

ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

*An International Society for the Advancement of Economic
Theory in its Relation to Statistics and Mathematics*

<http://www.econometricsociety.org/>

Econometrica, Vol. 77, No. 5 (September, 2009), 1683–1701

ASYMPTOTICS FOR STATISTICAL TREATMENT RULES

KEISUKE HIRANO

University of Arizona, Tucson, AZ 85721-0108, U.S.A.

JACK R. PORTER

University of Wisconsin, Madison, WI 53706-1393, U.S.A.

The copyright to this Article is held by the Econometric Society. It may be downloaded, printed and reproduced only for educational or research purposes, including use in course packs. No downloading or copying may be done for any commercial purpose without the explicit permission of the Econometric Society. For such commercial purposes contact the Office of the Econometric Society (contact information may be found at the website <http://www.econometricsociety.org> or in the back cover of *Econometrica*). This statement must be included on all copies of this Article that are made available electronically or in any other format.

ASYMPTOTICS FOR STATISTICAL TREATMENT RULES

BY KEISUKE HIRANO AND JACK R. PORTER¹

This paper develops asymptotic optimality theory for statistical treatment rules in smooth parametric and semiparametric models. Manski (2000, 2002, 2004) and Dehejia (2005) have argued that the problem of choosing treatments to maximize social welfare is distinct from the point estimation and hypothesis testing problems usually considered in the treatment effects literature, and advocate formal analysis of decision procedures that map empirical data into treatment choices. We develop large-sample approximations to statistical treatment assignment problems using the limits of experiments framework. We then consider some different loss functions and derive treatment assignment rules that are asymptotically optimal under average and minmax risk criteria.

KEYWORDS: Statistical decision theory, treatment assignment, minmax, minmax regret, Bayes rules, semiparametric models.

1. INTRODUCTION

ONE MAJOR GOAL of treatment evaluation in the social and medical sciences is to provide guidance on how to assign individuals to treatments. For example, a number of studies have examined the problem of “profiling” individuals to identify those likely to benefit from a social program.² Manski (2000, 2002, 2004) and Dehejia (2005) suggested placing the problem within a decision-theoretic framework and specifying a loss function that quantifies the consequences of choosing different treatments under different states of nature. Schlag (2006) and Stoye (2006) derived exact minmax-regret rules for randomized experiments with a discrete covariate and a bounded continuous outcome. Despite these important results, it is difficult to obtain exact optimality results in many empirically relevant settings, in the same way that it is difficult to obtain exactly optimal estimators or hypothesis tests.

In this paper, we develop large sample results to compare treatment rules and show how to construct approximately optimal procedures from efficient estimates of treatment effect parameters. The data could come from a randomized experiment or an observational data source, and we allow for unrestricted outcome and covariate distributions (including continuously distributed covariates). The key requirements are a local asymptotic normality condition and a welfare contrast parameter that is point-identified. When social welfare contrasts are point-identified, there will typically exist many treatment rules that

¹We are grateful to Gary Chamberlain, Michael Jansson, Charles Manski, Ping Yu, numerous seminar participants, a co-editor, and the referees for comments. Porter thanks the National Science Foundation for research support under grant SES-0438123.

²See, for example, Worden (1993), O’Leary, Decker, and Wandner (1998, 2005), Berger, Black, and Smith (2001), and Black, Smith, Berger, and Noel (2003).

are consistent, in the sense that they assign the “better” treatment with probability approaching 1. Our goal here is to make finer comparisons among rules and to base these comparisons on risk rather than conventional statistical criteria that are not tightly connected to the underlying decision problem.

We first study regular parametric models, using a local parametrization so that the problem of determining whether to assign the treatment does not become trivial as the sample size increases. Using Le Cam’s limits of experiments framework (see Le Cam (1986)), we show that the treatment assignment problem is asymptotically equivalent to a simpler problem, in which one observes a single draw from a shifted Gaussian distribution and must decide whether a linear function of the mean vector is greater than zero. We solve the approximate version of the decision problem and then construct a sequence of decision rules in the original problem that asymptotically matches the solution. The forms of optimal rules will depend on the loss function and the way in which risk (expected loss) is aggregated over the parameter space—for example, we could work with average (Bayes) risk or minmax risk. We consider some specific loss functions and show that some simple rules based on efficient parameter estimates are asymptotically optimal under average and minmax risk criteria.

We then extend the results to a semiparametric setting, where the welfare gain of the treatment can be expressed as a regular functional of the unknown distribution. The analysis in this setting mirrors the parametric setting, but involves a Gaussian sequence model instead of a finite-dimensional Gaussian model. We obtain both minmax and average risk optimality results; to define average risk in the semiparametric case, we propose a class of prior weightings on the tangent space. We illustrate our results by showing that Manski’s conditional empirical success rules are asymptotically optimal under certain symmetric loss functions according to both average risk and minmax risk efficiency criteria. All proofs are contained in the online Supplemental Material (Hirano and Porter (2009)).

2. STATISTICAL TREATMENT ASSIGNMENT PROBLEM

Following Manski (2000, 2002, 2004), we consider a social planner, who assigns individuals to different treatments based on their observed background variables. Suppose that a randomly drawn individual has covariates X with probability distribution F_X on a space \mathcal{X} . The set of possible treatment values is $\mathcal{T} = \{0, 1\}$. The planner observes $X = x$ and assigns the individual to treatment 1 according to a treatment rule

$$\delta(x) = \Pr(T = 1|X = x).$$

Let Y_0 and Y_1 denote potential outcomes for the individual, with conditional probability distributions $F_0(\cdot|x)$ and $F_1(\cdot|x)$ on the same space \mathcal{Y} . Given a rule δ , the outcome distribution conditional on $X = x$ is

$$F_\delta(\cdot|x) = \delta(x)F_1(\cdot|x) + (1 - \delta(x))F_0(\cdot|x).$$

For a given outcome distribution F , let the *social welfare* be a functional $W(F)$. We define

$$W_0(x) = W(F_0(\cdot|x)), \quad W_1(x) = W(F_1(\cdot|x)).$$

One optimal rule is then $\delta^*(x) = 1(W_1(x) > W_0(x))$. Of course, this rule is generally infeasible since F_0 and F_1 (and hence W_0 and W_1) are unknown.

We suppose that F_0 and F_1 belong to families of distributions indexed by a parameter $\theta \in \Theta$, where the parameter space could be finite dimensional or infinite dimensional. Let $w_0(\theta, x)$ and $w_1(\theta, x)$ denote the values for $W_0(x)$ and $W_1(x)$ under θ . It will be convenient to work with the welfare contrast

$$g(\theta, x) = w_1(\theta, x) - w_0(\theta, x).$$

We assume that w_0 and g are continuously differentiable in θ for F_X -almost all x .³

Suppose we have some data that are informative about θ , such as data from a randomized experiment or an observational study. For simplicity, we assume that the data $Z^n = (Z_1, \dots, Z_n)$ are independent and identically distributed (i.i.d.) with $Z_i \sim P_\theta$ on some space \mathcal{Z} .⁴ We will consider below a sequence of experiments $\mathcal{E}_n = \{P_\theta^n, \theta \in \Theta\}$ as the sample size grows.

