

Optimal Decision Rules when Payoffs are Partially Identified*

Timothy Christensen[†] Hyungsik Roger Moon[‡] Frank Schorfheide[§]

January 23, 2026

Abstract

We derive asymptotically optimal statistical decision rules for discrete choice problems when payoffs depend on a partially-identified parameter θ and the decision maker can use a point-identified parameter μ to deduce restrictions on θ . Examples include treatment choice under partial identification and pricing with rich unobserved heterogeneity. Our notion of optimality combines a minimax approach to handle the ambiguity from partial identification of θ given μ with an average risk minimization approach for μ . We show how to implement optimal decision rules using the bootstrap and (quasi-)Bayesian methods in both parametric and semiparametric settings. We provide detailed applications to treatment choice and optimal pricing. Our asymptotic approach is well suited for realistic empirical settings in which the derivation of finite-sample optimal rules is intractable.

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, C. Manski, F. Molinari, J. Porter, Q. Vuong, E. Vytlačil, three anonymous referees, and participants in various seminars and conferences for helpful comments and suggestions. This paper 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).

[†]Yale University. timothy.christensen@yale.edu

[‡]University of Southern California. moonr@usc.edu

[§]University of Pennsylvania, CEPR, NBER, and PIER. 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. How should they proceed?

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

We propose a theory of optimal statistical decision making in this setting, building on a line of research going back to [Manski \(2000\)](#). We depart from this body of work by combining a minimax approach to handle the ambiguity that arises from partial identification of θ given μ and average (or integrated) risk minimization for μ . This asymmetric treatment of parameters is in the spirit of the generalized Bayes-minimax principle of [Hurwicz \(1951\)](#). The resulting optimal decision rules are “robust” in the sense that they minimize maximum risk or regret over the identified set $\Theta_0(\mu)$ for θ conditional on μ , and use the data to learn efficiently about features of μ germane to the choice problem.

Our optimal decision rules can be implemented very easily in realistic empirical settings. All we require is that, for every choice, the maximum risk (or regret) over $\theta \in \Theta_0(\mu)$ conditional on μ can be computed. The maximum risk is averaged across a bootstrap distribution for an efficient estimator $\hat{\mu}$ of μ , a posterior distribution for μ in parametric models, or a quasi-posterior based on a limited-information criterion for μ in semiparametric models. Our optimal decision is then simply to choose whatever choice has smallest average maximum risk. We refer to this rule as a *bootstrap decision* when the maximum risk is averaged over a bootstrap distribution, a *Bayes decision* when averaged over a posterior,¹ and a *quasi-Bayes decision* when averaged over a quasi-posterior. Despite its simplicity, we provide a formal

¹These are not Bayes decision rules in the usual sense, but rather hybrid Bayes-minimax rules which, as we explain below, also have an interpretation as a type of robust Bayes (or conditional Γ -minimax) rules. For simplicity, we refer to them as Bayes decisions and drop the robust/minimax qualifiers.

asymptotic optimality theory to justify our approach.

As running example, we show how to implement optimal decisions in the context of a treatment assignment problem under partial identification of the average treatment effect (ATE). In an empirical illustration, taken from [Ishihara and Kitagawa \(2021\)](#), a decision maker decides whether or not to adopt a job-training program based on several RCT estimates from other studies and their standard errors. The identified set $\Theta_0(\mu)$ is defined through intersection bounds derived by extrapolating multiple studies. In addition to treatment assignment problems, we discuss how our framework can be used for optimal pricing decisions in an environment with rich unobserved heterogeneity, where revealed preference arguments may be used to derive bounds on demand responses under counterfactual prices. In both contexts, the maximum risk of different choices is available in closed form or can be computed by solving a standard optimization problem, e.g., a linear program.

To elaborate on practicality a little, consider the two competing paradigms: Bayes and minimax. In our setting, the usual Bayes decision rule requires specifying a prior on the parameter space for (θ, μ) , computing the posterior for (θ, μ) having observed the data, then choosing the decision that minimizes posterior risk. A common criticism of Bayes decisions (see, e.g., [Manski \(2021\)](#)) is that they are only justified if one can elicit a credible subjective prior, which can be difficult in practice, and that the resulting decision will depend, to some extent, on the decision maker's choice of prior. While this is true under both point- and partial identification, the problem is more severe under the latter because the prior for θ is not updated by the data (e.g., [Moon and Schorfheide \(2012\)](#)). Thus, even asymptotically, the decision will depend on the prior for θ . By contrast, our (quasi-)Bayesian rules only require specifying a prior for μ and our decisions are asymptotically independent of this choice of prior. Moreover, our bootstrap rules sidestep the choice of a prior altogether.

Minimax decisions minimize the maximum risk over the parameter space for (θ, μ) . While this may be desirable, it is often not feasible. To derive minimax decisions one usually has to make strong assumptions on the data-generating process and restrict the dependence of payoffs on parameters to be of a very simple form. Indeed, we are not aware of any work deriving a 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 on the ATE must be estimated from data. Algorithms such as that of [Chamberlain \(2000\)](#) may be used to compute approximate minimax decisions, but the performance gap between these and the true minimax decision can be difficult to quantify. Adopting a minimax approach with respect to θ and a Bayes approach with respect to μ lends a great deal of tractability, allowing us to derive asymptotically optimal rules for a very broad class of empirically relevant settings where minimax rules are intractable.

We develop an asymptotic optimality theory for decisions based on parametric and semi-parametric models. For parametric models, our optimality criterion extends the asymptotic average risk criterion introduced by [Hirano and Porter \(2009\)](#) for point-identified settings to partially identified settings. We further extend this notion to semiparametric models via a least favorable parametric submodel. Both of these extensions represent new contributions to the literature on asymptotic optimality for statistical decision rules. Our main results show formally that the proposed (quasi-)Bayesian rule is asymptotically optimal. Moreover, any decision rule that is asymptotically equivalent to the (quasi-)Bayes rule is optimal as well. This includes an implementation that replaces averaging under a (quasi-)posterior by averaging under a bootstrap approximation of the sampling distribution of an efficient estimator $\hat{\mu}$ of μ , in cases in which these two distributions are asymptotically equivalent.

Importantly, we show asymptotic equivalence to the (quasi-)Bayes or bootstrap rules is *necessary* for optimality: any decision whose asymptotic behavior is different from these is sub-optimal. It follows that “plug-in” rules, which plug an efficient estimator $\hat{\mu}$ into the oracle decision rule if μ were known, can perform sub-optimally.² [Manski \(2021, 2023\)](#) shows in numerical experiments that plug-in rules may perform poorly under finite-sample minimax regret criteria. Our necessity result provides a complementary and quite general theoretical explanation for why plug-in rules may perform poorly, albeit under a different (but related) optimality criterion.³ In our empirical application to the adoption of a job-training program, our optimal decision produces different treatment recommendations than the plug-in rule for some sub-populations.

On the technical side, partial identification often leads to non-differentiability of key components of the maximum risk function in μ . We address this challenge by extending arguments in [Hirano and Porter \(2009\)](#) for models with smooth welfare contrasts to settings with directional (but not full) differentiability. An important technical building-block is the derivation of the asymptotic distribution of the (quasi-)posterior mean of directionally differentiable functions in both parametric and semiparametric models (see Propositions 2 and 3 in Appendix C). These results are of independent interest. While [Kitagawa, Montiel Olea, Payne, and Velez \(2020\)](#) derived the asymptotic behavior of the posterior distribution of directionally differentiable functions, we instead characterize the large-sample (frequentist) distribution of the posterior mean of such functions. Our results offer a novel contribution to

²[Manski \(2021, 2023\)](#) refers to plug-in rules as “as-if” optimization, because estimators $\hat{\mu}$ are treated as if they are the true parameters.

³For the intuition, consider a treatment assignment problem under partial identification of the ATE. Oracle rules depend on a robust welfare contrast $b(\mu)$ formed from bounds on the ATE. The bounds are non-smooth functions of μ in many empirically relevant settings reviewed in Section 6. This non-smoothness leads to a failure of the δ -method that breaks the asymptotic equivalence between plug-in rules, which depend on $b(\hat{\mu})$, and optimal rules, which depend on the average of $b(\cdot)$ across a bootstrap or posterior distribution.

the literature on asymptotics for non-smooth functions (Dümbgen, 1993; Fang and Santos, 2019). We also introduce the concept of σ -optimality to handle settings in which the average (under Lebesgue measure) maximum risk is infinite at the optimal decision.

Our paper complements prior work on treatment assignment under partial identification, including Manski (2000, 2007a, 2020, 2021, 2023), 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 asymptotic framework allows us to relax restrictive parametric assumptions and accommodate a much broader class of data-generating processes, including semiparametric models. 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 finite-sample distribution is unknown.⁴ Second, except for Chamberlain (2011), these works use optimality criteria that (in our notation) are minimax over (θ, μ) , whereas our criterion is minimax over the partially-identified parameter θ and averages over the point-identified parameter μ , reflecting the asymmetric parameterization of the problem.

Our approach is related to the multiple priors framework of Gilboa and Schmeidler (1989) and a particular form of robust Bayes decision making, called conditional Γ -minimax; see DasGupta and Studden (1989), Betro and Ruggeri (1992), and the recent survey by Giacomini, Kitagawa, and Read (2021). We show how, under specific informational assumptions, our decision rule arises from a two-player zero-sum game between a decision maker and an adversarial nature. In the conditional Γ -minimax version of this game, nature chooses a prior for $\theta \in \Theta_0(\mu)$ conditional on the reduced-form parameter μ and the data that are available for estimation of μ .

In the optimal pricing application we consider an environment in which a monopolist is choosing to set a price but faces ambiguity about the true distribution of demand. As in Bergemann and Schlag (2011), we assume the decision maker has a preference for robustness. However, we use revealed preference demand theory to derive bounds on demand, rather than assuming true demand is in a neighborhood of a pre-specified distribution. Moreover, in our setting the decision maker must estimate the bounds from data. 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

⁴Finite-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.

Stoye (2015), Manski (2007b, 2014), and Kitamura and Stoye (2018, 2019). These works are primarily concerned with testing rationality or deriving bounds on demand. We instead focus on using bounds to solve an optimal pricing problem, and using demand data efficiently in that context.

The remainder of the paper is structured as follows. Section 2 presents an application to treatment assignment based on extrapolating meta-analyses. We first discuss the oracle decision based on a known μ , then describe our proposed optimal decision rule in this context, and finally provide an empirical illustration. Section 3 outlines our general framework for optimal decision making, discusses Bayesian and bootstrap implementations, and provides comparisons to minimax and conditional Γ -minimax decisions. The large sample optimality theory is presented in Section 4 for parametric models and Section 5 for semiparametric models. Further applications to other treatment assignment problems and optimal pricing are discussed in Sections 6 and 7, respectively. Finally, Section 8 concludes. Proofs and additional results and lemmas are relegated to the Appendix.

2 Treatment Assignment

To set the stage, we will present a specific treatment assignment application in Section 2.1, describe the proposed decision rule in Sections 2.2 and 2.3, and implement it numerically for the application in Section 2.4.

2.1 The Basic Problem and an Application

Suppose a social planner is choosing whether or not to introduce a treatment in a target population. The social planner does not know the target population’s ATE, defined as $\theta = \mathbb{E}[Y]$, where $Y := Y_1 - Y_0$ is the treatment effect associated with any individual, Y_1 is the outcome under treatment and Y_0 the outcome in the absence of treatment. Instead, the social planner observes a meta-analysis consisting of estimates $\hat{\mu}_k$ of the ATE μ_k in several related populations $k = 1, \dots, K$, their standard errors s_k , characteristics x_k of the related populations, and characteristics x_0 of the target population.

For concreteness, we revisit Example 2 of Ishihara and Kitagawa (2021). The authors 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 consisting of gender (males, females, or both), age (youths, adults, or both), OECD membership status, GDP growth (standardized) and unemployment (standardized).

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 growth of 3.48% and an unemployment rate of 9.45%).

To connect the RCT studies to the ATE θ associated with the target population, suppose that treatment effects are Lipschitz in the distance between size-specific covariates:

$$|\mu_k - \theta| \leq C\|x_k - x_0\|, \quad k = 1, \dots, K, \quad (1)$$

where C is a pre-specified constant. We refer to $\mu := (\mu_1, \dots, \mu_k)$ as the reduced-form parameter. It is assumed to be consistently estimable from the data and its estimate is denoted by $\hat{\mu}$. We obtain the following lower (L) and upper (U) bounds for θ :

$$b_L(\mu) := \max_{1 \leq k \leq K} (\mu_k - C\|x_0 - x_k\|) \leq \theta \leq \min_{1 \leq k \leq K} (\mu_k + C\|x_0 - x_k\|) =: b_U(\mu). \quad (2)$$

The bounds in (2) define an identified set for θ as a function of μ :

$$\Theta_0(\mu) := [b_L(\mu), b_U(\mu)]. \quad (3)$$

2.2 Oracle Decision

For now we will assume that μ is known. Following Manski (2000, 2004), it is common to derive treatment rules under a utilitarian social welfare function that is linear in the target population's ATE. Interpreting negative utility as loss, we write

$$l(d, \theta) = -d\theta,$$

where $d \in \{0, 1\}$ indicates treatment. The loss function is minimized by the treatment decision $\mathbb{I}[\theta \geq 0]$, where $\mathbb{I}[\cdot]$ is the indicator function that is equal to one if its argument is true and zero otherwise. Following Manski (2000, 2004), we subsequently use the regret criterion

$$r(d, \theta) = l(d, \theta) - \min_{d'} l(d', \theta) = (\mathbb{I}[\theta \geq 0] - d)\theta. \quad (4)$$

Failure to treat ($d = 0$) incurs zero regret when $\theta < 0$, otherwise the regret is θ . Similarly, treating ($d = 1$) incurs zero regret when $\theta \geq 0$, otherwise the regret is $-\theta$.

If the social planner knows μ , they can choose the decision that minimizes maximum regret over $\Theta_0(\mu)$, defined as

$$R(d, \mu) := \sup_{\theta \in \Theta_0(\mu)} r(d, \theta).$$

In our specific application we obtain

$$R(0, \mu) = (b_U(\mu))_+ \quad \text{and} \quad R(1, \mu) = -(b_L(\mu))_-, \quad (5)$$

where $(a)_+ := \max\{a, 0\}$ and $(a)_- := \min\{a, 0\}$. This leads to the *oracle decision*

$$\delta^o = \mathbb{I}[b(\mu) \geq 0] \quad (6)$$

that assigns treatment if the *robust welfare contrast*

$$b(\mu) := (b_U(\mu))_+ + (b_L(\mu))_-$$

is non-negative, and non-treatment otherwise.

Notice from plugging (2) into (5) that the max and min operators make $R(0, \mu)$ and $R(1, \mu)$ only directionally differentiable with respect to μ . Directional differentiability of the maximum risk function is a generic feature of all the applications discussed in this paper. It generates technical challenges that we tackle in our large sample analysis and distinguishes our analysis from that in [Hirano and Porter \(2009\)](#).

2.3 Optimal Decision Rule

The oracle rule is an infeasible first-best as it requires knowledge of the true μ . In practice μ is unknown and any practical rule must also confront sampling uncertainty about μ . One option is to simply plug $\hat{\mu}$ into (6), which leads to the plug-in rule

$$\delta_n^{plug} = \mathbb{I}[b(\hat{\mu}) \geq 0]. \quad (7)$$

However, this rule does not account for the precision of $\hat{\mu}$. [Manski \(2023\)](#) explores how different estimates of parameters affect the performance of plug-in rules.

We develop an optimality theory in Sections 3 to 5 below to guide decision-making in situations such as these where a payoff-relevant parameter θ is partially identified and its identified set $\Theta_0(\mu)$ is indexed by a parameter μ which can be estimated from data. In the context of the treatment assignment application, the procedure amounts to replacing $b(\mu)$ in (6) by its average \bar{b}_n under a posterior for μ (in parametric models) or quasi-posterior for μ (in semiparametric models):

$$\delta_n^* = \mathbb{I}[\bar{b}_n \geq 0]. \quad (8)$$

We show formally in Sections 4 and 5 that this decision is more efficient than the plug-in




Figure 1: Distribution of the robust welfare contrast $b(\mu)$ for $\mu \sim N(\hat{\mu}, \hat{\Sigma})$, its quasi-posterior mean \bar{b}_n , and the plug-in value $b(\hat{\mu})$.

rule, as it takes into account the sampling uncertainty in $\hat{\mu}$.

2.4 Empirical Illustration

We use a limited-information quasi-posterior distribution for μ that interprets the sampling distribution $\hat{\mu}_k \stackrel{\text{approx}}{\sim} N(\mu_k, s_k^2)$ as a quasi-likelihood. The estimates $\hat{\mu}_k$ are independent as they come from independent RCTs. Following Doksum and Lo (1990), Kim (2002), and Müller (2013), we combine the $N(\mu, \hat{\Sigma})$ quasi-likelihood for $\hat{\mu}$ with a flat prior for μ to obtain a $N(\hat{\mu}, \hat{\Sigma})$ quasi-posterior for μ , where $\hat{\Sigma}$ denotes a diagonal matrix with $(s_k^2)_{k=1}^K$ down its diagonal. The quasi-posterior mean of $b(\mu)$, denoted by \bar{b}_n , is computed by drawing independent $\mu \sim N(\hat{\mu}, \hat{\Sigma})$, computing $b(\mu)$, then averaging across a large number of draws. The efficient decision from (8) is then to treat if \bar{b}_n is non-negative, and not treat otherwise.

We compute the optimal decision as described above using data from Ishihara and Kitagawa (2021). We increase the Lipschitz constant in (1) from $C = 0.025$ in their paper to $C = 0.25$ so that the bounds on θ are non-empty. Figures 1a and 1b plot the quasi-posterior distribution of $b(\mu)$ under $\mu \sim N(\hat{\mu}, \hat{\Sigma})$ for males and females, respectively. Both figures display the mean \bar{b}_n whose sign determines the optimal treatment decision δ_n^* in (8). We also display the plug-in value $b(\hat{\mu})$ whose sign determines the plug-in rule δ_n^{plug} in (7).

In Figure 1a (male youths) we see that the distribution has a pronounced right-skew. Recall that the lower bound is $\max_{1 \leq k \leq K} (\mu_k - C \|x_0 - x_k\|)$. When evaluated at $\hat{\mu}$, the largest value of $\hat{\mu}_k - C \|x_0 - x_k\|$ is -0.3190 (corresponding to a US study) and the second-largest is -0.3298 (corresponding to a Brazilian study). As the maximum is not well separated relative to sampling uncertainty (the average s_k across the studies is 0.034), the distribution

of the lower bound is right-skewed because it behaves like a maximum of two Gaussians. The upper bound is less skewed because the minimum (corresponding to a Colombian study) is better separated. As a consequence of this asymmetry, the mean \bar{b}_n is positive and the optimal decision is to treat. By contrast, the plug-in value $b(\hat{\mu})$ is negative so the plug-in decision is to not treat. In Section 4 we show this difference between the two rules, where the optimal rule predicts treatment more aggressively than the plug-in rule, is as predicted by our optimality theory.

The situation is different for female youths (Figure 1b). Here the minima and maxima are relatively better separated, so the distribution of $b(\mu)$ for $\mu \sim N(\hat{\mu}, \hat{\Sigma})$ is close to Gaussian. In consequence, the mean \bar{b}_n and the plug-in value $b(\hat{\mu})$ are almost identical, and the optimal and plug-in decisions are both to treat.

3 Optimal Decisions and Their Implementation

The treatment assignment application in the previous section imposed a lot of structure which we will now relax. We begin by stating the general decision problem in Section 3.1. Many empirical examples are covered by our framework: see Sections 6 and 7 for applications to treatment assignment and optimal pricing and Christensen, Moon, and Schorfheide (2020) for an application to forecasting. We develop our concept of optimal decision rules in Section 3.2. Bayesian and bootstrap implementations are presented in Sections 3.3 and 3.4. Finally, in Section 3.5 we discuss a two-player zero-sum game representation of our setup and its relationship to minimax and robust Bayesian decision making.

3.1 Decision Problem

A decision maker (DM) observes $X^n \sim P_{n,\mu}$ taking values in a sample space \mathcal{X}^n . Often X^n may be a sample (X_1, \dots, X_n) of size n with joint distribution $P_{n,\mu}$. But X^n may also be a vector of estimators (e.g., OLS or IV estimators) with sampling distribution $P_{n,\mu}$. The distribution $P_{n,\mu}$ is indexed by a *reduced-form parameter* $\mu \in \mathcal{M}$.

After observing X^n , the DM chooses an action d in the action space $\mathcal{D} = \{0, 1, \dots, D\}$. A *statistical decision rule* $\delta_n : \mathcal{X}^n \rightarrow \mathcal{D}$ maps realizations of the data into actions.⁵ The DM's utility $u(d, Y, \theta, \mu)$ from choosing d depends on a random variable $Y \sim G_\theta$ independent of X^n , a *structural parameter* $\theta \in \Theta$, and μ . In a treatment assignment problem, X^n may represent data or summary statistics from an experiment, Y may represent the treatment effect for an

⁵We consider nonrandomized rules as opposed to randomized rules where the action space is the set of all distributions over $\{0, 1, \dots, D\}$.

individual sampled at random from the remaining population (hence independent of X^n), and $\mathbb{E}_\theta[Y]$ may represent the average treatment effect. Both G_θ and $\mathbb{E}_\theta[\cdot]$ can be conditional on covariates but we suppress this to simplify notation.

We interpret negative utility as loss and write $l(d, Y, \theta, \mu) \equiv -u(d, Y, \theta, \mu)$. The DM incurs this loss once Y is realized. Here it is helpful to draw a distinction between the *ex-post problem* that the DM faces after observing X^n but before Y is realized, and the *ex-ante problem* the DM faces before X^n is realized. In the ex-post problem, the DM faces risk

$$r(d, \theta, \mu) := \mathbb{E}_\theta[l(d, Y, \theta, \mu)]$$

from choosing $d \in \mathcal{D}$, where the expectation is taken with respect to $Y \sim G_\theta$.⁶ We refer to r as “risk” because it involves taking an expectation with respect to the random outcome Y in the ex-post problem. One can equally view r as a loss function in the DM’s ex-ante problem of choosing a decision rule δ_n before X^n is observed.

An important aspect of this paper is that we allow the payoff-relevant parameter θ to be set-identified. That is, the most that the DM can infer about θ is that $\theta \in \Theta_0(\mu) \subseteq \Theta$, where $\Theta_0(\cdot)$ is a known set-valued mapping from \mathcal{M} to Θ . While X^n can be used to learn about μ , it contains no information about the identity of the true θ within $\Theta_0(\mu)$.

Example: Treatment Assignment. Consider the setup in Section 2: $d \in \{0, 1\}$ is a binary treatment indicator, θ is the ATE, and $r(d, \theta, \mu) \equiv r(d, \theta)$ is given in display (4). The DM has bounds $b_L(\mu)$ and $b_U(\mu)$ on θ as a function of μ and observes data or summary statistics $X^n \sim P_{n, \mu}$ that may be used to learn μ . In the application from Section 2, the statistics were the vector of ATE estimators $X^n = (\hat{\mu}_1, \dots, \hat{\mu}_k)$, $P_{n, \mu}$ is their sampling distribution, the reduced-form parameter is the vector of ATEs $\mu = (\mu_1, \dots, \mu_k)$, and $b_L(\mu)$ and $b_U(\mu)$ are given in (2). We review other methods for constructing bounds in different empirical settings in Section 6. The identified set is $\Theta_0(\mu) = [b_L(\mu), b_U(\mu)]$. \square

3.2 Optimality

Our setup features two types of parameters: a point-identified parameter μ and a set-identified parameter θ . We adopt a minimax approach to handle the ambiguity that arises from the partial identification of θ conditional on μ and average (or integrated) risk minimization to estimate μ . Our asymmetric treatment of parameters is in the spirit of Hurwicz (1951); we discuss the connection in more detail in Section 3.3. We first consider the case in

⁶This notation nests welfare regret by setting $l(d, y, \theta, \mu) = \max_{d' \in \mathcal{D}} \mathbb{E}_\theta[W(d', Y, \theta, \mu)] - W(d, y, \theta, \mu)$ for a welfare function W . Hence, we do not distinguish between “risk” and “regret” in what follows.

which μ is known and then extend the analysis to the case of unknown μ , which will lead us to the definition of our optimality concept.

