

Manski, Charles

Working Paper

Statistical treatment rules for heterogeneous populations

cemmap working paper, No. CWP03/03

Provided in Cooperation with:

Institute for Fiscal Studies (IFS), London

Suggested Citation: Manski, Charles (2003) : Statistical treatment rules for heterogeneous populations, cemmap working paper, No. CWP03/03, Centre for Microdata Methods and Practice (cemmap), London,
<https://doi.org/10.1920/wp.cem.2003.0303>

This Version is available at:

<https://hdl.handle.net/10419/79304>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.


Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



STATISTICAL TREATMENT RULES FOR HETEROGENEOUS POPULATIONS

Charles Manski

THE INSTITUTE FOR FISCAL STUDIES
DEPARTMENT OF ECONOMICS, UCL
cemmap working paper CWP03/03

STATISTICAL TREATMENT RULES FOR HETEROGENEOUS POPULATIONS

Charles F. Manski
Department of Economics and Institute for Policy Research
Northwestern University

Revised: May 2003

The research presented here was supported in part by National Science Foundation grant SES-0001436. An early version of this paper was circulated in May 1999 as National Bureau of Economic Research Technical Working Paper No. 242 and a subsequent version was circulated in October 2001. I am grateful for the opportunity to present aspects of this work in seminars at Duke University, New York University, Northwestern University, Stanford University, UCLA, the University of Bristol, University College London, the University of Chicago, the University of Uppsala, the University of Virginia, and the University of Wisconsin-Madison. I have benefitted from the comments of Hidehiko Ichimura and Francesca Molinari. I am grateful to Jörg Stoye for able research assistance and for comments.

Abstract

An important objective of empirical research on treatment response is to provide decision makers with information useful in choosing treatments. Manski (2000, 2002, 2003) showed how identification problems generate ambiguity about the identity of optimal treatment choices. This paper studies treatment choice using sample data. I consider a planner who must choose among alternative *statistical treatment rules*, these being functions that map observed covariates of population members and sample data on treatment response into a treatment allocation. I study the use of *risk* (Wald, 1950) to evaluate the performance of alternative rules and, more particularly, the *minimax-regret* criterion to choose a treatment rule. These concepts may also be used to choose a sample design. Wald's development of statistical decision theory directly confronts the problem of finite-sample inference without recourse to the approximations of asymptotic theory. However, it is computationally challenging to implement. The main original work of this paper is to study implementation using data from a classical randomized experiment. Analysis of a simple problem of evaluation of an innovation yields a concise description of the set of undominated treatment rules and tractable computation of the minimax-regret rule. Analysis of a more complex problem of treatment choice using covariate information yields computable bounds on the maximum regret of alternative *conditional empirical success* rules, and consequent *sufficient sample sizes* for the beneficial use of covariate information. Numerical findings indicate that prevailing practices in the use of covariate information in treatment choice are too conservative.

1. Introduction

An important objective of empirical research on treatment response is to provide decision makers with information useful in choosing treatments. Manski (2000, 2002, 2003) showed how identification problems generate ambiguity about the identity of optimal treatment choices. This paper applies the Wald (1950) development of statistical decision theory to study treatment choice using sample data.

I consider a planner who must choose treatments for a heterogeneous population. The planner observes some covariates for each member of the treatment population. The observed covariates determine the set of treatment rules that are feasible to implement: the feasible rules are functions that assign all persons with the same observed covariates to one treatment or, more generally, that randomly allocate these persons across the different treatments. Each member of the population has a response function which maps treatments into an outcome. I assume that the planner wants to choose a treatment rule that maximizes population mean welfare, where welfare depends on the realized outcome; thus, the planner wants to maximize a utilitarian social welfare function.

To illustrate, consider a physician choosing medical treatments for a population of patients. The physician may observe each patient's demographic attributes, medical history, and the results of diagnostic tests. He may then choose a rule that makes treatment a function of these covariates. If the physician acts on behalf of his patients, the outcome of interest may be patient health status and the cost of treatment. Or consider a judge choosing sentences for a population of convicted offenders. The judge may observe each offender's past criminal record, demeanor in court, and other attributes. Subject to legislated sentencing guidelines, she may consider these covariates when choosing sentences. If the judge acts on behalf of society, the outcome of interest may be recidivism and the cost of carrying out a sentence.

As a prelude to consideration of treatment choice with sample data, Section 2 poses an idealized planning problem in which the planner has the information needed to choose an optimal treatment rule; this problem slightly generalizes the one examined in Manski (2000). Section 3 sets forth the planning problem

of interest, which is to make treatment choices using sample data. Now the planner must choose among the feasible *statistical treatment rules*, these being functions that map the observed covariates and sample data into a treatment allocation. The term *statistical treatment rule*, or *STR* for short, recalls Wald (1950), who used the term *statistical decision function* to describe functions that map sample data into decisions.