EXAMPLE 2.1: Dehejia (2005) used data from a randomized evaluation comparing the Greater Avenues for Independence (GAIN) program to the standard Aid to Families with Dependent Children (AFDC) program for welfare recipients in Alameda County, California. The outcome of interest is individual earnings in various quarters after the program. Since many welfare recipients had zero earnings, Dehejia used a Tobit model

$$Y_i = \max\{0, \alpha'_1 X_i + \alpha_2 T_i + \alpha'_3 X_i \cdot T_i + \epsilon_i\},$$

where $T_i = 1$ denotes receipt of the experimental program and $\epsilon_i | X_i, T_i \stackrel{\text{i.i.d.}}{\sim} N(0, \sigma^2)$. Dehejia computed posterior distributions based on observation of the n experimental subjects and then produced predictive distributions for a hypothetical $(n + 1)$ th subject to assess different treatment assignment rules. In our notation, $\theta = (\alpha_1, \alpha_2, \alpha_3, \sigma)$ and $Z^n = \{(T_i, X_i, Y_i) : i = 1, \dots, n\}$.

A *randomized statistical treatment rule* is a mapping $\delta: \mathcal{Z}^n \times \mathcal{X} \rightarrow [0, 1]$. We interpret it as the probability of assigning a (future) individual with covariate X to treatment, given past data Z^n :

$$\delta(z^n, x) = \Pr(T = 1 | Z^n = z^n, X = x).$$

³For a discussion of the relationship between the net social welfare and traditional measures of effects of treatments, such as the average treatment effect, see Dehejia (2003).

⁴The i.i.d. assumption could be weakened to allow for dependent data satisfying local asymptotic normality, but at the cost of complicating the arguments below.

Let $L(\delta, \theta, x)$ be some loss function, which specifies penalties for using the rule δ when the true parameter is θ and the future individual's covariate is $X = x$. We will discuss some specific choices for loss below.

Given a loss L , the risk of a rule $\delta(z^n, x)$ under θ is

$$R(\delta, \theta, x) = E_{\theta} L(\delta(Z^n, x), \theta, x) = \int L(\delta(z^n, x), \theta, x) dP_{\theta}^n(z^n).$$

We evaluate risk pointwise in x . In principle, we could integrate the risk over the marginal distribution of X , but this pointwise form fits most naturally with our local asymptotic approximations.

The risk of a decision rule can vary with θ and, typically, there will not exist a rule that uniformly dominates all other rules unless one restricts the class of rules substantially. There are two classic ways to define an ordering over risk functions: one can average the risk of a rule with respect to some prior measure Π on Θ , obtaining a Bayes risk:

$$\int R(\delta, \theta, x) d\Pi(\theta) = \int \int L(\delta(z^n, x), \theta, x) dP_{\theta}^n(z^n) d\Pi(\theta).$$

Alternatively, one can focus on worst-case risk:

$$\sup_{\theta \in \Theta} R(\delta, \theta, x) = \sup_{\theta \in \Theta} \int L(\delta(z^n, x), \theta, x) dP_{\theta}^n(z^n).$$

3. REGULAR PARAMETRIC MODELS

3.1. *Limit Experiment*

In this section we consider regular parametric models, where the likelihood is smooth in a finite-dimensional parameter. To develop asymptotic approximations, we adopt a local parametrization, as is standard in the literature on efficiency of estimators and test statistics. The local parametrization is used to derive an asymptotic description of the treatment assignment problem using the limits of experiments framework (Le Cam (1986)). Although this framework is typically applied to study point estimation and hypothesis testing, it applies much more broadly to general statistical decision problems. In regular parametric models, a simple Gaussian shift model provides an approximation to the original decision problem.

We first reparametrize the model in terms of local alternatives.⁵ The idea is to consider values for θ such that $g(\theta, x)$ is “close” to 0 and there is a nontrivial

⁵Alternatively, we could use large-deviations asymptotics, in analogy with Bahadur efficiency of hypothesis tests. Manski (2004) used finite-sample large-deviations results to bound the risk properties of certain types of treatment assignment rules in a binary-outcome randomized experiment. Puhalskii and Spokoiny (1998) developed a large-deviations version of asymptotic statistical decision theory and applied it to estimation and hypothesis testing.

difficulty distinguishing between the effects of the two treatments as sample size grows. In our setting, for a given value of x , we center the localization around θ_0 such that

$$(3.1) \quad g(\theta_0, x) = 0,$$

and consider parameter sequences of the form $\theta_0 + \frac{h}{\sqrt{n}}$ for $h \in \mathbb{R}^k$. This is the same localization device used in local asymptotic power calculations and efficiency bounds, although here the centering value θ_0 is tied to a particular covariate value x .

Equation (3.1) is not the only case of interest, but for establishing asymptotic optimality, it is the key case to focus on. For combinations of (θ_0, x) such that $g(\theta_0, x) \neq 0$, the treatment that is better at θ_0 will be better for all local alternatives $\theta_0 + h/\sqrt{n}$ asymptotically, and many rules, including the rules we will propose below, will select the better treatment with probability approaching 1. Our localization around a centering value satisfying (3.1) ensures that we are looking at the hardest cases, where it is difficult to determine the best treatment even with large sample sizes.

To simplify the notation, we will suppress the dependence on x in the remainder of the analysis, writing $g(\theta, x)$ as $g(\theta)$ and similarly for other quantities. All results should be interpreted as being stated for a fixed x .⁶

Let Θ be an open subset of \mathbb{R}^k and suppose $\theta_0 \in \Theta$ satisfies Equation (3.1). We assume that the sequence of experiments $\mathcal{E}_n = \{P_\theta^n, \theta \in \Theta\}$ satisfies differentiability in quadratic mean (DQM) at θ_0 : there exists a function $s: \mathcal{Z} \rightarrow \mathbb{R}^k$ such that

$$\begin{aligned} & \int \left[dP_{\theta_0+h}^{1/2}(z) - dP_{\theta_0}^{1/2}(z) - \frac{1}{2}h's(z) dP_{\theta_0}^{1/2}(z) \right]^2 \\ & = o(\|h\|^2) \quad \text{as } h \rightarrow 0. \end{aligned}$$

The function s is the score function associated with the statistical model \mathcal{E}_1 and can usually be calculated as the derivative of the log-likelihood function. Let $I_0 = E_{\theta_0}[ss']$.

The DQM assumption implies that the log-likelihood ratios of the original model converge weakly to the log-likelihood ratios of a multivariate normal experiment and is the basis for the following result, which specializes Theorems 7.2 and 15.1 of Van der Vaart (1998).