Known μ . If the DM knows μ , then the data are irrelevant and the germane problem is the ex-post problem in which d is chosen. To handle the ambiguity about $\theta \in \Theta_0(\mu)$, the DM evaluates actions by their maximum risk

$$R(d, \mu) := \sup_{\theta \in \Theta_0(\mu)} r(d, \theta, \mu). \quad (9)$$

Define

$$\delta^o(\mu) = \operatorname{argmin}_{d \in \mathcal{D}} R(d, \mu),$$

if the argmin is unique, otherwise $\delta^o(\mu)$ is chosen randomly from $\operatorname{argmin}_{d \in \mathcal{D}} R(d, \mu)$. This choice can be interpreted as the equilibrium in a two-player zero-sum game in which an adversarial nature chooses θ in response to the DM's choice of d to minimize $\mathbb{E}_\theta[u(d, Y, \theta, \mu)]$. We call $\delta^o(\mu)$ the *oracle* decision: it represents the DM's optimal choice if μ was known. The oracle decision is infeasible in any practical application because μ will need to be estimated. Nevertheless, it serves as a useful benchmark because $R(\delta_n(X^n), \mu) \geq R(\delta^o(\mu), \mu) \equiv \min_{d \in \mathcal{D}} R(d, \mu)$ for any data-dependent decision rule $\delta_n(X^n)$.

Unknown μ . Here the DM uses the data X^n to learn about μ before choosing $d \in \mathcal{D}$. We treat $R(d, \mu)$ as a loss function for the DM's ex-ante problem. For a decision rule δ_n , the DM incurs ex-ante risk

$$\mathbb{E}_\mu[R(\delta_n(X^n), \mu)], \quad (10)$$

where $\mathbb{E}_\mu[\cdot]$ denotes expectation with respect to $X^n \sim P_{n,\mu}$. Our goal is to construct sequences $\{\delta_n\}$ of decision rules that use the data efficiently, so that criterion (10) is minimized in large samples over a range of data-generating processes.

To simplify exposition, we consider parametric models in the rest of this section and defer discussion of semiparametric models to Section 5. Following Hirano and Porter (2009), we work in a local asymptotic framework centered at a fixed $\mu_0 \in \mathcal{M} \subseteq \mathbb{R}^K$ and parameterized by a local parameter $h = \sqrt{n}(\mu - \mu_0)$ ranging over \mathbb{R}^K . As h is not consistently estimable, this asymptotic framework is designed to mimic the finite-sample problem faced by the researcher, where μ is not known with certainty. We let $P_{n,h}$ denote the distribution of X^n when $\mu = \mu_0 + h/\sqrt{n}$ and let $\mathbb{E}_{n,h}$ denote expectation under $P_{n,h}$.

We rank decision rules by their ex-ante risk in excess of the oracle. We scale excess risk

by \sqrt{n} so that the large-sample limit is not degenerate. This leads to the criterion

$$\mathcal{R}(\{\delta_n\}; \mu_0, h) = \limsup_{n \rightarrow \infty} \sqrt{n} \left(\mathbb{E}_{n,h} [R(\delta_n(X^n), \mu_0 + h/\sqrt{n})] - \min_{d \in \mathcal{D}} R(d, \mu_0 + h/\sqrt{n}) \right). \quad (11)$$

As there is no a priori reason to assign more weight to one local parameter than another, we integrate with respect to Lebesgue measure over h to arrive at the criterion

$$\mathcal{R}(\{\delta_n\}; \mu_0) = \int \mathcal{R}(\{\delta_n\}; \mu_0, h) dh. \quad (12)$$

We rank sequences of decision rules by criterion (12). For now we assume $\mathcal{R}(\{\delta_n\}; \mu_0)$ is finite for at least one sequence of decisions $\{\delta_n\}$, so that criterion (12) provides a meaningful ranking. Section 4.4 discusses extensions to settings where this criterion is infinite and the improper Lebesgue prior on h is approximated by a sequence of proper priors.

Before introducing our definition of optimality, we first define a class \mathbb{D} of sequences of decision rules for which the above criteria are well defined. Let

$$\underline{\mathcal{D}}_{\mu_0} = \operatorname{argmin}_{d \in \mathcal{D}} R(d, \mu_0)$$

denote the set of optimal choices at μ_0 . Note that $\underline{\mathcal{D}}_{\mu_0}$ will be a non-singleton at values of μ_0 where there are multiple choices that minimize $R(d, \mu_0)$.

Definition 1 \mathbb{D} consists of all sequences $\{\delta_n\}$ such that, for all $\mu_0 \in \mathcal{M}$,

- (i) $\lim_{n \rightarrow \infty} P_{n,h}(\delta_n(X^n) = d)$ exists for all $d \in \mathcal{D}$ and $h \in \mathbb{R}^K$;
- (ii) $\lim_{n \rightarrow \infty} \sqrt{n} P_{n,h}(\delta_n(X^n) \notin \underline{\mathcal{D}}_{\mu_0}) = 0$ for all $h \in \mathbb{R}^K$.

Definition 1(i) requires that the sequence of random variables $\{\delta_n(X^n)\}$ converges in distribution along $\{P_{n,h}\}$. Definition 1(ii) is a minor technical condition requiring δ_n to choose among $\underline{\mathcal{D}}_{\mu_0}$ with high probability under small perturbations of μ around μ_0 . For the intuition, by the maximum theorem we know that optimal actions under μ near μ_0 should belong to $\underline{\mathcal{D}}_{\mu_0}$. If d_n chooses actions outside this set, then it makes mistakes: it chooses actions that should be easy to identify as sub-optimal. Condition (ii) requires the probability of these mistakes vanishes sufficiently fast. Bayes decisions satisfy this condition: their mistake probabilities vanish at rate n^{-1} (see Lemma 11 in Appendix D).

We refer to sequences $\{\delta_n\}$ as optimal if they satisfy the following definition:

Definition 2 Say $\{\delta_n\} \in \mathbb{D}$ is **optimal** if $\mathcal{R}(\{\delta_n\}; \mu_0) = \inf_{\{\delta'_n\} \in \mathbb{D}} \mathcal{R}(\{\delta'_n\}; \mu_0)$ for all $\mu_0 \in \mathcal{M}$.

This is a non-standard problem and there are different notions of optimality that may lead to different rankings over sequences of decision rules, as we discuss in Section 3.5 below. An attractive feature of our optimality criterion is that the DM only needs to be able to compute $R(d, \mu)$ in order to implement optimal decisions. This makes our approach broadly applicable including in scenarios when $R(d, \mu)$ has no closed-form expression, such as those in Sections 6 and 7. In the next subsections we present Bayesian and bootstrap implementations.

3.3 Bayesian Interpretation and Implementation

Integrating criterion (10) with respect to a prior π yields the *integrated maximum risk* criterion

$$\int \mathbb{E}_\mu[R(\delta_n(X^n), \mu)] d\pi(\mu). \quad (13)$$

After observing X^n , the DM can form a posterior $\pi_n(\mu) = \pi_n(\mu|X^n)$ for μ . Standard arguments (e.g. Wald, 1950, Chapter 5.1) imply that (13) can be minimized by minimizing the posterior maximum risk

$$\bar{R}_n(d) := \int R(d, \mu) d\pi_n(\mu) \quad (14)$$

with respect to $d \in \mathcal{D}$ for almost every realization of the data. This leads to the Bayes rule

$$\delta_n^*(X^n; \pi) = \operatorname{argmin}_{d \in \mathcal{D}} \bar{R}_n(d) \quad (15)$$

if the argmin is unique, otherwise $\delta_n^*(X^n; \pi)$ is chosen randomly from $\operatorname{argmin}_{d \in \mathcal{D}} \bar{R}_n(d)$. We include π as an argument to indicate that the decision depends on π in finite samples.

Bayes decisions $\delta_n^*(X^n; \pi)$ are optimal under criterion (12) for any prior π whose density is positive, bounded, and continuous (Theorem 1). To see the connection between (12) and (13), use a change-of-variables to express the prior for h as $\propto \pi(\mu_0 + h/\sqrt{n})$. This prior becomes uniform as $n \rightarrow \infty$, which is the weight function underlying (12).

Example: Treatment Assignment. Recall from Section 2 that the maximum risks of $d \in \{0, 1\}$ are the positive and negative parts of the bounds on the ATE:

$$R(0, \mu) = (b_U(\mu))_+, \quad R(1, \mu) = -(b_L(\mu))_-.$$

Averaging $R(1, \mu)$ and $R(0, \mu)$ across π_n yields

$$\bar{R}_n(0) = \int (b_U(\mu))_+ d\pi_n(\mu), \quad \bar{R}_n(1) = \int -(b_L(\mu))_- d\pi_n(\mu).$$

We have $\bar{R}_n(1) \leq \bar{R}_n(0)$ if and only if $\bar{b}_n := \int b(\mu) d\pi_n(\mu) \geq 0$ with $b(\mu) := (b_U(\mu))_+ + (b_L(\mu))_-$. Hence, the Bayes rule is $\delta_n^*(X^n; \pi) = \mathbb{I}[\bar{b}_n \geq 0]$.⁷ We discuss implementation in a number of different empirical settings in Section 6. \square

3.4 Bootstrap Implementation

Let $\hat{\mu}$ denote an efficient estimator of μ and let \mathbb{E}_n^* denote expectation with respect to the bootstrap version $\hat{\mu}^*$ of $\hat{\mu}$ conditional on X^n . Define the bootstrap average maximum risk

$$R_n^*(d) = \mathbb{E}_n^* [R(d, \hat{\mu}^*)].$$

The bootstrap rule is

$$\delta_n^{**}(X^n) = \operatorname{argmin}_{d \in \mathcal{D}} R_n^*(d)$$

if the argmin is unique (with a random selection from the argmin otherwise). We show in Theorem 1 below that as long as δ_n^{**} behaves asymptotically like δ_n^* it inherits the optimality property of the Bayes decision.

Example: Treatment assignment. Here we have

$$R_n^*(1) = \mathbb{E}_n^* [-(b_L(\hat{\mu}^*))_-], \quad R_n^*(0) = \mathbb{E}_n^* [(b_U(\hat{\mu}^*))_+].$$

Therefore, the bootstrap rule is $\delta_n^{**}(X^n) = \mathbb{I}[\mathbb{E}_n^* [b(\hat{\mu}^*)] \geq 0]$. \square

3.5 Further Discussion

The first distinguishing feature of our analysis is that we separate two groups of parameters, θ and μ , and apply the minimax reasoning only to θ because only it is partially identified. The approach of treating groups of parameters differently dates back to [Hurwicz \(1951\)](#). He argued that one might consider multiple priors for some parameters and referred to it as a generalized Bayes-minimax principle. He provided a two-parameter example in which the marginal prior for one of the parameters, μ in our notation, is fixed, whereas a family of priors is considered for the conditional distribution of the second parameter θ given μ .

Recall that $\delta_n^*(\cdot; \pi)$ minimizes criterion (13) over all decision rules $\delta_n : \mathcal{X}^n \rightarrow \mathcal{D}$. By the

⁷This rule deterministically chooses treatment when there is a tie (i.e., when $\bar{b}_n = 0$). Any (possibly randomized) tie-breaking rule leads to decision which is optimal under our optimality criteria.

definition of $R(d, \mu)$ and the discussion in Section 3.3, we see that $\delta_n^*(x^n; \pi)$ solves

$$\min_{d \in \mathcal{D}} \int_{\mathcal{M}} \left(\sup_{\theta \in \Theta_0(\mu)} r(d, \theta, \mu) \right) d\pi_n(\mu|x^n) \quad (16)$$

for almost every realization x^n of X^n , where $r(d, \theta, \mu) = \mathbb{E}_\theta[-u(d, Y, \theta, \mu)]$. Unlike the usual minimax framework, we allow the adversary (“nature”) to choose θ conditional on X^n . This is the second distinguishing feature of our analysis. It creates a more adversarial setting than usual, but it is disciplined by restricting nature’s choice to the identified set $\Theta_0(\mu)$ rather than the whole parameter space Θ .

In this regard, our approach is closely related to the derivation of conditional Γ -minimax decision rules in the Bayesian literature, e.g., [DasGupta and Studden \(1989\)](#), [Betro and Ruggeri \(1992\)](#), and [Giacomini, Kitagawa, and Read \(2021\)](#). One difference is that we keep the marginal prior distribution of μ fixed, which alleviates concerns about the conservativeness of the approach. Suppose one combines the unique prior π for μ with a family Λ of conditional priors $\lambda := \{\lambda(\cdot|\mu) : \mu \in \mathcal{M}\}$ for θ , where $\lambda(\theta|\mu)$ has support contained in $\Theta_0(\mu)$ for each $\mu \in \mathcal{M}$. Then one can define $\Gamma = \{\pi \otimes \lambda : \lambda \in \Lambda\}$, where each $\gamma \in \Gamma$ is a prior over (θ, μ) .⁸ The conditional Γ -minimax decision solves the following min-max problem for almost every realization x^n of X^n :

$$\min_{d \in \mathcal{D}} \sup_{\gamma \in \Gamma} \int_{\Theta \times \mathcal{M}} r(d, \theta, \mu) d\gamma_n(\theta, \mu|x^n), \quad (17)$$

where $\gamma_n(\theta, \mu|x^n)$ is the posterior for (θ, μ) given $X^n = x^n$ for the prior $\gamma \in \Gamma$. This is a robust Bayes criterion for the DM’s ex-post problem faced after observing X^n . Accordingly, X^n is treated as fixed: the DM and nature know X^n when choosing their actions.

In responding to the DM’s choice of d in (17), nature has to choose a set of conditional priors $\{\lambda(\theta|\mu) : \mu \in \mathcal{M}\}$ where each $\lambda(\cdot|\mu)$ has support contained in $\Theta_0(\mu)$. For each $\mu \in \mathcal{M}$, define $\theta_*(\mu; d) \in \arg \sup_{\theta \in \Theta_0(\mu)} r(d, \theta, \mu)$, assuming that the arg sup is non-empty, and let $\lambda_*(\theta|\mu, d)$ be a point mass at $\theta_*(\mu; d)$. In combination, π and $\{\lambda_*(\theta|\mu, d) : \mu \in \mathcal{M}\}$ define a joint prior $\gamma_* = \pi \otimes \lambda_*$ over (θ, μ) . Then, by construction, the posterior risk of d under the γ_* prior equals the posterior maximum risk

$$\int_{\Theta \times \mathcal{M}} r(d, \theta, \mu) d\gamma_{*n}(\theta, \mu|x^n) = \int_{\mathcal{M}} \left(\sup_{\theta \in \Theta_0(\mu)} r(d, \theta, \mu) \right) d\pi_n(\mu|x^n).$$

⁸We use the notation $\pi \otimes \lambda$ to denote the conditional probability measures $\lambda = \{\lambda(\cdot|\mu) : \mu \in \mathcal{M}\}$ integrated against a marginal probability measure π for μ . Thus, for $\gamma = \pi \otimes \lambda$ and $g : \Theta \times \mathcal{M} \rightarrow \mathbb{R}$, we define $\int_{\Theta \times \mathcal{M}} g(\theta, \mu) d\gamma(\theta, \mu) = \int_{\mathcal{M}} \int_{\Theta} g(\theta, \mu) d\lambda(\theta|\mu) d\pi(\mu)$.

Because the supremum over $\theta \in \Theta_0(\mu)$ is weakly larger than the average under any $\lambda(\theta|\mu)$, we can deduce that, provided $\gamma_* \in \Gamma$,

$$\sup_{\gamma \in \Gamma} \int_{\Theta \times \mathcal{M}} r(d, \theta, \mu) d\gamma_n(\theta, \mu|x^n) = \int_{\mathcal{M}} \left(\sup_{\theta \in \Theta_0(\mu)} r(d, \theta, \mu) \right) d\pi_n(\mu|x^n). \quad (18)$$

Combining (16) and (18), we see that the decision rules that are considered optimal under Definition 2 can be viewed as solutions to a conditional Γ -minimax problem.

As is well known, criterion (17) can lead to different optimal decision than the worst-case Bayes risk (or unconditional Γ -minimax) criterion for the DM's ex-ante problem faced before X^n is realized:

$$\inf_{\delta_n} \sup_{\gamma \in \Gamma} \int_{\Theta \times \mathcal{M}} \mathbb{E}_\mu [r(\delta_n(X^n), \theta, \mu)] d\gamma(\theta, \mu), \quad (19)$$

where the infimum is over all measurable $\delta_n : \mathcal{X}^n \rightarrow \mathcal{D}$; see [Giacomini, Kitagawa, and Read \(2021\)](#) and [Aradillas Fernández, Montiel Olea, Qiu, Stoye, and Tinda \(2024\)](#) for discussions. Thus, decision rules that are optimal under criterion (19) may not align with optimal actions in the conditional Γ -minimax problem (17). Further, optimal decisions under criterion (19) are, in general, intractable (see, e.g., [Giacomini, Kitagawa, and Read \(2021\)](#)).

Our proposal—namely, to use criterion (13) for the ex-ante problem—circumvents both of these problems. Optimal decisions under criterion (13) are tractable, align with optimal actions under the conditional Γ -minimax criterion (17), and are asymptotically efficient in a frequentist sense that we formalize in Section 4.

Finally, we note our approach departs from the textbook Wald approach in two regards. The Wald approach would seek a minimax decision rule that solves

$$\inf_{\delta_n} \sup_{\mu \in \mathcal{M}} \sup_{\theta \in \Theta_0(\mu)} \mathbb{E}_\mu [r(\delta_n(X^n), \theta, \mu)].$$

Comparing with criterion (13), we see that our approach moves the supremum over θ inside the expectation, potentially making the criterion more conservative. But as we discussed above, this conservativeness is constrained by the identified set $\Theta_0(\mu)$. Moreover, this conservativeness is offset, at least in part, by the fact that we average across a prior π for μ rather than taking a supremum over μ . These departures seem reasonable given the asymmetric nature of the identification problem. They also buy tractability, allowing the derivation of optimal decisions in a much broader class of problems than the standard Wald approach.

4 Optimality Theory for Parametric Models

We now present the main optimality results for parametric models. We first outline the assumptions in Section 4.1. Section 4.2 presents two main results. Theorem 1 establishes that Bayes decision rules $\delta_n^*(\cdot; \pi)$, defined in (15), are optimal. Theorem 2 shows that any decision whose asymptotic behavior is different from $\delta_n^*(\cdot; \pi)$ is sub-optimal. In particular, we prove in Section 4.3 that plug-in decisions δ_n^{plug} are sub-optimal when the maximum risk $R(d, \mu)$ is a non-smooth function of μ . As we documented in the empirical illustration in Section 2.4, this finding can have important implications for treatment assignment under partial identification. Finally, in Section 4.4 we provide a refinement of our analysis for settings in which the average maximum risk at the optimal decision not finite.

4.1 Assumptions

We first place some assumptions on the risk functions. Say $f : \mathcal{M} \rightarrow \mathbb{R}^k$ is *directionally differentiable* at μ_0 if there is a continuous map $\dot{f}_{\mu_0}[\cdot] : \mathbb{R}^K \rightarrow \mathbb{R}^k$ such that

$$\lim_{n \rightarrow \infty} \frac{f(\mu_0 + t_n h_n) - f(\mu_0)}{t_n} = \dot{f}_{\mu_0}[h]$$

for all sequences $\{t_n\} \subset \mathbb{R}_+$ and $\{h_n\} \subset \mathbb{R}^K$ with $t_n \downarrow 0$ and $h_n \rightarrow h \in \mathbb{R}^K$. If so, we refer to $\dot{f}_{\mu_0}[\cdot]$ as the directional derivative of f at μ_0 . Note that $\dot{f}_{\mu_0}[\cdot]$ is positively homogeneous of degree one but not necessarily linear. If $\dot{f}_{\mu_0}[\cdot]$ is linear, then f is (*fully*) *differentiable* at μ_0 .

Assumption 1 (i) $R(d, \cdot)$ is bounded and continuous on \mathcal{M} for each $d \in \mathcal{D}$;
(ii) For every $\mu_0 \in \mathcal{M}$ for which $|\underline{\mathcal{D}}_{\mu_0}| > 1$, $R(d, \cdot)$ is directionally differentiable at μ_0 for each $d \in \underline{\mathcal{D}}_{\mu_0}$.

We note that $R(d, \cdot)$ may be directionally but not fully differentiable in many empirically relevant cases of treatment assignment under partial identification (see Sections 2 and 6). We only require directional differentiability to hold when $|\underline{\mathcal{D}}_{\mu_0}| > 1$. In this case, optimal actions can differ for different μ close to μ_0 and the problem of distinguishing between optimal actions remains nontrivial as the sample size increases. In these scenarios, we use directional differentiability to derive asymptotic approximations for different decision rules.

We also place some regularity conditions on the statistical model $\mathcal{P} = \{P_\mu : \mu \in \mathcal{M}\}$. We present assumptions for a random sample to simplify exposition — so $X^n = (X_1, \dots, X_n)$ where each X_i is an independent draw from P_μ — though it is straightforward to extend our theory to weakly dependent data. We assume implicitly that each P_μ admits a density p_μ

with respect to a common dominating measure ν . Say that \mathcal{P} is *differentiable in quadratic mean* (DQM) at μ if there exists a vector of measurable functions $\dot{\ell}_\mu$ such that

$$\int \left(\sqrt{p_{\mu+h}} - \sqrt{p_\mu} - \frac{1}{2} h^T \dot{\ell}_\mu \sqrt{p_\mu} \right)^2 d\nu = o(\|h\|^2)$$

as $h \rightarrow 0$. Under DQM, the Fisher information matrix is $I_\mu := \int \dot{\ell}_\mu \dot{\ell}_\mu^T dP_\mu$. Let $D_{KL}(p_\mu \| p_{\mu'}) = \int p_\mu \log(p_\mu/p_{\mu'}) d\nu$ if $P_\mu(p_{\mu'}(X) = 0) = 0$ and $D_{KL}(p_\mu \| p_{\mu'}) = +\infty$ if $P_\mu(p_{\mu'}(X) = 0) > 0$. Say that \mathcal{P} is *locally quadratic* for D_{KL} if for all $\mu_0 \in \mathcal{M}$,

$$D_{KL}(p_\mu \| p_{\mu'}) \leq 2(\mu - \mu')^T I_{\mu_0} (\mu - \mu')$$

holds for all μ, μ' in a neighborhood of μ_0 . Following [Clarke and Barron \(1990\)](#), we say that \mathcal{P} is *sound* if weak convergence of P_μ to $P_{\mu'}$ is equivalent to convergence of μ to μ' .

- Assumption 2**
- (i) \mathcal{M} is an open subset of \mathbb{R}^K ;
 - (ii) \mathcal{P} is differentiable in quadratic mean at each $\mu_0 \in \mathcal{M}$;
 - (iii) I_{μ_0} is finite and nonsingular at each $\mu_0 \in \mathcal{M}$;
 - (iv) \mathcal{P} is locally quadratic for D_{KL} ;
 - (v) \mathcal{P} is sound.

Assumptions 2(i)-(iii) are similar to the conditions for parametric models in [Hirano and Porter \(2009\)](#). DQM and local quadraticity hold under standard smoothness and integrability conditions. These conditions rule out certain “irregular” models, such as those with parameter-dependent support and parameters on the boundary. Assumption 2(iv) holds under conditions similar to those used to establish asymptotic normality of maximum likelihood estimators. Soundness is a minimal identifiability condition. Following [Schwartz \(1965\)](#), posterior consistency is typically established in the weak topology. Soundness converts this to consistency for the parameter μ . This condition trivially holds: (i) in exponential families ([Clarke and Barron, 1990](#), p. 470); or (ii) if μ is identified (i.e., $P_\mu = P_{\mu'}$ if and only if $\mu = \mu'$), \mathcal{M} is precompact, and $\mu \mapsto P_\mu$ is weakly continuous on the closure of \mathcal{M} .⁹

Example: Treatment Assignment. Consider Assumption 1. In the empirical application, the lower bound for male youths behaved approximately like the maximum of two

⁹Alternatively, one could assume the existence of uniformly consistent tests of $H_0 : \mu = \mu_0$ against $H_1 : \|\mu - \mu_0\| \geq \epsilon$ ([Schwartz, 1965](#); [van der Vaart, 1998](#)). Soundness is a weak sufficient condition for the existence of such tests ([Clarke and Barron, 1990](#)). It also allows us to control the error rate of tests along $\{P_{n,h}\}$, for which the usual classical testing condition seemed inadequate.

independent Gaussian random variables (corresponding to estimates from US and Brazilian RCTs) while the upper bound behaved approximately like a third independent Gaussian random variable (corresponding to the estimate from a Colombian RCT). We can mimic this setting with the following stylized example. Let

$$\mu = (\mu_1, \mu_2, \mu_3), \quad b_L(\mu) = (\mu_1 \vee \mu_2), \quad b_U(\mu) = \mu_3. \quad (20)$$

As the bounds from the three studies are roughly the same magnitude, let $\mu_0 = (-c, -c, c)$ for $c > 0$. Then $R(0, \mu) = (\mu_3)_+$ and $R(1, \mu) = -(\mu_1 \vee \mu_2)_-$, and so $R(0, \mu_0) = R(1, \mu_0) = c$ and $\underline{\mathcal{D}}_{\mu_0} = \{0, 1\}$. We deduce that $R(d, \mu)$ is bounded (provided $\mathcal{M} \subset \mathbb{R}^3$ is bounded) and continuous in μ . Here $R(0, \mu)$ is smooth at μ_0 with $\dot{R}_{0, \mu_0}[h] = h_3$. However, $R(1, \mu)$ is only directionally differentiable at μ_0 with $\dot{R}_{1, \mu_0}[h] = -(h_1 \vee h_2)$. Overall, $R(d, \mu)$ satisfies Assumption 1.