Wald proposed that alternative statistical decision functions be compared in terms of *risk*; that is, expected performance as the sampling process is engaged repeatedly to draw independent data samples. Risk (here called *expected welfare*) generally depends on quantities whose values are not known, so an operational decision criterion must somehow cope with these unknown quantities. Clearly, one should not choose a *dominated* (aka *inadmissible*) action; that is, one which is weakly inferior to another rule in all states of nature and strictly inferior in some state of nature. There is no consensus on how to choose among undominated actions, but statistical decision theorists have studied various approaches, most prominently the Bayes rules, which subjectively average across the feasible values of the unknowns, and the maximin and minimax-regret rules, which aim in different ways to behave uniformly well across the feasible values of the unknowns. Section 3 describes these venerable ideas in the context of treatment choice.

Section 3 presents a conceptually appealing prescription for treatment choice using sample data, but this prescription is computationally challenging to implement. The main original work of this paper is to study implementation in settings of interest. This is done in Sections 4 and 5, which examine treatment choice using experimental data.

Section 4 considers what may be the simplest non-trivial case of treatment choice using sample data. Outcomes are binary, there are no observed covariates, and there are two treatments, one being the status quo and the other being an innovation. The planner knows the response distribution of the status quo treatment, but not that of the innovation. To learn about the innovation, an experiment with randomized treatment assignment is performed. The objective is to use the experimental data to inform treatment choice.

A theorem of Karlin and Rubin (1956) shows that, in this setting, the undominated treatment rules

coincide with the *monotone treatment rules*. These are rules which assign all persons to the status quo treatment if the experimental success rate is below some threshold and all to the innovation if the success rate is above the threshold. Observing that the class of monotone treatment rules is mathematically “small” but substantively “large,” I go on to examine various rules within this class. The minimax-regret rule is particularly appealing. This rule solves a sensible, objective, optimization problem and makes judicious use of sample data.

Section 5 addresses a question of considerable analytical interest and practical importance, this being the use of covariate information in treatment choice. Suppose that outcomes are real-valued with bounded support, the covariate space is finite, and there are two treatments. To learn about response to these treatments, an experiment with randomized treatment assignment is performed. I consider designs that draw subjects at random within groups stratified by covariates and treatments, and ones that use simple random sampling to draw subjects. An experiment of unlimited size would enable the planner to implement the idealized optimal treatment rule, which conditions treatment choice on all observed covariates. The question of interest is the appropriate use of covariate information when the available data are from an experiment with a finite sample of subjects.

A complete analysis would first determine the set of undominated treatment rules and then consider choice among the undominated rules. This is not practical, but restriction of attention to a tractable class of treatment rules makes progress possible. In particular, I consider the family of *conditional empirical success (CES)* rules; these rules select treatments that maximize sample-average outcomes conditional on specified subsets of the observed covariates. A large-deviations theorem of Hoeffding (1963) implies a closed-form bound on the maximum regret of any CES rule. These bounds imply *sufficient sample sizes* for the beneficial use of covariate information. When the available sample size exceeds the sufficiency boundary, a planner can be certain that conditioning treatment choice on the observed covariates is preferable (in terms of minimax regret) to not conditioning.

Numerical findings in the simple case of a binary covariate indicate that prevailing practices in the use of covariate information are too conservative. It is common to condition treatment choice on covariates only if treatment response varies in a statistically significant manner across covariates values. Showing statistical significance typically requires sample sizes that exceed the sufficiency boundaries determined here.

2. The Idealized Planning Problem

This section poses the decision problem that, I presume, the planner wants to solve. This problem has a simple solution—the optimal treatment rule assigns to each member of the population a treatment that maximizes mean welfare conditional on the person’s observed covariates. I refer to this as the *idealized planning problem* because planners usually do not have the knowledge of treatment response needed to implement the optimal rule. Thus, this section is preparatory to consideration of treatment choice using sample data.

2.1. The Choice Set

I suppose that there is a finite set T of mutually exclusive and exhaustive treatments. Each member j of the treatment population, denoted J , has a predetermined response function $y_j(\cdot): T \rightarrow Y$ mapping treatments $t \in T$ into outcomes $y_j(t) \in Y$. The population is a probability space (J, Ω, P) , and the probability distribution $P[y(\cdot)]$ of the random function $y(\cdot): T \rightarrow Y$ describes treatment response across the population. The population is “large,” in the formal sense that J is uncountable and $P(j) = 0, j \in J$.