PROPOSITION 3.1: *Let Θ be an open subset of \mathbb{R}^k and suppose $\theta_0 \in \Theta$. Let $\mathcal{E}_n = \{P_\theta^n, \theta \in \Theta\}$ satisfy DQM with I_0 nonsingular. Consider a sequence of treatment assignment rules $\delta_n(z^n)$ in the experiments \mathcal{E}_n and let $\beta_n(h) = E_{\theta_0+h/\sqrt{n}}[\delta_n(Z^n)]$. Suppose $\beta_n(h) \rightarrow \beta(h)$ for every h . Then there exists a function $\delta: \mathbb{R}^k \rightarrow [0, 1]$*

⁶We will reintroduce the dependence on x in the last example of the paper, where we examine a specific treatment assignment rule based on conditional sample averages.

such that for every $h \in \mathbb{R}^k$,

$$\beta(h) = \int \delta(\Delta) dN(\Delta|h, I_0^{-1}),$$

where $N(\Delta|h, I_0^{-1})$ is the multivariate normal distribution with mean h and variance I_0^{-1} .

Proposition 3.1 shows that any converging sequence of treatment rules in the original problem is matched by some treatment rule in a simpler experiment where Δ has a shifted normal distribution with known variance. In this sense, the $N(h, I_0^{-1})$ model is a “limit experiment” for the original problem.

3.2. Loss Functions

Having obtained an asymptotic version of the statistical experiment, we need to complete the specification of the decision problem by choosing a loss function, and then examine the limiting forms of loss, risk, and Bayes risk. Generally, we will need to normalize, or modify somewhat, the loss function in the original problem so that limiting versions of risk and Bayes risk are well defined and lead to useful comparisons of treatment rules. A key component of the loss and risk functions that we will consider is the welfare contrast $g(\theta)$. Let \dot{g} be the vector of partial derivatives of g at θ_0 . Then, since $g(\theta_0) = 0$, we have $\sqrt{ng}(\theta_0 + h/\sqrt{n}) \rightarrow \dot{g}'h$ as $n \rightarrow \infty$.

3.2.1. Asymmetric Losses

We consider two loss functions that penalize differently for two types of errors—assigning to treatment 1 when treatment 0 is better and vice versa. The first loss has been used in the literature on hypothesis testing and gives fixed penalties for the two types of errors:

Hypothesis Testing Loss:

$$L^H(\delta, \theta) = \begin{cases} (1 - \delta), & \text{if } g(\theta) > 0, \\ K \cdot \delta, & \text{if } g(\theta) \leq 0, \end{cases}$$

where $K > 0$.

For a rule δ , the loss can be written

$$\begin{aligned} L^H\left(\delta, \theta_0 + \frac{h}{\sqrt{n}}\right) &= 1\left(g\left(\theta_0 + \frac{h}{\sqrt{n}}\right) > 0\right) \\ &\quad + \delta\left[K \cdot 1\left(g\left(\theta_0 + \frac{h}{\sqrt{n}}\right) \leq 0\right) \right. \\ &\quad \left. - 1\left(g\left(\theta_0 + \frac{h}{\sqrt{n}}\right) > 0\right)\right]. \end{aligned}$$

This converges as $n \rightarrow \infty$ for values of h such that $\dot{g}'h \neq 0$. Due to the discontinuity in the loss L^H , the case $\dot{g}'h = 0$ presents a problem for taking limits, but we can define a lower bound limit as

$$L_\infty^H(\delta, h) = 1(\dot{g}'h > 0) + \delta[K \cdot 1(\dot{g}'h < 0) - 1(\dot{g}'h > 0)].$$

For a converging sequence of rules δ_n with $\beta_n(h)$ and $\beta(h)$ as defined in Proposition 3.1, we can define risk and its limiting lower bound as

$$\begin{aligned} R_n^H\left(\delta_n, \theta_0 + \frac{h}{\sqrt{n}}\right) &= 1\left(g\left(\theta_0 + \frac{h}{\sqrt{n}}\right) > 0\right) \\ &\quad + \beta_n(h)\left[K \cdot 1\left(g\left(\theta_0 + \frac{h}{\sqrt{n}}\right) \leq 0\right) \right. \\ &\quad \left. - 1\left(g\left(\theta_0 + \frac{h}{\sqrt{n}}\right) > 0\right)\right], \end{aligned}$$

$$R_\infty^H(\delta, h) = 1(\dot{g}'h > 0) + \beta(h)[K \cdot 1(\dot{g}'h < 0) - 1(\dot{g}'h > 0)].$$

For Bayes risk, let Π be a prior on the parameter space Θ , with a Lebesgue density $\pi(\theta)$ that is positive and continuous at θ_0 . Define

$$B_n^H(\delta_n, \Pi) = \int R_n^H\left(\delta_n, \theta_0 + \frac{h}{\sqrt{n}}\right)\pi\left(\theta_0 + \frac{h}{\sqrt{n}}\right) dh.$$

Since the prior is smooth around θ_0 and places mass zero on h such that $\dot{g}'h = 0$, its limit is

$$B_\infty^H(\delta) = \pi(\theta_0) \int R_\infty^H(\delta, h) dh.$$

Asymmetric Welfare Regret: Tetenov (2007) proposed a loss that extends loss H by penalizing based on the amount of welfare lost by choosing the worse treatment:

$$L^T(\delta, \theta) = \begin{cases} g(\theta)(1 - \delta), & \text{if } g(\theta) > 0, \\ -K\delta g(\theta), & \text{if } g(\theta) \leq 0, \end{cases}$$

with $K > 0$. The case $K = 1$ corresponds to “welfare regret,” which we will discuss further below.

The loss can be written as

$$\begin{aligned} L^T\left(\delta, \theta_0 + \frac{h}{\sqrt{n}}\right) &= g\left(\theta_0 + \frac{h}{\sqrt{n}}\right)1\left(g\left(\theta_0 + \frac{h}{\sqrt{n}}\right) > 0\right)(1 - \delta) \\ &\quad - K \cdot g\left(\theta_0 + \frac{h}{\sqrt{n}}\right)1\left(g\left(\theta_0 + \frac{h}{\sqrt{n}}\right) \leq 0\right)\delta. \end{aligned}$$

Here, the case $\dot{g}'h = 0$ does not present a problem for taking limits, but we do need to normalize by \sqrt{n} so that the limit is nondegenerate:

$$\begin{aligned} &\sqrt{n}L^T\left(\delta, \theta_0 + \frac{h}{\sqrt{n}}\right) \\ &\rightarrow L^T_\infty(\delta, h) = (\dot{g}'h)1(\dot{g}'h > 0)(1 - \delta) - K \cdot (\dot{g}'h)1(\dot{g}'h \leq 0)\delta. \end{aligned}$$

We then define risk and Bayes risk analogously to loss H .