Assumption 2 places restrictions on the family of probability distributions $\{P_\mu : \mu \in \mathcal{M}\}$. One formalization of the application described in Section 2.4 is to regard the published estimates, rather than the observations upon which these estimates are based, as data. So $X^n = \hat{\mu}$. Using the Normality assumption underlying our empirical illustration, P_μ is a multivariate Normal distribution with mean μ and a diagonal covariance matrix with elements s_k^2 . The Normal family clearly satisfies Assumption 2. In Section 5 we will relax the strict Normality assumption and cast the paper in a semiparametric framework. \square

4.2 Main Results

We first introduce some terminology. Let Π denote the class of all priors on \mathcal{M} with positive, continuous, and bounded Lebesgue density. We say that $\{\delta_n\}, \{\delta'_n\} \in \mathbb{D}$ are *asymptotically equivalent* if

$$\lim_{n \rightarrow \infty} P_{n,h}(\delta_n(X^n) = d) = \lim_{n \rightarrow \infty} P_{n,h}(\delta'_n(X^n) = d)$$

for all $d \in \mathcal{D}$, $h \in \mathbb{R}^K$, and $\mu_0 \in \mathcal{M}$. Proposition 2 in Appendix C shows that $\sqrt{n}(\bar{R}_n(d) - R(d, \mu_0))$ converges in distribution along $P_{n,h}$ to $\mathbb{E}^*[\dot{R}_{d, \mu_0}[Z^* + Z]|Z]$, where \dot{R}_{d, μ_0} is the directional derivative of $R(d, \mu)$ at μ_0 , $Z \sim N(h, I_{\mu_0}^{-1})$ represents the asymptotic behavior of an efficient estimator of μ , $Z^* \sim N(0, I_{\mu_0}^{-1})$ independently of Z represents asymptotic uncertainty about the local parameter h , and \mathbb{E}^* denotes expectation with respect to Z^* , which integrates over the uncertainty in the local parameter. We say there are *no first-order ties* if for each $\mu_0 \in \mathcal{M}$ the minimizer of $d \mapsto \mathbb{E}^*[\dot{R}_{d, \mu_0}[Z^* + z]]$ over $\underline{\mathcal{D}}_{\mu_0}$ is unique for almost every z . This condition trivially holds if $\underline{\mathcal{D}}_{\mu_0}$ is a singleton. To interpret the condition when $|\underline{\mathcal{D}}_{\mu_0}| > 1$, suppose that R is differentiable at μ_0 , which is analogous to the case of smooth

welfare contrasts studied in [Hirano and Porter \(2009\)](#). Then $\mathbb{E}^*[\dot{R}_{d,\mu_0}[Z^* + z]] = g_{d,\mu_0}^T z$ with $g_{d,\mu_0} = \frac{\partial R(d,\mu_0)}{\partial \mu}$. In that case, a sufficient condition for no first-order ties is that the derivatives are different: $g_{d,\mu_0} \neq g_{d',\mu_0}$ for $d, d' \in \underline{\mathcal{D}}_{\mu_0}$.

Our first main result establishes optimality of Bayes decisions $\{\delta_n^*(\cdot; \pi)\}$ for $\pi \in \Pi$.

Theorem 1 *Suppose that Assumptions 1 and 2 hold, $\mathcal{R}(\{\delta_n\}; \mu_0)$ is finite for at least one $\{\delta_n\} \in \mathbb{D}$, and there are no first-order ties. Then*

- (i) $\{\delta_n^*(\cdot; \pi)\}$ is optimal for any $\pi \in \Pi$;
- (ii) $\{\delta_n^*(\cdot; \pi)\}$ and $\{\delta_n^*(\cdot; \pi')\}$ are asymptotically equivalent for all $\pi, \pi' \in \Pi$;
- (iii) $\{\delta_n\} \in \mathbb{D}$ is optimal if it is asymptotically equivalent to $\{\delta_n^*(\cdot; \pi)\}$ for some $\pi \in \Pi$.

According to [Theorem 1](#), Bayes decisions with priors $\pi \in \Pi$ are asymptotically equivalent and optimal. All such Bayes decisions are asymptotically independent of π . Part (iii) implies that any decision that is asymptotically equivalent to a Bayes decision under a prior $\pi \in \Pi$ is optimal. Optimality of the bootstrap implementation discussed in [Section 3.4](#) follows from (iii) under suitable regularity conditions.

We now show that asymptotic equivalence to a Bayes decision is *necessary* for optimality. Say $\{\delta_n\}, \{\delta'_n\} \in \mathbb{D}$ fail to be asymptotically equivalent at μ_0 if

$$\lim_{n \rightarrow \infty} P_{n,h}(\delta_n(X^n) = d) \neq \lim_{n \rightarrow \infty} P_{n,h}(\delta'_n(X^n) = d)$$

for some $h \in \mathbb{R}^K$ and some $d \in \mathcal{D}$.

Theorem 2 *Suppose that Assumptions 1 and 2 hold, $\mathcal{R}(\{\delta_n\}; \mu_0)$ is finite for at least one $\{\delta_n\} \in \mathbb{D}$, and there are no first-order ties. Suppose that $\{\delta_n\} \in \mathbb{D}$ is not asymptotically equivalent to $\{\delta_n^*(\cdot; \pi)\}$ for some (equivalently, all) $\pi \in \Pi$. Then at any μ_0 at which asymptotic equivalence fails,*

$$\mathcal{R}(\{\delta_n\}; \mu_0) > \mathcal{R}(\{\delta_n^*(\cdot; \pi)\}; \mu_0) \quad \text{for all } \pi \in \Pi.$$

In the next subsection we discuss further implications of [Theorem 2](#) with respect to decision rules based on plugging-in efficient estimators of μ .

4.3 Plug-in Rules

A natural alternative to $\delta_n^*(X^n; \pi)$ is to plug in an efficient estimator $\hat{\mu} = \hat{\mu}(X^n)$ of μ into the oracle decision rule, yielding the “plug-in” rule $\delta_n^{plug}(X^n) = \delta^o(\hat{\mu})$. [Manski \(2021, 2023\)](#)

refers to this approach as “as-if” optimization: the estimator $\hat{\mu}$ is treated “as if” it is the true parameter.¹⁰ We show that asymptotic equivalence of the Bayes and plug-in rules can fail when the maximum risk does not depend smoothly on μ ; if so, Theorem 2 implies plug-in rules are sub-optimal. We illustrate this difference within the context of the application to treatment assignment from Section 2.

To understand when asymptotic equivalence holds, we turn to Lemma 1 in Appendix B, which characterizes the asymptotic behavior of sequences of decision rules. Lemma 1 implies

$$\lim_{n \rightarrow \infty} P_{n,h}(\delta_n^*(X^n; \pi) = d) = \begin{cases} \mathbb{P}_h\left(\operatorname{argmin}_{i \in \mathcal{D}_{\mu_0}} \mathbb{E}^* \left[\dot{R}_{i,\mu_0}[Z^* + Z] \middle| Z \right] = d \right), & d \in \mathcal{D}_{\mu_0}, \\ 0, & \text{otherwise,} \end{cases}$$

where \mathbb{P}_h is the probability measure of $Z \sim N(h, I_{\mu_0}^{-1})$. One can similarly show

$$\lim_{n \rightarrow \infty} P_{n,h}(\delta_n^{plug}(X^n) = d) = \begin{cases} \mathbb{P}_h\left(\operatorname{argmin}_{i \in \mathcal{D}_{\mu_0}} \dot{R}_{i,\mu_0}[Z] = d \right), & d \in \mathcal{D}_{\mu_0}, \\ 0, & \text{otherwise.} \end{cases} \quad (21)$$

Asymptotic equivalence of the Bayes and plug-in rules holds when the right-hand side probabilities in the above two displays are equal.

Suppose $R(d, \mu)$ is fully differentiable at μ_0 for all $d \in \mathcal{D}_{\mu_0}$. Then each $\dot{R}_{d,\mu_0}[\cdot]$ is linear, and

$$\mathbb{E}^* \left[\dot{R}_{d,\mu_0}[Z^* + Z] \middle| Z \right] = \dot{R}_{d,\mu_0}[\mathbb{E}^*[Z^*] + Z] \equiv \dot{R}_{d,\mu_0}[Z].$$

Hence, the plug-in and Bayes rules are asymptotically equivalent, and therefore optimal by Theorem 1. Formally:

Corollary 1 *Let the conditions of Theorem 1 hold, let the plug-in rule δ_n^{plug} satisfy (21), and let $R(d, \mu)$ be fully differentiable at μ_0 for all $d \in \mathcal{D}_{\mu_0}$ and all $\mu_0 \in \mathcal{M}$. Then $\{\delta_n^{plug}\}$ is optimal.*

Corollary 1 is consistent with Theorem 3.2 of Hirano and Porter (2009), which shows plug-in rules are optimal when welfare contrasts depend smoothly on a point-identified, regularly estimable parameter.

If $R(d, \mu)$ is only directionally differentiable at μ_0 , then $\dot{R}_{d,\mu_0}[\cdot]$ is not linear and the above reasoning no longer applies. In this case, asymptotic equivalence of the Bayes and plug-in rules cannot be guaranteed. If it fails, Theorem 2 implies the plug-in rule is sub-optimal.

¹⁰This approach also has connections to anticipated utility (Kreps, 1998; Cogley and Sargent, 2008).

Example: Treatment Assignment. We continue the calculations for the stylized example introduced in Section 4.1. By Jensen’s inequality,

$$\begin{aligned} \mathbb{P}_h\left(\operatorname{argmin}_{d \in \{0,1\}} \mathbb{E}^* \left[\dot{R}_{d,\mu_0}[Z^* + Z] \middle| Z \right] = 1\right) &= \mathbb{P}_h\left(\mathbb{E}^*[(Z_1^* + Z_1) \vee (Z_2^* + Z_2) | Z] + Z_3 \geq 0\right) \\ &> \mathbb{P}_h\left((Z_1 \vee Z_2) + Z_3 \geq 0\right) \\ &= \mathbb{P}_h\left(\operatorname{argmin}_{d \in \{0,1\}} \dot{R}_{d,\mu_0}[Z] = 1\right). \end{aligned}$$

Hence, asymptotic equivalence fails and the plug-in rule is sub-optimal: the Bayes rule recommends treatment more aggressively than the plug-in rule. This is reflected in the empirical application, where the optimal rule recommends treatment but the plug-in rule does not. \square

4.4 Finiteness of the Criterion

Criterion $\mathcal{R}(\{\delta_n\}; \mu_0)$ in (12) is formed by integrating $\mathcal{R}(\{\delta_n\}; \mu_0, h)$ with respect to Lebesgue measure on \mathbb{R}^K . This raises the possibility that $\mathcal{R}(\{\delta_n\}; \mu_0) = +\infty$ for all $\{\delta_n\} \in \mathbb{D}$, in which case it does not produce a meaningful ranking. We first note this is not possible if $K = 1$:

Proposition 1 *Suppose that $K = 1$ and Assumptions 1 and 2 hold. Fix any $\mu_0 \in \mathcal{M}$ for which there are no first-order ties. Then $\mathcal{R}(\{\delta_n^*(\cdot; \pi)\}; \mu_0) < \infty$ for all $\pi \in \Pi$.*

If $K > 1$, however, we may have $\mathcal{R}(\{\delta_n\}; \mu_0) = +\infty$ for all $\{\delta_n\} \in \mathbb{D}$. For instance, one could take the $K = 1$ case and augment the parameter space with redundant parameters. To handle cases where criterion (12) is infinite, we approximate the improper Lebesgue prior on h by a sequence of proper priors. We used Lebesgue measure on \mathbb{R}^K in criterion (12) because there is no a priori reason to view one local parameter as more likely than another. A similar effect is achieved for a $N(0, \sigma I)$ prior with large σ . However, for any finite σ the integrated risk under this proper prior is also finite because $R(d, \cdot)$ is bounded according to Assumption 1(i). Let π_σ denote the $N(0, \sigma I)$ density. Define

$$\mathcal{R}_\sigma(\{\delta_n\}; \mu_0) = \int \mathcal{R}(\{\delta_n\}; \mu_0, h) \pi_\sigma(h) dh. \quad (22)$$

Lemma 1 implies $\mathcal{R}_\sigma(\{\delta_n\}; \mu_0)$ is finite for all $\{\delta_n\} \in \mathbb{D}$. Analogously to Definition 2, we define the concept of σ -optimality:

Definition 3 *Say $\{\delta_n\} \in \mathbb{D}$ is σ -optimal if $\mathcal{R}_\sigma(\{\delta_n\}; \mu_0) = \inf_{\{\delta'_n\} \in \mathbb{D}} \mathcal{R}_\sigma(\{\delta'_n\}; \mu_0)$ for all $\mu_0 \in \mathcal{M}$.*

As we are using the large- σ limit of the $N(0, \sigma I)$ prior to approximate Lebesgue measure, we are really interested in the behavior of σ -optimal decisions as $\sigma \rightarrow \infty$. The following result shows Bayes and σ -optimal decisions lead to the same choices in large samples.

Theorem 3 *Suppose that Assumptions 1 and 2 hold and there are no first-order ties. Let $\{\delta_{n,\sigma}^*\} \in \mathbb{D}$ be σ -optimal at μ_0 . Then for all $d \in \mathcal{D}$ and $h \in \mathbb{R}^K$,*

$$\lim_{n \rightarrow \infty} P_{n,h}(\delta_n^*(X^n; \pi) = d) = \lim_{\sigma \rightarrow \infty} \lim_{n \rightarrow \infty} P_{n,h}(\delta_{n,\sigma}^*(X^n) = d).$$

The proof of Theorem 3 shows that Bayes and σ -optimal decisions are *identical* in large samples, for almost every realization of the data, provided σ is sufficiently large. We omit the details here to avoid introducing additional technicalities, and instead illustrate the idea within the context of our running example.¹¹

Example: Treatment Assignment. We continue with our discussion of the stylized example from (20) that mimics the empirical application. We have $\sqrt{n}(\hat{\mu} - \mu) \rightarrow_d Z \sim N(h, \Sigma)$ along $P_{n,h}$, where $\Sigma = \text{diag}(s_1, s_2, s_3)$, say, since the RCTs are independent. Lemma 1 in Appendix B implies that asymptotically, any Bayes decision behaves like

$$\delta_\infty^*(z) = \mathbb{I}[\mathbb{E}^*[(Z_1^* + z_1) \vee (Z_2^* + z_2)] + z_3 \geq 0],$$

for almost every realization z of Z , where \mathbb{E}^* denotes expectation taken with respect to $Z^* \sim N(0, \Sigma)$. The expectation can be computed in closed-form (see, e.g., [Nadarajah and Kotz \(2008\)](#)) to give

$$\delta_\infty^*(z) = \mathbb{I}[f(z; s_1, s_2) \geq 0],$$

$$\text{where } f(z; s_1, s_2) = z_1 \Phi\left(\frac{z_1 - z_2}{\sqrt{s_1 + s_2}}\right) + z_2 \Phi\left(\frac{z_2 - z_1}{\sqrt{s_1 + s_2}}\right) + \sqrt{s_1 + s_2} \phi\left(\frac{z_1 - z_2}{\sqrt{s_1 + s_2}}\right) + z_3,$$

and Φ and ϕ are the standard normal CDF and PDF, respectively. Similarly, Lemma 2 in Appendix B implies that asymptotically, the σ -optimal decision behaves like

$$\delta_\sigma^*(z) = \mathbb{I}[f_\sigma(z; s_1, s_2) \geq 0],$$

$$\text{where } f_\sigma(z; s_1, s_2) = f\left((I + \sigma^{-1}\Sigma)^{-1}z; \frac{\sigma s_1}{\sigma + s_1}, \frac{\sigma s_2}{\sigma + s_2}\right).$$

¹¹Here is some intuition: in many Bayesian settings posterior distributions of parameters are proper even under improper priors. But integrated risk is typically only finite under a proper prior. As long as the likelihood function asymptotically dominates the prior, posteriors, and decisions derived from them, under an improper prior are very close to those obtained from a prior with a large variance.

The difference between the two arises because in the σ -optimal decision the posterior mean is shrunk from z to $(I + \sigma^{-1}\Sigma)^{-1}z$ and the posterior variance is shrunk from Σ to $(\Sigma^{-1} + \sigma^{-1}I)^{-1}$. There are no first-order ties because the set of z values for which $f(z; s_1, s_2) = 0$ has measure zero. For any z outside this negligible set, we have $\text{sign}(f(z; s_1, s_2)) = \text{sign}(f_\sigma(z; s_1, s_2))$ for σ sufficiently large. Hence, for almost every realization z of Z , there is a minimal value of σ above which the Bayes and σ -optimal decisions are identical. \square

5 Optimality Theory for Semiparametric Models

This section extends our approach to semiparametric models, which is relevant in many empirical contexts. For example, the data may not follow a specific parametric model, or X^n may be a vector of summary statistics with unknown finite-sample distribution, as in the empirical application in Section 2. We present the model in Section 5.1 then generalize our optimality concept in Section 5.2. Section 5.3 describes a quasi-Bayesian implementation of optimal decisions. Section 5.4 presents the main optimality results. Theorem 4 shows that quasi-Bayes decision rules are optimal, while Theorem 5 shows that any decision whose asymptotic behavior differs from quasi-Bayes rules is sub-optimal.

5.1 Model

Let $X^n \sim P_{n,(\mu,\eta)}$ with $\mu \in \mathcal{M} \subseteq \mathbb{R}^K$ and $\eta \in \mathcal{H}$, an infinite-dimensional space. In a GMM model, μ is a finite-dimensional parameter vector, η is the marginal distribution of each observation $X_i \in \mathcal{X}$, and $\mathcal{H} = \{\eta \in \mathbb{M}(\mathcal{X}) : \int g(x, \mu) d\eta(x) = 0 \text{ for some } \mu \in \mathcal{M}\}$ for a vector of moment functions g , where $\mathbb{M}(\mathcal{X})$ is the set of all probability measures on \mathcal{X} . We again assume that, given μ , the structural parameter θ takes values in a set $\Theta_0(\mu)$. Therefore, the set of payoff distributions is indexed only by the parametric component μ and the nonparametric component η is a nuisance parameter. For instance, μ may be a vector of population moments used to construct bounds on θ . The nuisance parameter η represents other features of the distribution of X^n that are irrelevant for the DM's decision problem.

5.2 Optimality Criterion

Our optimality criterion (12) integrates the excess maximum risk of $\delta_n(X^n)$ relative to the oracle using Lebesgue measure on the local perturbations $h \in \mathbb{R}^K$ of μ_0 . This approach does not extend to perturbations of (μ_0, η_0) due to measure-theoretic complications in infinite-dimensional spaces. We therefore form our optimality criterion using local perturbations of μ_0 within a *least favorable submodel* in which X^n carries the least information about μ of all

parametric submodels. The problem of parameter estimation in the least favorable submodel is asymptotically equivalent to the problem of estimating μ in the full semiparametric model.

To simplify exposition, we present the following discussion and results within the context of a random sample — so $X^n = (X_1, \dots, X_n)$ where each X_i is an independent draw from $P_{(\mu, \eta)}$ — though it is straightforward to extend our theory to weakly dependent data. We say that $\mathcal{P} = \{P_{(\mu, \eta)} : \mu \in \mathcal{M}, \eta \in \mathcal{H}\}$ has a *least favorable submodel* at (μ, η) if there exists an open neighborhood $\mathcal{M}_{(\mu, \eta)}$ of μ and a map $t \mapsto \eta_t$ from $\mathcal{M}_{(\mu, \eta)}$ into \mathcal{H} such that $\{P_{\beta(t)} : t \in \mathcal{M}_{(\mu, \eta)}\}$ with $\beta(t) = (t, \eta_t)$ have densities $p_{\beta(t)}$ with respect to a common dominating measure ν that satisfy the DQM condition

$$\int \left(\sqrt{p_{\beta(\mu+h)}} - \sqrt{p_{\beta(\mu)}} - \frac{1}{2} h^T \dot{\ell}_{(\mu, \eta)} \sqrt{p_{\beta(\mu+h)}} \right)^2 d\nu = o(\|h\|^2)$$

as $h \rightarrow 0$, where $\dot{\ell}_{(\mu, \eta)} : \mathcal{X} \rightarrow \mathbb{R}^K$ is the efficient score for μ and $I_{(\mu, \eta)} := \int \dot{\ell}_{(\mu, \eta)} \dot{\ell}_{(\mu, \eta)}^T dP_{(\mu, \eta)}$ is the semiparametric information bound. In other words, there is a regular parametric submodel $\{P_{\beta(t)} : t \in \mathcal{M}_{(\mu, \eta)}\}$ whose information matrix at $t = \mu$ is the semiparametric information bound at (μ, η) . The path $t \mapsto \eta_t$ and dominating measure ν can depend on (μ, η) , but we suppress this to simplify notation. We note that our approach only requires the least favorable model to exist: the researcher doesn't need to derive it in order to implement optimal decisions. The least favorable model also needn't be unique, but any such model satisfying the above conditions will lead to the same optimality criterion.

Our optimality criterion is analogous to the parametric case. For each $(\mu_0, \eta_0) \in \mathcal{M} \times \mathcal{H}$, we reparametrize the least favorable submodel $\{P_{\beta(t)} : t \in \mathcal{M}_{(\mu_0, \eta_0)}\}$ using $t = \mu_0 + h/\sqrt{n}$ for $h \in \mathbb{R}^K$. Let $P_{n, h}$ denote the distribution of X^n under $P_{\beta(\mu_0 + h/\sqrt{n})}$ with $\beta(\mu_0 + h/\sqrt{n}) = (\mu_0 + h/\sqrt{n}, \eta_{\mu_0 + h/\sqrt{n}})$ and $\mathbb{E}_{n, h}$ denote expectation under $P_{n, h}$. Define

$$\begin{aligned} \mathcal{R}(\{\delta_n\}; (\mu_0, \eta_0), h) \\ = \limsup_{n \rightarrow \infty} \sqrt{n} \left(\mathbb{E}_{n, h} [R(\delta_n(X^n), \mu_0 + h/\sqrt{n})] - \min_{d \in \mathcal{D}} R(d, \mu_0 + h/\sqrt{n}) \right), \end{aligned} \quad (23)$$

and let

$$\mathcal{R}(\{\delta_n\}; (\mu_0, \eta_0)) = \int \mathcal{R}(\{\delta_n\}; (\mu_0, \eta_0), h) dh. \quad (24)$$

The class \mathbb{D} is defined analogously to the parametric case (cf. Definition 1). Recall that $\underline{\mathcal{D}}_{\mu_0} = \operatorname{argmin}_{d \in \mathcal{D}} R(d, \mu_0)$ denotes the set of optimal choices at μ_0 .

¹²Note that $\mathcal{R}(\{\delta_n\}; (\mu_0, \eta_0), h)$ and $\mathcal{R}(\{\delta_n\}; (\mu_0, \eta_0))$ remain well defined despite the fact that $\mu_0 + h/\sqrt{n}$ may be outside $\mathcal{M}_{(\mu_0, \eta_0)}$ for small n .

Definition 4 \mathbb{D} consists of all sequences $\{\delta_n\}$ such that, for all $(\mu_0, \eta_0) \in \mathcal{M} \times \mathcal{H}$,

- (i) $\lim_{n \rightarrow \infty} P_{n,h}(\delta_n(X^n) = d)$ exists for all $d \in \mathcal{D}$ and $h \in \mathbb{R}^K$;
- (ii) $\lim_{n \rightarrow \infty} \sqrt{n} P_{n,h}(\delta_n(X^n) \notin \underline{\mathcal{D}}_{\mu_0}) = 0$ for all $h \in \mathbb{R}^K$.

Finally, we say a sequence of decision rules $\{\delta_n\}$ is optimal if it minimizes the criterion $\mathcal{R}(\cdot; (\mu_0, \eta_0))$ in (24) over \mathbb{D} for all $(\mu_0, \eta_0) \in \mathcal{M} \times \mathcal{H}$ (cf. Definition 2):

Definition 5 Say $\{\delta_n\} \in \mathbb{D}$ is **optimal** if $\mathcal{R}(\{\delta_n\}; (\mu_0, \eta_0)) = \inf_{\{\delta'_n\} \in \mathbb{D}} \mathcal{R}(\{\delta'_n\}; (\mu_0, \eta_0))$ for all $(\mu_0, \eta_0) \in \mathcal{M} \times \mathcal{H}$.

If $\mathcal{R}(\{\delta_n\}; (\mu_0, \eta_0))$ is infinite for all $\{\delta_n\} \in \mathbb{D}$, then a similar approach to Section 4.4 can be followed whereby the Lebesgue prior on h is approximated by a sequence of proper priors.

