepsilon_{it}$  is a vector of disturbances, which are independent across individuals and time. Consider an intervention that changes in covariates from  $x^0$  to  $x^1$ . The ATE associated with the intervention is

$$\int g(x^1, \alpha, \varepsilon) - g(x^0, \alpha, \varepsilon) dQ(\alpha, \varepsilon),$$

where  $Q$  is a distribution over  $(\alpha, \varepsilon)$ . With discrete regressors  $X_{it}$  and discrete outcomes  $Y_{it}$ , parametric restrictions on the distribution of  $\varepsilon$  and functional form restrictions on  $g$  are generally insufficient to point identify the ATE without parametric restrictions on the distribution of  $\alpha$ . A leading example is dynamic panel data models in which  $X_{it}$  collects lagged values of a discrete outcome  $Y_{it}$ —see, e.g., [Honoré and Tamer \(2006\)](#) and [Torgovitsky \(2019a\)](#). Building on [Honoré and Tamer \(2006\)](#), [Chernozhukov et al. \(2013\)](#) and [Torgovitsky \(2019b\)](#) derive sharp bounds on the ATE without parametric assumptions on  $Q$ . Their bounds may be expressed as the value of optimization problems (linear programs) parameterized by a finite-dimensional vector of choice probabilities  $P$  for individual histories  $(X_{it})_{t=1}^T$ . As in the previous example,  $b_L(P)$  and  $b_U(P)$  may therefore only be directionally differentiable in  $P$ .

**Implementation.** For the first two examples,  $X^n$  may represent the data (if the DM has access to it) or the vector of estimates  $\hat{P}$  of  $P$ . The data-generating process for  $X^n$  will

typically be semiparametric.<sup>18</sup> If the DM has a vector of (efficient) estimates  $\hat{P}$  of  $P$  and a consistent estimate  $\hat{I}^{-1}$  of the asymptotic variance of  $\hat{P}$ , then the efficient-robust decision can be implemented based on a  $N(\hat{P}, (n\hat{I})^{-1})$  quasi-posterior. Alternatively, if the researcher observes the data  $X^n$ , then a quasi-posterior based on an efficient GMM objective function can be used.

In the nonseparable panel data case with discrete  $Y_{it}$  and  $X_{it}$ , the distribution  $F_{n,P}$  of the data  $((X_{it})_{t=1}^T)_{i=1}^n$  can be identified with a multinomial distribution parameterized by  $P$ . In this case our parametric theory applies and the efficient-robust can easily be implemented using either the bootstrap or Bayesian methods. For the latter, a Dirichlet prior for  $P$  leads to a Dirichlet posterior for  $P$ , which is trivial to sample from.  $\square$

## 6 Optimal Pricing with Rich Unobserved Heterogeneity

This section expands on our running example of optimal pricing with rich unobserved heterogeneity. Section 6.1 presents techniques for computing sharp bounds on functionals of counterfactual demand using linear programming. Section 6.2 then discusses how to implement our methods.

### 6.1 Bounds on Functionals of Counterfactual Demand

Kitamura and Stoye (2019) present a general approach using linear programming to derive bounds on functionals of counterfactual demands. We introduce their approach by way of Example 2 from Section 2.

Consider consumer behavior on budgets  $\mathcal{B}_1$  and  $\mathcal{B}_2$  in Figure 1. There are 9 potential combinations of patches consumers may choose from among  $\mathcal{B}_1$  and  $\mathcal{B}_2$ :  $(s_{11}, s_{12})$ ,  $(s_{11}, s_{22})$ ,  $(s_{11}, s_{32})$ ,  $(s_{21}, s_{12})$ ,  $(s_{21}, s_{22})$ ,  $(s_{21}, s_{32})$ ,  $(s_{31}, s_{12})$ ,  $(s_{31}, s_{22})$ ,  $(s_{31}, s_{23})$ . By revealed preference we know a consumer will never choose  $(s_{21}, s_{12})$  or  $(s_{31}, s_{12})$ . This leaves a total of 7 rational types of consumer. Let

$$A = \begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 & 1 \end{bmatrix}.$$

The rows of  $A$  correspond to choosing  $s_{11}, s_{21}, s_{31}, s_{12}, s_{22}, s_{32}$  and the columns of  $A$  cor-

---

<sup>18</sup>An exception is the intersection bounds example, where in each study we observe a binary treatment indicator and a binary outcome across individuals. That case can be handled by forming a multinomial likelihood parametrized by the reduced-form parameters.

respond to the 7 rational types. Let  $p = (p_1(s_{11}), p_1(s_{21}), p_1(s_{31}), p_2(s_{12}), p_2(s_{22}), p_2(s_{32}))'$  collect the corresponding choice probabilities. Kitamura and Stoye (2018) showed that the stochastic demand system  $p$  is rationalizable if and only if

$$p = A\pi^v,$$

for some  $\pi^v \in \Delta^6$ , the unit simplex in  $\mathbb{R}^7$ . Each entry of  $\pi^v$  corresponds to probabilities of the 7 rational types. While these probabilities constrain the set of potential distributions  $\Pi^v$  of utilities, they are insufficient to point-identify  $\Pi^v$  without further restrictions.

Now suppose we consider the implications of these restrictions for choice behavior on the counterfactual budget set  $\mathcal{B}_0$ . Each of the 7 rational types of consumer may choose a counterfactual demand in patch  $s_{10}$ ,  $s_{20}$ , or  $s_{30}$ , for a total of 21 potential types. We can again use revealed preference to deduce that a consumer who chose  $s_{31}$  must choose  $s_{20}$  or  $s_{30}$ , and that a consumer who chose  $s_{32}$  must choose  $s_{30}$ . This leaves a total of 16 rational types. We may represent the system as

$$\begin{bmatrix} p \\ p^* \end{bmatrix} = A^* \pi^{v*}, \quad (21)$$

where  $p^* = (p_0(s_{10}), p_0(s_{20}), p_0(s_{30}))'$  collects the counterfactual choice probabilities for patches  $s_{10}$ ,  $s_{20}$ , and  $s_{30}$ ,  $\pi^{v*} \in \Delta^{15}$  collects the probabilities of observing each rational type, and

$$A^* = \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 \\ 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 1 & 0 & 1 & 1 & 1 \end{bmatrix},$$

where the rows correspond to choosing  $s_{11}, s_{21}, s_{31}, s_{12}, s_{22}, s_{32}, s_{10}, s_{20}, s_{30}$ . We may partition  $A^*$  as

$$A^* = \begin{bmatrix} A_{obs}^* \\ A_{unobs}^* \end{bmatrix},$$

where  $A_{obs}^*$  collects the first 6 rows of  $A^*$ , which correspond to the patches for the observed budget sets, while  $A_{unobs}^*$  collects the final 3 rows of  $A^*$  corresponding to the patches  $s_{10}$ ,  $s_{20}$ , and  $s_{30}$  for the counterfactual budget set.