3.2.2. *Welfare Loss*

Since $w_0(\theta) + \delta g(\theta)$ is social welfare, it is natural to use its negative as a loss function:

$$L^W(\delta, \theta) = -w_0(\theta) - \delta \cdot g(\theta).$$

To keep the second term on the right nondegenerate when sample size increases, we would typically scale the loss by \sqrt{n} . However, this will lead the first term, involving w_0 , to diverge.⁷ To obtain nondegenerate limits, we could recenter welfare loss as

$$\begin{aligned} L^R(\delta, \theta) &= L^W(\delta, \theta) - [-w_0(\theta) - 1(g(\theta) > 0)g(\theta)] \\ &= g(\theta)[1(g(\theta) > 0) - \delta]. \end{aligned}$$

This subtracts off the loss associated with the infeasible optimal rule $\delta^* = 1(g(\theta) > 0)$ and leads to what is called regret loss.⁸ Then

$$\sqrt{n}L^R\left(\delta, \theta_0 + \frac{h}{\sqrt{n}}\right) \rightarrow L^R_\infty(\delta, h) = \dot{g}'h[1(\dot{g}'h > 0) - \delta].$$

Note that regret loss is equal to Tetenov’s loss with $K = 1$, so we will not need to treat regret loss separately in the formal results to follow.

The recentering does not affect Bayes rules. Consider a rule δ_n that minimizes

$$\begin{aligned} &B_n^W(\delta_n, \Pi) \\ &= \int \int L^W\left(\delta_n(z^n), \theta_0 + \frac{h}{\sqrt{n}}\right) dP^n_{\theta_0+h/\sqrt{n}}(z^n)\pi\left(\theta_0 + \frac{h}{\sqrt{n}}\right) dh. \end{aligned}$$

Clearly, adding any function $f(\theta)$ of the parameter to the loss function does not change the solution. Hence the minimizer of B_n^W will also minimize Bayes

⁷This is the motivation for Assumption 1 in the Supplemental Material.

⁸Savage (1951) suggested this type of centering in a discussion of the minmax criterion. The label “regret” is standard in the decision theory literature.

risk for regret loss. However, the recentering does affect the minmax solution. In this sense, only certain combinations of loss function and optimality criterion have nontrivial asymptotic approximations.

3.3. Asymptotic Optimality

We have approximated the sequence of statistical experiments by a Gaussian one and we have approximated loss functions by certain asymptotic versions. If we can find an optimal rule (according to some criterion) in this limiting version of the decision problem, then it will serve as a benchmark for how well any sequence of decision rules can perform in the original problem. Typically, the solution will also suggest the form of a sequence of decision rules that asymptotically match the optimal rule. We develop results for two standard optimality concepts: average and minmax risk.

3.3.1. Average Risk Optimality

First, consider the average (Bayes) risk criterion in the limiting Gaussian model. When the limiting loss function is given by $L_\infty(\delta, h)$, we wish to find the rule δ that minimizes

$$B_\infty(\delta) = \int R_\infty(\delta, h) dh = \int \int L_\infty(\delta(\Delta), h) dN(\Delta|h, I_0^{-1}) dh.$$

Directly minimizing this expression would involve searching over the space of decision rules, but the problem can be simplified by reversing the order of integration and noting that, for each Δ , the solution will minimize

$$\int L_\infty(\delta(\Delta), h) \exp(-(h - \Delta)'I_0(h - \Delta)/2) dh,$$

which is equivalent to minimizing posterior expected risk, where the posterior has $h|\Delta \sim N(\Delta, I_0^{-1})$.

Let $\sigma_g^2 = \dot{g}'I_0^{-1}\dot{g}$. In the Supplemental Material, we show that the average risk optimal rule for hypothesis testing loss L_∞^H is

$$\delta(\Delta) = 1\left(\frac{\dot{g}'\Delta}{\sigma_g} > c^{H,B}\right), \quad \text{where} \quad c^{H,B} = \Phi^{-1}\left(\frac{K}{1+K}\right).$$

For asymmetric welfare regret loss, average risk is minimized by the rule

$$\delta(\Delta) = 1\left(\frac{\dot{g}'\Delta}{\sigma_g} > c^{T,B}\right), \quad \text{where} \quad c^{T,B} \text{ solves } c = \frac{(K-1)\phi(c)}{\Phi(c) + K\Phi(-c)}.$$

For both L_∞^H and L_∞^T , the cutoff is equal to 0 when $K = 1$.

Both optimal rules in the limiting Gaussian model have a simple cutoff form, which suggests how to construct rules in the original problem that are asymptotically equivalent. Let $\hat{\theta}_n$ be an estimator in the original sequence of models that is best regular:

$$(3.2) \quad \sqrt{n}(\hat{\theta}_n - \theta_0 - h/\sqrt{n}) \overset{h}{\rightsquigarrow} N(0, I_0^{-1}) \quad \forall h \in \mathbb{R}^k,$$

where $\overset{h}{\rightsquigarrow}$ denotes convergence in distribution under the sequence of probability measures $P_{\theta_0+h/\sqrt{n}}^n$. Both the maximum likelihood estimator and the Bayesian posterior mean would usually satisfy this condition. If we also have a consistent estimator $\hat{\sigma}_g$ of σ_g , then the feasible decision rule

$$(3.3) \quad \delta_n^{H,B}(Z^n) = 1\left(\frac{\sqrt{n}g(\hat{\theta}_n)}{\hat{\sigma}_g} > c^{H,B}\right)$$

will have limiting distributions that match $1(\dot{g}'\Delta/\sigma_g > c^{H,B})$ for every h . For Tetenov's loss, we define $\delta_n^{T,B}$ analogously. These two decision rule are asymptotically optimal for average risk:

THEOREM 3.2: *Suppose the conditions of Proposition 3.1 are satisfied, $g(\theta_0) = 0$, $g(\theta)$ is differentiable at θ_0 , and the prior measure Π admits a density π with respect to Lebesgue measure that is continuous and positive at θ_0 . Suppose $\hat{\theta}_n$ is a best regular estimator satisfying Equation (3.2) and $\hat{\sigma}_g \xrightarrow{p} \sigma_g$ under θ_0 . Then*

$$\lim_{n \rightarrow \infty} B_n^H(\delta_n^{H,B}, \Pi) = \inf_{\delta_n \in \mathcal{D}} \liminf_{n \rightarrow \infty} B_n^H(\delta_n, \Pi)$$

and

$$\lim_{n \rightarrow \infty} \sqrt{n}B_n^T(\delta_n^{T,B}, \Pi) = \inf_{\delta_n \in \mathcal{D}} \liminf_{n \rightarrow \infty} \sqrt{n}B_n^T(\delta_n, \Pi),$$

where \mathcal{D} denotes the set of all sequences of decision rules that converge in the sense of Proposition 3.1.

REMARK: This result focuses on the case where $g(\theta_0) = 0$. When $g(\theta_0) \neq 0$, one treatment is always preferred under all local parameters, so any rule choosing the appropriate treatment with probability approaching 1 will be asymptotically optimal. In particular, the rules given in Theorem 3.2 remain optimal when $g(\theta_0) \neq 0$ under suitable rate normalizations for the Bayes risk. Similar remarks apply to the other asymptotic optimality results to follow.