5.3 Quasi-Bayesian Implementation

Optimal decisions are formed similarly to the parametric case, but we replace the posterior with a *quasi-posterior* formed from a limited-information criterion for μ . Following Doksum and Lo (1990), Kim (2002), and Müller (2013), consider a limited information $N(\hat{\mu}, (n\hat{I})^{-1})$ quasi-likelihood for μ , where $\hat{\mu}$ is an efficient estimator of μ and \hat{I}^{-1} is a consistent estimator of its asymptotic variance, namely $I_{(\mu,\eta)}^{-1}$. Combining the quasi-likelihood with a prior π on \mathcal{M} yields the *quasi-posterior*

$$\pi_n(\mu|X^n) \propto e^{-\frac{1}{2}(\mu-\hat{\mu})^T(n\hat{I})(\mu-\hat{\mu})} \pi(\mu). \quad (25)$$

The *quasi-posterior maximum risk* $\bar{R}_n(d)$ is calculated by averaging $R(d, \mu)$ across the quasi-posterior, as in (14). The *quasi-Bayes decision* $\delta_n^*(X^n; \pi)$ is chosen to minimize the quasi-posterior maximum risk $\bar{R}_n(d)$, as in (15).

Unlike the parametric case, here the optimal decision cannot be justified on the basis of robust Bayes analysis. A formal Bayesian approach would require specifying a prior on $\mathcal{M} \times \mathcal{H}$ then forming a marginal posterior for μ .¹³

Example: Treatment Assignment. In Section 4.1 we rationalized the empirical illustration in Section 2.4 through the assumption that the vector $\hat{\mu}$ is exactly Normally distributed. The semiparametric extension allows us to relax this assumption. Treating the estimates $\hat{\mu}$ as data, the semiparametric model allows for more general distributions of the form $\hat{\mu} \sim P_{n,(\mu,\eta)}$,

¹³See, for instance, the Bayesian exponentially tilted empirical likelihood approach of Schennach (2005) or the Bayesian GMM approaches of Shin (2015) and Walker (2025). Our approach is computationally simple and avoids the delicate issue of specifying priors in infinite-dimensional parameter spaces.

satisfying the moment restriction $\int(\hat{\mu} - \mu)dP_{n,(\mu,\eta)} = 0$. The density in (25) is identical to the one implied by the Normal model in Section 4.1, but the interpretation changes from an exact posterior to a quasi-posterior. \square

5.4 Optimality of Quasi-Bayes Decisions

We first state and discuss regularity conditions, then present our optimality results. Let $\hat{\lambda}_{\min}$ and $\hat{\lambda}_{\max}$ denote the smallest and largest eigenvalues of \hat{I} . Let $\xrightarrow{P_{n,(\mu_0,\eta_0)}}$ denote convergence in probability under $\{P_{n,(\mu_0,\eta_0)}\}$. Let $\xrightarrow{P_{n,h}}$ denote convergence in distribution under $\{P_{n,h}\}$.

Assumption 3 (i) \mathcal{M} is an open subset of \mathbb{R}^K ;

(ii) \mathcal{P} has a least favorable submodel at each $(\mu_0, \eta_0) \in \mathcal{M} \times \mathcal{H}$;

(iii) $I_{(\mu_0,\eta_0)}$ is finite and nonsingular at each $(\mu_0, \eta_0) \in \mathcal{M} \times \mathcal{H}$;

(iv) For each $(\mu_0, \eta_0) \in \mathcal{M} \times \mathcal{H}$ and $h \in \mathbb{R}^K$:

(a) $\sqrt{n}P_{n,h}(\|\hat{\mu} - \mu_0\| > \epsilon) \rightarrow 0$ for each $\epsilon > 0$;

(b) there exists $c \in (0, 1)$ such that $\sqrt{n}P_{n,h}(c \leq \hat{\lambda}_{\min}, \hat{\lambda}_{\max} \leq c^{-1}) \rightarrow 0$;

(v) For each $(\mu_0, \eta_0) \in \mathcal{M} \times \mathcal{H}$:

(a) $\sqrt{n}(\hat{\mu} - \mu_0) \xrightarrow{P_{n,h}} Z$ with $Z \sim N(h, I_{(\mu_0,\eta_0)}^{-1})$ for all $h \in \mathbb{R}^K$;

(b) $\hat{I} \xrightarrow{P_{n,(\mu_0,\eta_0)}} I_{(\mu_0,\eta_0)}$.

Assumptions 3(i)-(iii) are analogous to Assumptions 2(i)-(iii). We do not require versions of Assumptions 2(iv) and 2(v) because the quasi-likelihood here has a particular quadratic structure. But we need to ensure that the estimators $\hat{\mu}$ and \hat{I} used in the quasi-likelihood are sufficiently well behaved, which is the role of Assumptions 3(iv) and 3(v). These may be verified for a wide variety of estimators $\hat{\mu}$ and \hat{I} under suitable regularity conditions that cover the RCT estimators $\hat{\mu}$ of the empirical illustration in Section 2.4.

Example: Treatment Assignment. Each subpopulation corresponds to an independent randomized experiment. In each subpopulation k we observe a random sample of (X_i, Y_i) of size n drawn from a distribution η_k ,¹⁴ where X_i is a binary treatment indicator and $Y_i = X_i Y_{i1} + (1 - X_i) Y_{i0}$ with Y_{i1} and Y_{i0} denoting the treated and untreated outcomes for individual i . Recall $\mu = (\mu_k)_{k=1}^K$ and $\hat{\mu} = (\hat{\mu}_k)_{k=1}^K$, where we estimate the subpopulation- k ATE μ_k by, for example, using the slope coefficient $\hat{\mu}_k$ in an OLS regression of Y_i on X_i . Let s_k denote the heteroskedasticity-robust standard error of $\hat{\mu}_k$. Then \hat{I} is given by $(n\hat{I})^{-1} = \text{diag}(s_1^2, \dots, s_K^2)$. The nuisance parameter is the joint distribution $\eta = \times_{k=1}^K \eta_k$

¹⁴It is straightforward to extend the analysis to have different sample sizes n_k in each subpopulation, provided the n_k all grow at the same rate.

across the independent subpopulations. We have $P_{(\mu,\eta)} \equiv P_\eta = \eta$. We similarly drop dependence of μ in quantities such as the score and information matrix that follow.

Each $\hat{\mu}_k$ is semiparametrically efficient (Hahn, 1998) with influence function

$$\psi_k(x, y) = \frac{x}{p_k}(y - \bar{\mu}_{1,k}) - \frac{1-x}{1-p_k}(y - \bar{\mu}_{0,k}),$$

where p_k , $\bar{\mu}_{1,k}$, and $\bar{\mu}_{0,k}$ represent the means of X_i , Y_{i1} , and Y_{i0} in subpopulation k and are available from the moments of X , XY and $(1-X)Y$ under η_k . We construct a least favorable submodel as follows. For each subpopulation k , fix any $\eta_{0,k}$ with $0 < p_k < 1$ and $0 < \int \psi_k^2 d\eta_{0,k} < \infty$. Let

$$\dot{\ell}_{k,\eta_0}(x, y) = \frac{\psi_k(x, y)}{\int \psi_k^2 d\eta_{0,k}}.$$

Without confusion we also let $\eta_{0,k}$ denote the density of $\eta_{0,k}$ with respect to some dominating measure ν . We define a smooth parametric model $\eta_{t_k,k}$ passing through $\eta_{0,k}$ at $t_k = \mu_{0,k}$ with score $\dot{\ell}_{k,\eta_0}$ by

$$\eta_{t_k,k}(x, y) = \eta_{0,k}(x, y) \frac{v((t_k - \mu_{0,k})\dot{\ell}_{k,(\mu_0,\eta_0)}(x, y))}{\int v((t_k - \mu_{0,k})\dot{\ell}_{k,(\mu_0,\eta_0)}(x, y))d\eta_{0,k}(x, y)},$$

where $v(u) = 2(1 + e^{-2u})^{-1}$ satisfies $v(0) = v'(0) = 1$ (van der Vaart, 1998, Example 25.16). Letting $\eta_t = \times_{k=1}^K \eta_{t_k,k}$ for $t = (t_k)_{k=1}^K$, we have a smooth parametric family passing through η_0 at $t = \mu_0$ whose score $\dot{\ell}_{\eta_0} = (\dot{\ell}_{k,\eta_0})_{k=1}^K$ satisfies $\int \dot{\ell}_{\eta_0} \dot{\ell}_{\eta_0}^T d\eta_0 = \text{diag}(\sigma_1^{-2}, \dots, \sigma_K^{-2}) \equiv I_{\eta_0}^{-1}$, where $\sigma_k^2 = \int \psi_k^2 d\eta_{0,k}$ is the (efficient) asymptotic variance of $\hat{\mu}_k$.

By Le Cam's third lemma (van der Vaart, 1998, Example 6.7), for any $h \in \mathbb{R}^K$ we may deduce that $\sqrt{n}(\hat{\mu} - \mu_0)$ converges in distribution to a $N(h, I_{\eta_0}^{-1})$ random vector under $\{P_{n,h}\}$, where $P_{n,h}$ is the product measure formed by drawing n copies of (X, Y) under $\eta_{\mu_{0,k} + h_k/\sqrt{n},k}$ for each subpopulation. This verifies Assumption 3(v)a. Assumption 3(v)b holds by standard consistency arguments for heteroskedasticity-robust standard errors. It is also straightforward to verify, for instance by using suitable concentration inequalities, that Assumption 3(iv) holds. \square

We now present analogues of Theorems 1 and 2 for the semiparametric case. Recall that Π denotes the class of all priors on \mathcal{M} with positive, continuous, and bounded Lebesgue density. Similar to the parametric case, we say that $\{\delta_n\}, \{\delta'_n\} \in \mathbb{D}$ are *asymptotically equivalent* if

$$\lim_{n \rightarrow \infty} P_{n,h}(\delta_n(X^n) = d) = \lim_{n \rightarrow \infty} P_{n,h}(\delta'_n(X^n) = d)$$

for all $d \in \mathcal{D}$, $h \in \mathbb{R}^K$, and $(\mu_0, \eta_0) \in \mathcal{M} \times \mathcal{H}$.

Theorem 4 *Suppose that Assumptions 1 and 3 hold, $\mathcal{R}(\{\delta_n\}; (\mu_0, \eta_0))$ is finite for at least one $\{\delta_n\} \in \mathbb{D}$, and there are no first-order ties. Then*

- (i) $\{\delta_n^*(\cdot; \pi)\}$ is optimal for any $\pi \in \Pi$;
- (ii) $\{\delta_n^*(\cdot; \pi)\}$ and $\{\delta_n^*(\cdot; \pi')\}$ are asymptotically equivalent for all $\pi, \pi' \in \Pi$;
- (iii) $\{\delta_n\} \in \mathbb{D}$ is optimal if it is asymptotically equivalent to $\{\delta_n^*(\cdot; \pi)\}$ for some $\pi \in \Pi$.

It follows from Theorem 4 that quasi-Bayes decisions with priors $\pi \in \Pi$ are asymptotically equivalent and optimal. Moreover, any decision that is asymptotically equivalent to a quasi-Bayes decision under a prior $\pi \in \Pi$ is optimal. In particular, bootstrap rules will be asymptotically equivalent to quasi-Bayes rules (and hence optimal) provided the bootstrap distribution for an efficient estimator $\hat{\mu}$ of μ and the quasi-posterior are sufficiently close.

Similar to the parametric case, we say $\{\delta_n\}, \{\delta'_n\} \in \mathbb{D}$ fail to be asymptotically equivalent at (μ_0, η_0) if

$$\lim_{n \rightarrow \infty} P_{n,h}(\delta_n(X^n) = d) \neq \lim_{n \rightarrow \infty} P_{n,h}(\delta'_n(X^n) = d)$$

for some $h \in \mathbb{R}^K$ and some $d \in \mathcal{D}$.

Theorem 5 *Suppose that Assumptions 1 and 3 hold, $\mathcal{R}(\{\delta_n\}; (\mu_0, \eta_0))$ is finite for at least one $\{\delta_n\} \in \mathbb{D}$, and there are no first-order ties. Suppose that $\{\delta_n\} \in \mathbb{D}$ is not asymptotically equivalent to $\{\delta_n^*(\cdot; \pi)\}$ for some (equivalently, all) $\pi \in \Pi$. Then at any (μ_0, η_0) at which asymptotic equivalence fails,*

$$\mathcal{R}(\{\delta_n\}; (\mu_0, \eta_0)) > \mathcal{R}(\{\delta_n^*(\cdot; \pi)\}; (\mu_0, \eta_0)) \quad \text{for all } \pi \in \Pi.$$

As in the parametric case, asymptotic equivalence to $\{\delta_n^*(\cdot; \pi)\}$ for some $\pi \in \Pi$ is necessary for optimality whenever there are no first-order ties. Thus, as discussed in Section 4.3, plug-in rules may fail to be optimal in settings where the maximum risks $R(d, \mu)$ is only directionally differentiable at μ_0 . This includes the empirical application from Section 2 and the further examples to treatment assignment and optimal pricing that we discuss in the next two sections.

6 Treatment Rules: Further Applications

This section expands on our running example of treatment assignment under partial identification. We first review several empirically relevant approaches for constructing bounds on the ATE in Section 6.1. Each of these constructions leads to bounds $b_L(\mu)$ and $b_U(\mu)$ that will in general be only directionally differentiable in a vector of reduced-form parameters μ .

We then discuss implementation of the optimal decision rules in these settings in Section 6.2. As far as we are aware, ours is the first work to propose optimal decisions for these realistic empirical settings.

6.1 Examples of Treatment Effect Bounds

Intersection Bounds. This setting generalizes the empirical application from Section 2. Suppose X^n consists of data from K observational studies. In each study k we can consistently estimate lower and upper bounds $b_{L,k}(\mu_k)$ and $b_{U,k}(\mu_k)$ on the ATE as a function of population moments μ_k . We then obtain the intersection bounds

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

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

Bounds via IV-like Estimands. Mogstad, Santos, and Torgovitsky (2018) present an approach for bounding the ATE and other causal effects using IV-like estimands from observational studies. Suppose treatment is determined by $D = \mathbb{I}[U \leq v(Z)]$ 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, Santos, and Torgovitsky (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, consider 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 this case, the bounds of [Mogstad, Santos, and Torgovitsky \(2018\)](#) are

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

where \mathcal{S} is a class of functions and $\mu = (\text{Cov}(Y, Z_0), \text{Cov}(D, Z_0), \mathbb{E}[Z_0], p)$, which is finite-dimensional if Z_0 has finite support (e.g. binary Z_0). They show that $b_L(\mu)$ and $b_U(\mu)$ may be expressed as the optimal values of linear programs parameterized by μ . 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 only be 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, α_i is a latent individual effect, and ε_{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 (α, ε) . When outcomes and covariates are discrete, parametric restrictions on the distribution of ε and functional form restrictions on g are generally insufficient to point identify the ATE without parametric restrictions on the distribution of α . 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, Fernández-Val, Hahn, and Newey \(2013\)](#) and [Torgovitsky \(2019b\)](#) derive 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 μ . As in the previous example, $b_L(\mu)$ and $b_U(\mu)$ may therefore only be directionally differentiable in μ .

6.2 Implementation

The first two examples are semiparametric and we can take X^n to be a vector of estimators $\hat{\mu}$ of μ . Assuming $\hat{\mu}$ is asymptotically efficient and the DM has available a consistent estimator \hat{I}^{-1} of the asymptotic variance of $\hat{\mu}$, then the optimal decision can be implemented based on a $N(\hat{\mu}, (n\hat{I})^{-1})$ quasi-posterior as in (25).

In the third example with discrete outcomes and covariates, the distribution $P_{n,\mu}$ of the data $((Y_{it}, X_{it})_{t=1}^T)_{i=1}^n$ can be identified with a multinomial distribution parameterized by μ . In this case our parametric theory applies and the optimal decision can easily be implemented using either the bootstrap or Bayesian methods.

7 Optimal Pricing with Unobserved Heterogeneity

Our methods can also be applied to make optimal pricing decisions in models with rich unobserved heterogeneity using revealed-preference demand theory. Section 7.1 provides the model and empirical setting. Section 7.2 presents techniques for computing sharp bounds on functionals of counterfactual demand using linear programming. Finally, we discuss how to implement our methods in Section 7.3.

7.1 Model

The DM observes repeated cross sections of individual demands $X^n = (X_{b,1}, \dots, X_{b,n})_{b=1}^B$, where each $X_{b,i} \in \mathbb{R}^m$ is the demand of individual i for m goods under prices q_b :

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

where $\mathcal{B}_b = \{x \in \mathbb{R}^m : x'q_b = 1\}$ is the budget set (expenditure is normalized to one). Individuals are heterogeneous in their utility functions u_i . We assume the demand system is rationalized by a random utility model with a probability distribution F over utility functions u . The demand under q_b of a randomly selected individual may therefore be interpreted as stochastic. The mass $p_b(s)$ of individuals whose demand is in a set $s \subset \mathcal{B}_b$ at price q_b is

$$p_b(s) = \int \mathbb{I}[\arg \max_{x \in \mathcal{B}_b} u(x) \in s] dF(u), \quad s \in \mathcal{B}_b, \quad b = 1, \dots, B \quad (26)$$

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

The DM wishes to choose a new price vector q_d for $d \in \mathcal{D} = \mathcal{O} \cup \mathcal{C}$. The price vectors $\{q_d : d \in \mathcal{O}\}$ are a subset of the observed prices q_1, \dots, q_B , while each $\{q_d : d \in \mathcal{C}\}$ is a set




Figure 2: Choices and Budget Lines

of counterfactual price vectors. In principle, the set \mathcal{C} of new price vectors could be large, representing prices rounded to nearest currency units or tax rates rounded to the nearest percentage. Let $w_d(X_d)$ represent a functional of demand (e.g., revenue or welfare) under prices q_d . The DM’s goal is to choose $d \in \mathcal{D}$ to maximize the average of $w_d(X_d)$. For $d \in \mathcal{O}$ the average demand $\mathbb{E}[w_d(X_d)]$ is identified from observed choice behavior under q_d . However, the observed choice behavior is only sufficient to bound, but not point-identify, $\mathbb{E}[w_d(X_d)]$ for counterfactual prices q_d , $d \in \mathcal{C}$.

7.2 Bounds on Functionals of Counterfactual Demand

Kitamura and Stoye (2019) present a general approach using linear programming to bound functionals of counterfactual demand. We introduce their approach with an example. Consider Figure 2. There are two goods, and the DM has observed the demand for two price vectors q_1 and q_2 , which generate the budget sets \mathcal{B}_1 and \mathcal{B}_2 . The counterfactual price vector is q_0 which generates the counterfactual budget set \mathcal{B}_0 .

The budget lines in Figure 2 are divided 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 between \mathcal{B}_2 and \mathcal{B}_0 , and so forth.¹⁵ 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_{32}) . Each combination corresponds to a consumer type. By revealed preference we know a consumer will never choose (s_{21}, s_{12})

¹⁵Like KS18, 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.

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 $s_{11}, s_{21}, s_{31}, s_{12}, s_{22}, s_{32}$ and the columns of A correspond 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. KS18 showed that the demand system p is rationalizable if and only if

$$p = Af,$$

for some $f \in \Delta^6$, the unit simplex in \mathbb{R}^7 , representing the probabilities of the 7 rational types. These probabilities constrain the distribution F of utilities.

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

$$\begin{bmatrix} p \\ p^* \end{bmatrix} = A^* f^*, \quad (27)$$

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} , $f^* \in \Delta^{15}$ collects the probabilities of observing each rational type, and

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

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

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

where A_{obs}^* collects the first 6 rows of A^* (corresponding to the observed budget sets) while A_{unobs}^* collects the rows of A^* corresponding to the patches s_{10}, s_{20} , and s_{30} .

Following [Kitamura and Stoye \(2019\)](#), we may deduce sharp bounds on $\mathbb{E}[w_0(X_0)]$ as follows. Let $(\underline{w}_1, \underline{w}_2, \underline{w}_3)$ and $(\bar{w}_1, \bar{w}_2, \bar{w}_3)$ denote row vectors which collect the smallest and largest values of $w_0(x)$ for x in s_{10}, s_{20} and s_{30} . Then the bounds on $\mathbb{E}[w_0(X_0)]$ are

$$\begin{aligned} w_L(p) &= \min \left\{ (\underline{w}_1, \underline{w}_2, \underline{w}_3)^T A_{unobs}^* f^* : p = A_{obs}^* f^*, f^* \in \Delta^{15} \right\}, \\ w_U(p) &= \max \left\{ (\bar{w}_1, \bar{w}_2, \bar{w}_3)^T A_{unobs}^* f^* : p = A_{obs}^* f^*, f^* \in \Delta^{15} \right\}. \end{aligned}$$

More generally, the preceding argument applies with a collection of B observed budget sets and multiple goods. In that case, representation (27) holds 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 \mathcal{B}_d generated by q_d for $d \in \mathcal{C}$. Suppose there are J patches across \mathcal{B}_d and T rational types. Then letting row vectors \underline{w}_d and \bar{w}_d collect the smallest and largest values of $w_d(x)$ for x in each of the J patches, sharp bounds on $\mathbb{E}[w_d(X_d)]$ are

$$\begin{aligned} w_{d,L}(p) &= \min \left\{ \underline{w}_d^T A_{unobs}^* f^* : p = A_{obs}^* f^*, f^* \in \Delta^{T-1} \right\}, \\ w_{d,U}(p) &= \max \left\{ \bar{w}_d^T A_{unobs}^* f^* : p = A_{obs}^* f^*, f^* \in \Delta^{T-1} \right\}, \end{aligned} \tag{28}$$

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

7.3 Implementation

As in [Bergemann and Schlag \(2011\)](#), we assume the DM has a preference for robustness. That is, the DM wishes to choose d to minimize maximum risk, where the maximum is taking over the identified set of counterfactual demand responses.

The DM's problem maps into our framework as follows. For each $d \in \mathcal{O}$, the value $\mathbb{E}[w_d(X_d)]$ is identified from observed choice behavior under q_d . The remaining reduced-form parameters are the patch probabilities p . Thus, $\mu = (p, (\mathbb{E}[w_d(X_d)])_{d \in \mathcal{O}})$. Here $\theta = F$ and $\Theta_0(\mu)$ is all F that rationalize the patch probabilities p consistent with revealed preference. The vector $Y = (w_d(X_d))_{d \in \mathcal{C}}$ collects functionals of demand under the counterfactual prices. Each $\theta \equiv F \in \Theta_0(\mu)$ induces a distribution G_θ for Y . Although we suppressed it in the previous subsections, we now write $\mathbb{E}_\theta[w_d(X_d)]$ for $d \in \mathcal{C}$ to denote that the average counterfactual

demand functional depends on the structural parameter θ .

For $d \in \mathcal{O}$, the DM incurs risk

$$r(d, \theta, \mu) = -\mathbb{E}[w_d(X_d)],$$

which is point identified from $\mu \equiv (p, (\mathbb{E}[w_d(X_d)])_{d \in \mathcal{O}})$. For $d \in \mathcal{C}$, the DM incurs risk

$$r(d, \theta, \mu) = -\mathbb{E}_\theta[w_d(X_d)],$$

which is set-identified and may be bounded as described in the previous subsection. The maximum risk is

$$R(d, \mu) = \begin{cases} -\mathbb{E}[w_d(X_d)] & d \in \mathcal{O}, \\ -w_{d,L}(p) & d \in \mathcal{C}, \end{cases}$$

where $w_{d,L}(p)$ was defined in (28).

This is a semiparametric model and our implementation follows the steps described in Section 5.3. Partition $\mu = (\mu_1, \dots, \mu_B)$ where μ_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 μ_b by just-identified GMM. As we observe repeated cross sections under B different price vectors, we can simply take $\eta = \times_{b=1}^B \eta_b$, where η_b is the distribution of $X_{b,i}$. We can then form a quasi-posterior $\pi_n(\mu|X^n)$ based on a limited-information Gaussian quasi-likelihood as in (25). For $d \in \mathcal{O}$, the expected value $\mathbb{E}[w_d(X_d)]$ is an element of μ and so $\bar{R}_n(d)$ is simply its quasi-posterior mean:

$$\bar{R}_n(d) = - \int \mathbb{E}[w_d(X_d)] d\pi_n(\mu).$$

For $d \in \mathcal{C}$, the posterior maximum risk is

$$\bar{R}_n(d) = - \int w_{d,L}(p) d\pi_n(\mu).$$

This may be computed by sampling p from the quasi-posterior, solving the linear program (28) defining $w_{d,L}(p)$ for each draw of p , then taking the average across draws.

The optimal decision minimizes $\bar{R}_n(d)$ for $d \in \mathcal{D} = \mathcal{O} \cup \mathcal{C}$. As the value of a linear program is typically directionally differentiable, the asymptotic distribution of the posterior mean $\bar{R}_n(d)$ for $d \in \mathcal{C}$ may be different from that of the corresponding plug-in values $-w_{d,L}(\hat{p})$. In this case, the optimal decision δ_n^* may outperform the plug-in rule.