Suppose as in Example 2 from Section 2 that  $\mathbb{E}_\theta[u(0, Y, \theta, P)]$  depends on the expected value of a functional  $Y = h(X_0)$  of counterfactual demand, where the expectation is with respect to the distribution  $G_\theta$  of  $Y$  induced by  $\Pi^v$ . Following Kitamura and Stoye (2019), we may deduce sharp bounds on  $\mathbb{E}_\theta[h(X_0)]$  as follows. Let  $(\underline{h}_1, \underline{h}_2, \underline{h}_3)$  and  $(\bar{h}_1, \bar{h}_2, \bar{h}_3)$  denote row vectors which collect the smallest and largest values of  $h(x)$  for  $x$  in  $s_{10}$ ,  $s_{20}$  and  $s_{30}$ . Then the lower bound on  $\mathbb{E}_\theta[h(X_0)]$  is the value of a linear program:

$$h_L(p) = \min \{ (\underline{h}_1, \underline{h}_2, \underline{h}_3) A_{unobs}^* \pi^{v*} : p = A_{obs}^* \pi^{v*}, \pi^{v*} \in \Delta^{15} \}.$$

Similarly, the corresponding upper bound is

$$h_U(p) = \max \{ (\bar{h}_1, \bar{h}_2, \bar{h}_3) A_{unobs}^* \pi^{v*} : p = A_{obs}^* \pi^{v*}, \pi^{v*} \in \Delta^{15} \}.$$

The preceding argument applies more generally with an arbitrary collection of  $B$  observed budget sets and multiple goods. In that case, we again have the representation (21) for observed choice probabilities  $p$  of patches on the  $B$  observed budget sets and counterfactual choice probabilities  $p^*$  of patches on the counterfactual budget set. Suppose there are  $J$  patches across the counterfactual budget set and a total of  $H$  rational types. Then letting  $\underline{h}$  and  $\bar{h}$  be row vectors collecting the smallest and largest values of  $h(x)$  for  $x$  in each of the  $J$  patches, bounds on  $\mathbb{E}_\theta[h(X_0)]$  are

$$\begin{aligned} h_L(p) &= \min \{ \underline{h} A_{unobs}^* \pi^{v*} : p = A_{obs}^* \pi^{v*}, \pi^{v*} \in \Delta^{H-1} \}, \\ h_U(p) &= \max \{ \bar{h} A_{unobs}^* \pi^{v*} : p = A_{obs}^* \pi^{v*}, \pi^{v*} \in \Delta^{H-1} \}, \end{aligned}$$

where we have partitioned the  $A^*$  matrix analogously to the simple 3 budget example.

## 6.2 Implementation

We present an implementation when  $\mathbb{E}_\theta[u(d, Y, \theta, P)]$  is the expected value of a functional  $h_d(X_d)$  of demand (e.g. revenue) at prices  $q_d$ . Suppose the DM wishes to choose a price vector  $q_d$  for  $d$  in a finite set  $\mathcal{D} = \mathcal{O} \cup \mathcal{C}$ . Each  $q_d$  for  $d \in \mathcal{O}$  is a price vector for which demand is observed. The expected value  $\mathbb{E}[h_d(X_d)]$  is identified from observed choice behavior on the budget set  $\mathcal{B}_d$  and forms an element of  $P$ . Each  $q_d$  for  $d \in \mathcal{C}$  is a counterfactual price vector for which demand is not observed. As argued in the previous subsection, we may deduce sharp bounds on the counterfactual  $\mathbb{E}_\theta[h_d(X_d)]$  as functions of the observed patch probabilities  $p$ . The reduced form parameter is therefore  $P = (p, (\mathbb{E}[h(X_d)])_{d \in \mathcal{O}})$ .

This is a semiparametric model and our implementation follows the steps described in Section 4.3. Partition  $P = (P_1, \dots, P_B)$  where  $P_b$  collects the parameters corresponding to budget set  $\mathcal{B}_b$ . As we have assumed that each of the  $B$  observed budgets is sampled in

a repeated cross section, we can estimate each  $P_b$  by just-identified GMM. We can then form a quasi-posterior  $\pi_n(P|X^n)$  based on a limited-information Gaussian quasi-likelihood, exponentiating the GMM objective function, or by the Bayesian bootstrap. For  $d \in \mathcal{O}$ , the expected value  $\mathbb{E}[h_d(X_d)]$  is an element of  $P$  and so  $\bar{R}(d|X^n)$  is simply its posterior mean

$$\bar{R}(d|X^n) = - \int \mathbb{E}[h_d(X_d)] d\pi(P|X^n).$$

For  $d \in \mathcal{C}$ , a lower bound  $h_{d,L}(p)$  can be deduced as a function of the observed patch probabilities  $p$  using linear programming. The posterior maximum risk

$$\bar{R}(d|X^n) = - \int h_{d,L}(p) d\pi(P|X^n)$$

is computed by resampling  $p$  from the posterior, solving the linear program for each draw, then taking the average across draws. The efficient-robust decision  $d_n^*(X^n)$  is the choice  $d$  that minimizes  $\bar{R}(d|X^n)$  over the choice set  $\mathcal{D}$ . As is evident from the intersection bound derived in Example 2 in Section 2 and the general treatment via linear programming in the previous subsection, the posterior mean  $\bar{R}(d|X^n)$  for  $d \in \mathcal{C}$  may be different, even asymptotically, from the bound  $-h_{d,L}(\hat{p})$  deduced by plugging in an estimate  $\hat{p}$  of the patch probabilities. As such, the efficient-robust decision may dominate a plug-in rule under our optimality criterion.

Efficient-robust decisions can also be implemented replacing the posterior distribution for  $P$  by the bootstrap distribution of an efficient estimator  $\hat{P}$  of  $P$ . For observed prices  $d \in \mathcal{O}$ , the bootstrap average maximum risk  $R_n^*(d)$  is simply the average of the estimate of  $-\mathbb{E}[h_d(X_d)]$  across bootstrap distributions. Typically, this will just correspond (or be asymptotically equivalent to) to the estimate of  $-\mathbb{E}[h_d(X_d)]$  in  $\hat{P}$ . For counterfactual prices  $d \in \mathcal{C}$ , the bootstrap average maximum risk is computed by averaging  $-h_{d,L}(\hat{p}^*)$  across bootstrap draws  $\hat{p}^*$  of the estimated patch probabilities  $\hat{p}$ :

$$R_n^*(d|X^n) = -\mathbb{E}_n^*[h_{d,L}(\hat{p}^*)].$$

Note that this quantity may differ, even asymptotically, from the plug-in value  $-h_{d,L}(\hat{p})$ . The bootstrap decision  $d_n^{**}(X^n)$  is the choice  $d$  that minimizes  $R_n^*(d|X^n)$  over the choice set  $\mathcal{D}$ .

## 7 Conclusion

We derived optimal statistical decision rules for discrete choice problems when the distribution of payoffs depends on a set-identified parameter  $\theta$  and the decision maker can use a point-identified parameter  $P$  to deduce restrictions on  $\theta$ . Our proposed efficient-robust decision rules minimize maximum risk or regret over the identified set for  $\theta$  conditional on  $P$ , and

use the data efficiently to learn about  $P$ . In many empirically relevant applications, the maximum risk depends non-smoothly on  $P$ . In those cases, plug-in rules are suboptimal under our asymptotic efficiency criterion. We provided detailed applications to optimal treatment choice under partial identification and optimal pricing with rich unobserved heterogeneity. Our asymptotic approach is well suited for empirical settings in which the derivation of finite-sample optimal rules is intractable. While our asymptotic optimality theory was developed for discrete decisions, our general approach can be used for continuous decisions as well. We conjecture that similar optimality results also apply in that case.