Theorem 3.2 shows that a simple rule, which replaces θ by an efficient estimator $\hat{\theta}_n$ and replaces σ_g by a consistent estimator, is approximately Bayes

optimal. If the posterior distribution is tractable, we could also solve the finite-sample Bayes problem directly. Not surprisingly, this will also be asymptotically optimal:

COROLLARY 3.3: *For $j = H, T$, let $\delta_n^{j, \text{Bayes}} = \arg \min_{\delta} B_n^j(\delta, \Pi)$. If the arg min does not exist for any n , let $\delta_n^{j, \text{Bayes}}$ be any rule such that $B_n^j(\delta_n^{j, \text{Bayes}}, \Pi) \leq B_n^j(\delta_n^{j, B}, \Pi)$. Then Theorem 3.2 also holds with $\delta_n^{j, B}$ replaced by $\delta_n^{j, \text{Bayes}}$.*

3.3.2. Minmax Optimality

Next, consider the minmax criterion in the limiting Gaussian model. We wish to solve the functional minimization problem $\inf_{\delta} \sup_h R_{\infty}(\delta, h)$ over the class of all decision rules, a difficult task in general. However, the structure of our problem can be used to simplify the solution. We consider “slices” of the parameter space constructed in the following way: fix an h_0 such that $\dot{g}'h_0 = 0$ and for any $b \in \mathbb{R}$, define

$$h_1(b, h_0) = h_0 + \frac{b}{\dot{g}'I_0^{-1}\dot{g}} I_0^{-1} \dot{g}.$$

Along each slice, the quantity $b = \dot{g}'h_1$ gives the welfare contrast. It turns out that for many loss functions of interest, rules of the form $\delta_c = 1(\dot{g}'\Delta > c)$ for $c \in \mathbb{R}$ form an essential complete class on each slice, so that it is sufficient to search among cutoff rules to solve the minmax problem along a slice. Furthermore, when the loss function only depends on $\dot{g}'h$, the same cutoff value c solves the minmax problem along each slice and leads to a minmax rule over the entire parameter space.

THEOREM 3.4: *Suppose that $\Delta \sim N(h, I_0^{-1})$ for $h \in \mathbb{R}^k$, and consider a decision problem with action space $\{0, 1\}$ and loss $L(a, h)$ such that for all h with $\dot{g}'h \neq 0$,*

$$[L(1, h) - L(0, h)](\dot{g}'h) < 0.$$

(i) *For any randomized decision rule $\tilde{\delta}(\Delta)$ and any fixed $h_0 \in \mathbb{R}^k$, there exists a rule of the form $\delta_c(\Delta) = 1(\dot{g}'\Delta > c)$ which is at least as good as $\tilde{\delta}$ on the subspace $\{h_1(b, h_0) : b \in \mathbb{R}\}$.*

(ii) *Additionally, suppose $L(a, h)$ depends on h only through $\dot{g}'h$.⁹ If a minmax decision rule exists, then $\delta_{c^*}(\Delta)$ is minmax for some c^* . Moreover, the optimal value c^* can be obtained by solving $\inf_c \sup_b E_{h_1(b, 0)} L(\delta_c, h_1(b, 0))$.*

The condition $[L(h, 1) - L(h, 0)](\dot{g}'h) < 0$ requires that the loss impose greater penalties for incorrect assignment, and is satisfied by losses L_{∞}^H and L_{∞}^T .

⁹There exists a function L_g such that $L(a, h) = L_g(a, \dot{g}'h)$ for all a, h .

Part (i) of this result is a mild extension of the essential complete class theorem of Karlin and Rubin (1956). Part (ii) provides a simple method for constructing a minmax rule.

Using (ii), the minmax rule for loss L_∞^H is derived in the Supplemental Material and is a cutoff rule $1(\hat{g}'\Delta/\sigma_g > c^{H,M})$, where $c^{H,M} = c^{H,B}$. In the case of loss L_∞^T , Tetenov (2007) provided a solution in the scalar normal case with known variance, that extends to our multivariate setting in light of Theorem 3.4(ii). The minmax rule for loss L_∞^T is a cutoff rule where the cutoff $c^{T,M}$ is the solution to

$$\sup_{b \leq 0} (-K \cdot b \cdot \Phi(b - c^{T,M})) = \sup_{b > 0} b \cdot \Phi(c^{T,M} - b).$$

As with the Bayes criterion, if either loss is symmetric ($K = 1$), then the optimal cutoff is zero. For general K , the minmax criterion and the Bayes risk criterion lead to the same optimal decision rules for loss L_∞^H , but interestingly not for loss L_∞^T . These limit experiment solutions lead to the following asymptotic minmax result:

THEOREM 3.5: *Suppose the conditions of Proposition 3.1 are satisfied, $g(\theta_0) = 0$, and $g(\theta)$ is differentiable at θ_0 . Suppose $\hat{\theta}_n$ is a best regular estimator satisfying Equation (3.2) and $\hat{\sigma}_g \xrightarrow{p} \sigma_g$ under θ_0 . Let $\delta_n^{H,M}$ and $\delta_n^{T,M}$ be defined analogously to Equation (3.3). Then these rules are locally asymptotically minmax:*

$$\begin{aligned} & \sup_J \liminf_{n \rightarrow \infty} \sup_{h \in J} R_n^H \left(\delta_n^{H,M}, \theta_0 + \frac{h}{\sqrt{n}} \right) \\ &= \inf_{\delta_n \in \mathcal{D}} \sup_J \liminf_{n \rightarrow \infty} \sup_{h \in J} R_n^H \left(\delta_n, \theta_0 + \frac{h}{\sqrt{n}} \right) \end{aligned}$$

and

$$\begin{aligned} & \sup_J \liminf_{n \rightarrow \infty} \sup_{h \in J} \sqrt{n} R_n^T \left(\delta_n^{T,M}, \theta_0 + \frac{h}{\sqrt{n}} \right) \\ &= \inf_{\delta_n \in \mathcal{D}} \sup_J \liminf_{n \rightarrow \infty} \sup_{h \in J} \sqrt{n} R_n^T \left(\delta_n, \theta_0 + \frac{h}{\sqrt{n}} \right), \end{aligned}$$

where the outer supremum is over all finite subsets J of \mathbb{R}^k .

REMARK: Although the average and minmax asymptotic risk comparisons allow h , and hence $\theta_0 + h/\sqrt{n}$, to take on arbitrary values, this reparametrization has important consequences. The localization reduces the statistical information to have an approximately Gaussian form (a type of asymptotic sufficiency) and leads to limiting risk functions that depend on the parameter only

through $\dot{g}'h$. These two effects lead to the simplification of the general decision problem, with the trade-off that only certain combinations of loss function and optimality notion lead to nontrivial comparisons of decision procedures.

4. SEMIPARAMETRIC MODELS

4.1. Gaussian Sequence Limit Experiment

Empirical studies of treatment effects often use nonparametric or semiparametric specifications to allow for more flexibility in the modeling of treatment effects. In this section, we extend the results from the previous section to models with an infinite-dimensional parameter space.