A planner must choose a treatment rule assigning a treatment to each member of J . A fully specified *treatment rule* is a function $\tau(\cdot): J \rightarrow T$ that assigns a treatment to each person. Person j 's outcome under rule

$\tau(\cdot)$ is $y_j[\tau(j)]$. Treatment is individualistic; that is, a person's outcome may depend on the treatment he is assigned, but not on the treatments assigned to others.

The planner observes certain predetermined covariates $x_j \in X$ for each member of the population; thus, $x: J \rightarrow X$ is the random variable mapping persons into their observable covariates. To simplify analysis, I suppose that the covariate space X is finite and that $P(x = \xi) > 0, \forall \xi \in X$. The planner can systematically differentiate persons with different observed covariates, but cannot distinguish among persons with the same observed covariates. Hence, a feasible treatment rule is a function that assigns all persons with the same observed covariates to one treatment or, more generally, a function that randomly allocates such persons across the different treatments.¹

Formally, let Z denote the space of functions that map $T \times X$ into the unit interval and that satisfy the adding-up conditions: $z(\cdot, \cdot) \in Z \Rightarrow \sum_{t \in T} z(t, \xi) = 1, \forall \xi \in X$. Then the feasible treatment rules are the elements of Z . An important subclass of Z are the *singleton* rules that assign all persons with the same observed covariates to one treatment; that is, $z(\cdot, \cdot)$ is a singleton rule if, for each $\xi \in X$, $z(t, \xi) = 1$ for some $t \in T$ and $z(s, \xi) = 0$ for all $s \neq t$. Non-singleton rules randomly allocate persons with covariates ξ across multiple treatments, with assignment shares $[z(t, \xi), t \in T]$. This definition of non-singleton rules does not specify which persons with covariates x receive each treatment, only the assignment shares. Designation of the particular persons receiving each treatment is immaterial because assignment is random, the population is large, and the planner's objective (see Section 2.2) is to maximize population mean welfare.

¹ In practice, institutional constraints may restrict the feasible treatment rules to a proper subset of these functions. For example, the planner may be precluded from using certain covariates (say race or gender) to assign treatments. The present analysis continues to hold if x is defined to be the covariates that the planner is permitted to consider, rather than the full vector of covariates that the planner observes.

2.2. The Objective Function and the Optimal Treatment Rule

I suppose that the planner wants to choose a feasible treatment rule that maximizes population mean welfare. The welfare from assigning treatment t to person j is $u[y_j(t), t, x_j]$, where $u(\cdot, \cdot, \cdot): Y \times T \times X \rightarrow R$ is the welfare function. For each feasible treatment rule z , the population mean welfare that would be realized if the planner were to choose rule z is

$$(1) \quad U(z, P) \equiv \sum_{\xi \in X} P(x = \xi) \sum_{t \in T} z(t, \xi) \cdot E\{u[y(t), t, \xi] \mid x = \xi\}.$$

The planner wants to solve the problem ²

$$(2) \quad \max_{z \in Z} U(z, P).$$

Let $S^{[T]}$ denote the unit simplex in $R^{[T]}$. The maximum in (2) is achieved if, for each $\xi \in X$, the planner chooses the treatment allocation $z(\cdot, \xi)$ to solve the problem

² Stafford (1985, pp. 112-114) appears to have been the first to explicitly view the objective of econometric evaluation of social programs as solution of a social planning problem. The specific social planning problem (2) has normative and analytical appeal. Maximization of a population mean outcome, or perhaps some weighted average outcome, is the standard utilitarian criterion of the public economics literature on social planning; the outcome of interest measures the social benefits minus costs of a treatment. The linearity of the expectation operator yields substantial analytical simplifications, particularly through use of the law of iterated expectations.

Of course, planning problems that are not of the form (2) may be of interest. For example, Heckman, Smith, and Clements (1997) suggest that a social preference for equity in outcomes may motivate a non-utilitarian criterion for choice between social programs. These authors also consider situations in which society uses a voting process to choose between programs. Then, one program is “better” than another if a majority of the population achieves better outcomes under the former program than under the latter, and hence votes for the former program.

$$(3) \quad \max_{z(\cdot, \xi) \in S^{|T|}} \sum_{t \in T} z(t, \xi) \cdot E\{u[y(t), t, \xi] \mid x = \xi\}.$$

The maximum in (3) is achieved by a singleton rule that allocates all persons with covariates ξ to a treatment that solves the problem³

$$(4) \quad \max_{t \in T} E\{u[y(t), t, \xi] \mid x = \xi\}.$$

The population mean welfare achieved by an optimal rule is

$$(5) \quad U^*(P) \equiv \sum_{\xi \in X} P(x = \xi) \left\{ \max_{t \in T} E\{u[y(t), t, \xi] \mid x = \xi\} \right\}.$$

3. The Planner's Problem, with Sample Data

3.1. The Expected Welfare (Risk) of a Statistical Treatment Rule

A planner facing the decision problem of Section 2 can choose an optimal treatment rule if he knows the conditional response distributions $\{P[y(t) \mid x = \xi], (t, \xi) \in T \times X\}$. Full knowledge of these distributions is rare, but sample data on treatment response may be available. The question is: How may planners use sample data to make treatment choices?