8 Conclusion

We derived optimal statistical decision rules for discrete choice problems when payoffs depend on a set-identified parameter θ and the decision maker can use a point-identified parameter μ to deduce restrictions on θ . Our notion of optimality combines a minimax approach to handle the ambiguity from partial identification of θ given μ with average risk minimization for μ . In many empirically relevant applications, the maximum risk depends non-smoothly on μ , making plug-in rules sub-optimal. 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 continuous decisions fall outside the scope of our theory and analysis, it would be interesting to study asymptotic efficiency in the continuous case as well.

Data Availability Statement

The data and code underlying this article is available on Zenodo at <https://dx.doi.org/10.5281/zenodo.18157514>.

References

- ARADILLAS FERNÁNDEZ, A., J. L. MONTIEL OLEA, C. QIU, J. STOYE, AND S. TINDA (2024): “Robust Bayes Treatment Choice with Partial Identification,” *arXiv preprint arXiv:2408.11621*.
- BERGEMANN, D., AND K. SCHLAG (2011): “Robust monopoly pricing,” *Journal of Economic Theory*, 146(6), 2527–2543.
- BETRO, B., AND F. RUGGERI (1992): “Conditional Γ -Minimax Actions Under Convex Losses,” *Communications in Statistics: Theory and Methods*, 21(4), 1051–1066.
- BLUNDELL, R., M. BROWNING, AND I. CRAWFORD (2007): “Improving revealed preference bounds on demand responses,” *International Economic Review*, 48(4), 1227–1244.
- (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 coun-

- terfactuals with multidimensional unobserved heterogeneity,” Discussion paper, cemmap working paper.
- 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. (2000): “Econometric Applications of Maxmin Expected Utility,” *Journal of Applied Econometrics*, 15(6), 625–644.
- (2011): “Bayesian Aspects of Treatment Choice,” in *The Oxford Handbook of Bayesian Econometrics*. Oxford University Press.
- CHERNOZHUKOV, V., I. FERNÁNDEZ-VAL, J. HAHN, AND W. NEWEY (2013): “Average and quantile effects in nonseparable panel models,” *Econometrica*, 81(2), 535–580.
- CHRISTENSEN, T., H. R. MOON, AND F. SCHORFHEIDE (2020): “Robust Forecasting,” *arXiv preprint arXiv:2011.03153*.
- CLARKE, B., AND A. BARRON (1990): “Information-theoretic asymptotics of Bayes methods,” *IEEE Transactions on Information Theory*, 36(3), 453–471.
- 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.
- DASGUPTA, A., AND W. J. STUDDEN (1989): “Frequentist Behavior of Robust Bayes Estimates of Normal Means,” *Statistics & Decisions*, 7, 333–361.
- 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.
- DÜMBGEN, L. (1993): “On nondifferentiable functions and the bootstrap,” *Probability Theory and Related Fields*, 95, 125–140.
- FANG, Z., AND A. SANTOS (2019): “Inference on Directionally Differentiable Functions,” *The Review of Economic Studies*, 86(1), 377–412.
- GHOSAL, S., AND A. VAN DER VAART (2017): *Fundamentals of Nonparametric Bayesian Inference*. Cambridge University Press.
- 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.
- HAHN, J. (1998): “On the Role of the Propensity Score in Efficient Semiparametric Estimation of Average Treatment Effects,” *Econometrica*, 66(2), 315–331.
- 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.
- KITAGAWA, T., J. L. MONTIEL OLEA, J. PAYNE, AND A. VELEZ (2020): “Posterior distribution of nondifferentiable functions,” *Journal of Econometrics*, 217(1), 161–175.
- 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 *Frontiers of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures, 1983-1997*, ed. by D. P. Jacobs, E. Kalai, M. I. Kamien, and N. L. Schwartz, Econometric Society Monographs, pp. 242–274. Cambridge University Press.
- 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.
- (2004): “Statistical Treatment Rules for Heterogeneous Populations,” *Econometrica*, 72(4), 1221–1246.
- (2007a): “Minimax-regret treatment choice with missing outcome data,” *Journal of Econometrics*, 139(1), 105–115, Endogeneity, instruments and identification.
- (2007b): “Partial Identification of Counterfactual Choice Probabilities,” *International Economic Review*, 48(4), 1393–1410.
- (2014): “Identification of income-leisure preferences and evaluation of income tax policy,” *Quantitative Economics*, 5(1), 145–174.

- (2020): “Toward Credible Patient-centered Meta-analysis,” *Epidemiology*, 31(3), 345–352.
- MANSKI, C. F. (2021): “Econometrics for Decision Making: Building Foundations Sketched by Haavelmo and Wald,” *Econometrica*, 89(6), 2827–2853.
- (2023): “Probabilistic Prediction for Binary Treatment Choice: with focus on personalized medicine,” *Journal of Econometrics*, 234, 647–663.
- 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.
- MOON, H. R., AND F. SCHORFHEIDE (2012): “Bayesian and frequentist inference in partially identified models,” *Econometrica*, 80(2), 755–782.
- MÜLLER, U. K. (2013): “Risk of Bayesian Inference in Misspecified Models, and the Sandwich Covariance Matrix,” *Econometrica*, 81(5), 1805–1849.
- NADARAJAH, S., AND S. KOTZ (2008): “Exact Distribution of the Max/Min of Two Gaussian Random Variables,” *IEEE Transactions on Very Large Scale Integration Systems*, 6(2), 210–212.
- 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.
- SCHWARTZ, L. (1965): “On Bayes procedures,” *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 4(1), 10–26.
- SHAPIRO, A. (1990): “On concepts of directional differentiability,” *Journal of Optimization Theory and Applications*, 66(3), 477–487.
- 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.
- (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.
- (1998): *Asymptotic Statistics*. Cambridge University Press, Cambridge.
- WALD, A. (1950): *Statistical Decision Functions*. John Wiley, New York.
- WALKER, C. (2025): “Semiparametric Bayesian Inference for a Conditional Moment Equality Model,” *Manuscript, Duke University*.
- 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

A Preliminaries

We first introduce relevant concepts on limit experiments and asymptotic representations. We refer the reader to [van der Vaart \(1991, 1998\)](#) for comprehensive treatments.

An *experiment* is a measurable space equipped with a family of probability measures. Let $\mathcal{E}_n = (\mathcal{X}^n, \mathcal{B}_n, P_{n,h} : h \in \mathbb{R}^K)$ where \mathcal{B}_n denotes the Borel σ -algebra on \mathcal{X}^n . Each \mathcal{E}_n corresponds to observing a random sample X_1, \dots, X_n from $P_{\mu_0+h/\sqrt{n}}$ with parameter $h \in \mathbb{R}^K$. Also let $\mathcal{E} = (\mathbb{R}^K, \mathcal{B}_K, N(h, I_{\mu_0}^{-1}) : h \in \mathbb{R}^K)$ where \mathcal{B}_K denotes the Borel σ -algebra on \mathbb{R}^K . This experiment corresponds to observing a single $Z \sim N(h, I_{\mu_0}^{-1})$ with unknown h . We suppress dependence of \mathcal{E}_n , \mathcal{E} , $P_{n,h}$, and N_h on μ_0 to simplify notation.

The sequence of experiments $\{\mathcal{E}_n\}$ is said to converge to the *limit experiment* \mathcal{E} if for all finite subsets H of \mathbb{R}^K and every $h_0 \in \mathbb{R}^K$, we have

$$\left(\frac{dP_{n,h}}{dP_{n,h_0}}(X^n) \right)_{h \in H} \overset{P_{n,h_0}}{\rightsquigarrow} \left(\frac{dN(h, I_{\mu_0}^{-1})}{dN(h_0, I_{\mu_0}^{-1})}(Z) \right)_{h \in H},$$

where $\overset{P_{n,h_0}}{\rightsquigarrow}$ denotes convergence in distribution along $\{P_{n,h}\}$. This convergence holds under Assumption 2(i)-(iii) ([van der Vaart, 1998](#), Example 7.15 and Corollary 9.5).

Because \mathcal{E} is dominated,¹⁶ we can invoke an asymptotic representation theorem of [van der Vaart \(1991\)](#). For each $\{\delta_n\} \in \mathbb{D}$ there exists a *matching decision* $\delta_\infty(Z, U)$ given by a measurable map $\delta_\infty : \mathbb{R}^K \times [0, 1] \rightarrow \mathcal{D}$, the observation $Z \sim N(h, I_{\mu_0}^{-1})$, and an independent randomization $U \sim \text{Uniform}[0, 1]$, with the property that for each $h \in \mathbb{R}^K$,

$$\lim_{n \rightarrow \infty} P_{n,h}(\delta_n(X^n) = d) = \mathbb{P}_h(\delta_\infty(Z, U) = d), \quad d \in \mathcal{D},$$

where \mathbb{P}_h is the product of the $N(h, I_{\mu_0}^{-1})$ and $U[0, 1]$ distributions. Although $\delta_\infty(Z, U)$ depends on the randomization U in general, under our regularity conditions this randomization is unnecessary and any optimal decision depends only on Z .

¹⁶In particular, each $N(h, I_{\mu_0}^{-1})$ distribution has density with respect to Lebesgue measure.

B Proof of Main Results

B.1 Proofs for Section 4

We first state and prove a lemma which is used to prove Theorems 1 and 2. The first part shows that if $\pi \in \Pi$, then $\{\delta_n^*(\cdot; \pi)\} \in \mathbb{D}$ and its matching decision is independent of the prior π and nonrandomized. Thus, we drop dependence of the matching decision on π and U and simply write $\delta_\infty^*(Z)$. In the second part, we derive expressions for criteria (11) and (12) in terms of the limit experiment. Finally, in the third part, we show that our optimality criterion (12) is minimized for matching decisions corresponding to $\{\delta_n^*(\cdot; \pi)\}$ for $\pi \in \Pi$.

Recall that \dot{R}_{d,μ_0} denotes the directional derivative of $R(d, \mu)$ at μ_0 , $Z \sim N(h, I_{\mu_0}^{-1})$ represents the asymptotic behavior of an efficient estimator of μ , $Z^* \sim N(0, I_{\mu_0}^{-1})$ independently of Z represents asymptotic uncertainty about the local parameter h , and \mathbb{E}^* denotes expectation with respect to Z^* .

Lemma 1 *Suppose that Assumptions 1 and 2 hold. Fix any $\mu_0 \in \mathcal{M}$ for which there are no first-order ties.*

(i) *If $\pi \in \Pi$, then $\{\delta_n^*(\cdot; \pi)\} \in \mathbb{D}$ and the matching decision of $\{\delta_n^*(\cdot; \pi)\}$ is $\delta_\infty^*(Z)$ with*

$$\delta_\infty^*(z) = \operatorname{argmin}_{d \in \underline{\mathcal{D}}_{\mu_0}} \mathbb{E}^* \left[\dot{R}_{d,\mu_0}[Z^* + z] \right]$$

for almost every $z \in \mathbb{R}^K$.

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

$$\begin{aligned} \mathcal{R}(\{\delta_n\}; \mu_0, h) &= \sum_{d \in \underline{\mathcal{D}}_{\mu_0}} \left(\dot{R}_{d,\mu_0}[h] - \min_{i \in \underline{\mathcal{D}}_{\mu_0}} \dot{R}_{i,\mu_0}[h] \right) \mathbb{P}_h(\delta_\infty(Z, U) = d), \\ \mathcal{R}(\{\delta_n\}; \mu_0) &= \int \left(\sum_{d \in \underline{\mathcal{D}}_{\mu_0}} \left(\dot{R}_{d,\mu_0}[h] - \min_{i \in \underline{\mathcal{D}}_{\mu_0}} \dot{R}_{i,\mu_0}[h] \right) \mathbb{P}_h(\delta_\infty(Z, U) = d) \right) dh, \end{aligned} \quad (29)$$

where $\delta_\infty(Z, U)$ is the matching decision of $\{\delta_n\}$.

(iii) *The matching decision $\delta_\infty^*(Z)$ in (i) minimizes the right-hand side of (29) over all matching decisions $\delta_\infty(Z, U)$ of sequences $\{\delta_n\} \in \mathbb{D}$.*

Proof of Lemma 1. Part (i): Let A_n denote the event $\delta_n^*(X^n; \pi) \in \underline{\mathcal{D}}_{\mu_0}$. By Lemma 11, we have $\lim_{n \rightarrow \infty} \sqrt{n} P_{n,h}(A_n^c) = 0$ for all $h \in \mathbb{R}^K$. Hence, condition (ii) in Definition 1 holds. In particular, if $\underline{\mathcal{D}}_{\mu_0} \subsetneq \mathcal{D}$, then

$$\lim_{n \rightarrow \infty} P_{n,h}(\delta_n^*(X^n; \pi) = d) = 0 \quad \text{all } h \in \mathbb{R}^K \text{ and } d \notin \underline{\mathcal{D}}_{\mu_0}. \quad (30)$$

Hence, if $\mathcal{D}_{\mu_0} = \{d\}$ we have $\lim_{n \rightarrow \infty} P_{n,h}(\delta_n^*(X^n; \pi) = d) = 1$.

It remains to characterize $\lim_{n \rightarrow \infty} P_{n,h}(\delta_n^*(X^n; \pi) = d)$ for $d \in \mathcal{D}_{\mu_0}$ when $|\mathcal{D}_{\mu_0}| \geq 2$. We require a rule to break ties when $\operatorname{argmin}_{d \in \mathcal{D}} \bar{R}_n(d)$ is not a singleton. We shall take the smallest index d among the set of minimizers, though it will be seen below that the matching decision doesn't depend on the tie-breaking rule. Thus, any (possibly randomized) tie-breaking rule would lead to the same matching decision. Under this tie-breaking rule, we have

$$\mathbb{I}[\delta_n^*(X^n; \pi) = d] = \mathbb{I} \left[\bar{R}_n(d) < \min_{i \in \mathcal{D}_{\mu_0}: i < d} \bar{R}_n(i), \text{ and } \bar{R}_n(d) \leq \min_{i \in \mathcal{D}_{\mu_0}: i \geq d} \bar{R}_n(i) \right] \text{ on } A_n,$$

where the minimum over an empty set is $+\infty$. Since $|\Pr(A) - \Pr(B)| \leq 2\Pr(C^c)$ if $A = B$ holds on C , we obtain

$$\begin{aligned} \left| P_{n,h}(\delta_n^*(X^n; \pi) = d) - P_{n,h} \left(\bar{R}_n(d) < \min_{i \in \mathcal{D}_{\mu_0}: i < d} \bar{R}_n(i), \text{ and } \bar{R}_n(d) \leq \min_{i \in \mathcal{D}_{\mu_0}: i \geq d} \bar{R}_n(i) \right) \right| \\ \leq 2P_{n,h}(A_n^c) \rightarrow 0. \end{aligned}$$

As $R(d, \mu_0) = R(d', \mu_0)$ for all $d, d' \in \mathcal{D}_{\mu_0}$, we have

$$\begin{aligned} & P_{n,h} \left(\bar{R}_n(d) < \min_{i \in \mathcal{D}_{\mu_0}: i < d} \bar{R}_n(i), \text{ and } \bar{R}_n(d) \leq \min_{i \in \mathcal{D}_{\mu_0}: i \geq d} \bar{R}_n(i) \right) \\ &= P_{n,h} \left(\int \sqrt{n}(R(d, \mu) - R(d, \mu_0)) d\pi_n(\mu) < \min_{i \in \mathcal{D}_{\mu_0}: i < d} \int \sqrt{n}(R(i, \mu) - R(i, \mu_0)) d\pi_n(\mu), \right. \\ & \quad \left. \text{and } \int \sqrt{n}(R(d, \mu) - R(d, \mu_0)) d\pi_n(\mu) \leq \min_{i \in \mathcal{D}_{\mu_0}: i \geq d} \int \sqrt{n}(R(i, \mu) - R(i, \mu_0)) d\pi_n(\mu) \right). \end{aligned}$$

Now by Proposition 2, we have

$$\left(\int \sqrt{n}(R(d, \mu) - R(d, \mu_0)) d\pi_n(\mu) \right)_{d \in \mathcal{D}_{\mu_0}} \stackrel{P_{n,h}}{\rightsquigarrow} \left(\mathbb{E}^* [\dot{R}_{d,\mu_0} [Z^* + Z] | Z] \right)_{d \in \mathcal{D}_{\mu_0}}.$$

Define $U_d = \{(u_i)_{i \in \mathcal{D}_{\mu_0}} : u_d < \min_{i \in \mathcal{D}_{\mu_0}: i < d} u_i, \text{ and } u_d \leq \min_{i \in \mathcal{D}_{\mu_0}: i \geq d} u_i\} \subset \mathbb{R}^{|\mathcal{D}_{\mu_0}|}$ for $d \in \mathcal{D}_{\mu_0}$. These sets form a partition of $\mathbb{R}^{|\mathcal{D}_{\mu_0}|}$. Evidently, $\operatorname{int}(U_d) \supseteq \{(u_i)_{i \in \mathcal{D}_{\mu_0}} : u_d < \min_{i \in \mathcal{D}_{\mu_0}: i \neq d} u_i\}$. Thus, boundary points of U_d are points at which $\operatorname{argmin}_{i \in \mathcal{D}_{\mu_0}} u_i$ is not a singleton. Since $\operatorname{argmin}_{d \in \mathcal{D}_{\mu_0}} \mathbb{E}^* [\dot{R}_{d,\mu_0} [Z^* + z]]$ is a singleton for almost every z , it follows that

$$\mathbb{P}_h \left(\left(\mathbb{E}^* [\dot{R}_{d,\mu_0} [Z^* + Z] | Z] \right)_{d \in \mathcal{D}_{\mu_0}} \in \partial U_d \right) = 0.$$

Hence, by the Portmanteau theorem and the preceding four displays, for all $d \in \underline{\mathcal{D}}_{\mu_0}$ we have

$$\lim_{n \rightarrow \infty} P_{n,h}(\delta_n^*(X^n; \pi) = d) = \mathbb{P}_h \left(\left(\mathbb{E}^*[\dot{R}_{d,\mu_0}[Z^* + Z]|Z] < \min_{i \in \underline{\mathcal{D}}_{\mu_0}: i < d} \mathbb{E}^*[\dot{R}_{i,\mu_0}[Z^* + Z]|Z] \right), \text{ and} \right. \\ \left. \left(\mathbb{E}^*[\dot{R}_{d,\mu_0}[Z^* + Z]|Z] \leq \min_{i \in \underline{\mathcal{D}}_{\mu_0}: i \geq d} \mathbb{E}^*[\dot{R}_{i,\mu_0}[Z^* + Z]|Z] \right) \right). \quad (31)$$

Thus, we have shown that $\{\delta_n^*(\cdot; \pi)\} \in \mathbb{D}$.

We can now invoke Theorem 3.1 of [van der Vaart \(1991\)](#) to guarantee existence of the matching form $\delta_\infty^*(Z, U; \pi)$ of $\{\delta_n^*(\cdot; \pi)\}$. By (30), we see that $\delta_\infty^*(z, u; \pi)$ must take values in $\underline{\mathcal{D}}_{\mu_0}$ for almost every (z, u) . Moreover, as $\operatorname{argmin}_{d \in \underline{\mathcal{D}}_{\mu_0}} \mathbb{E}^*[\dot{R}_{d,\mu_0}[Z^* + z]]$ is a singleton for almost every z , we may restate (31) as

$$\lim_{n \rightarrow \infty} P_{n,h}(\delta_n^*(X^n; \pi) = d) = \mathbb{P}_h \left(d = \operatorname{argmin}_{i \in \underline{\mathcal{D}}_{\mu_0}} \mathbb{E}^*[\dot{R}_{i,\mu_0}[Z^* + Z]|Z] \right).$$

Hence, we may take $\delta_\infty^*(z, u; \pi) = \operatorname{argmin}_{i \in \underline{\mathcal{D}}_{\mu_0}} \mathbb{E}^*[\dot{R}_{i,\mu_0}[Z^* + z]]$ for almost every (z, u) . As $\delta_\infty^*(z, u; \pi)$ doesn't depend on u or π , we simply write $\delta_\infty^*(z)$.

Part (ii): By Assumption 1(i), $\min_{d \in \mathcal{D}} R(d, \mu_0 + h/\sqrt{n}) = \min_{d \in \underline{\mathcal{D}}_{\mu_0}} R(d, \mu_0 + h/\sqrt{n})$ holds for all n sufficiently large. Hence, for all n sufficiently large,

$$\sqrt{n} \left(\mathbb{E}_{n,h} \left[R(\delta_n(X^n), \mu_0 + h/\sqrt{n}) \right] - \min_{i \in \mathcal{D}} R(i, \mu_0 + h/\sqrt{n}) \right) \\ = \sum_{d \in \underline{\mathcal{D}}_{\mu_0}} P_{n,h}(\delta_n(X^n) = d) \sqrt{n} \left(R(d, \mu_0 + h/\sqrt{n}) - \min_{i \in \underline{\mathcal{D}}_{\mu_0}} R(i, \mu_0 + h/\sqrt{n}) \right) \\ + \sum_{d \notin \underline{\mathcal{D}}_{\mu_0}} \sqrt{n} P_{n,h}(\delta_n(X^n) = d) \left(R(d, \mu_0 + h/\sqrt{n}) - \min_{i \in \underline{\mathcal{D}}_{\mu_0}} R(i, \mu_0 + h/\sqrt{n}) \right),$$

where the second sum is zero if $\mathcal{D} = \underline{\mathcal{D}}_{\mu_0}$; otherwise, it converges to zero as $n \rightarrow \infty$ by Definition 1(ii) and Assumption 1(i). For the first sum, Theorem 3.1 of [van der Vaart \(1991\)](#) implies

$$\lim_{n \rightarrow \infty} P_{n,h}(\delta_n(X^n) = d) = \mathbb{P}_h(\delta_\infty(Z, U) = d), \quad d \in \underline{\mathcal{D}}_{\mu_0},$$

where $\delta_\infty(Z, U)$ is the matching decision of $\{\delta_n\}$. As $R(d, \mu_0) = R(d', \mu_0)$ for all $d, d' \in \underline{\mathcal{D}}_{\mu_0}$,

we have

$$\begin{aligned}
& \sqrt{n} \left(R(d, \mu_0 + h/\sqrt{n}) - \min_{i \in \mathcal{D}_{\mu_0}} R(i, \mu_0 + h/\sqrt{n}) \right) \\
&= \sqrt{n} \left(R(d, \mu_0 + h/\sqrt{n}) - R(d, \mu_0) \right) - \sqrt{n} \left(\min_{i \in \mathcal{D}_{\mu_0}} R(i, \mu_0 + h/\sqrt{n}) - \min_{i \in \mathcal{D}_{\mu_0}} R(i, \mu_0) \right). \\
&= \dot{R}_{d, \mu_0}[h] - \min_{i \in \mathcal{D}_{\mu_0}} \dot{R}_{i, \mu_0}[h],
\end{aligned}$$

where the final line follows from the fact that if $\{f_i\}_{i \in \mathcal{I}}$ is a finite set of functions that are directionally differentiable at μ_0 with $f_i(\mu_0) = f_{i'}(\mu_0)$ for all $i, i' \in \mathcal{I}$, then $\min_{i \in \mathcal{I}} f_i(\mu)$ has directional derivative $\min_{i \in \mathcal{I}} \dot{f}_{i, \mu_0}[h]$ at μ_0 .

Part (iii): Let $g_{d, \mu_0}[h] = \dot{R}_{d, \mu_0}[h] - \min_{i \in \mathcal{D}_{\mu_0}} \dot{R}_{i, \mu_0}[h]$. By part (ii), we have

$$\begin{aligned}
\mathcal{R}(\{\delta_n\}; \mu_0) &= \int \left(\sum_{d \in \mathcal{D}_{\mu_0}} g_{d, \mu_0}[h] \mathbb{P}_h(\delta_\infty(Z, U) = d) \right) dh \\
&\propto \sum_{d \in \mathcal{D}_{\mu_0}} \int \int \int_0^1 \mathbb{I}[\delta_\infty(z, u) = d] e^{-\frac{1}{2}(z-h)^T I_{\mu_0}(z-h)} du dz (g_{d, \mu_0}[h]) dh \\
&= \int_0^1 \int \sum_{d \in \mathcal{D}_{\mu_0}} \mathbb{I}[\delta_\infty(z, u) = d] \left(\int g_{d, \mu_0}[h] e^{-\frac{1}{2}(z-h)^T I_{\mu_0}(z-h)} dh \right) dz du.
\end{aligned}$$

Changing the order of integration in the final line above is justified by Tonelli's theorem, which only requires non-negativity of the integrand. Define

$$\mathcal{I}(z) = \operatorname{argmin}_{i \in \mathcal{D}_{\mu_0}} \int \dot{R}_{i, \mu_0}[h] e^{-\frac{1}{2}(z-h)^T I_{\mu_0}(z-h)} dh \equiv \operatorname{argmin}_{i \in \mathcal{D}_{\mu_0}} \mathbb{E}^* \left[\dot{R}_{i, \mu_0}[Z^* + z] \right]. \quad (32)$$

Evidently, any δ_∞ with $\delta_\infty(z, u) \in \mathcal{I}(z)$ for almost every (z, u) minimizes the criterion. Since there are no first-order ties at μ_0 , the set $\mathcal{I}(z)$ is a singleton for almost every z and the criterion is minimized by $\delta_\infty^*(z)$ derived in part (i). ■

Proof of Theorem 1. By Lemma 1, for any $\{\delta_n\} \in \mathbb{D}$, we have

$$\begin{aligned}
\mathcal{R}(\{\delta_n\}; \mu_0) &= \int \left(\sum_{d \in \mathcal{D}_{\mu_0}} \left(\dot{R}_{d, \mu_0}[h] - \min_{i \in \mathcal{D}_{\mu_0}} \dot{R}_{i, \mu_0}[h] \right) \mathbb{P}_h(\delta_\infty(Z, U) = d) \right) dh \\
&\geq \int \left(\sum_{d \in \mathcal{D}_{\mu_0}} \left(\dot{R}_{d, \mu_0}[h] - \min_{i \in \mathcal{D}_{\mu_0}} \dot{R}_{i, \mu_0}[h] \right) \mathbb{P}_h(\delta_\infty^*(Z) = d) \right) dh = \mathcal{R}(\{\delta_n^*(\cdot; \pi)\}; \mu_0).
\end{aligned}$$

Result (i) now follows by taking the infimum over $\{\delta_n\} \in \mathbb{D}$ and noting that $\{\delta_n^*(\cdot; \pi)\} \in \mathbb{D}$ by Lemma 1(i).