Suppose Z^n consists of an i.i.d. sample of size n drawn from a probability measure $P \in \mathcal{P}$, where \mathcal{P} is the set of probability measures defined by the underlying semiparametric model. In some cases, the set \mathcal{P} will include all distributions satisfying certain weak conditions (so that the model is nonparametric); in other cases, the form of the semiparametric model may restrict the feasible distributions in \mathcal{P} .

We fix $P_0 \in \mathcal{P}$, and define local alternatives to P_0 in a standard way following Van der Vaart (1991a). Consider subsets of the form $\{P_{t,h} : t \in (0, \eta)\} \subset \mathcal{P}$, where $\eta > 0$ and h is a real-valued measurable function on \mathcal{Z} satisfying

$$(4.1) \quad \int \left[\frac{1}{t} (dP_{t,h}^{1/2} - dP_0^{1/2}) - \frac{1}{2} h dP_0^{1/2} \right]^2 \rightarrow 0 \quad \text{as } t \downarrow 0.$$

Each subset $\{P_{t,h} : t \in (0, \eta)\}$ is then a smooth one-dimensional submodel (or path) of \mathcal{P} . Given P_0 , the collection of such paths will be denoted $\mathcal{P}(P_0)$. The function h is the score function associated with the submodel and satisfies $\int h dP_0 = 0$ and $\int h^2 dP_0 < \infty$.

For fixed t and h , $P_{t/\sqrt{n},h}$ is a sequence of measures that approaches P_0 as $n \rightarrow \infty$. It will be enough to consider the sequences $P_{1/\sqrt{n},h}$, so we can view each h as a local parameter, in analogy with the parametric case. Let the tangent set $T(P_0) \subset L_2(P_0)$ be the set of (equivalence classes of) functions h satisfying Equation (4.1).

We will assume that $T(P_0)$ is a separable linear space, so that $\overline{T(P_0)}$ is a separable Hilbert space with the usual inner product and norm for $L_2(P_0)$. Let ϕ_1, ϕ_2, \dots denote any orthonormal basis of $\overline{T(P_0)}$. We identify $\overline{T(P_0)}$ with an l_2 space in the usual way, through the isomorphism $h \mapsto (h_1, h_2, \dots)$ with $h_j = \langle h, \phi_j \rangle$, so that $h(\cdot) = \sum_j h_j \phi_j(\cdot)$.

Again, we use g to denote the difference in social welfare $W_1(x) - W_0(x)$. For a probability measure $P \in \mathcal{P}$, we denote this welfare contrast by $g(P, x)$, or $g(P)$ for short. We assume that there exists a continuous linear map $\dot{g} : T(P_0) \rightarrow \mathbb{R}$ such that

$$(4.2) \quad \frac{1}{t} (g(P_{t,h}) - g(P_0)) \rightarrow \dot{g}(h) \quad \text{as } t \downarrow 0$$

for every path in $\mathcal{P}(P_0)$.¹⁰ This implies

$$\sqrt{n}(g(P_{1/\sqrt{n},h}) - g(P_0)) \rightarrow \dot{g}(h).$$

By the Riesz representation theorem, the functional $\dot{g}(\cdot)$ can be associated with an element $\dot{g} \in \overline{T(P_0)}$ such that $\dot{g}(h) = \langle \dot{g}, h \rangle$ for all $h \in T(P_0)$. (This parallels the notation $\dot{g}'h$ in the parametric case.) Assume $\|\dot{g}\|^2 = \langle \dot{g}, \dot{g} \rangle > 0$.

Van der Vaart (1991a) showed that an asymptotic representation theorem similar to the parametric case holds, where the shifted multivariate Gaussian limit experiment is replaced by an infinite (shifted) Gaussian sequence. This leads to the following result for treatment rules:

PROPOSITION 4.1: *Let $\mathcal{E}_n = \{P_{1/\sqrt{n},h}^n : h \in T(P_0)\}$ satisfy Equation (4.1). Consider a sequence of treatment rules $\delta_n(z^n)$ in the experiments \mathcal{E}_n and let $\beta_n(h) = \int \delta_n dP_{1/\sqrt{n},h}$. Suppose $\beta_n(h) \rightarrow \beta(h)$ for every h . Then there exists a function δ such that $\beta(h) = E_h[\delta(\Delta_1, \Delta_2, \dots)]$, where $(\Delta_1, \Delta_2, \dots)$ is a sequence of independent random variables with $\Delta_j \stackrel{h}{\sim} N(h_j, 1)$.*

4.2. Semiparametric Optimality

We consider loss functions that are analogs of losses H and T in the parametric case, where the parameter θ is replaced by $P \in \mathcal{P}$. We denote these by L^H and L^T as before, in a slight abuse of notation. Then the limiting versions of the loss functions will have the same form as before, with $\langle \dot{g}, h \rangle$ replacing $\dot{g}'h$.

Defining and working with average risk is more complicated in infinite-dimensional models. In the parametric case, simple conditions ensure that a prior measure behaves locally like Lebesgue measure. In infinite product spaces, however, there is no natural analog of Lebesgue measure, and the asymptotic properties of Bayes procedures can be quite sensitive to the choice of prior and the specific model at hand (see, for example, Diaconis and Freedman (1986)). Instead of working with some fixed prior on the space \mathcal{P} , we define a prior on the tangent space and compare procedures by their average risk with respect to this prior.¹¹

It is useful to choose the orthonormal basis ϕ_1, ϕ_2, \dots so that the welfare contrast $\langle \dot{g}, h \rangle$ is attached to the leading term. Let $\phi_1 = \dot{g}/\|\dot{g}\|$ and let ϕ_2, ϕ_3, \dots be an orthonormal basis for the orthocomplement of the space spanned by ϕ_1 . We can view the Gaussian sequence $\Delta_1, \Delta_2, \dots$ in Proposition 4.1 as being defined relative to this choice of orthonormal basis. In particular, under h , $\Delta_1 \sim N(\langle \dot{g}, h \rangle / \|\dot{g}\|, 1)$ and $\Delta_2, \Delta_3, \dots$ have distributions that do not depend on the value of $\langle \dot{g}, h \rangle$.

¹⁰Van der Vaart (1991b) provided a thorough discussion of this differentiability notion, which is related to Hadamard differentiability.

¹¹In a different setting, Andrews and Ploberger (1994) used priors on local parameters to define local average power optimality of tests.

Define a prior measure for (h_1, h_2, \dots) by $\Pi = \lambda \times \rho$, where λ is Lebesgue measure on the real line and ρ is some finite or σ -finite measure on \mathcal{I}^2 . Since Π is a finite product of σ -finite measures on separable spaces, it is well defined.