## References

- Berger, J. O. (1985). *Statistical Decision Theory and Bayesian Analysis*. Springer Verlag, New York.
- Blundell, R., M. Browning, and I. Crawford (2007). Improving revealed preference bounds on demand responses. *International Economic Review* 48(4), 1227–1244.
- Blundell, R., M. Browning, and I. Crawford (2008). Best nonparametric bounds on demand responses. *Econometrica* 76(6), 1227–1262.
- Blundell, R., D. Kristensen, and R. Matzkin (2014). Bounding quantile demand functions using revealed preference inequalities. *Journal of Econometrics* 179(2), 112–127.
- Blundell, R. W., D. Kristensen, and R. L. Matzkin (2017). Individual counterfactuals with multidimensional unobserved heterogeneity. Technical report, cemmap working paper.
- Brinch, C. N., M. Mogstad, and M. Wiswall (2017). Beyond LATE with a discrete instrument. *Journal of Political Economy* 125(4), 985–1039.
- Card, D., J. Kluve, and A. Weber (2017). What Works? A Meta Analysis of Recent Active Labor Market Program Evaluations. *Journal of the European Economic Association* 16(3), 894–931.
- Chamberlain, G. (2011). Bayesian aspects of treatment choice. In *The Oxford Handbook of Bayesian Econometrics*. Oxford University Press.
- Chamberlain, G. and G. W. Imbens (2003). Nonparametric applications of bayesian inference. *Journal of Business & Economic Statistics* 21(1), 12–18.
- Chen, X., T. M. Christensen, and E. Tamer (2018). Monte carlo confidence sets for identified sets. *Econometrica* 86(6), 1965–2018.
- Chernozhukov, V., I. Fernández-Val, J. Hahn, and W. Newey (2013). Average and quantile effects in nonseparable panel models. *Econometrica* 81(2), 535–580.
- Chernozhukov, V. and H. Hong (2003). An MCMC approach to classical estimation. *Journal of Econometrics* 115(2), 293–346.

- Christensen, T., H. R. Moon, and F. Schorfheide (2020). Robust forecasting. *arXiv preprint arXiv:2011.03153*.
- Cogley, T. and T. J. Sargent (2008). Anticipated utility and rational expectations as approximations of Bayesian decision making. *International Economic Review* 49(1), 185–221.
- Doksum, K. A. and A. Y. Lo (1990). Consistent and Robust Bayes Procedures for Location Based on Partial Information. *The Annals of Statistics* 18(1), 443–453.
- Ghosh, J. and R. Ramamoorthi (2003). *Bayesian Nonparametrics*. Springer Verlag, New York.
- Giacomini, R., T. Kitagawa, and M. Read (2021). Robust Bayesian analysis for econometrics. *Federal Reserve Bank of Chicago Working Paper 2021-11*.
- Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics* 18(2), 141–153.
- Hartigan, J. A. (1983). *Bayes Theory*. Springer Verlag, New York.
- Heckman, J. J. and E. Vytlacil (2005). Structural equations, treatment effects, and econometric policy evaluation. *Econometrica* 73(3), 669–738.
- Heckman, J. J. and E. J. Vytlacil (1999). Local instrumental variables and latent variable models for identifying and bounding treatment effects. *Proceedings of the National Academy of Sciences* 96(8), 4730–4734.
- Hirano, K. and J. R. Porter (2009). Asymptotics for statistical treatment rules. *Econometrica* 77(5), 1683–1701.
- Hoderlein, S. and J. Stoye (2015). Testing stochastic rationality and predicting stochastic demand: the case of two goods. *Economic Theory Bulletin* 3(2), 313–328.
- Honoré, B. E. and E. Tamer (2006). Bounds on parameters in panel dynamic discrete choice models. *Econometrica* 74(5), 611–629.
- Hurwicz, L. (1951). Some specification problems and applications to econometric models. *Econometrica* 19(3), 343–344.
- Ishihara, T. and T. Kitagawa (2021). Evidence aggregation for treatment choice. *arXiv preprint arXiv:2108.06473*.
- Kim, J.-Y. (2002). Limited information likelihood and bayesian analysis. *Journal of Econometrics* 107(1-2), 175–193.
- Kitamura, Y. and J. Stoye (2018). Nonparametric analysis of random utility models. *Econometrica* 86(6), 1883–1909.
- Kitamura, Y. and J. Stoye (2019). Nonparametric counterfactuals in random utility models. *arXiv preprint arXiv:1902.08350*.

- Kreps, D. M. (1998). Anticipated utility and dynamic choice. In D. P. Jacobs, E. Kalai, M. I. Kamien, and N. L. Schwartz (Eds.), *Frontiers of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures, 1983-1997*, Econometric Society Monographs, pp. 242–274. Cambridge University Press.
- Manski, C. F. (1990). Nonparametric bounds on treatment effects. *The American Economic Review* 80(2), 319–323.
- Manski, C. F. (2000). Identification problems and decisions under ambiguity: Empirical analysis of treatment response and normative analysis of treatment choice. *Journal of Econometrics* 95(2), 415–442.
- Manski, C. F. (2004). Statistical treatment rules for heterogeneous populations. *Econometrica* 72(4), 1221–1246.
- Manski, C. F. (2007a). Minimax-regret treatment choice with missing outcome data. *Journal of Econometrics* 139(1), 105–115. Endogeneity, instruments and identification.
- Manski, C. F. (2007b). Partial identification of counterfactual choice probabilities. *International Economic Review* 48(4), 1393–1410.
- Manski, C. F. (2014). Identification of income-leisure preferences and evaluation of income tax policy. *Quantitative Economics* 5(1), 145–174.
- Manski, C. F. (2020). Toward credible patient-centered meta-analysis. *Epidemiology* 31(3), 345–352.
- Manski, C. F. (2021a). Econometrics for decision making: Building foundations sketched by Haavelmo and Wald. *Econometrica* 89(6), 2827–2853.
- Manski, C. F. (2021b). Probabilistic prediction for binary treatment choice: with focus on personalized medicine. *NBER Working Paper 29358*.
- Milgrom, P. and I. Segal (2002). Envelope theorems for arbitrary choice sets. *Econometrica* 70(2), 583–601.
- Mills, H. D. (1956). Marginal values of matrix games and linear programs. In *Linear Inequalities and Related Systems.*, pp. 183–194. Princeton University Press.
- Mogstad, M., A. Santos, and A. Torgovitsky (2018). Using instrumental variables for inference about policy relevant treatment parameters. *Econometrica* 86(5), 1589–1619.
- Murphy, S. A. and A. W. van der Vaart (2000). On profile likelihood. *Journal of the American Statistical Association* 95(450), 449–465.
- Robbins, H. (1951). Asymptotically subminimax solutions of compound decision problems. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, Volume I. University of California Press, Berkeley and Los Angeles.
- Russell, T. M. (2020). Policy transforms and learning optimal policies. *arXiv preprint arXiv:2012.11046*.