³ There is a unique optimal rule if problem (4) has a unique solution for every $\xi \in X$. There are multiple optimal rules if problem (4) has multiple solutions for some $\xi \in X$. In the latter case, all rules that randomly allocate persons with the same covariates among their optimal treatments are optimal.

To address this question, we first need to generalize the concept of a treatment rule. *Statistical treatment rules* (STRs) are functions that map covariates and sample data into treatment allocations. Let Q denote the sampling process generating the available data and let Ψ denote the sample space; that is, Ψ is the set of data samples that may be drawn under Q . Let Z henceforth denote the space of functions that map $T \times X \times \Psi$ into the unit interval and that satisfy the adding-up conditions: $z \in Z \Rightarrow \sum_{t \in T} z(t, \xi, \psi) = 1$, $\forall (\xi, \psi) \in X \times \Psi$. Then each function $z \in Z$ defines a statistical treatment rule.

Wald (1950) recommended evaluation of statistical decision functions as *procedures* applied as the sampling process is engaged repeatedly to draw independent data samples. In particular, he proposed that the performance of a statistical decision function be measured by its expected result across realizations of the sampling process. This was termed *risk* by Wald in the context of loss-minimization but will be called *expected welfare* here, where the objective is to maximize population mean welfare.

Let z be any feasible STR. Repeated engagement of the sampling process Q to draw independent samples ψ makes population mean welfare a real random variable. The expected welfare yielded by z in repeated samples is

$$\begin{aligned} (6) \quad W(z, P, Q) &\equiv \int \left\{ \sum_{\xi \in X} P(x = \xi) \sum_{t \in T} z(t, \xi, \psi) \cdot E\{u[y(t), t, \xi \mid x = \xi]\} dQ(\psi) \right\} \\ &= \sum_{\xi \in X} P(x = \xi) \sum_{t \in T} E[z(t, \xi, \psi)] \cdot E\{u[y(t), t, \xi \mid x = \xi]\}. \end{aligned}$$

Here $E[z(t, \xi, \psi)] \equiv \int z(t, \xi, \psi) dQ(\psi)$ is the expected (across repeated samples) fraction of persons with covariates ξ who are assigned to treatment t . In the case of a singleton rule, the expected allocation to treatment t is the probability of drawing a sample in which z assigns all persons with covariates ξ to this treatment; that is, $E[z(t, \xi, \psi)] = Q[z(t, \xi, \psi) = 1]$.

Following Wald, I use expected welfare (risk) to measure the performance of a statistical treatment

rule. This is appropriate if the planner is risk neutral and, moreover, it is analytically convenient. It should be understood, however, that Wald's general recommendation that rules be evaluated as procedures does not require that one employ the specific criterion of risk. For example, a planner could reasonably measure the performance of an STR by the median value of population welfare across repeated samples.

3.2. Implementable Criteria for Treatment Choice

Having defined expected welfare, consider the optimization problem

$$(7) \quad \max_{z \in Z} W(z, P, Q).$$

This problem has the same solution as the idealized problem (2), namely the singleton rule

$$(8) \quad z^*(\xi, \Psi) \equiv \operatorname{argmax}_{t \in T} E\{u[y(t), t, \xi] \mid x = \xi\}, \quad \xi \in X, \quad \Psi \in \Psi.$$

Thus, maximization of expected welfare yields the optimal treatment rule. The difficulty, of course, is that determination of the optimal rule requires knowledge of the conditional response distributions $\{P[y(t) \mid x = \xi], (t, \xi) \in T \times X\}$. Hence, problem (7) is no more solvable in practice than is problem (2).

Although problem (7) is not an implementable criterion for treatment choice, it provides the foundation for development of criteria that are implementable. First, one eliminates treatment rules that are dominated. Then, one chooses among the undominated rules.

Let Γ index the set of feasible states of nature; that is, let $[(P_\gamma, Q_\gamma), \gamma \in \Gamma]$ be the set of (P, Q) pairs that the planner deems possible. A statistical treatment rule is dominated if there exists another feasible rule

that yields weakly larger expected welfare in all feasible states of nature and strictly larger expected welfare in some state of nature. Thus, rule z is dominated if there exists another rule, say z' , such that $W(z', P_\gamma, Q_\gamma) \geq W(z, P_\gamma, Q_\gamma)$ for all $\gamma \in \Gamma$ and $W(z', P_\gamma, Q_\gamma) > W(z, P_\gamma, Q_\gamma)$ for some $\gamma \in \Gamma$.