Define $\Delta = (\Delta_1, \Delta_2, \dots)$, let R_∞ be the associated risk function, and let

$$B_\infty(\delta(\Delta), \Pi) = \int R_\infty(\delta(\Delta), h) d\Pi(h).$$

Suppose that the loss function only depends on h through $\langle \dot{g}, h \rangle$, so we can write (with slight abuse of notation) $L_\infty(\delta(\Delta), \langle \dot{g}, h \rangle)$ (see the supposition of Theorem 3.4(ii)). By interchanging the order of integration, it follows that the Bayes rule can be obtained by minimizing, for each Δ ,

$$\begin{aligned} & \int L_\infty(\delta(\Delta), u) dN(u|\tilde{g}, \sigma_g^2) \\ &= \int L_\infty(0, u) dN(u|\tilde{g}, \sigma_g^2) \\ & \quad + \delta(\Delta) \int (L_\infty(1, u) - L_\infty(0, u)) dN(u|\tilde{g}, \sigma_g^2), \end{aligned}$$

where $\tilde{g} = \|\dot{g}\|\Delta_1$ and $\sigma_g^2 = \|\dot{g}\|^2$. The optimal rule does not depend on $\Delta_2, \Delta_3, \dots$ or on ρ . From here, the analysis is essentially the same as in the parametric case. The optimal Bayes rules in the Gaussian sequence model for L_∞^H and L_∞^T can be written

$$1(\Delta_1 > c^{H,B}), \quad 1(\Delta_1 > c^{T,B}),$$

where $c^{H,B}$ and $c^{T,B}$ are the same constants as in the parametric case.

To construct asymptotic approximation results, define

$$R_n^H(\delta_n, P_{1/\sqrt{n}, h}) = \int L^H(\delta_n(z^n), P_{1/\sqrt{n}, h}) dP_{1/\sqrt{n}, h}^n(z^n),$$

$$B_n^H(\delta_n, \Pi) = \int R_n^H(\delta_n, P_{1/\sqrt{n}, h}) d\Pi(h),$$

and similarly for loss L^T . Then we have the following theorem, which shows how to construct optimal rules from a best regular estimator:

THEOREM 4.2: *Suppose the conditions for Proposition 4.1 are satisfied, $g(P_0) = 0$, g satisfies Equation (4.2), $\hat{\sigma}_g \xrightarrow{P} \|\dot{g}\|$ under P_0 , and $\hat{g}_n(Z^n)$ satisfies:*

$$(4.3) \quad \sqrt{n}(\hat{g}_n(Z^n) - g(P_{1/\sqrt{n}, h})) \overset{h}{\rightsquigarrow} N(0, \|\dot{g}\|^2)$$

for all $h \in T(P_0)$, where $\overset{h}{\rightsquigarrow}$ denotes convergence in distribution under $P_{1/\sqrt{n},h}$. Let

$$\delta_n^{H,B}(Z^n) = 1\left(\frac{\sqrt{n}\hat{g}_n}{\hat{\sigma}_g} > c^{H,B}\right), \quad \delta_n^{T,B}(Z^n) = 1\left(\frac{\sqrt{n}\hat{g}_n}{\hat{\sigma}_g} > c^{T,B}\right).$$

Then

$$\lim_{n \rightarrow \infty} B_n^H(\delta_n^{H,B}, \Pi) = \inf_{\delta_n \in \mathcal{D}} \liminf_{n \rightarrow \infty} B_n^H(\delta_n, \Pi)$$

and

$$\lim_{n \rightarrow \infty} \sqrt{n}B_n^T(\delta_n^{T,B}, \Pi) = \inf_{\delta_n \in \mathcal{D}} \liminf_{n \rightarrow \infty} \sqrt{n}B_n^T(\delta_n, \Pi),$$

where \mathcal{D} denotes the set of all sequences of decision rules that converge in the sense of Proposition 4.1.

REMARK: Theorem 4.2 has a different interpretation than the parametric Bayes result in Theorem 3.2. Here, the prior Π is a weighting on the tangent space of P_0 , so its influence does not disappear as the sample size grows large. We use it here to compare different rules which are derived from asymptotic considerations. In general, a sequence of Bayes rules derived from some fixed prior on \mathcal{P} will not be optimal under our criterion.

The slicing argument used for the minmax analysis of the parametric case also extends to the infinite-dimensional case. Our choice of basis proves convenient, because $\langle \dot{g}, h \rangle$ depends on h only through its first term h_1 : $\langle \dot{g}, h \rangle = \|\dot{g}\|h_1$.

THEOREM 4.3: Let $\Delta = (\Delta_1, \Delta_2, \dots)$ have $\Delta_j \overset{\text{ind}}{\rightsquigarrow} N(h_j, 1)$ under $h = (h_1, h_2, \dots) \in T(P_0)$. Let the action space be $\{0, 1\}$ and for all h such that $\langle \dot{g}, h \rangle \neq 0$, let the loss $L(a, h)$ satisfy

$$[L(1, h) - L(0, h)]\langle \dot{g}, h \rangle < 0.$$

(i) Then for any randomized decision rule $\tilde{\delta}(\Delta)$ and $(0, h_2, h_3, \dots) \in T(P_0)$, there exists a rule of the form $\delta_c(\Delta) = 1(\Delta_1 > c)$ which is at least as good as $\tilde{\delta}$ on the one-dimensional subspace $\{(b, h_2, h_3, \dots) : b \in \mathbb{R}\}$.

(ii) Additionally, suppose $L(a, (h_1, h_2, \dots))$ only depends on (h_1, h_2, h_3, \dots) through $\langle \dot{g}, h \rangle$.¹² If a minmax decision rule exists, then $\delta_{c^*}(\Delta)$ is minmax for some c^* . Moreover, the optimal value c^* can be obtained by solving $\inf_c \sup_b E_{h=(b,0,0,\dots)} L(\delta_c, h)$.

¹²This supposition is equivalent to assuming that $L(a, (h_1, h_2, \dots))$ does not depend on h_2, h_3, \dots for $a = 0, 1$.

Theorem 4.3 leads to the optimal rules $1(\Delta_1 > c^{H,M})$ and $1(\Delta_1 > c^{T,M})$, where the constants are the same as in the parametric case. If we can match the distribution of Δ_1 asymptotically, we can obtain asymptotic minmax optimality:

THEOREM 4.4: *Suppose the conditions for Proposition 4.1 are satisfied, $g(P_0) = 0$, g satisfies Equation (4.2), $\hat{\sigma}_g \xrightarrow{P} \|\dot{g}\|$ under P_0 , and \hat{g}_n is a best regular estimator for $g(P)$ satisfying Equation (4.3). Let*

$$\delta_n^{H,M}(Z^n) = 1\left(\frac{\sqrt{n}\hat{g}_n}{\hat{\sigma}_g} > c^{H,M}\right), \quad \delta_n^{T,M}(Z^n) = 1\left(\frac{\sqrt{n}\hat{g}_n}{\hat{\sigma}_g} > c^{T,M}\right).$$

Then

$$\sup_J \liminf_{n \rightarrow \infty} \sup_{h \in J} R_n^H(\delta_n^{H,M}, P_{1/\sqrt{n},h}) = \inf_{\delta_n \in \mathcal{D}} \sup_J \liminf_{n \rightarrow \infty} \sup_{h \in J} R_n^H(\delta_n, P_{1/\sqrt{n},h})$$

and

$$\begin{aligned} &\sup_J \liminf_{n \rightarrow \infty} \sup_{h \in J} \sqrt{n}R_n^T(\delta_n^{T,M}, P_{1/\sqrt{n},h}) \\ &= \inf_{\delta_n \in \mathcal{D}} \sup_J \liminf_{n \rightarrow \infty} \sup_{h \in J} \sqrt{n}R_n^T(\delta_n, P_{1/\sqrt{n},h}), \end{aligned}$$

where the outer supremum is over finite subsets J of $T(P_0)$.