- Schennach, S. M. (2005). Bayesian exponentially tilted empirical likelihood. *Biometrika* 92(1), 31–46.
- Shin, M. (2015). *Bayesian GMM*. Ph. D. thesis.
- Stoye, J. (2012). Minimax regret treatment choice with covariates or with limited validity of experiments. *Journal of Econometrics* 166, 138–156.
- Torgovitsky, A. (2019a). Nonparametric inference on state dependence in unemployment. *Econometrica* 87(5), 1475–1505.
- Torgovitsky, A. (2019b). Partial identification by extending subdistributions. *Quantitative Economics* 10(1), 105–144.
- van der Vaart, A. (1991). An asymptotic representation theorem. *International Statistical Review* 59(1), 97–121.
- van der Vaart, A. W. (1998). *Asymptotic Statistics*. Cambridge University Press, Cambridge.
- Wald, A. (1950). *Statistical Decision Functions*. John Wiley, New York.
- Williams, A. C. (1963). Marginal values in linear programming. *Journal of the Society for Industrial and Applied Mathematics* 11(1), 82–94.
- Yata, K. (2021). Optimal decision rules under partial identification. *arXiv preprint arXiv:2111.04926*.

# Online Appendix: Optimal Discrete Decisions when Payoffs are Partially Identified

Timothy Christensen, Hyungsik Roger Moon, and Frank Schorfheide

This Appendix consists of the following sections:

- A. Assumptions
- B. Proofs

## A Assumptions

This section presents the main regularity conditions under which Theorems 1 and 2 are derived. We first place some assumptions on the maximum risk. Say  $f : \mathcal{P} \rightarrow \mathbb{R}^k$  is *directionally differentiable* at  $P_0$  if the limit

$$\lim_{t \downarrow 0} \frac{f(P_0 + th) - f(P_0)}{t} =: \dot{f}_{P_0}[h]$$

exists for every  $h \in \mathbb{R}^k$ , in which case  $\dot{f}_{P_0}[\cdot]$  denotes the directional derivative of  $f$  at  $P_0$ . If it exists, the directional derivative  $\dot{f}_{P_0}[\cdot]$  is positively homogeneous of degree one but not necessarily linear. If  $\dot{f}_{P_0}[\cdot]$  is linear then  $f$  is *fully* differentiable at  $P_0$ . Define  $\rho : \mathcal{P} \rightarrow \mathbb{R}^{D+1}$  by  $\rho(P) = (R(0, P), R(1, P), \dots, R(D, P))'$ . If  $\rho$  is directionally differentiable at  $P_0$ , we let  $\dot{\rho}_{P_0}[\cdot]$  denote its directional derivative.

- Assumption 1** (i) *The function  $\rho$  is bounded and continuous on  $\mathcal{P}$ ;*  
(ii) *The function  $\rho$  is directionally differentiable any  $P_0 \in \mathcal{P}$  for which  $\arg \min_{d \in \mathcal{D}} R(d, P_0)$  is not a singleton.*

The remaining assumptions are presented separately for parametric and semiparametric models.

### A.1 Parametric Models

We first impose some assumptions on the data-generating process. Let  $\overset{P_{n,h}}{\rightsquigarrow}$  and  $\overset{P_{n,h}}{\rightarrow}$  denote convergence in distribution and in probability under the sequence of measures  $\{F_{n,P_{n,h}}\}$  with  $P_{n,h} = P_0 + h/\sqrt{n}$  for  $P_0$  fixed and  $h$  ranging over  $\mathbb{R}^k$ . Say the model for  $X^n$  is *locally asymptotically normal* at  $P_0$  if for each  $h_0 \in \mathbb{R}^k$ , the likelihood ratio processes indexed by any finite subset  $H \subset \mathbb{R}^k$  converge weakly to the likelihood ratio in a shifted normal model:

$$\left( \frac{dF_{n,P_{n,h}}}{dF_{n,P_{n,h_0}}} \right)_{h \in H} \overset{P_{n,h_0}}{\rightsquigarrow} \left( \exp \left( (h - h_0)' Z - \frac{1}{2} (h - h_0)' I_0 (h - h_0) \right) \right)_{h \in H},$$

with  $Z \sim N(h_0, I_0^{-1})$  and the nonsingular matrix  $I_0 = I_0(P_0)$  is the (asymptotic) information matrix at  $P_0$ .

**Assumption 2** (i)  $\mathcal{P}$  is an open subset of  $\mathbb{R}^k$  with  $P_0 \in \mathcal{P}$ ;  
(ii) The model for  $X^n$  is locally asymptotically normal at each  $P_0 \in \mathcal{P}$ .

Finally, we impose some high-level assumptions on the large-sample behavior of the posterior. We say that  $\pi_n(P|X^n)$  is *consistent* if  $\pi_n(P \in N|X^n) \xrightarrow{P_0} 1$  for every neighborhood  $N$  containing  $P_0$ , for each  $P_0 \in \mathcal{P}$ . Let  $\mathbb{P}_h$  denote probability with respect to  $Z \sim N(h, I_0^{-1})$  for  $I_0 = I_0(P_0)$  and let  $\mathbb{E}^*$  denote expectation with respect to  $Z^*$  where  $Z^* \sim N(0, I_0^{-1})$  independent of  $Z$ .

**Assumption 3** (i) The posterior  $\pi_n(P|X^n)$  is consistent;  
(ii) At any  $P_0 \in \mathcal{P}$  with non-singleton  $\arg \min_{d \in \mathcal{D}} R(d, P_0)$ , for any Borel set  $A \subseteq \mathbb{R}^{D+1}$  and  $h \in \mathbb{R}^k$ , we have

$$\lim_{n \rightarrow \infty} F_{n, P_n, h} \left( \int \sqrt{n}(\rho(P) - \rho(P_0)) \, d\pi_n(P|X^n) \in A \right) = \mathbb{P}_h(\mathbb{E}^*[\dot{\rho}_{P_0}[Z^* + Z]|Z] \in A).$$

Assumption 3(i) is a mild consistency condition that is satisfied under minimal conditions. For the intuition behind Assumption 3(ii), suppose  $X^n = (X_1, \dots, X_n)$  where the  $X_i$  are independent  $N(P, I_0^{-1})$  random variables under  $F_{n, P}$ . Under a flat prior for  $P$ ,  $\pi_n(P|X^n)$  is  $N(\hat{P}, (nI_0)^{-1})$  where  $\hat{P} = \bar{X}_n$  is the sample mean, which itself has a  $N(P_{n, h}, (nI_0)^{-1})$  distribution under  $F_{n, P_{n, h}}$ . On a shrinking neighborhood of  $P_0$ , we then have

$$\sqrt{n}(\rho(P) - \rho(P_0)) \approx \dot{\rho}_{P_0}[\sqrt{n}(P - \hat{P}) + \sqrt{n}(\hat{P} - P_0)],$$

where  $\sqrt{n}(P - \hat{P}) \sim N(0, I_0^{-1})$  under  $\pi_n(P|X^n)$  and  $\sqrt{n}(\hat{P} - P_0) \sim N(h, I_0)$  under  $F_{n, P_{n, h}}$ . Similar reasoning applies in more general smooth parametric models by the Bernstein–von Mises theorem and standard asymptotic distribution theory for the MLE.<sup>19</sup>

## A.2 Semiparametric Models

Our conditions for the semiparametric case parallel those for the parametric case. First we place assumptions on the data-generating process. Similar to [Murphy and van der Vaart \(2000\)](#), we say the model for  $X^n$  has an *approximately least-favorable submodel* at  $(P_0, \eta_0)$  if

---

<sup>19</sup>Fundamental conditions for consistency and the asymptotic normality of posterior distributions in point-identified models can be found, for instance, in the textbook treatments of [Hartigan \(1983\)](#), [van der Vaart \(1998\)](#), or [Ghosh and Ramamoorthi \(2003\)](#).