There is essentially a consensus among statistical decision theorists that dominated actions should not be chosen.⁴ There is no similar consensus on choice among undominated actions, but two broad ideas have been prominent. Bayesian decision theorists recommend that a decision maker should place a subjective distribution on the states of nature and minimize *Bayes risk*; that is, the subjective expected value of risk with respect to this distribution. Other decision theorists recommend choice of an action that behaves uniformly well across the possible states of nature, in the sense of the minimax or minimax-regret criterion. Berger (1985) nicely exposit these decision criteria; see Section 1.5.2 and Chapter 5.

In the present treatment-choice context, where the objective is to maximize expected welfare rather than minimize risk, the Bayesian recommendation is that the planner place a subjective probability measure π on the set Γ and solve the optimization problem⁵

$$(9) \quad \sup_{z \in Z^*} \int W(z, P_\gamma, Q_\gamma) d\pi(\gamma),$$

where Z^* is the set of undominated STRs. The maximin criterion is

⁴ I say “essentially a consensus” because one may be able to prove that an action is dominated, but not be able to identify a specific alternative action that dominates it and is itself undominated. In such circumstances, statistical decision theory does not provide a clear prescription for decision making.

⁵ Problem (9), which is the Bayesian recommendation for evaluation of a treatment rule as a *procedure*, differs from the more familiar *conditional Bayes* prescription for decision making. The latter calls on the decision maker to form a posterior subjective distribution on Γ , conditional on the realized sample data, and to maximize the expected value of $U(z, P)$ with respect to this posterior distribution. Although the two problems differ, it turns out that solution of the latter problem at all points in the sample space yields the solution to the former one. See Berger (1985), Section 4.4.1.

$$(10) \quad \sup_{z \in Z^*} \inf_{\gamma \in \Gamma} W(z, P_\gamma, Q_\gamma).$$

The *minimax-regret* criterion, first proposed explicitly in Savage (1951), is

$$(11) \quad \inf_{z \in Z^*} \sup_{\gamma \in \Gamma} U^*(P_\gamma) - W(z, P_\gamma, Q_\gamma).$$

Here, $U^*(P_\gamma)$ is the optimal population mean welfare that would be achievable if it were known that $P = P_\gamma$; that is, by (5),

$$(12) \quad U^*(P_\gamma) \equiv \sum_{\xi \in X} P_\gamma(x = \xi) \left\{ \max_{t \in T} \int u[y(t), t, \xi] dP_\gamma[y(t) | x = \xi] \right\}.$$

The quantity $U^*(P_\gamma) - W(z, P_\gamma, Q_\gamma)$, called the *regret* of rule z in state of nature γ , is the loss in expected welfare that results from not knowing the true state of nature.

The Bayesian recommendation does not, per se, yield a specific treatment rule; the solution to problem (9) depends on the prior distribution π . The maximin and minimax-regret criteria interpret in two ways the idea that the planner should choose a rule that behaves uniformly well across the possible states of nature. A maximin treatment rule yields the greatest lower bound on expected welfare across all states of nature. A minimax-regret rule yields the least upper bound on the loss in expected welfare that results from not knowing the true state of nature.

Observe that these decision criteria address the problem of finite-sample statistical inference directly, without recourse to the large-sample approximations of asymptotic statistical theory. Indeed, the concept of regret provides an attractive decision theoretic foundation for the development of asymptotic theory. Consider a commensurate sequence of sampling processes and treatment rules $(z_N, Q_N; N = 1, \dots, \infty)$, where

N indexes sample size. This sequence is *pointwise consistent* if regret converges to zero in all states of nature and it is uniformly consistent if maximum regret converges to zero. Thus, sequence $(z_N, Q_N; N = 1, \dots, \infty)$ is pointwise consistent if $\lim_{N \rightarrow \infty} U^*(P_\gamma) - W(z_N, P_\gamma, Q_N) = 0$, all $\gamma \in \Gamma$. It is uniformly consistent if $\lim_{N \rightarrow \infty} \{\sup_{\gamma \in \Gamma} U^*(P_\gamma) - W(z_N, P_\gamma, Q_N)\} = 0$.

3.3. Discussion

Generality of the Prescription

Wald's development of statistical decision theory provides a three-step prescription for treatment choice with sample data: (1) use expected welfare (risk) to evaluate the performance in repeated samples of alternative statistical treatment rules; (2) eliminate dominated rules; (3) use some implementable criterion—perhaps Bayesian, maximin, or minimax regret—to choose among the undominated rules. The prescription is operational whenever the expected welfare function $W(\cdot, \cdot, \cdot)$ exists on its domain $Z \times [(P_\gamma, Q_\gamma), \gamma \in \Gamma]$ and a well-defined criterion is used to choose among the undominated rules.