Result (ii) follows from the fact that the matching forms δ_∞^* of $\{\delta_n^*(\cdot; \pi)\}$ derived in Lemma 1(i) do not depend on the prior $\pi \in \Pi$.

Finally, if $\{\delta_n\} \in \mathbb{D}$ is asymptotically equivalent to $\{\delta_n^*(\cdot; \pi)\}$, then $\{\delta_n\}$ must also have matching form $\delta_\infty^*(Z)$. Hence, by similar arguments to the above, we have $\mathcal{R}(\{\delta_n'\}; \mu_0) \geq \mathcal{R}(\{\delta_n\}; \mu_0)$ for any $\{\delta_n'\} \in \mathbb{D}$. Taking the infimum over $\{\delta_n'\} \in \mathbb{D}$ yields result (iii). ■

Proof of Theorem 2. First note that asymptotic equivalence cannot fail if $|\underline{\mathcal{D}}_{\mu_0}| = 1$, since if $\underline{\mathcal{D}}_{\mu_0} = \{d\}$, say, then $\lim_{n \rightarrow \infty} P_{n,h}(\delta_n(X^n) = d) = 1$ must hold for any $\{\delta_n\} \in \mathbb{D}$.

Now suppose $|\underline{\mathcal{D}}_{\mu_0}| \geq 2$. Let $g_{d,\mu_0}[h] = \dot{R}_{d,\mu_0}[h] - \min_{i \in \underline{\mathcal{D}}_{\mu_0}} \dot{R}_{i,\mu_0}[h]$. By Lemma 1(ii)-(iii), it suffices to show

$$\int \left(\sum_{d \in \underline{\mathcal{D}}_{\mu_0}} g_{d,\mu_0}[h] \mathbb{P}_h(\delta_\infty(Z, U) = d) \right) dh > \int \left(\sum_{d \in \underline{\mathcal{D}}_{\mu_0}} g_{d,\mu_0}[h] \mathbb{P}_h(\delta_\infty^*(Z) = d) \right) dh,$$

where $\delta_\infty(Z, U)$ is the matching form of $\{\delta_n\}$ and $\delta_\infty^*(Z)$ is the matching form of $\{\delta_n^*(\cdot; \pi)\}$ from Lemma 1(i). Since both integrands are non-negative, by Tonelli's theorem we may restate this inequality as

$$\begin{aligned} \int_0^1 \int \sum_{d \in \underline{\mathcal{D}}_{\mu_0}} \mathbb{I}[\delta_\infty(z, u) = d] \left(\int g_{d,\mu_0}[h] e^{-\frac{1}{2}(z-h)^T I_{\mu_0}(z-h)} dh \right) dz du \\ > \int_0^1 \int \sum_{d \in \underline{\mathcal{D}}_{\mu_0}} \mathbb{I}[\delta_\infty^*(z) = d] \left(\int g_{d,\mu_0}[h] e^{-\frac{1}{2}(z-h)^T I_{\mu_0}(z-h)} dh \right) dz du. \end{aligned}$$

As $\{\delta_n\}$ and $\{\delta_n^*(\cdot; \pi)\}$ are not asymptotically equivalent, we have

$$\lim_{n \rightarrow \infty} P_{n,h_0}(\delta_n(X^n) = d) \neq \lim_{n \rightarrow \infty} P_{n,h_0}(\delta_n^*(X^n; \pi) = d)$$

for some $h_0 \in \mathbb{R}^K$ and some $d \in \mathcal{D}$. As $\{\delta_n\}$ and $\{\delta_n^*(\cdot; \pi)\}$ are both elements of \mathbb{D} , we can use Theorem 3.1 of van der Vaart (1991) and restate this in terms of their matching forms $\delta_\infty(Z, U)$ and $\delta_\infty^*(Z)$:

$$\mathbb{P}_{h_0}(\delta_\infty(Z, U) = d) \neq \mathbb{P}_{h_0}(\delta_\infty^*(Z) = d).$$

Hence, it must be the case that $\delta_\infty(z, u)$ and $\delta_\infty^*(z)$ disagree on a set S of positive Lebesgue measure. As $\delta_\infty^*(z) \in \mathcal{I}(z)$ for almost every z , and $\mathcal{I}(z)$ is a singleton for almost every z as there are no first-order ties, we must have that $\delta_\infty(z, u) \notin \mathcal{I}(z)$ on a set S' of positive

Lebesgue measure. Now by definition of $\mathcal{I}(z)$, we also have that

$$\begin{aligned} \sum_{d \in \mathcal{D}_{\mu_0}} \mathbb{I}[\delta_\infty(z, u) = d] & \left(\int g_{d, \mu_0}[h] e^{-\frac{1}{2}(z-h)^T I_{\mu_0}(z-h)} dh \right) \\ & \geq \sum_{d \in \mathcal{D}_{\mu_0}} \mathbb{I}[\delta_\infty^*(z) = d] \left(\int g_{d, \mu_0}[h] e^{-\frac{1}{2}(z-h)^T I_{\mu_0}(z-h)} dh \right), \end{aligned}$$

holds for almost every (z, u) , with strict inequality for (z, u) in a set of positive measure. Integrating with respect to (z, u) yields the strict inequality above. ■

Proof of Corollary 1. Immediate from Theorem 1 and the argument given in the text preceding the statement of this corollary. ■

Proof of Proposition 1. Let $g_{d, \mu_0}[h] = \dot{R}_{d, \mu_0}[h] - \min_{i \in \mathcal{D}_{\mu_0}} \dot{R}_{i, \mu_0}[h]$. Note $g_{d, \mu_0}[h] \geq 0$. By Lemma 1, we have

$$\mathcal{R}(\{\delta_n^*(\cdot; \pi)\}; \mu_0) \leq \sum_{d \in \mathcal{D}_{\mu_0}} \int g_{d, \mu_0}[h] \mathbb{P}_h \left(\operatorname{argmin}_{i \in \mathcal{D}_{\mu_0}} \dot{R}_{i, \mu_0}[Z] = d \right) dh,$$

where interchanging summation and integration is justified by Tonelli's theorem. Since $h \in \mathbb{R}$, we split the integral into an integral over $[0, \infty)$ and an integral over $(-\infty, 0)$. Suppose $h \geq 0$. For each $d \in \mathcal{D}_{\mu_0}$, note $g_{d, \mu_0}[h] = h g_{d, \mu_0}[1]$ by positive homogeneity. Suppose $g_{d, \mu_0}[1] > 0$ for some $d \in \mathcal{D}_{\mu_0}$, else the integral is zero. Then $d = \operatorname{argmin}_{i \in \mathcal{D}_{\mu_0}} \dot{R}_{i, \mu_0}[z]$ implies $g_{d, \mu_0}[z] = 0$. And since $g_{d, \mu_0}[1] > 0$ and $g_{d, \mu_0}[\cdot]$ is positively homogeneous, $g_{d, \mu_0}[z] = 0$ in turn implies $z \leq 0$. Hence,

$$\mathbb{P}_h(\operatorname{argmin}_{i \in \mathcal{D}_{\mu_0}} \dot{R}_{i, \mu_0}[Z] = d) \leq \mathbb{P}_h(Z \leq 0) = \Phi(-h/\sigma),$$

where $\sigma = I_{\mu_0}^{-1/2} > 0$, and Φ is the CDF of the $N(0, 1)$ distribution. Then

$$\int_0^\infty g_{d, \mu_0}[h] \mathbb{P}_h \left(\operatorname{argmin}_{i \in \mathcal{D}_{\mu_0}} \dot{R}_{i, \mu_0}[Z] = d \right) dh \leq g_d[1] \int_0^\infty h \Phi(-h/\sigma) dh < \infty.$$

One may similarly show the integral over $(-\infty, 0)$ is finite. ■

We first present a preliminary result before turning to the proof of Theorem 3.

Lemma 2 *Suppose that Assumption 1 and 2 hold. Then $\{\delta_n\} \in \mathbb{D}$ is σ -optimal at μ_0 if its*

matching decision is $\delta_\sigma^*(Z, U)$ with

$$\delta_\sigma^*(z, u) \in \operatorname{argmin}_{d \in \mathcal{D}_{\mu_0}} \mathbb{E}_\sigma^* \left[\dot{R}_{d, \mu_0} [Z_\sigma^* + (I + (\sigma I_{\mu_0})^{-1})^{-1} z] \right]$$

for almost every $(z, u) \in \mathbb{R}^K \times [0, 1]$, where \mathbb{E}_σ^* denotes expectation with respect to $Z_\sigma^* \sim N(0, (I_{\mu_0} + \sigma^{-1}I)^{-1})$.

Proof of Lemma 2. By Lemma 1(ii), for any $\{\delta_n\} \in \mathbb{D}$ we have

$$\mathcal{R}_\sigma(\{\delta_n\}; \mu_0) \propto \int \sum_{d \in \mathcal{D}_{\mu_0}} g_{d, \mu_0}[h] \mathbb{P}_h(\delta_\infty(Z, U) = d) e^{-\frac{1}{2\sigma} \|h\|^2} dh,$$

where $g_{d, \mu_0}[h] = \dot{R}_{d, \mu_0}[h] - \min_{i \in \mathcal{D}_{\mu_0}} \dot{R}_{i, \mu_0}[h]$ and $\delta_\infty(Z, U)$ is the matching decision of $\{\delta_n\}$. By Tonelli's theorem, we have

$$\mathcal{R}_\sigma(\{\delta_n\}; \mu_0) \propto \int_0^1 \int \sum_{d \in \mathcal{D}_{\mu_0}} \mathbb{I}[\delta_\infty(z, u) = d] \left(\int g_{d, \mu_0}[h] e^{-\frac{1}{2}(z-h)^T I_{\mu_0} (z-h) - \frac{1}{2\sigma} \|h\|^2} dh \right) dz du.$$

Therefore, $\mathcal{R}_\sigma(\{\delta_n\}; \mu_0)$ is minimized when

$$\delta_\infty(z, u) \in \operatorname{argmin}_{i \in \mathcal{D}_{\mu_0}} \int \dot{R}_{i, \mu_0}[h] e^{-\frac{1}{2}(z-h)^T I_{\mu_0} (z-h) - \frac{1}{2\sigma} \|h\|^2} dh$$

for almost every (z, u) . The result follows by completing the square. ■

Proof of Theorem 3. Lemma 1(i) implies

$$\lim_{n \rightarrow \infty} P_{n, h}(\delta_n^*(X^n; \pi) = d) = \mathbb{P}_h(\delta^*(Z) = d),$$

for all $d \in \mathcal{D}$ and $h \in \mathbb{R}^K$, where

$$\delta^*(z) = \mathcal{I}(z) := \operatorname{argmin}_{i \in \mathcal{D}_{\mu_0}} \mathbb{E}^* \left[\dot{R}_{i, \mu_0} [Z^* + z] \right]$$

for almost every z , since there are no first-order ties. Similarly, Lemma 2 implies

$$\lim_{n \rightarrow \infty} P_{n, h}(\delta_{n, \sigma}^*(X^n) = d) = \mathbb{P}_h(\delta_\sigma^*(Z, U) = d),$$

for all $d \in \mathcal{D}$ and $h \in \mathbb{R}^K$, where

$$\delta_\sigma^*(z, u) \in \mathcal{I}_\sigma(z) := \operatorname{argmin}_{i \in \underline{\mathcal{D}}_{\mu_0}} \mathbb{E}_\sigma^* \left[\dot{R}_{i, \mu_0} [Z_\sigma^* + (I + (\sigma I_{\mu_0})^{-1})^{-1} z] \right]$$

for almost every $(z, u) \in \mathbb{R}^K \times [0, 1]$.

For almost every z , we have by dominated convergence that

$$\lim_{\sigma \rightarrow \infty} \mathbb{E}_\sigma^* \left[\dot{R}_{i, \mu_0} [Z_\sigma^* + (I + (\sigma I_{\mu_0})^{-1})^{-1} z] \right] = \mathbb{E}^* \left[\dot{R}_{i, \mu_0} [Z^* + z] \right].$$

Hence, $\lim_{\sigma \rightarrow \infty} \mathcal{I}_\sigma(z) = \mathcal{I}(z)$ for almost every z , and so $\lim_{\sigma \rightarrow \infty} \delta_\sigma^*(z, u) = \delta^*(z)$ for almost every (z, u) . Hence, by dominated convergence and Theorem 3.1 of [van der Vaart \(1991\)](#),

$$\begin{aligned} \lim_{n \rightarrow \infty} P_{n, h}(\delta_n^*(X^n; \pi) = d) &= \mathbb{P}_h(\delta^*(Z) = d) \\ &= \lim_{\sigma \rightarrow \infty} \mathbb{P}_h(\delta_\sigma^*(Z, U) = d) = \lim_{\sigma \rightarrow \infty} \lim_{n \rightarrow \infty} P_{n, h}(\delta_{n, \sigma}^*(X^n) = d), \end{aligned}$$

as required. ■

B.2 Proofs for Section 5

Once we fix (μ_0, η_0) , the least favorable model $\{P_{\beta(t)} : t \in \mathcal{M}_{(\mu_0, \eta_0)}\}$ is a parametric model. By Assumption 3(i)-(iii) and Example 7.15 and Corollary 9.5 of [van der Vaart \(1998\)](#), the sequence of experiments $\{\mathcal{E}_n\}$, with $\mathcal{E}_n = (\mathcal{X}^n, \mathcal{B}_n, P_{n, h} : h \in \mathbb{R}^K)$ and $P_{n, h}$ denoting the distribution of X^n under $P_{\beta(\mu_0 + h/\sqrt{n})}$, converges to a limit experiment $\mathcal{E} = (\mathbb{R}^K, \mathcal{B}_K, N(h, I_{(\mu_0, \eta_0)}^{-1}) : h \in \mathbb{R}^K)$. Because \mathcal{E} is dominated, the asymptotic representation theorem of [van der Vaart \(1991\)](#) implies that each $\{\delta_n\} \in \mathbb{D}$ is matched by a $\delta_\infty(Z, U)$ given by a measurable map $\delta_\infty : \mathbb{R}^K \times [0, 1] \rightarrow \mathcal{D}$, the observation $Z \sim N(h, I_{(\mu_0, \eta_0)}^{-1})$, and an independent randomization $U \sim \text{Uniform}[0, 1]$, with the property that for each $h \in \mathbb{R}^K$,

$$\lim_{n \rightarrow \infty} P_{n, h}(\delta_n(X^n) = d) = \mathbb{P}_h(\delta_\infty(Z, U) = d), \quad d \in \mathcal{D},$$

where \mathbb{P}_h is the product of the $N(h, I_{(\mu_0, \eta_0)}^{-1})$ and $U[0, 1]$ distributions

Once these modifications are made, the proofs of Theorems 4 and 5 follow identical arguments to the proofs of Theorems 1 and 2, using Lemma 12 in place of Lemma 11 and Proposition 3 in place of Proposition 2.

C Asymptotic Distribution of the Posterior Mean of Directionally Differentiable Functions

In this section we derive the asymptotic distribution of the posterior mean of directionally differentiable functions. Let $f : \mathcal{M} \rightarrow \mathbb{R}^k$ be a directionally differentiable function of μ . The first main result is Proposition 2, which presents the asymptotic distribution of the centered and scaled posterior mean

$$\bar{f}_n := \int \sqrt{n}(f(\mu) - f(\mu_0))d\pi_n(\mu)$$

along sequences of local parameters $P_{n,h}$. In Proposition 3 we then extend this result to semiparametric models. Both results are of independent interest. The most closely related work is Kitagawa, Montiel Olea, Payne, and Velez (2020), which characterizes the asymptotic behavior of the posterior distribution of directionally differentiable functions. We instead characterize the (frequentist) asymptotic distribution of its posterior mean. For both results, we assume the following:

Assumption 4 (i) f is integrable with respect to the prior π for μ ;
(ii) f is directionally differentiable at μ_0 .

C.1 Parametric Models

We first consider parametric models as described in Section 4. We present some regularity conditions and preliminary lemmas before giving the main result.

Recall the notation from Section 4. Define $Z_{n,\mu_0} = \frac{1}{\sqrt{n}} \sum_{i=1}^n I_{\mu_0}^{-1} \dot{\ell}_{\mu_0}(X_i)$. Let $\tilde{\pi}_n$ denote the posterior induced by $\pi_n(\mu)$ under a change of variables from μ to $h = \sqrt{n}(\mu - \mu_0) - Z_{n,\mu_0}$. Let P^* denote the distribution of $Z^* \sim N(0, I_{\mu_0}^{-1})$ and let $\|\cdot\|_{TV}$ denote total variation distance.

Lemma 3 *Suppose that Assumptions 2(i)-(iii) and (v) hold, Assumption 4(ii) holds, and $\pi \in \Pi$. Then there exists a sequence $\{B_n\} \subset \mathbb{R}_+$ with $B_n \uparrow +\infty$, $B_n = o(\sqrt{n})$, such that for all n sufficiently large,*

$$\sup_{h: \|h\| \leq 2B_n} \left\| \sqrt{n} (f(\mu_0 + h/\sqrt{n}) - f(\mu_0)) - \dot{f}_{\mu_0}[h] \right\| \leq \frac{1}{B_n},$$

and

$$P_{n,\mu_0} \left(\|\tilde{\pi}_n - P^*\|_{TV} > \frac{1}{B_n^2} \right) \leq \frac{1}{B_n}.$$

Proof of Lemma 3. By Assumption 4(ii) and Proposition 3.3 of Shapiro (1990), for each $a \in \mathbb{N}$ there exists $n_{a,f} \in \mathbb{N}$ such that $n \geq n_{a,f}$ implies

$$\sup_{h: \|h\| \leq 2a} \left\| \sqrt{n} (f(\mu_0 + h/\sqrt{n}) - f(\mu_0)) - \dot{f}_{\mu_0}[h] \right\| \leq \frac{1}{a}.$$

Under Assumptions 2(i)-(iii) and since $\pi \in \Pi$, we may invoke Theorem 10.1 of van der Vaart (1998), using Assumption 2(v) and Proposition 6.2 of Clarke and Barron (1990) to deduce existence of tests with the desired properties. Further, we may use invariance of total variation under location changes to deduce that for all $a \in \mathbb{N}$ there exists $n_{a,\pi} \in \mathbb{N}$ such that $n \geq n_{a,\pi}$ implies

$$P_{n,\mu_0} \left(\|\tilde{\pi}_n - P^*\|_{TV} > \frac{1}{a^2} \right) \leq \frac{1}{a}.$$

The result follows by setting $B_n = \sup\{a \in \mathbb{N} : n \geq (n_{a,f} \vee n_{a,\pi})\}$ for $n \geq (n_{1,f} \vee n_{1,\pi})$, then letting B_n increase slower if necessary so that $B_n = o(\sqrt{n})$. ■

Let $\xrightarrow{P_{n,\mu_0}}$ denote convergence in probability under $\{P_{n,\mu_0}\}$. Let \mathbb{E}^* denote expectation with respect to $Z^* \sim N(0, I_{\mu_0}^{-1})$ independent of Z_{n,μ_0} .

Lemma 4 *Suppose Assumptions 2 and 4 hold and $\pi \in \Pi$. Then*

$$\left\| \bar{f}_n - \mathbb{E}^* \left[\dot{f}_{\mu_0} [Z^* + Z_{n,\mu_0}] \middle| Z_{n,\mu_0} \right] \right\| \xrightarrow{P_{n,\mu_0}} 0$$

as $n \rightarrow \infty$.

Proof of Lemma 4. Let $\{B_n\}$ be as in Lemma 3. By a change of variables, we have

$$\begin{aligned} \bar{f}_n &= \int_{h: \|h\| \leq B_n} \sqrt{n} (f(\mu_0 + n^{-1/2}(Z_{n,\mu_0} + h)) - f(\mu_0)) d\tilde{\pi}_n(h) \\ &\quad + \int_{\mu: \|\sqrt{n}(\mu - \mu_0) - Z_{n,\mu_0}\| > B_n} \sqrt{n} (f(\mu) - f(\mu_0)) d\pi_n(\mu) =: \mathcal{J}_{1,n} + \mathcal{J}_{2,n}. \end{aligned}$$

Let A_n denote the event $\|Z_{n,\mu_0}\| \leq B_n$ and $\|\tilde{\pi}_n - P^*\|_{TV} \leq B_n^{-2}$. Assumption 2(iii) and Lemma 3 imply $\lim_{n \rightarrow \infty} P_{n,\mu_0}(A_n^c) = 0$. Again by Lemma 3, on A_n we have

$$\left\| \mathcal{J}_{1,n} - \int_{h: \|h\| \leq B_n} \dot{f}_{\mu_0} [Z_{n,\mu_0} + h] d\tilde{\pi}_n(h) \right\| \leq \frac{1}{B_n} \rightarrow 0.$$

Note that $\dot{f}_{\mu_0}[\cdot]$ is continuous under Assumption 4(ii) (Shapiro, 1990, Proposition 3.1) and

hence that

$$C_f := \sup_{u: \|u\|=1} \|\dot{f}_{\mu_0}[u]\| < \infty. \quad (33)$$

Moreover, $\dot{f}_{\mu_0}[\cdot]$ is positively homogeneous of degree one (Shapiro, 1990, p. 478). It therefore follows by Lemma 3 that on A_n we have

$$\begin{aligned} \left\| \int_{h: \|h\| \leq B_n} \dot{f}_{\mu_0}[Z_{n,\mu_0} + h] d\tilde{\pi}_n(h) - \int_{h: \|h\| \leq B_n} \dot{f}_{\mu_0}[Z_{n,\mu_0} + h] dP^*(h) \right\| \\ \leq 2B_n C_f \|\tilde{\pi}_n - P^*\|_{TV} \leq \frac{2C_f}{B_n} \rightarrow 0. \end{aligned}$$

Now, as $\|Z_{n,\mu_0}\| \leq B_n$ holds on A_n , we have by the triangle and Chebyshev inequalities that

$$\begin{aligned} \left\| \int_{h: \|h\| > B_n} \dot{f}_{\mu_0}[Z_{n,\mu_0} + h] dP^*(h) \right\| &\leq C_f \int_{h: \|h\| > B_n} (\|Z_{n,\mu_0}\| + \|h\|) dP^*(h) \\ &\leq C_f \left(B_n P^*(\|Z^*\| > B_n) + \int_{h: \|h\| > B_n} \|h\| dP^*(h) \right) \\ &\leq \frac{2C_f \text{tr}(I_{\mu_0}^{-1})}{B_n} \rightarrow 0 \end{aligned}$$

by Assumption 2(iii). Combining the preceding three displays, we see that

$$\left\| \mathcal{J}_{1,n} - \mathbb{E}^* \left[\dot{f}_{\mu_0}[Z^* + Z_{n,\mu_0}] \middle| Z_{n,\mu_0} \right] \right\| \rightarrow 0 \quad \text{on } A_n.$$

It follows that $\|\mathcal{J}_{1,n} - \mathbb{E}^*[\dot{f}_{\mu_0}[Z^* + Z_{n,\mu_0}]|Z_{n,\mu_0}]\| \xrightarrow{P_{n,\mu_0}} 0$ because $P_{n,\mu_0}(A_n^c) \rightarrow 0$.