(i) there exists a map  $t \mapsto \eta_{(P_0, \eta_0)}(t)$  from an open neighborhood  $\mathcal{P}(P_0, \eta_0)$  of  $P_0$  into  $\mathcal{H}$  for which  $\eta_{(P_0, \eta_0)}(P_0) = \eta_0$ , and (ii) with  $\beta_{(P_0, \eta_0)}(P_{n,h}) = (P_{n,h}, \eta_{(P_0, \eta_0)}(P_{n,h}))$ , for each  $h_0 \in \mathbb{R}^k$  and any finite subset  $H \subset \mathbb{R}^k$ , we have

$$\left( \frac{dF_{n, \beta_{(P_0, \eta_0)}(P_{n,h})}}{dF_{n, \beta_{(P_0, \eta_0)}(P_{n,h_0})}} \right)_{h \in H} \overset{P_{n,h_0}}{\rightsquigarrow} \left( \exp \left( (h - h_0)' Z - \frac{1}{2} (h - h_0)' I_0 (h - h_0) \right) \right)_{h \in H},$$

with  $\overset{P_{n,h_0}}{\rightsquigarrow}$  and  $\overset{P_{n,h_0}}{\rightarrow}$  denoting convergence in distribution and in probability under the sequence of measures  $\{F_{n, \beta_{(P_0, \eta_0)}(P_{n,h})}\}$  with  $P_{n,h} = P_0 + h/\sqrt{n}$  for  $h \in \mathbb{R}^k$ , and where  $Z \sim N(h_0, I_0^{-1})$  with  $I_0 = I_0(P_0, \eta_0)$  the (nonsingular) semiparametric information bound at  $(P_0, \eta_0)$ .

**Assumption 4** (i)  $\mathcal{P}$  is an open subset of  $\mathbb{R}^k$  with  $P_0 \in \mathcal{P}$ ;

(ii) The model for  $X^n$  has an approximately least-favorable submodel at each  $(P_0, \eta_0) \in \mathcal{P} \times \mathcal{H}$ .

We then impose some high-level assumptions on the large-sample behavior of the quasi-posterior. In this case, we say that  $\pi_n(P|X^n)$  is *consistent* if  $\pi_n(P \in N|X^n) \xrightarrow{P_0} 1$  for every neighborhood  $N$  containing  $P_0$ , for each  $(P_0, \eta_0) \in \mathcal{P} \times \mathcal{H}$ . Let  $\mathbb{P}_h$  denote probability with respect to  $Z \sim N(h, I_0^{-1})$  where  $I_0 = I_0(P_0, \eta_0)$  is the semiparametric information bound and let  $\mathbb{E}^*$  denote expectation with respect to  $Z^*$  where  $Z^* \sim N(0, I_0^{-1})$  independent of  $Z$ .

**Assumption 5** (i) The quasi-posterior  $\pi_n(P|X^n)$  is consistent;

(ii) At any  $P_0 \in \mathcal{P}$  with non-singleton  $\arg \min_{d \in \mathcal{D}} R(d, P_0)$ , for any Borel set  $A \subseteq \mathbb{R}^{D+1}$  and  $h \in \mathbb{R}^k$ , we have

$$\begin{aligned} \lim_{n \rightarrow \infty} F_{n, \beta_{(P_0, \eta_0)}(P_{n,h})} \left( \int \sqrt{n}(\rho(P) - \rho(P_0)) d\pi_n(P|X^n) \in A \right) \\ = \mathbb{P}_h(\mathbb{E}^*[\dot{\rho}_{P_0}[Z^* + Z]|Z] \in A) \end{aligned}$$

for all  $\eta_0 \in \mathcal{H}$ .

## B Proofs

To simplify notation, throughout we write  $\pi_n(P)$  for the posterior  $\pi_n(P|X^n)$ . We adopt the convention that  $+\infty \times 0 = 0$ . We also suppress dependence of  $d_n^*$  on  $\pi$ , since any  $\pi \in \Pi$  leads to a decision rule with the same asymptotic behavior.

## B.1 Proof of Theorems 1 and 2 and Corollary 1

### B.1.1 Preliminaries

Our proof of Theorem 1 follows the approach of Hirano and Porter (2009), with appropriate modifications to handle non-binary problems. By an asymptotic representation theorem of van der Vaart (1991), Assumption 2 implies that for any  $\{d_n\} \in \mathbb{D}$  there exists a function  $d_\infty(z, u; P_0)$  with  $d_\infty(Z, U; P_0) \sim Q_{P_0, h}$  where  $Z \sim N(h, I_0^{-1})$  and  $U \sim \text{Uniform}[0, 1]$  with  $U$  and  $Z$  independent. To summarize the argument, suppose that

$$\lim_{n \rightarrow \infty} \sqrt{n} \left( \mathbb{E}_{P_{n,h}} [R(d_n(X^n), P_{n,h})] - \min_{d \in \mathcal{D}} R(d, P_{n,h}) \right) = L_\infty(d_\infty; P_0, h) \quad (\text{A.1})$$

for a suitable function  $L_\infty$ . We may therefore rewrite the criterion (9) as

$$\mathcal{R}(\{d_n\}; P_0) = \int L_\infty(d_\infty; P_0, h) dh.$$

We show that the asymptotic representation  $d_\infty^*$  of  $\{d_n^*\}$  is optimal in the limit experiment: for each  $P_0 \in \mathcal{P}$ ,

$$\int L_\infty(d_\infty^*; P_0, h) dh = \inf_{d_\infty} \int L_\infty(d_\infty; P_0, h) dh,$$

where the infimum is over all such (possibly randomized)  $\mathcal{D}$ -valued decisions  $d_\infty(Z, U; P_0)$  in the limit experiment. It follows that any  $\{d_n\} \in \mathbb{D}$  which is asymptotically equivalent to  $d_n^*$  must be asymptotically optimal.

Deriving the asymptotic representation of  $\{d_n\} \in \mathbb{D}$  requires a tie-breaking rule when the argmin is not a singleton. In what follows, we take the smallest index  $d$  among the set of minimizers, though our optimality result does not rely on this and any (possibly randomized) tie-breaking rule is optimal. Because we use a non-randomized rule, asymptotic representations are of the form  $d_\infty(Z; P_0)$ . Recall that  $\mathbb{P}_h$  denotes probability with respect to  $Z \sim N(h, I_0^{-1})$  and  $\mathbb{E}^*$  denote expectation with respect to  $Z^*$  where  $Z^* \sim N(0, I_0^{-1})$  independently of  $Z$ . Recall  $\rho(P) = (R(0, P), R(1, P), \dots, R(D, P))'$ . Let  $\rho_d(P) = R(d, P)$  and let  $\dot{\rho}_{d, P_0}[\cdot]$  denote its directional derivative at  $P_0$ .