Subject to these regularity conditions, the prescription is entirely general. It enables comparison of all feasible statistical treatment rules. It applies whatever the sampling process may be and whatever information the planner may have about the true state of nature and the sampling process. The set $[(P_\gamma, Q_\gamma), \gamma \in \Gamma]$ of feasible states of nature expresses the assumptions about treatment response and the sampling process that the planner is willing to maintain. This set may be finite-dimensional (parametric) or infinite-dimensional (nonparametric). The maintained assumptions may point-identify or only partially identify mean treatment response.

When one can influence sample design, Wald's prescription enables comparison of alternative sampling processes. Consider a two-period world, with sample data collected in the first period and treatment choices made in the second. A planner may want to jointly choose a sampling process and a treatment rule

that uses the data generated by this process. Let $C(Q)$ denote the cost of sampling process Q , with cost measured in the same units as population welfare. Then the expected welfare of any feasible (STR, sampling process) pair (z, Q) is $W(z, P, Q) - C(Q)$. A pair (z, Q) is optimal if it maximizes $W(z, P, Q) - C(Q)$.

Conceptual and Computational Issues

The primary conceptual criticism of Wald's statistical decision theory has been the conditional Bayes argument that statistical inference should be based only on observed data and not on frequentist thought experiments that contemplate how a procedure would perform in repeated sampling. Thus, decision theorists of the conditional Bayes school argue that risk is not an appropriate measure of the performance of a statistical decision function; see, for example, Berger (1985, Chapter 1). These decision theorists do not put forward an alternative proposal to measure performance. Rather, they argue on grounds of *procedural rationality* that conditional Bayes thinking is coherent and does not rest on frequentist thought experiments that Bayesians consider irrelevant to decision making.⁶ The continuing conceptual debates between and within the frequentist and conditional Bayes schools of thought will not be settled here.

Apart from the conditional Bayes criticism of frequentist statistics, the primary impediment to use of Wald's statistical decision theory has been computational rather than conceptual. Early on, statistical decision theorists found that it can be difficult to determine what actions are dominated and to solve the optimization problems yielding a Bayes, minimax or minimax-regret action. As a consequence, the surge of decision theoretic research that immediately followed publication of Wald (1950) did not endure. Of course, conclusions about computational tractability drawn fifty years ago may not be entirely relevant today. Hence, the practicality of Wald's theory warrants a fresh appraisal.

⁶ The conditional Bayes criticism of Wald's frequentist statistical decision theory applies only to decision making using given sample data, not to sample design. Conditional Bayesians recognize that frequentist thinking is necessary to choose among alternative sampling processes.

Application to Treatment Choice

Whether for computational or other reasons, treatment choice appears not to have been studied from the Wald perspective. Instead, treatment choice has been strongly influenced by the frequentist theory of hypothesis testing; see the discussion in Sections 4.2 and 5.5 below. To the extent that statistical decision theory has been applied to treatment choice, it has been the conditional Bayes paradigm rather than that of Wald; see, for example, Rubin (1978) and Dehejia (2003).

The remainder of this paper applies Wald's statistical decision theory to the problem of treatment choice with sample data. Because this is an initial exploration of a potentially vast subject, I focus on the leading case in which the sample data are the outcomes of a classical randomized experiment. Moreover, I restrict attention to problems of choice between two treatments. Section 4 analyzes a simple problem of evaluation of an innovation. Section 5 examines a more complex problem of treatment choice using covariate information.

4. Using a Randomized Experiment to Evaluate an Innovation

4.1. The Setting

In this section there are two treatments, a binary outcome, and no observed covariates; thus, $T = \{0, 1\}$, $Y = \{0, 1\}$ and $X = \phi$. One treatment, say $t = 0$, is the status quo and the other, $t = 1$, is an innovation. The planner knows the response distribution $P[y(0)]$ of the status quo treatment, perhaps through observation of historical experience. The planner does not know the response distribution $P[y(1)]$ of the innovation. I suppose for convenience that welfare is the outcome of a treatment; that is, $u[y_j(t), t, x_j] = y_j(t)$, $j \in J$, $t \in T$. In the present setting — with a binary outcome, no covariates, and $P[y(0)]$ known—this is a normalization

rather than a substantive assumption.

An experiment is performed to learn about outcomes under the innovation, with N subjects randomly drawn from the population and assigned to treatment 1. Of these subjects, a number n realize outcome $y = 1$ and the remaining $N - n$ realize outcome $y = 0$. I assume that the outcomes of all subjects are observed.