It remains to show that $\|\mathcal{J}_{2,n}\| \xrightarrow{P_{n,\mu_0}} 0$. First note that by Assumptions 2(i)-(iii) we may invoke Lemma 10.3 of van der Vaart (1998), using Assumption 2(v) and Proposition 6.2 of Clarke and Barron (1990) to deduce existence of uniformly consistent tests as required in Theorem 10.1 of van der Vaart (1998). Hence, there exists a sequence of tests $\{\phi_n\}$ and a constant $c > 0$ such that for every sufficiently large n and every $\mu \in \mathcal{M}$ with $\|\mu - \mu_0\| \geq B_n/(2\sqrt{n})$, we have $\mathbb{E}_{n,\mu}[\phi_n(X^n)] \rightarrow 0$ and $\mathbb{E}_{n,\mu}[1 - \phi_n(X^n)] \leq e^{-cn(\|\mu - \mu_0\|^2 \wedge 1)}$. Without loss of generality we may take the tests to be nonrandomized (van der Vaart, 1998, p. 148). Then

$$\mathcal{J}_{2,n} = \phi_n(X^n) \mathcal{J}_{2,n} + (1 - \phi_n(X^n)) \mathcal{J}_{2,n},$$

where $P_{n,\mu_0}(\|\phi_n(X^n) \mathcal{J}_{2,n}\| > 0) \leq P_{n,\mu_0}(\phi_n(X^n) = 1) \rightarrow 0$. The remaining term is

$$(1 - \phi_n(X^n)) \frac{\int_{\mu: B_n < \|\sqrt{n}(\mu - \mu_0) - Z_{n,\mu_0}\|} \sqrt{n} (f(\mu) - f(\mu_0)) \prod_{i=1}^n p_\mu(X_i) / p_{\mu_0}(X_i) d\pi(\mu)}{\mathcal{J}_n},$$

where $\mathcal{J}_n = \int \prod_{i=1}^n p_\mu(X_i)/p_{\mu_0}(X_i) d\pi(\mu)$. By Lemma 6.26 of Ghosal and van der Vaart (2017) and Lemma 10, there exists $C > 0$ such that for any $C_n > 0$, the inequality

$$P_{n,\mu_0}(\mathcal{J}_n < Cn^{-K/2}e^{-2C_n}) \leq \frac{1}{C_n}$$

holds for all n sufficiently large. Below we shall specify a sequence $\{C_n\} \subset \mathbb{R}_+$ with $C_n \uparrow \infty$. Letting A'_n denote the event upon which $\|Z_{n,\mu_0}\| \leq \frac{1}{2}B_n$ and $\mathcal{J}_n \geq Cn^{-K/2}e^{-2C_n}$ both hold, we have by the above display and Assumption 2(iii) that $P_{n,\mu_0}(A_n^c) \rightarrow 0$. It follows by the union bound, Markov's inequality, and Fubini's theorem that for any $\epsilon > 0$,

$$\begin{aligned} P_{n,\mu_0}(\|(1 - \phi_n(X^n))\mathcal{J}_{2,n}\| > \epsilon) \\ \leq \frac{\int_{\mu: \|\sqrt{n}(\mu - \mu_0)\| > \frac{1}{2}B_n} \|\sqrt{n}(f(\mu) - f(\mu_0))\| e^{-cn(\|\mu - \mu_0\|^2 \wedge 1)} d\pi(\mu)}{\epsilon Cn^{-K/2}e^{-2C_n}} + P_{n,\mu_0}(A_n^c). \end{aligned}$$

Split the numerator into integrals over $N_{1,n} = \{\mu : \frac{1}{2}B_n < \|\sqrt{n}(\mu - \mu_0)\| \leq \epsilon' \sqrt{n}\}$ and $N_{2,n} = \{\mu : \|\mu - \mu_0\| > \epsilon'\}$ for $\epsilon' \in (0, 1)$ sufficiently small that $\pi(\mu) \leq C_2$ on $\{\mu : \|\mu - \mu_0\| \leq \epsilon'\}$. Finally, Proposition 3.3 of Shapiro (1990) implies $\sup_{h: \|h\| \leq 1} \|f(\mu_0 + th) - f(\mu_0) - \dot{f}_{\mu_0}[th]\| = o(t)$ as $t \downarrow 0$. Hence, we may choose $\bar{t} > 0$ such that $\|f(\mu_0 + th) - f(\mu_0) - \dot{f}_{\mu_0}[th]\| \leq t$ holds for all $t \leq \bar{t}$ and all $\|h\| \leq 1$. Let $\epsilon' = \min\{\epsilon', \bar{t}\}$. It then follows for any $\mu \in N_{1,n}$ that, by setting $t = \|\mu - \mu_0\|$ and $h = \frac{\mu - \mu_0}{\|\mu - \mu_0\|}$, we have $\|f(\mu) - f(\mu_0)\| \leq \|\dot{f}_{\mu_0}[\mu - \mu_0]\| + \|\mu - \mu_0\|$. Hence, by (33), we have that $\|f(\mu) - f(\mu_0)\| \leq (1 + C_f)\|\mu - \mu_0\|$ holds on $N_{1,n}$. Then

$$\begin{aligned} & \int_{N_{1,n}} \|\sqrt{n}(f(\mu) - f(\mu_0))\| e^{-cn(\|\mu - \mu_0\|^2 \wedge 1)} d\pi(\mu) \\ & \leq C_2(1 + C_f) \int_{N_{1,n}} \|\sqrt{n}(\mu - \mu_0)\| e^{-c\|\sqrt{n}(\mu - \mu_0)\|^2} d\mu \\ & \leq C_2(1 + C_f)n^{-K/2} \int_{h: \|h\| > \frac{1}{2}B_n} \|h\| e^{-c\|h\|^2} dh, \end{aligned}$$

where the final line is by a change of variables from μ to $h = \sqrt{n}(\mu - \mu_0)$. The final integral vanishes as $n \rightarrow \infty$ by dominated convergence. Further, by Assumption 4(i), we have

$$\int_{N_{2,n}} \|\sqrt{n}(f(\mu) - f(\mu_0))\| e^{-cn(\|\mu - \mu_0\|^2 \wedge 1)} d\pi(\mu) \leq \sqrt{n}e^{-nc(\epsilon')^2} \int \|f(\mu) - f(\mu_0)\| d\pi(\mu).$$

Choosing $\{C_n\} \subset \mathbb{R}_+$ with $C_n \uparrow \infty$ sufficiently slowly that

$$e^{2C_n} \int_{h: \|h\| > \frac{1}{2}B_n} \|h\| e^{-c\|h\|^2} dh \rightarrow 0, \quad \text{and} \quad n^{(K+1)/2} e^{2C_n - nc(\epsilon')^2} \rightarrow 0,$$

it follows from the preceding four displays that $P_{n,\mu_0} (\|(1 - \phi_n(X^n))\mathcal{J}_{2,n}\| > \epsilon) \rightarrow 0$. ■

For the next result, recall that $Z \sim N(h, I_{\mu_0}^{-1})$, \mathbb{E}^* denotes expectation with respect to $Z^* \sim N(0, I_{\mu_0}^{-1})$ independent of Z , and $\overset{P_{n,h}}{\rightsquigarrow}$ denotes convergence in distribution along $\{P_{n,h}\}$.

Proposition 2 *Suppose that Assumptions 2 and 4 hold and $\pi \in \Pi$. Then for every $h \in \mathbb{R}^K$,*

$$\bar{f}_n \overset{P_{n,h}}{\rightsquigarrow} \mathbb{E}^* \left[\dot{f}_{\mu_0} [Z^* + Z] \middle| Z \right].$$

Proof of Proposition 2. Assumption 2(i)-(iii) imply $\mathcal{P} = \{P_\mu : \mu \in \mathcal{M}\}$ is locally asymptotically normal (van der Vaart, 1998, Example 7.15) and so $\{P_{n,\mu_0}\}$ and $\{P_{n,h}\}$ are mutually contiguous (van der Vaart, 1998, Example 6.5). Hence, with $\overset{P_{n,h}}{\rightarrow}$ denoting convergence in probability under $\{P_{n,h}\}$, we have by Lemma 4 that

$$\left\| \bar{f}_n - \mathbb{E}^* \left[\dot{f}_{\mu_0} [Z^* + Z_{n,\mu_0}] \middle| Z_{n,\mu_0} \right] \right\| \overset{P_{n,h}}{\rightarrow} 0.$$

Moreover, under Assumption 2(i)-(iii) we have by Le Cam's third lemma (van der Vaart, 1998, Example 6.7) that

$$Z_{n,\mu_0} \overset{P_{n,h}}{\rightsquigarrow} Z \sim N(h, I_{\mu_0}^{-1}).$$

The result follows by the continuous mapping theorem, noting $z \mapsto \mathbb{E}^*[\dot{f}_{\mu_0}[Z^* + z]]$ is everywhere continuous by Assumption 4(ii) and Proposition 3.1 of Shapiro (1990). ■

C.2 Semiparametric Models

We now extend Proposition 2 to semiparametric models. In this case,

$$\bar{f}_n = \int \sqrt{n}(f(\mu) - f(\mu_0))d\pi_n(\mu)$$

with π_n denoting the quasi-posterior in (25) based on a Gaussian quasi-likelihood. Before giving the main result we first state and prove several lemmas. The first is a Bernstein–von Mises result for the quasi-posterior. Let $Z_{n,(\mu_0,\eta_0)} = \sqrt{n}(\hat{\mu} - \mu_0)$. Let $\tilde{\pi}_n(h) \propto e^{-\frac{1}{2}h^T \hat{I}h} \pi(\mu_0 + n^{-1/2}(h + Z_{n,(\mu_0,\eta_0)}))$ denote the quasi-posterior for h under a change of variables from μ to $h = \sqrt{n}(\mu - \mu_0) - Z_{n,(\mu_0,\eta_0)}$. Let P^* denote the distribution of $Z^* \sim N(0, I_{(\mu_0,\eta_0)}^{-1})$.

Lemma 5 *Suppose that Assumption 3(i),(iii),(v) holds and $\pi \in \Pi$. Then*

$$\|\tilde{\pi}_n - P^*\|_{TV} \overset{P_{n,(\mu_0,\eta_0)}}{\rightarrow} 0.$$

Proof of Lemma 5. Fix any Borel set $B \subset \mathbb{R}^K$. We have

$$\tilde{\pi}_n(B) = \frac{\int_{h: \mu_0 + n^{-1/2}(h + Z_{n,(\mu_0, \eta_0)}) \in \mathcal{M}} \mathbb{I}[h \in B] e^{-\frac{1}{2}h^T \hat{I} h} \pi(\mu_0 + n^{-1/2}(h + Z_{n,(\mu_0, \eta_0)})) dh}{\int_{h: \mu_0 + n^{-1/2}(h + Z_{n,(\mu_0, \eta_0)}) \in \mathcal{M}} e^{-\frac{1}{2}h^T \hat{I} h} \pi(\mu_0 + n^{-1/2}(h + Z_{n,(\mu_0, \eta_0)})) dh}.$$

Let $\{M_n\} \subset \mathbb{R}_+$ be a sequence with $M_n \uparrow +\infty$, $M_n = o(n^{1/2})$. Let A_n denote the event upon which $\|Z_{n,(\mu_0, \eta_0)}\| \leq M_n$ and $c \leq \hat{\lambda}_{\min}$, $\hat{\lambda}_{\max} \leq c^{-1}$ hold, where $\hat{\lambda}_{\min}$ and $\hat{\lambda}_{\max}$ are the smallest and largest eigenvalues of \hat{I} , and $c \in (0, 1)$ is chosen such that the minimum and maximum eigenvalues of $I_{(\mu_0, \eta_0)}$ are contained in (c, c^{-1}) (such a c exists by Assumption 3(iii)). We have $P_{n,(\mu_0, \eta_0)}(A_n) \rightarrow 1$ by Assumptions 3(iii) and 3(v). Moreover, by Assumption 3(i),

$$\{\mu_0 + n^{-1/2}(h + Z_{n,(\mu_0, \eta_0)}) : \|h\| \leq M_n\} \subset \{\mu_0 + n^{-1/2}h : \|h\| \leq 2M_n\} \subset \mathcal{M}$$

holds on A_n for all n sufficiently large. It follows that for n sufficiently large, we have

$$\sup_{\|h\| \leq M_n} |\pi(\mu_0 + n^{-1/2}(h + Z_{n,(\mu_0, \eta_0)})) - \pi(\mu_0)| \leq \sup_{\|h\| \leq 2M_n} |\pi(\mu_0 + n^{-1/2}h) - \pi(\mu_0)| =: \varepsilon_n \pi(\mu_0)$$

on A_n , where $\varepsilon_n \downarrow 0$ because $\pi \in \Pi$. The density $\pi(\mu)$ is uniformly bounded above by some $\bar{\pi} < +\infty$ because $\pi \in \Pi$. Hence, on A_n , for all n sufficiently large (independent of B),

$$\begin{aligned} \tilde{\pi}_n(B) &\leq \frac{\int_{h: \mu_0 + n^{-1/2}(h + Z_{n,(\mu_0, \eta_0)}) \in \mathcal{M}} \mathbb{I}[h \in B] e^{-\frac{1}{2}h^T \hat{I} h} \pi(\mu_0 + n^{-1/2}(h + Z_{n,(\mu_0, \eta_0)})) dh}{(1 - \varepsilon_n) \pi(\mu_0) \int_{h: \|h\| \leq M_n} e^{-\frac{1}{2}h^T \hat{I} h} dh} \\ &\leq \frac{(1 + \varepsilon_n) \int_{h: \|h\| \leq M_n} \mathbb{I}[h \in B] e^{-\frac{1}{2}h^T \hat{I} h} dh}{(1 - \varepsilon_n) \int_{h: \|h\| \leq M_n} e^{-\frac{1}{2}h^T \hat{I} h} dh} + \frac{\bar{\pi} \int_{h: \|h\| > M_n} e^{-\frac{1}{2}h^T \hat{I} h} dh}{(1 - \varepsilon_n) \pi(\mu_0) \int_{h: \|h\| \leq M_n} e^{-\frac{1}{2}h^T \hat{I} h} dh} \\ &\leq \frac{(1 + \varepsilon_n) \int \mathbb{I}[h \in B] e^{-\frac{1}{2}h^T \hat{I} h} dh}{(1 - \varepsilon_n) \int_{h: \|h\| \leq M_n} e^{-\frac{1}{2}h^T \hat{I} h} dh} + \frac{\bar{\pi} \int_{h: \|h\| > M_n} e^{-\frac{c}{2}\|h\|^2} dh}{(1 - \varepsilon_n) \pi(\mu_0) \int_{h: \|h\| \leq M_n} e^{-\frac{1}{2c}\|h\|^2} dh}, \end{aligned}$$

where the second term on the right-hand side, say $t_{1,n}$, converges to zero by dominated convergence. Moreover,

$$\delta_{1,n} := \frac{\int e^{-\frac{1}{2}h^T \hat{I} h} dh}{\int_{h: \|h\| \leq M_n} e^{-\frac{1}{2}h^T \hat{I} h} dh} - 1 = \frac{\int_{h: \|h\| > M_n} e^{-\frac{1}{2}h^T \hat{I} h} dh}{\int_{h: \|h\| \leq M_n} e^{-\frac{1}{2}h^T \hat{I} h} dh} \leq \frac{\int_{h: \|h\| > M_n} e^{-\frac{c}{2}\|h\|^2} dh}{\int_{h: \|h\| \leq M_n} e^{-\frac{1}{2c}\|h\|^2} dh} \rightarrow 0,$$

on A_n , again by dominated convergence. Hence,

$$\begin{aligned}\tilde{\pi}_n(B) &\leq \frac{(1 + \varepsilon_n)}{(1 - \varepsilon_n)} \frac{\int e^{-\frac{1}{2}h^T \hat{I}h} dh}{\int_{h:\|h\|\leq M_n} e^{-\frac{1}{2}h^T \hat{I}h} dh} \frac{\int \mathbb{I}[h \in B] e^{-\frac{1}{2}h^T \hat{I}h} dh}{\int e^{-\frac{1}{2}h^T \hat{I}h} dh} + t_{1,n} \\ &\leq \frac{(1 + \varepsilon_n)(1 + \delta_{1,n})}{1 - \varepsilon_n} \hat{P}(B) + t_{1,n},\end{aligned}\quad (34)$$

holds on A_n for n sufficiently large (independent of B), where \hat{P} denotes the distribution of $Z^* \sim N(0, \hat{I}^{-1})$, and ε_n , $\delta_{1,n}$ and $t_{1,n}$ are all independent of B and converge (either deterministically or in probability) to zero.

Similarly, for all n sufficiently large (independently of B), on A_n , we have

$$\begin{aligned}\tilde{\pi}_n(B) &\geq \frac{\int_{h:\|h\|\leq M_n} \mathbb{I}[h \in B] e^{-\frac{1}{2}h^T \hat{I}h} \pi(\mu_0 + n^{-1/2}(h + Z_{n,(\mu_0, \eta_0)})) dh}{\int_{h:\mu_0 + n^{-1/2}(h + Z_{n,(\mu_0, \eta_0)}) \in \mathcal{M}} e^{-\frac{1}{2}h^T \hat{I}h} \pi(\mu_0 + n^{-1/2}(h + Z_{n,(\mu_0, \eta_0)})) dh} \\ &\geq \frac{(1 - \varepsilon_n) \pi(\mu_0) \int_{h:\|h\|\leq M_n} \mathbb{I}[h \in B] e^{-\frac{1}{2}h^T \hat{I}h} dh}{\int_{h:\mu_0 + n^{-1/2}(h + Z_{n,(\mu_0, \eta_0)}) \in \mathcal{M}} e^{-\frac{1}{2}h^T \hat{I}h} \pi(\mu_0 + n^{-1/2}(h + Z_{n,(\mu_0, \eta_0)})) dh} \\ &\geq \frac{(1 - \varepsilon_n) \pi(\mu_0) \int_{h:\|h\|\leq M_n} \mathbb{I}[h \in B] e^{-\frac{1}{2}h^T \hat{I}h} dh}{(1 + \delta_{2,n}) \int_{h:\|h\|\leq M_n} e^{-\frac{1}{2}h^T \hat{I}h} \pi(\mu_0 + n^{-1/2}(h + Z_{n,(\mu_0, \eta_0)})) dh},\end{aligned}$$

where

$$\begin{aligned}\delta_{2,n} &:= \left| \frac{\int_{h:\mu_0 + n^{-1/2}(h + Z_{n,(\mu_0, \eta_0)}) \in \mathcal{M}} e^{-\frac{1}{2}h^T \hat{I}h} \pi(\mu_0 + n^{-1/2}(h + Z_{n,(\mu_0, \eta_0)})) dh}{\int_{h:\|h\|\leq M_n} e^{-\frac{1}{2}h^T \hat{I}h} \pi(\mu_0 + n^{-1/2}(h + Z_{n,(\mu_0, \eta_0)})) dh} - 1 \right| \\ &\leq \frac{\bar{\pi} \int_{h:\|h\|>M_n} e^{-\frac{1}{2}h^T \hat{I}h} dh}{(1 - \varepsilon_n) \pi(\mu_0) \int_{h:\|h\|\leq M_n} e^{-\frac{1}{2}h^T \hat{I}h} dh} \leq \frac{\bar{\pi} \int_{h:\|h\|>M_n} e^{-\frac{c}{2}\|h\|^2} dh}{(1 - \varepsilon_n) \pi(\mu_0) \int_{h:\|h\|\leq M_n} e^{-\frac{1}{2c}\|h\|^2} dh} \rightarrow 0.\end{aligned}$$

Hence, on A_n , for all n sufficiently large (independent of B),

$$\begin{aligned}\tilde{\pi}_n(B) &\geq \frac{(1 - \varepsilon_n)}{(1 + \delta_{2,n})(1 + \varepsilon_n)} \frac{\int e^{-\frac{1}{2}h^T \hat{I}h} dh}{\int_{h:\|h\|\leq M_n} e^{-\frac{1}{2}h^T \hat{I}h} dh} \frac{\int_{h:\|h\|\leq M_n} \mathbb{I}[h \in B] e^{-\frac{1}{2}h^T \hat{I}h} dh}{\int e^{-\frac{1}{2}h^T \hat{I}h} dh} \\ &\geq \frac{(1 - \varepsilon_n)(1 - \delta_{1,n})}{(1 + \delta_{2,n})(1 + \varepsilon_n)} \frac{\int_{h:\|h\|\leq M_n} \mathbb{I}[h \in B] e^{-\frac{1}{2}h^T \hat{I}h} dh}{\int e^{-\frac{1}{2}h^T \hat{I}h} dh} \\ &\geq \frac{(1 - \varepsilon_n)(1 - \delta_{1,n})}{(1 + \delta_{2,n})(1 + \varepsilon_n)} \hat{P}(B) - \frac{(1 - \varepsilon_n)(1 - \delta_{1,n})}{(1 + \delta_{2,n})(1 + \varepsilon_n)} \frac{\int_{h:\|h\|>M_n} e^{-\frac{c}{2}\|h\|^2} dh}{\int e^{-\frac{1}{2c}\|h\|^2} dh},\end{aligned}$$

where the second term on the right-hand side, say $t_{2,n}$, converges to zero by dominated convergence. Hence,

$$\tilde{\pi}_n(B) \geq \frac{(1 - \varepsilon_n)(1 - \delta_{1,n})}{(1 + \delta_{2,n})(1 + \varepsilon_n)} \hat{P}(B) - t_{2,n} \quad (35)$$

holds on A_n for n sufficiently large (independent of B), where ε_n , $\delta_{1,n}$, $\delta_{2,n}$, and $t_{2,n}$ are all independent of B and converge (either deterministically or in probability) to zero.

Combining (34) and (35) and using the fact that $P_{n,(\mu_0,\eta_0)}(A_n) \rightarrow 1$, we may deduce

$$\|\tilde{\pi}_n - \hat{P}\|_{TV} \equiv \sup_B \left| \tilde{\pi}_n(B) - \hat{P}(B) \right| \xrightarrow{P_{n,(\mu_0,\eta_0)}} 0,$$

where the supremum is over all Borel sets. Finally, it is easy to deduce (e.g., by using Pinsker's inequality) that the total variation distance between the $N(0, \Sigma_1)$ and $N(0, \Sigma_2)$ distributions converges to zero as $\Sigma_1 \rightarrow \Sigma_2$ for Σ_2 positive definite. It therefore follows by Assumptions 3(iii) and 3(v)b that $\|\hat{P} - P^*\|_{TV} \xrightarrow{P_{n,(\mu_0,\eta_0)}} 0$. ■

Lemma 6 *Suppose that Assumptions 3(i), (iii), (v) and 4(ii) hold and $\pi \in \Pi$. Then there exists a sequence $\{B_n\} \subset \mathbb{R}_+$ with $B_n \uparrow +\infty$, $B_n = o(\sqrt{n})$, such that for all n sufficiently large,*

$$\sup_{h: \|h\| \leq 2B_n} \left\| \sqrt{n} (f(\mu_0 + h/\sqrt{n}) - f(\mu_0)) - \dot{f}_{\mu_0}[h] \right\| \leq \frac{1}{B_n},$$

and

$$P_{n,(\mu_0,\eta_0)} \left(\|\tilde{\pi}_n - P^*\|_{TV} > \frac{1}{B_n^2} \right) \leq \frac{1}{B_n}.$$

Proof of Lemma 6. The result follows by similar arguments to the proof of Lemma 3, using Lemma 5 in place of the (parametric) Bernstein–von Mises theorem. ■

For the next result we invoke Assumption 3(v)a, but we only require $\|\sqrt{n}(\hat{\mu} - \mu_0)\|$ to be bounded in probability under $\{P_{n,(\mu_0,\eta_0)}\}$. Let \mathbb{E}^* denote expectation with respect to $Z^* \sim N(0, I_{(\mu_0,\eta_0)}^{-1})$ independent of $Z_{n,(\mu_0,\eta_0)}$.

Lemma 7 *Suppose Assumptions 3(i), (iii), (v) and 4 hold and $\pi \in \Pi$. Then*

$$\left\| \bar{f}_n - \mathbb{E}^* \left[\dot{f}_{\mu_0} [Z^* + Z_{n,(\mu_0,\eta_0)}] \middle| Z_{n,(\mu_0,\eta_0)} \right] \right\| \xrightarrow{P_{n,(\mu_0,\eta_0)}} 0$$

as $n \rightarrow \infty$.