**Lemma 1** *Suppose that Assumptions 1, 2, and 3 hold. Fix any  $P_0 \in \mathcal{P}$ . Without loss of generality, reorder the elements of  $\mathcal{D}$  so that  $R(0, P_0) \leq R(1, P_0) \leq \dots \leq R(D, P_0)$ . Let  $k \in \mathcal{D}$  satisfy  $R(0, P_0) = \dots = R(k, P_0)$  with  $R(k, P_0) < R(k+1, P_0)$  if  $k < D$ . Then:*

(i) The sequence  $\{d_n^*\}$  belongs to  $\mathbb{D}$ , and its asymptotic representation at  $P_0$  is

$$d_\infty^*(z; P_0) = \sum_{i=0}^k i \times \mathbb{I} \left[ \mathbb{E}^* [\dot{\rho}_{i, P_0} [Z^* + z]] < \min_{j < i} \mathbb{E}^* [\dot{\rho}_{j, P_0} [Z^* + z]] \right] \\ \times \left[ \mathbb{E}^* [\dot{\rho}_{i, P_0} [Z^* + z]] \leq \min_{i < j \leq k} \mathbb{E}^* [\dot{\rho}_{j, P_0} [Z^* + z]] \right],$$

where the minimum over an empty index is  $+\infty$ ;

(ii) For any  $\{d_n\} \in \mathbb{D}$ ,

$$L_\infty(d_\infty; P_0, h) \\ = \sum_{i=0}^k (\dot{\rho}_{i, P_0} [h] - \min_{j \leq k} \dot{\rho}_{j, P_0} [h]) \times \mathbb{P}_h(d_\infty(Z; P_0) = i) + \infty \times \mathbb{P}_h(d_\infty(Z; P_0) > k),$$

where  $d_\infty(Z; P_0)$  is the asymptotic representation of  $\{d_n\}$  at  $P_0$ ;

(iii) The decision  $d_\infty^*$  is optimal in the limit experiment.

**Proof of Lemma 1.** Part (i). As  $d_n$  takes values in a finite set  $\mathcal{D}$ , establishing convergence in distribution under  $\{F_{n, P_{n, h}}\}$  is equivalent to showing  $\lim_{n \rightarrow \infty} F_{n, P_{n, h}}(d_n(X^n) = 1)$  exists. Suppose  $k < D$ . Then by Assumption 1(i), choose  $\varepsilon > 0$  and a neighborhood  $N$  of  $P_0$  upon which  $\max_{i \leq k} R(i, P) \leq \min_{j \geq k} R(j, P) - \varepsilon$ . By posterior consistency (Assumption 3(i)),

$$\max_{i \leq k} \bar{R}(i|X^n) - \min_{j > k} \bar{R}(j|X^n) \leq -\varepsilon \pi_n(P \in N) + 2\|\rho\|_\infty \pi_n(P \notin N) \xrightarrow{P_0} -\varepsilon,$$

where  $\|\rho\|_\infty = \max_{d \in \mathcal{D}} \sup_{P \in \mathcal{P}} |R(d, P)|$ , which is finite under Assumption 1(i). The convergence in probability in the above display holds along  $P_{n, h}$  for any  $h$  by Le Cam's first lemma and Assumption 2. Therefore,

$$\mathbb{I}[d_n^*(X^n) > k] \xrightarrow{P_{n, h}} 0.$$

It follows that the asymptotic distribution  $Q_{P_0, h}^*$  of  $d_n^*(X^n)$  along  $P_{n, h}$  assigns zero probability to  $\{k+1, \dots, D\}$  for all  $h \in \mathbb{R}^k$ . We may therefore take  $d_\infty^*(z; P_0) \in \{0, 1, \dots, k\}$  for all  $z \in \mathbb{R}^k$ . In particular,  $d_\infty^*(z; P_0) \equiv 0$  when  $k = 0$ . We therefore consider only  $k \geq 1$  in what follows.

As  $\max_{i \leq k} \bar{R}(i|X^n) < \min_{j > k} \bar{R}(j|X^n)$  wpa1 along  $P_{n, h}$ , in what follows it is without loss of generality to work on the sequence of events upon which  $\max_{i \leq k} \bar{R}(i|X^n) < \min_{j > k} \bar{R}(j|X^n)$ . Under the above tie-breaking rule, for any  $i \leq k$  we have

$$\mathbb{I}[d_n^*(X^n) = i] = \mathbb{I} \left[ \bar{R}(i|X^n) < \min_{j < i} \bar{R}(j|X^n) \text{ and } \bar{R}(i|X^n) \leq \min_{i < j \leq k} \bar{R}(j|X^n) \right],$$

where the minimum over an empty set is  $+\infty$ . As  $R(0, P_0) = \dots = R(k, P_0)$ , we may rewrite the previous expression in terms of  $\rho_{i0}(P) = R(i, P) - R(i, P_0)$ :

$$\begin{aligned} \mathbb{I}[d_n^*(X^n) = i] &= \mathbb{I} \left[ \int \sqrt{n} \rho_{i0}(P) d\pi_n(P) < \min_{j < i} \int \sqrt{n} \rho_{j0}(P) d\pi_n(P) \right] \\ &\quad \times \mathbb{I} \left[ \int \sqrt{n} \rho_{i0}(P) d\pi_n(P) \leq \min_{i < j \leq k} \int \sqrt{n} \rho_{j0}(P) d\pi_n(P) \right]. \end{aligned}$$

By Assumption 3(ii) with  $A = \{(x_0, x_1, \dots, x_D) : x_i < \min_{j < i} x_j \text{ and } x_i \leq \min_{i < j \leq k} x_j\}$ ,

$$\begin{aligned} \lim_{n \rightarrow \infty} F_{n, P_{n, h}}(d_n^*(X^n) = i) &= \mathbb{P}_h \left( \left( \mathbb{E}^*[\dot{\rho}_{i, P_0}[Z^* + Z]|Z] < \min_{j < i} \mathbb{E}^*[\dot{\rho}_{j, P_0}[Z^* + Z]|Z] \right) \right. \\ &\quad \left. \text{and } \left( \mathbb{E}^*[\dot{\rho}_{i, P_0}[Z^* + Z]|Z] \leq \min_{i < j \leq k} \mathbb{E}^*[\dot{\rho}_{j, P_0}[Z^* + Z]|Z] \right) \right). \end{aligned}$$

Therefore,  $\{d_n^*\} \in \mathbb{D}$ . The asymptotic representation of  $\{d_n^*\}$  follows from the preceding display.

Part (ii): First write

$$R(d, P_{n, h}) - \min_{d \in \mathcal{D}} R(d, P_{n, h}) = \sum_{i=0}^D \mathbb{I}[d = i] \left( R(i, P_{n, h}) - \min_{d \in \mathcal{D}} R(d, P_{n, h}) \right).$$