In this setting, the sample size N indexes the sampling process and the number n of experimental successes is a sufficient statistic for the sample data. The feasible statistical treatment rules are functions $z: T \times [0, \dots, N] \rightarrow [0, 1]$ that map the number of experimental successes into a treatment allocation; that is, for each value of n , rule z randomly allocates a fraction $z(1, n)$ of the population to treatment 1 and the remaining $z(0, n) = 1 - z(1, n)$ to treatment 0.

Let $p(t) \equiv P[y(t) = 1]$, $t \in T$. By (6), the expected welfare of rule z is

$$(13) \quad W(z, P, N) = p(0) \cdot E[z(0, n)] + p(1) \cdot E[z(1, n)] = p(0) + [p(1) - p(0)] \cdot E[z(1, n)].$$

The number of experimental successes is distributed binomial $\mathbf{B}[p(1), N]$, so

$$(14) \quad E[z(1, n)] = \sum_{i=0}^N z(1, i) \cdot b[n = i; p(1), N],$$

where $b[i; p(1), N] \equiv N!/[i!(N - i)!]^{-1} p(1)^i [1 - p(1)]^{N-i}$ is the Binomial probability of i successes.

The only unknown determinant of expected welfare is $p(1)$. Hence, Γ indexes the set of feasible values of $p(1)$. In what follows, I presume that $[p_\gamma(1), \gamma \in \Gamma]$ contains values that are smaller and larger than $p(0)$; otherwise, the choice problem is trivial.

4.2. Undominated Treatment Rules

In the setting described above, larger values of n are evidence for larger values of $p(1)$. Hence, it is reasonable to conjecture that undominated treatment rules should be ones in which the share $z(1, \cdot)$ of the population allocated to treatment 1 increases with n . In fact, the undominated treatment rules are a simple subclass of these rules.

Define a *monotone treatment rule* to be one of the form

$$(15) \quad \begin{aligned} z(1, n) &= 0 && \text{for } n < n_0, \\ z(1, n) &= \lambda && \text{for } n = n_0, \\ z(1, n) &= 1 && \text{for } n > n_0, \end{aligned}$$

where $0 \leq n_0 \leq N$ and $0 \leq \lambda \leq 1$. Thus, a monotone rule allocates all persons to treatment 0 if n is smaller than the specified threshold n_0 , a fraction λ to treatment 1 if $n = n_0$, and all persons to treatment 1 if n is larger than n_0 . Observe that monotone rules have simple expressions for the expected treatment allocation. By (15),

$$(16) \quad E[z(1, n)] = b[n > n_0; p(1), N] + \lambda \cdot b[n = n_0; p(1), N].$$

The result is

Proposition 1: Let $0 < p(0) < 1$ and let the feasible set $[p_\gamma(1), \gamma \in \Gamma]$ exclude the values 0 and 1. Then the collection of monotone treatment rules is the set of undominated rules. □

Proof: This is an application of Karlin and Rubin (1956, Theorem 4). The Karlin and Rubin theorem

concerns the class of decision problems with three features. First, there are two feasible actions. Second, the unknown state of nature is real-valued and there exists a unique interior *change point* such that one action has better risk under states of nature lower than the change point, and the other has better risk under states of nature higher than the change point. Third, the sample data is the realization of a real random variable whose density function (with respect to some σ -finite measure) has the strict form of the monotone likelihood ratio property.

In the present setting, the two actions are choice of treatment 0 and 1. The unknown state of nature is $p(1)$ and the change point is $p(0)$; treatment 0 is better in state of nature γ if $p_\gamma(1) < p(0)$ and treatment 1 is better if $p_\gamma(1) > p(0)$. The sample data is the realization of n , whose binomial distribution has density with respect to counting measure. This density function, namely $b[n; p_\gamma(1), N]$, has the strict form of the monotone likelihood ratio property. That is,

$$i > i' \text{ and } p_\gamma(1) > p_{\gamma'}(1) \Rightarrow \frac{b[n = i; p_\gamma(1), N]}{b[n = i'; p_\gamma(1), N]} > \frac{b[n = i; p_{\gamma'}(1), N]}{b[n = i'; p_{\gamma'}(1), N]}.$$

Thus, the theorem applies here.

Q. E. D.

A weaker version of the theorem applies if $p(0)$ or $p(1)$ can take the value 0 or 1. Then, Karlin and Rubin (1956, Theorem 1) shows that the collection of monotone treatment rules is *essentially complete*. That is, given any non-monotone rule z' , there exists some monotone rule z such that $W(z, P_\gamma, N) \geq W(z', P_\gamma, N)$, $\forall \gamma \in \Gamma$. See Berger (1985, Section 8.3) for exposition of the Karlin-Rubin theorems and related findings.