To close, we illustrate how this result applies to Manski’s conditional empirical success rule.

EXAMPLE 4.5: Suppose that $W_0(x) = 0$ and that we observe a random sample (X_i, Y_i) , $i = 1, \dots, n$, where X_i has a finitely supported distribution and $Y_i|X_i$ has conditional distribution $F_1(y|x)$. The social welfare contrast is the functional $g(x, F_1) = \int w(y) dF_1(y|x)$. The conditional distribution function F_1 is unknown and the set of possible cumulative distribution functions \mathcal{P} is the largest set satisfying

$$\sup_{F_1 \in \mathcal{P}} E[|w(Y)|^2 | X = x] < \infty.$$

The conditional empirical success rule of Manski (2004) can be expressed as

$$\hat{\delta}_n(x) = 1(\hat{g}_n(x) > 0),$$

where

$$\hat{g}_n(x) := \frac{\sum_{i=1}^n w(Y_i) \cdot 1(X_i = x)}{\sum_{i=1}^n 1(X_i = x)}.$$

The estimator $\hat{g}_n(x)$ is an asymptotically efficient estimator of $g(x, F_1)$ (Bickel, Klaassen, Ritov, and Wellner (1993, pp. 67–68)). Therefore, $\hat{\delta}_n$ is asymptotically minmax and asymptotically Bayes optimal for both losses H and T under $K = 1$.

This result extends easily to the case where $W_0(x)$ is not known; then $\hat{g}_n(x)$ would be a difference of conditional mean estimates for outcomes under treatments 1 and 0.

REFERENCES

- ANDREWS, D. W. K., AND W. PLOBERGER (1994): “Optimal Tests When a Nuisance Parameter Is Present Only Under an Alternative,” *Econometrica*, 62 (6), 1383–1414. [1696]
- BERGER, M. C., D. BLACK, AND J. SMITH (2001): “Evaluating Profiling as a Means of Allocating Government Services,” in *Econometric Evaluation of Labour Market Policies*, ed. by M. Lechner and F. Pfeiffer. Heidelberg: Physica-Verlag, 59–84. [1683]
- BICKEL, P. J., C. A. KLAASEN, Y. RITOV, AND J. A. WELLNER (1993): *Efficient and Adaptive Estimation for Semiparametric Models*. New York: Springer-Verlag. [1700]
- BLACK, D., J. SMITH, M. BERGER, AND B. NOEL (2003): “Is the Threat of Training More Effective Than Training Itself? Experimental Evidence From the UI System,” *American Economic Review*, 93 (4), 1313–1327. [1683]
- DEHEJIA, R. (2003): “When Is ATE Enough? Risk Aversion and Inequality Aversion in Evaluating Training Programs,” Working Paper, Columbia University. [1685]
- (2005): “Program Evaluation as a Decision Problem,” *Journal of Econometrics*, 125, 141–173. [1683,1685]
- DIACONIS, P., AND D. FREEDMAN (1986): “On the Consistency of Bayes Estimates,” *The Annals of Statistics*, 14 (1), 1–26. [1696]
- HIRANO, K., AND J. PORTER (2009): “Supplement to ‘Asymptotics for Statistical Treatment Rules,’” *Econometrica Supplemental Material*, 77, http://www.econometricsociety.org/ecta/Supmat/6630_Proofs.pdf. [1684]
- KARLIN, S., AND H. RUBIN (1956): “The Theory of Decision Procedures for Distributions With Monotone Likelihood Ratio,” *Annals of Mathematical Statistics*, 27, 272–299. [1694]
- LE CAM, L. (1986): *Asymptotic Methods in Statistical Decision Theory*. New York: Springer-Verlag. [1684,1686]
- MANSKI, C. F. (2000): “Identification Problems and Decisions Under Ambiguity: Empirical Analysis of Treatment Response and Normative Analysis of Treatment Choice,” *Journal of Econometrics*, 95, 415–442. [1683,1684]
- (2002): “Treatment Choice Under Ambiguity Induced by Inferential Problems,” *Journal of Statistical Planning and Inference*, 105, 67–82. [1683,1684]
- (2004): “Statistical Treatment Rules for Heterogeneous Populations,” *Econometrica*, 72 (4), 1221–1246. [1683,1684,1686,1699]
- O’LEARY, C. J., P. T. DECKER, AND S. A. WANDNER (1998): “Reemployment Bonuses and Profiling,” Discussion Paper 98-51, W. E. Upjohn Institute for Employment Research. [1683]
- (2005): “Cost-Effectiveness of Targeted Reemployment Bonuses,” *The Journal of Human Resources*, 40 (1), 270–279. [1683]
- PUHALSKI, A., AND V. SPOKOINY (1998): “On Large Deviation Efficiency in Statistical Inference,” *Bernoulli*, 4 (2), 203–272. [1686]
- SAVAGE, L. (1951): “The Theory of Statistical Decision,” *Journal of the American Statistical Association*, 46, 55–67. [1690]
- SCHLAG, K. (2006): “Eleven-Tests Needed for a Recommendation,” Working Paper ECO 2006-2, EUI. [1683]

- STOYE, J. (2006): "Minimax Regret Treatment Choice With Finite Samples," Working Paper, New York University. [1683]
- TETENOV, A. (2007): "Statistical Treatment Choice Based on Asymmetric Minmax Regret Criteria," Working Paper, Northwestern University. [1689,1694]
- VAN DER VAART, A. W. (1991a): "An Asymptotic Representation Theorem," *International Statistical Review*, 59, 99–121. [1695,1696]
- (1991b): "On Differentiable Functionals," *The Annals of Statistics*, 19, 178–204. [1696]
- (1998): *Asymptotic Statistics*. New York: Cambridge University Press. [1687]
- WORDEN, K. (1993): "Profiling Dislocated Workers for Early Referral to Reemployment Services," Unpublished Manuscript, U.S. Department of Labor. [1683]

Dept. of Economics, University of Arizona, 401 McClelland Hall, 1130 East Helen Street, Tucson, AZ 85721-0108, U.S.A.; hirano@u.arizona.edu
and

Dept. of Economics, University of Wisconsin, 6448 Social Science Building, 1180 Observatory Drive, Madison, WI 53706-1393, U.S.A.; jrporter@ssc.wisc.edu.

Manuscript received August, 2006; final revision received August, 2008.