Proof of Lemma 7. Let $\{B_n\}$ be as in Lemma 6. By a change of variables, we have

$$\begin{aligned}\bar{f}_n &= \int_{h: \|h\| \leq B_n} \sqrt{n} (f(\mu_0 + n^{-1/2}(Z_{n,(\mu_0, \eta_0)} + h)) - f(\mu_0)) d\tilde{\pi}_n(h) \\ &\quad + \int_{\mu: \|\sqrt{n}(\mu - \hat{\mu})\| > B_n} \sqrt{n} (f(\mu) - f(\mu_0)) d\pi_n(\mu) =: \mathcal{J}_{1,n} + \mathcal{J}_{2,n}.\end{aligned}$$

Let A_n denote the event $\|Z_{n,(\mu_0, \eta_0)}\| \leq B_n$, $\|\tilde{\pi}_n - P^*\|_{TV} \leq B_n^{-2}$, and $c \leq \hat{\lambda}_{\min}$, $\hat{\lambda}_{\max} \leq c^{-1}$ for some $c \in (0, 1)$. Assumptions 3(iii) and 3(v) and Lemma 6 imply $\lim_{n \rightarrow \infty} P_{n,(\mu_0, \eta_0)}(A_n^c) = 0$. Again by Lemma 6, on A_n we have

$$\left\| \mathcal{J}_{1,n} - \int_{h: \|h\| \leq B_n} \dot{f}_{\mu_0} [Z_{n,(\mu_0, \eta_0)} + h] d\tilde{\pi}_n(h) \right\| \leq \frac{1}{B_n} \rightarrow 0,$$

where, as in the proof of Lemma 4, we may deduce that

$$\left\| \int_{h: \|h\| \leq B_n} \dot{f}_{\mu_0} [Z_{n,(\mu_0, \eta_0)} + h] d\tilde{\pi}_n(h) - \int_{h: \|h\| \leq B_n} \dot{f}_{\mu_0} [Z_{n,(\mu_0, \eta_0)} + h] dP^*(h) \right\| \leq \frac{2C_f}{B_n} \rightarrow 0,$$

on A_n , and

$$\left\| \int_{h: \|h\| > B_n} \dot{f}_{\mu_0} [Z_{n,(\mu_0, \eta_0)} + h] dP^*(h) \right\| \leq \frac{2C_f \text{tr}(I_{(\mu_0, \eta_0)}^{-1})}{B_n} \rightarrow 0,$$

on A_n , where the latter convergence holds in view of Assumption 3(iii). It follows that

$$\|\mathcal{J}_{1,n} - \mathbb{E}^*[\dot{f}_{\mu_0}[Z^* + Z_{n,(\mu_0, \eta_0)}] | Z_{n,(\mu_0, \eta_0)}]\| \xrightarrow{P_{n,(\mu_0, \eta_0)}} 0.$$

It remains to show that $\|\mathcal{J}_{2,n}\| \xrightarrow{P_{n,(\mu_0, \eta_0)}} 0$. To this end, write

$$\begin{aligned}\mathcal{J}_{2,n} &= \int_{\mu: B_n < \|\sqrt{n}(\mu - \hat{\mu})\| \leq \epsilon\sqrt{n}} \sqrt{n} (f(\mu) - f(\mu_0)) d\pi_n(\mu) \\ &\quad + \int_{\mu: \|\sqrt{n}(\mu - \hat{\mu})\| > \epsilon\sqrt{n}} \sqrt{n} (f(\mu) - f(\mu_0)) d\pi_n(\mu) =: \mathcal{J}_{2,n,a} + \mathcal{J}_{2,n,b},\end{aligned}$$

where, as in the proof of Lemma 4, we choose $\epsilon > 0$ sufficiently small that $\|f(\mu) - f(\mu_0)\| \leq (1 + C_f)\|\mu - \mu_0\|$ holds for μ with $\|\mu - \mu_0\| \leq 2\epsilon$. Then on A_n , we have

$$\|\mu - \mu_0\| \leq \|\mu - \hat{\mu}\| + \|\hat{\mu} - \mu_0\| \leq \|\mu - \hat{\mu}\| + \frac{B_n}{\sqrt{n}},$$

so $\{\mu : \|\mu - \hat{\mu}\| \leq \epsilon\} \subseteq \{\mu : \|\mu - \mu_0\| \leq 2\epsilon\}$ for all n sufficiently large, in which case

$$\|\mathcal{I}_{2,n,a}\| \leq (1+C_f) \left(\int_{\mu: B_n < \|\sqrt{n}(\mu - \hat{\mu})\| \leq \epsilon\sqrt{n}} \|\sqrt{n}(\mu - \hat{\mu})\| d\pi_n(\mu) + B_n \pi_n(\{\mu : B_n < \sqrt{n}(\mu - \hat{\mu})\}) \right).$$

Taking ϵ smaller if necessary, we have $\{\mu : \|\mu - \mu_0\| \leq 2\epsilon\} \subset \mathcal{M}$ in view of Assumption 3(i). Then by dominated convergence and the fact that $c \leq \hat{\lambda}_{\min}$, $\hat{\lambda}_{\max} \leq c^{-1}$ holds on A_n , we have

$$\int_{\mu: B_n < \|\sqrt{n}(\mu - \hat{\mu})\| \leq \epsilon\sqrt{n}} \|\sqrt{n}(\mu - \hat{\mu})\| d\pi_n(\mu) \leq \frac{\int_{h: \|h\| > B_n} \|h\| e^{-\frac{\epsilon}{2}\|h\|^2} \pi(\mu_0 + n^{-1/2}(h + Z_{n,(\mu_0, \eta_0)})) dh}{\int e^{-\frac{1}{2c}\|h\|^2} \pi(\mu_0 + n^{-1/2}(h + Z_{n,(\mu_0, \eta_0)})) dh}$$

and

$$B_n \pi_n(\{\mu : B_n < \sqrt{n}(\mu - \hat{\mu})\}) \leq B_n \frac{\int_{h: \|h\| > B_n} e^{-\frac{\epsilon}{2}\|h\|^2} \pi(\mu_0 + n^{-1/2}(h + Z_{n,(\mu_0, \eta_0)})) dh}{\int e^{-\frac{1}{2c}\|h\|^2} \pi(\mu_0 + n^{-1/2}(h + Z_{n,(\mu_0, \eta_0)})) dh}.$$

It then follows by similar arguments to the proof of Lemma 5 that $\|\mathcal{J}_{2,n,a}\| \xrightarrow{P_{n,(\mu_0, \eta_0)}} 0$.

Finally, for $\mathcal{J}_{2,n,b}$, on A_n we have

$$\|\mathcal{J}_{2,n,b}\| \leq \frac{\sqrt{n} e^{-\frac{cn\epsilon^2}{2}} \int f(\mu) - f(\mu_0) \pi(\mu) d\mu}{n^{-K/2} \int e^{-\frac{1}{2c}\|h\|^2} \pi(\mu_0 + n^{-1/2}(h + Z_{n,(\mu_0, \eta_0)})) dh}.$$

The integral in the numerator is finite by Assumption 4(i) and the integral in the denominator may be shown to be bounded below by a positive constant on A_n by similar arguments to the proof of Lemma 5, from which it follows that $\|\mathcal{J}_{2,n,b}\| \xrightarrow{P_{n,(\mu_0, \eta_0)}} 0$. ■

For the next result, let $Z \sim N(h, I_{\mu_0}^{-1})$ and let \mathbb{E}^* denote expectation with respect to $Z^* \sim N(0, I_{\mu_0}^{-1})$ independent of Z .

Proposition 3 *Suppose that Assumptions 3 and 4 hold and $\pi \in \Pi$. Then for every $h \in \mathbb{R}^K$,*

$$\bar{f}_n \xrightarrow{P_{n,h}} \mathbb{E}^* \left[\dot{f}_{\mu_0} [Z^* + Z] \middle| Z \right].$$

Proof of Proposition 3. Assumption 3(i)-(iii) imply $\mathcal{P} = \{P_{\beta(t)} : t \in \mathcal{M}_{(\mu_0, \eta_0)}\}$ is locally asymptotically normal (van der Vaart, 1998, Example 7.15) and so $\{P_{n,(\mu_0, \eta_0)}\}$ and $\{P_{n,h}\}$ are mutually contiguous (van der Vaart, 1998, Example 6.5). Hence, with $\xrightarrow{P_{n,h}}$ denoting convergence in probability under $\{P_{n,h}\}$, we have by Lemma 7 that

$$\left\| \bar{f}_n - \mathbb{E}^* \left[\dot{f}_{\mu_0} [Z^* + Z_{n,(\mu_0, \eta_0)}] \middle| Z_{n,(\mu_0, \eta_0)} \right] \right\| \xrightarrow{P_{n,h}} 0.$$

Moreover, under Assumption 3(v)a we have

$$Z_{n,(\mu_0,\eta_0)} \stackrel{P_{n,h}}{\rightsquigarrow} Z \sim N(h, I_{(\mu_0,\eta_0)}^{-1}).$$

The result follows by the continuous mapping theorem, noting $z \mapsto \mathbb{E}^*[f_{\mu_0}[Z^* + z]]$ is everywhere continuous by Assumption 4(ii) and Proposition 3.1 of Shapiro (1990). ■

D Bounds on Mistake Probabilities

Definition 1(ii) requires mistake probabilities to vanish faster than $n^{-1/2}$. The main results of this section are Lemma 11, which shows that mistake probabilities for Bayes decisions vanish at rate n^{-1} , and Lemma 12, which presents a result for quasi-Bayes decisions.

D.1 Parametric Models

We first present some preliminary lemmas before proving Lemma 11.

Lemma 8 *Suppose that Assumption 1(i) holds. Then for any $\mu_0 \in \mathcal{M}$ there exists $\epsilon, \epsilon' > 0$ such that*

$$\pi_n(\{\mu : \|\mu - \mu_0\| \geq \epsilon\}) \leq 2\epsilon' \quad (36)$$

implies $\delta_n^(X^n; \pi) \in \underline{\mathcal{D}}_{\mu_0}$.*

Proof of Lemma 8. If $\underline{\mathcal{D}}_{\mu_0} = \mathcal{D}$ then the result is trivial. Now suppose $\underline{\mathcal{D}}_{\mu_0} \subsetneq \mathcal{D}$. By Assumption 1(i), we may choose $\epsilon, \epsilon' > 0$ such that

$$\min_{i \notin \underline{\mathcal{D}}_{\mu_0}} R(i, \mu) - \max_{i \in \underline{\mathcal{D}}_{\mu_0}} R(i, \mu) \geq \epsilon, \quad \mu \in N := \{\mu : \|\mu - \mu_0\| < \epsilon\}.$$

We have

$$\begin{aligned} \max_{i \in \underline{\mathcal{D}}_{\mu_0}} \bar{R}_n(i) &\leq \max_{i \in \underline{\mathcal{D}}_{\mu_0}} \int_N R(i, \mu) d\pi_n(\mu) + \max_{i \in \underline{\mathcal{D}}_{\mu_0}} \int_{N^c} R(i, \mu) d\pi_n(\mu) \\ &\leq \int_N \max_{i \in \underline{\mathcal{D}}_{\mu_0}} R(i, \mu) d\pi_n(\mu) + \pi_n(N^c) \|R\|_\infty, \end{aligned}$$

where $\|R\|_\infty := \sup_{(d,\mu) \in \mathcal{D} \times \mathcal{M}} |R(d, \mu)| < \infty$ by Assumption 1(i), and

$$\begin{aligned} \min_{i \notin \underline{\mathcal{D}}_{\mu_0}} \bar{R}_n(i) &\geq \min_{i \notin \underline{\mathcal{D}}_{\mu_0}} \int_N R(i, \mu) d\pi_n(\mu) + \min_{i \notin \underline{\mathcal{D}}_{\mu_0}} \int_{N^c} R(i, \mu) d\pi_n(\mu) \\ &\geq \int_N \min_{i \notin \underline{\mathcal{D}}_{\mu_0}} R(i, \mu) d\pi_n(\mu) - \pi_n(N^c) \|R\|_\infty. \end{aligned}$$

Hence,

$$\begin{aligned}
\min_{i \notin \underline{\mathcal{D}}_{\mu_0}} \bar{R}_n(i) - \max_{i \in \underline{\mathcal{D}}_{\mu_0}} \bar{R}_n(i) &\geq \int_N \left(\min_{i \notin \underline{\mathcal{D}}_{\mu_0}} R(i, \mu) - \max_{i \in \underline{\mathcal{D}}_{\mu_0}} R(i, \mu) \right) d\pi_n(\mu) - 2\pi_n(N^c) \|R\|_\infty \\
&\geq \pi_n(N)\varepsilon - 2\pi_n(N^c) \|R\|_\infty \\
&= \varepsilon - \pi_n(N^c)(\varepsilon + 2\|R\|_\infty).
\end{aligned}$$

If $\pi_n(N^c) \leq \frac{\varepsilon}{2(\varepsilon + 2\|R\|_\infty)} =: 2\varepsilon'$ holds, then

$$\min_{i \notin \underline{\mathcal{D}}_{\mu_0}} \bar{R}_n(i) - \max_{i \in \underline{\mathcal{D}}_{\mu_0}} \bar{R}_n(i) \geq \varepsilon - \frac{\varepsilon}{2} = \frac{\varepsilon}{2},$$

which implies

$$\max_{i \in \underline{\mathcal{D}}_{\mu_0}} \bar{R}_n(i) \leq \min_{i \notin \underline{\mathcal{D}}_{\mu_0}} \bar{R}_n(i) - \frac{\varepsilon}{2} < \min_{i \notin \underline{\mathcal{D}}_{\mu_0}} \bar{R}_n(i),$$

and hence that $\delta_n^*(X^n; \pi) \in \underline{\mathcal{D}}_{\mu_0}$. ■

Lemma 9 *Suppose that Assumption 1(i) holds. Take any $\mu_0 \in \mathcal{M}$ and $h \in \mathbb{R}^K$, and let $\varepsilon, \varepsilon'$ be as in (36) of Lemma 8. Suppose that there exist tests $\{\phi_n\}$ with the property that $\mathbb{E}_{n,h}[\phi_n(X^n)] \leq C'e^{-nC}$ and $\sup_{\mu: \|\mu - \mu_0\| \geq \varepsilon} \mathbb{E}_{n,\mu}[1 - \phi_n(X^n)] \leq C'e^{-nC}$ for $C, C' > 0$. Then*

$$P_{n,h}(\delta_n^*(X^n; \pi) \notin \underline{\mathcal{D}}_{\mu_0}) \leq \frac{C'e^{-nC}}{\varepsilon'} + \frac{C'e^{-nC/2}}{\varepsilon' \pi(\{\mu : D_{KL}(p_{\mu_0+h/\sqrt{n}} \| p_\mu) < n^{-1}\})} + \frac{4}{nC}.$$

Proof of Lemma 9. In view of Lemma 8, we have

$$P_{n,h}(\delta_n^*(X^n; \pi) \notin \underline{\mathcal{D}}_{\mu_0}) \leq P_{n,h}(\pi_n(\{\mu : \|\mu - \mu_0\| \geq \varepsilon\}) > 2\varepsilon').$$

We proceed by similar arguments to the consistency theorem of [Schwartz \(1965\)](#) but centering along $P_{n,h}$. Write

$$\begin{aligned}
P_{n,h}(\pi_n(\{\mu : \|\mu - \mu_0\| \geq \varepsilon\}) > 2\varepsilon') &\leq P_{n,h}(\phi_n(X^n) > \varepsilon') \\
&\quad + P_{n,h} \left(\frac{(1 - \phi_n(X^n)) \int_{N^c} \prod_{i=1}^n p_\mu(X_i) / p_{\mu_0+h/\sqrt{n}}(X_i) d\pi(\mu)}{\mathcal{J}_n} > \varepsilon' \right),
\end{aligned}$$

where $\mathcal{J}_n = \int \prod_{i=1}^n p_\mu(X_i) / p_{\mu_0+h/\sqrt{n}}(X_i) d\pi(\mu)$ and $N^c = \{\mu : \|\mu - \mu_0\| > \varepsilon\}$. The first probability on the right-hand side is bounded by $C'e^{-nC}/\varepsilon'$ by Markov's inequality. For the second, Lemma 6.26 of [Ghosal and van der Vaart \(2017\)](#) implies that for any $B > 0$ and

$\varepsilon \geq n^{-1}$,¹⁷

$$P_{n,h}(\mathcal{J}_n < \pi(\{\mu : D_{KL}(p_{\mu_0+h/\sqrt{n}}\|p_\mu) < \varepsilon\})e^{-2nB\varepsilon}) \leq \frac{1}{B}.$$

Hence we may deduce

$$\begin{aligned} P_{n,h} \left(\frac{(1 - \phi'_n(X^n)) \int_{N^c} \prod_{i=1}^n p_\mu(X_i)/p_{\mu_0+h/\sqrt{n}}(X_i) d\pi(\mu)}{\mathcal{J}_n} > \epsilon' \right) \\ \leq \frac{\mathbb{E}_{n,h} \left[(1 - \phi'_n(X^n)) \int_{N^c} \prod_{i=1}^n p_\mu(X_i)/p_{\mu_0+h/\sqrt{n}}(X_i) d\pi(\mu) \right]}{\epsilon' \pi(\{\mu : D_{KL}(p_{\mu_0+h/\sqrt{n}}\|p_\mu) < \varepsilon\})e^{-2nB\varepsilon}} + \frac{1}{B}. \end{aligned}$$

The numerator in the first term on the right-hand side may be bounded by $C'e^{-Cn}$ using Fubini's theorem (see, e.g., the proof of Lemma 6.17 of Ghosal and van der Vaart (2017)). Setting $B = Cn/4$ and $\varepsilon = n^{-1}$, the right-hand side of the above display is bounded by

$$\frac{C'e^{-nC/2}}{\epsilon' \pi(\{\mu : D_{KL}(p_{\mu_0+h/\sqrt{n}}\|p_\mu) < n^{-1}\})} + \frac{4}{nC},$$

as required. ■

Lemma 10 *Suppose that Assumption 2(i), (iii)-(iv) hold and $\pi \in \Pi$. Then for $\mu_0 \in \mathcal{M}$ and $h \in \mathbb{R}^K$,*

$$\liminf_{n \rightarrow \infty} n^{K/2} \pi(\{\mu : D_{KL}(p_{\mu_0+h/\sqrt{n}}\|p_\mu) < n^{-1}\}) > 0.$$

Proof of Lemma 10. Assumption 2(iv) implies $D_{KL}(p_{\mu_0+h/\sqrt{n}}\|p_\mu) \leq B\|\mu - \mu_0 - h/\sqrt{n}\|^2$ on a neighborhood of μ_0 for all sufficiently large n , with $B = 2\lambda_{\max}(I_{\mu_0}) < \infty$ by Assumption 2(iii). Hence, for all n sufficiently large, $\|\mu - \mu_0 - h/\sqrt{n}\|^2 < (Bn)^{-1}$ implies $D_{KL}(p_{\mu_0+h/\sqrt{n}}\|p_\mu) < n^{-1}$. Assumption 2(i) and the positivity and continuity of π imply that $\pi(\{\mu : \|\mu - \mu_0 - h/\sqrt{n}\|^2 < (Bn)^{-1}\})$ is bounded below by a multiple of $n^{-K/2}$ for all n sufficiently large, and hence that $\pi(\{\mu : D_{KL}(p_{\mu_0+h/\sqrt{n}}\|p_\mu) < n^{-1}\})$, is bounded below by a multiple of $n^{-K/2}$ for all n sufficiently large. ■

Lemma 11 *Suppose that Assumptions 1(i) and 2 hold and $\pi \in \Pi$. Then there exists $C > 0$ such that*

$$P_{n,h}(\delta_n^*(X^n; \pi) \notin \mathcal{D}_{\mu_0}) \leq \frac{C}{n}.$$

Proof of Lemma 11. We will prove the result by applying Lemma 9. We first establish existence of tests as required by Lemma 9. We proceed as in Propositions 6.1 and 6.2 of Clarke and Barron (1990). Let d_G denote the metric they construct to metrize weak convergence.

¹⁷The result is stated requiring $B > 1$, but inspection of the proof shows any $B > 0$ is valid: restricting to $B > 1$ is simply so the probability bound of $1/B$ is not vacuous.

Let $\hat{P}_n \equiv \hat{P}_n(X^n)$ denote the empirical measure. As shown in the proof of Proposition 6.2 of [Clarke and Barron \(1990\)](#), for any $\varepsilon > 0$ there exists an integer k depending only on ε such that

$$P_{n,\mu} \left(d_G(\hat{P}_n, P_\mu) \geq \varepsilon \right) \leq 2ke^{-n\varepsilon^2}, \quad \mu \in \mathcal{M}. \quad (37)$$

By Assumption 2(v), there exists $\varepsilon' > 0$ such that $\|\mu - \mu_0\| \geq \varepsilon$ implies $d_G(P_\mu, P_{\mu_0}) \geq \varepsilon'$. Consider the test $\phi_n(X^n) = \mathbb{I}[d_G(\hat{P}_n, P_{\mu_0}) > \varepsilon'/2]$. Again using Assumption 2(v), take n large enough that $d_G(P_{\mu_0+h/\sqrt{n}}, P_{\mu_0}) < \varepsilon'/4$. Then by (37), we have

$$\mathbb{E}_{n,h}[\phi_n(X^n)] \leq P_{n,h} \left(d_G(\hat{P}_n, P_{\mu_0+h/\sqrt{n}}) \geq \varepsilon'/4 \right) \leq C'e^{-nC}$$

for suitable C, C' . We may similarly deduce as in the proof of Proposition 6.1 of [Clarke and Barron \(1990\)](#) that for any μ with $\|\mu - \mu_0\| \geq \varepsilon$ we have $\mathbb{E}_{n,\mu}[1 - \phi_n(X^n)] \leq C'e^{-nC}$ for suitable C, C' not depending on μ . Thus, we have verified existence of a sequence of tests with the desired properties.

By Lemma 9, it suffices to show $e^{nc}\pi(\{\mu : D_{KL}(p_{\mu_0+h/\sqrt{n}}\|p_\mu) < n^{-1}\}) \rightarrow +\infty$ for any $c > 0$. This follows immediately from Lemma 10. ■

D.2 Semiparametric Models

Lemma 12 *Suppose that Assumptions 1(i) and 3 hold and $\pi \in \Pi$. Then for any $(\mu_0, \eta_0) \in \mathcal{M} \times \mathcal{H}$,*

$$\sqrt{n} P_{n,h} \left(\delta_n^*(X^n; \pi) \notin \underline{\mathcal{D}}_{\mu_0} \right) \rightarrow 0$$

for all $h \in \mathbb{R}^K$.

Proof of Lemma 12. In view of Lemma 8, we have

$$P_{n,h} \left(\delta_n^*(X^n; \pi) \notin \underline{\mathcal{D}}_{\mu_0} \right) \leq P_{n,h} \left(\pi_n(\{\mu : \|\mu - \mu_0\| \geq \varepsilon\}) > 2\varepsilon' \right)$$

for some $\varepsilon, \varepsilon' > 0$. By Assumption 3(i), choose $C > 0$ such that $\{\mu : \|\mu - \mu_0\| \leq C\} \subset \mathcal{M}$. Taking ε smaller if needed, we may assume $\varepsilon \leq C$. Recall that $\hat{\lambda}_{\min}$ and $\hat{\lambda}_{\max}$ denote the smallest and largest eigenvalues of \hat{I} . Let A_n denote the event upon which $\|\hat{\mu} - \mu_0\| \leq \varepsilon/2$ and $c < \hat{\lambda}_{\min} \leq \hat{\lambda}_{\max} < c^{-1}$ both hold, where the constant $c \in (0, 1)$ is from Assumption 3(iv)b. Note π has a positive, continuous, and bounded Lebesgue density $\pi(\mu)$ on \mathcal{M} , hence $\bar{\pi} :=$

$\sup_{\mu \in \mathcal{M}} \pi(\mu) < \infty$ and $\underline{\pi} := \inf_{\mu: \|\mu - \mu_0\| \leq C} \pi(\mu) > 0$. It follows that on A_n , we have

$$\begin{aligned} \pi_n(\{\mu : \|\mu - \mu_0\| \geq \epsilon\}) &\leq \pi_n(\{\mu : \|\mu - \hat{\mu}\| \geq \epsilon/2\}) \\ &\leq \frac{\int_{\mu: \|\mu - \hat{\mu}\| \geq \epsilon/2} e^{-\frac{1}{2}(\mu - \hat{\mu})^T (n\hat{I})(\mu - \hat{\mu})} d\pi(\mu)}{\int_{\mu: \|\mu - \hat{\mu}\| \leq C - \epsilon/2} e^{-\frac{1}{2}(\mu - \hat{\mu})^T (n\hat{I})(\mu - \hat{\mu})} d\pi(\mu)} \\ &\leq \frac{\bar{\pi} \int_{h: \|h\| \geq \sqrt{n}\epsilon/2} e^{-\frac{\epsilon}{2}\|h\|^2} dh}{\underline{\pi} \int_{h: \|h\| \leq \sqrt{n}(C - \epsilon/2)} e^{-\frac{1}{2\epsilon}\|h\|^2} dh} \rightarrow 0, \end{aligned}$$

where the first inequality is by the reverse triangle inequality, the final inequality is by a change of variables from μ to $h = \sqrt{n}(\mu - \hat{\mu})$, and convergence to zero is by the dominated convergence theorem. We therefore have that $\pi_n(\{\mu : \|\mu - \mu_0\| \geq \epsilon\}) < 2\epsilon'$ holds on A_n for all n sufficiently large, in which case

$$P_{n,h}(\pi_n(\{\mu : \|\mu - \mu_0\| \geq \epsilon\}) > 2\epsilon') \leq P_{n,h}(A_n^c).$$

Finally, $\sqrt{n}P_{n,h}(A_n^c) \rightarrow 0$ by Assumption 3(iv). ■