By Assumption 1(i),  $\min_{d \in \mathcal{D}} R(d, P_{n, h}) = \min_{j \leq k} R(j, P_{n, h})$  for all  $n$  sufficiently large, and so

$$\begin{aligned} L_n(d_n; P_0, h) &= \sum_{i=0}^k \mathbb{E}_{P_{n, h}} [\mathbb{I}[d_n(X^n) = i]] \sqrt{n} \left( R(i, P_{n, h}) - \min_{j \leq k} R(j, P_{n, h}) \right) \\ &\quad + \sum_{i=k+1}^D \mathbb{E}_{P_{n, h}} [\mathbb{I}[d_n(X^n) = i]] \sqrt{n} \left( R(i, P_{n, h}) - \min_{j \leq k} R(j, P_{n, h}) \right), \end{aligned}$$

for all  $n$  sufficiently large, where the second sum is zero if  $k = D$ . As  $\mathbb{E}_{P_{n, h}} [\mathbb{I}[d_n(X^n) = i]] \rightarrow \mathbb{P}_h[d_\infty(Z; P_0) = i]$  for all  $i \in \mathcal{D}$  and  $\liminf_{n \rightarrow \infty} (R(i, P_{n, h}) - \min_{j \leq k} R(j, P_{n, h})) > 0$  for  $i > k$ , we have

$$\sum_{i=k+1}^D \mathbb{E}_{P_{n, h}} [\mathbb{I}[d_n(X^n) = i]] \sqrt{n} \left( R(i, P_{n, h}) - \min_{j \leq k} R(j, P_{n, h}) \right) \rightarrow +\infty \times \mathbb{P}_h(d_\infty(Z; P_0) > k)$$

whenever  $k < D$ . Moreover, as  $R(0, P_0) = \dots = R(k, P_0)$ , for any  $i \leq k$  we have

$$\begin{aligned} \sqrt{n} \left( R(i, P_{n, h}) - \min_{j \leq k} R(j, P_{n, h}) \right) &= \left( \sqrt{n} \rho_{i0}(P_{n, h}) - \min_{j \leq k} \sqrt{n} \rho_{j0}(P_{n, h}) \right) \\ &\rightarrow \dot{\rho}_{i, P_0}[h] - \min_{j \leq k} \dot{\rho}_{j, P_0}[h], \end{aligned}$$

where the final assertion is by Assumption 1(ii) and the fact that if  $f_i$ ,  $i = 1, 2$  are directionally differentiable at  $P_0$  and  $f_1(P_0) = f_2(P_0)$ , then  $f(P) := \min_{i=1,2} f_i(P)$  has directional derivative  $\dot{f}_{P_0}[h] = \min_{i=1,2} \dot{f}_{i,P_0}[h]$ .

Part (iii): By part (ii), we see that  $d_\infty(z; P_0) \in \{0, 1, \dots, k\}$  for almost every  $z$  is necessary for optimality in the limit experiment. For any such  $d_\infty$ , we have by part (ii) that

$$\begin{aligned} \int L_\infty(d_\infty; P_0, h) dh &= \sum_{i=0}^k \int \mathbb{P}_h[d_\infty(Z; P_0) = i] \left( \dot{\rho}_{i,P_0}[h] - \min_{j \leq k} \dot{\rho}_{j,P_0}[h] \right) dh \\ &\propto \sum_{i=0}^k \int \int \mathbb{I}[d_\infty(z; P_0) = i] \left( \dot{\rho}_{i,P_0}[h] - \min_{j \leq k} \dot{\rho}_{j,P_0}[h] \right) e^{-\frac{1}{2}(z-h)'I_0(z-h)} dz dh \\ &= \sum_{i=0}^k \int \int \mathbb{I}[d_\infty(z; P_0) = i] \left( \dot{\rho}_{i,P_0}[h] - \min_{j \leq k} \dot{\rho}_{j,P_0}[h] \right) e^{-\frac{1}{2}(z-h)'I_0(z-h)} dh dz. \end{aligned}$$

As the integrand is non-negative, changing the order of integration in the final line is justified by Tonelli's theorem. Minimizing pointwise in  $z$ , we see that if

$$\mathcal{I}(z) = \arg \min_{i \leq k} \int \dot{\rho}_{i,P_0}[h] e^{-\frac{1}{2}(z-h)'I_0(z-h)} dh \equiv \arg \min_{i \leq k} \mathbb{E}^*[\dot{\rho}_{i,P_0}[Z + z]], \quad (\text{A.2})$$

then choosing  $d_\infty(z; P_0)$  to be any element of  $\mathcal{I}(z)$  for each  $z$  is optimal in the limit experiment. The tie-breaking rule used in part (i) is a special case with  $d_\infty^*(z; P_0) = \min \mathcal{I}(z)$ . ■

### B.1.2 Proof of Main Results

**Proof of Theorem 1.** Part (i) is immediate from parts (i) and (iii) of Lemma 1, which together imply

$$\mathcal{R}(\{d_n^*\}, P_0) = \int L_\infty(d_\infty^*; P_0, h) dh = \inf_{d_\infty} \int L_\infty(d_\infty; P_0, h) dh = \inf_{\{d'_n\} \in \mathbb{D}} \mathcal{R}(\{d'_n\}; P_0). \quad (\text{A.3})$$

Part (ii) follows from the fact that the asymptotic representation of  $d_n^*$  derived in part (i) of Lemma 1 does not depend on the prior  $\pi \in \Pi$ . For part (iii), note if  $\{d_n\}$  is asymptotically equivalent to  $\{d_n^*\}$ , then  $\{d_n\}$  must also have asymptotic representation  $d_\infty^*$ . The result now follows from (A.3). ■

**Proof of Theorem 2.** By Lemma 1(ii) and the proof of Theorem 1, it suffices to show that

$$\int L_\infty(d_\infty; P_0, h) dh > \int L_\infty(d_\infty^*; P_0, h) dh, \quad (\text{A.4})$$

where  $d_\infty$  is the asymptotic representation of  $\{d_n\}$ .

First suppose that  $\arg \min_{d \in \mathcal{D}} R(d, P_0)$  is a singleton. Without loss of generality let  $d = 0$

be the minimizing value. In view of Lemma 1(i), we must have

$$\lim_{n \rightarrow \infty} F_{n, P_n, h_0}(d_n(X^n) = 0) < \lim_{n \rightarrow \infty} F_{n, P_n, h_0}(d_n^*(X^n) = 0) = 1$$

for some  $h_0$  in  $\mathbb{R}^k$ . In terms of the asymptotic representation  $d_\infty$  of  $\{d_n\}$ , this means

$$\mathbb{P}_{h_0}(d_\infty(Z; P_0) = 0) < 1.$$

It follows by continuity of  $h \mapsto \mathbb{P}_h(d_\infty(Z; P_0) = 0)$  that there exists a neighborhood  $H \subset \mathbb{R}^k$  of  $h_0$  upon which  $\mathbb{P}_h(d_\infty(Z) = 0) < 1$ . In view of Lemma 1(ii), this implies  $L_\infty(d_\infty; P_0, h) = +\infty$  for all  $h \in H$  whereas  $L_\infty(d_\infty^*; P_0, h) = 0$  for all  $h \in \mathbb{R}^k$ . Therefore, inequality (A.4) holds.

Now suppose that  $\arg \min_{d \in \mathcal{D}} R(d, P_0)$  is not a singleton. As in Lemma 1, it is without loss of generality to reorder elements of  $\mathcal{D}$  so that  $R(0, P_0) \leq R(1, P_0) \leq \dots \leq R(D, P_0)$ . Let  $k \in \{1, \dots, D\}$  satisfy  $R(0, P_0) = \dots = R(k, P_0)$  with  $R(k, P_0) < R(k+1, P_0)$  if  $k < D$ . If the asymptotic representation  $d_\infty$  of  $\{d_n\}$  satisfies

$$\mathbb{P}_{h_0}(d_\infty(Z; P_0) \leq k) < 1$$

for any  $h_0 \in \mathbb{R}^k$ , then by similar arguments to the above we have  $L_\infty(d_\infty; P_0, h) = +\infty$  for all  $h$  in a neighborhood  $H$  of  $h_0$  and so inequality (A.4) holds. For the remainder of the proof we therefore suppose that

$$\mathbb{P}_h(d_\infty(Z; P_0) \leq k) = 1$$

for all  $h \in \mathbb{R}^k$ . By Lemma 1(ii) it suffices to show

$$\begin{aligned} \sum_{i=0}^k \int (\dot{\rho}_{i, P_0}[h] - \min_{j \leq k} \dot{\rho}_{j, P_0}[h]) \times \mathbb{P}_h(d_\infty(Z; P_0) = i) dh \\ > \sum_{i=0}^k \int (\dot{\rho}_{i, P_0}[h] - \min_{j \leq k} \dot{\rho}_{j, P_0}[h]) \times \mathbb{P}_h(d_\infty^*(Z; P_0) = i) dh. \end{aligned}$$

Since both integrands are non-negative, we may use Tonelli's theorem to restate the inequality as

$$\begin{aligned} \sum_{i=0}^k \int \mathbb{I}[d_\infty(z; P_0) = i] \left( \int (\dot{\rho}_{i, P_0}[h] - \min_{j \leq k} \dot{\rho}_{j, P_0}[h]) e^{-\frac{1}{2}(z-h)' I_0(z-h)} dh \right) dz \\ > \sum_{i=0}^k \int \mathbb{I}[d_\infty^*(z; P_0) = i] \left( \int (\dot{\rho}_{i, P_0}[h] - \min_{j \leq k} \dot{\rho}_{j, P_0}[h]) e^{-\frac{1}{2}(z-h)' I_0(z-h)} dh \right) dz. \end{aligned}$$

Recall the set  $\mathcal{I}(z)$  from (A.2). The function  $d_\infty^*(z; P_0)$  takes values in  $\mathcal{I}(z)$  for each  $z$ , so the

preceding inequality must hold weakly. To establish a strict inequality, note  $d_\infty(z; P_0)$  and  $d_\infty^*(z; P_0)$  must disagree on a set of positive Lebesgue measure, say  $\mathcal{Z}$  (otherwise  $d_\infty(Z; P_0)$  and  $d_\infty^*(Z; P_0)$  would have the same distribution for  $Z \sim N(h_0, I_0)$ ). For each  $z \in \mathcal{Z}$ , one of the following must hold:

- (a)  $\mathcal{I}(z)$  is a singleton and  $d_\infty(z, I) \notin \mathcal{I}(z)$ , or
- (b)  $\mathcal{I}(z)$  is not a singleton and  $d_\infty(z, P_0) \neq d_\infty^*(z; P_0)$ .

Condition (ii) in the statement of the theorem implies that  $\mathcal{I}(z)$  is non-singleton on a set of zero Lebesgue measure. Therefore, (a) must hold at almost every  $z \in \mathcal{Z}$ , and for any such  $z$  we have

$$\begin{aligned} \sum_{i=0}^k \mathbb{I}[d_\infty(z; P_0) = i] & \left( \int (\dot{\rho}_{i, P_0}[h] - \min_{j \leq k} \dot{\rho}_{j, P_0}[h]) e^{-\frac{1}{2}(z-h)' I_0(z-h)} dh \right) \\ & > \sum_{i=0}^k \mathbb{I}[d_\infty^*(z; P_0) = i] \left( \int (\dot{\rho}_{i, P_0}[h] - \min_{j \leq k} \dot{\rho}_{j, P_0}[h]) e^{-\frac{1}{2}(z-h)' I_0(z-h)} dh \right), \end{aligned}$$

proving (A.4). ■

Finally, we prove Corollary 1. For this we require  $\{d_n\}$  to satisfy a condition permitting the exchange of limits and integration, namely

$$\limsup_{n \rightarrow \infty} \mathcal{R}_n(d_n; P_0) \leq \int \limsup_{n \rightarrow \infty} \sqrt{n} \left( \mathbb{E}_{P_{n,h}} [R(d_n(X^n), P_{n,h})] - \min_{d \in \mathcal{D}} R(d, P_{n,h}) \right) w(P_{n,h}) dh \quad (\text{A.5})$$

for all  $P_0 \in \mathcal{P}$ . By the reverse Fatou lemma, condition (A.5) holds if for each  $P_0 \in \mathcal{P}$  there exists an  $h$ -integrable function  $g(h; P_0)$  with

$$\sqrt{n} \left( \mathbb{E}_{P_{n,h}} [R(d_n(X^n), P_{n,h})] - \min_{d \in \mathcal{D}} R(d, P_{n,h}) \right) w(P_{n,h}) \leq g(h; P_0)$$

for all  $n$ .

**Proof of Corollary 1.** The first part of our proof is similar to Lemma 1 of [Hirano and Porter \(2009\)](#). By Fatou's lemma and Lemma 1(ii), for any  $\{d'_n\} \in \mathbb{D}$  we have

$$\liminf_{n \rightarrow \infty} \mathcal{R}_n(d'_n; P_0) \geq w(P_0) \cdot \int L_\infty(d'_\infty; P_0, h) dh,$$

where  $d'_\infty$  is the asymptotic representation of  $\{d'_n\}$ . In particular this holds for  $\{d_n\}$ . In view of inequality (A.5) and Lemma 1(ii), we have

$$\liminf_{n \rightarrow \infty} \mathcal{R}_n(d_n; P_0) = w(P_0) \cdot \int L_\infty(d_\infty; P_0, h) dh = w(P_0) \cdot \inf_{d_\infty} \int L_\infty(d_\infty; P_0, h) dh,$$

where the final equality is by asymptotic equivalence of  $\{d_n\}$  and  $\{d_n^*\}$  and Lemma 1(iii). Therefore,

$$\lim_{n \rightarrow \infty} \mathcal{R}_n(d_n; P_0) = \lim_{n \rightarrow \infty} \mathcal{R}_n(d_n^*; P_0) = \inf_{\{d'_n\} \in \mathbb{D}} \liminf_{n \rightarrow \infty} \mathcal{R}_n(d'_n; P_0).$$

Finally, note that

$$\inf_{\{d'_n\} \in \mathbb{D}} \liminf_{n \rightarrow \infty} \mathcal{R}_n(d'_n; P_0) = \liminf_{n \rightarrow \infty} \inf_{d'_n} \mathcal{R}_n(d'_n; P_0).$$

because  $\mathcal{R}_n(d_n^*; P_0) = \inf_{d'_n} \mathcal{R}_n(d'_n; P_0)$  for each  $n$ . ■

## B.2 Proof of Theorems 3 and 4

Note that once we fix  $(P_0, \eta_0)$ , the approximately least-favorable model  $t \mapsto F_{n, \beta_{(P_0, \eta_0)}}(t)$  is a parametric model. In particular, under Assumption 4 any sequence  $\{d_n\} \in \mathbb{D}$  has an asymptotic representation in the limit experiment associated with  $\{F_{n, \beta_{(P_0, \eta_0)}}(P_{n, h}) : h \in \mathbb{R}^k\}$ , namely  $d_\infty(Z, U; (P_0, \eta_0)) \sim Q_{(P_0, \eta_0), h}$  for  $Z \sim N(h, I_0^{-1})$  with  $I_0 = I_0(P_0, \eta_0) > 0$  denoting the semiparametric information bound and  $U \sim \text{Uniform}[0, 1]$  with  $U$  and  $Z$  independent.<sup>20</sup> Once these modifications are made, the proofs of Theorems 3 and 4 follow identical arguments to the proofs of Theorems 1 and 2.

---

<sup>20</sup>As before, it suffices to consider asymptotic representations depending only on  $Z$ .