Some Monotone Rules

The collection of monotone treatment rules is a mathematically “small” but substantively “large”

subset of the space of all feasible treatment rules. This collection of rules is mathematically small in that it excludes almost all feasible rules, in the sense of Lebesgue measure. The space of all feasible rules is the Cartesian Product of N replicates of the unit simplex in \mathbb{R}^2 . Whereas this space has positive Lebesgue measure, the collection of monotone rules has Lebesgue measure zero.

To demonstrate that the collection of monotone rules is substantively large, I describe briefly various forms that such a rule may take. Another monotone rule, the minimax-regret rule, will be examined more closely in Section 4.3.

Data-Invariant Rules: These are the rules $z(1, \cdot) = 0$ and $z(1, \cdot) = 1$, which assign all persons to treatment 0 or 1 respectively, whatever the realization of n may be.

Maximin Rule: The maximin criterion yields a specific data-invariant rule. By (13), the infimum of expected welfare for rule z is

$$(17) \quad \inf_{\gamma \in \Gamma} W(z, P_\gamma, N) = p(0) + \min_{\gamma \in \Gamma} [p_\gamma(1) - p(0)] \cdot E_\gamma[z(1, n)],$$

where $E_\gamma[z(1, n)]$ is the expression in (14) with $p_\gamma(1)$ replacing $p(1)$. By assumption, $[p_\gamma(1), \gamma \in \Gamma]$ contains values that are smaller than $p(0)$. Moreover, $E_\gamma[z(1, n)] > 0$ for all $p_\gamma(1) > 0$ and all monotone treatment rules except for $z(1, \cdot) = 0$, the rule that always choose treatment 0. Hence, the maximin rule is $z(1, \cdot) = 0$.

Empirical Success Rules: Recall that an optimal treatment rule allocates all persons to treatment 0 if $p(1) < p(0)$ and all to treatment 1 if $p(1) > p(0)$. An empirical success rule emulates the optimal rule by replacing $p(1)$ with its sample analog, the empirical success rate n/N . Thus, an empirical success rule has the form

$$\begin{aligned}
(18) \quad z(1, n) &= 0 \quad \text{for } n < p(0)N, \\
z(1, n) &= \lambda \quad \text{for } n = p(0)N, \quad \text{where } 0 \leq \lambda \leq 1, \\
z(1, n) &= 1 \quad \text{for } n > p(0)N.
\end{aligned}$$

Statistical Significance Rules: These rules use a one-sided hypothesis test to choose between the status quo treatment and the innovation. The null hypothesis is that both treatments yield the same mean welfare; that is, $p(1) = p(0)$. The alternative is that treatment 1 is superior to treatment 0; that is, $p(1) > p(0)$. Treatment 1 is chosen if the null is rejected, and treatment 0 is chosen otherwise. Thus, the rule is

$$\begin{aligned}
(19) \quad z(1, n) &= 0 \quad \text{for } n \leq d[s, p(0), N], \\
z(1, n) &= 1 \quad \text{for } n > d[s, p(0), N],
\end{aligned}$$

where s is the specified size of the test and $d[s, p(0), N]$ is the associated critical value. Given that n is binomial, $d[s, p(0), N] = \min i: b[n > i; p(0), N] \leq s$.

The use of one-sided hypothesis tests to make treatment choices is institutionalized in the U. S. Food and Drug Administration pre-market drug approval process, which calls for comparison of a new drug under study ($t = 1$) with an approved treatment if one exists, or a placebo otherwise ($t = 0$). Approval of the new treatment normally requires a one-sided rejection of the null hypothesis in two independent randomized clinical trials (see Fisher and Moyé, 1999).

Although statistical significance rules are monotone treatment rules, the conventional practice of hypothesis testing is remote from the problem of treatment choice with sample data. If the null hypothesis $H_0: [p(1) = p(0)]$ is correct, all feasible treatment rules yield the same expected welfare. If not, alternative STRs may yield different expected welfare. A statistical test of H_0 indicates only whether the sample data are inconsistent (in the usual sense of having low probability of being realized under the null) with the

hypothesis that all feasible STRs are equivalent in expected welfare.

Bayes Rules: A planner who commits to use a Bayes rule has enormous discretion because the form of the Bayes rule depends critically on the prior subjective distribution placed on $p(1)$. To see this, consider the class of Beta priors, which form the conjugate family for a Binomial likelihood. Let $[p_\gamma(1), \gamma \in \Gamma] = (0, 1)$ and let the prior be Beta with parameters (α, β) . Then the posterior mean for $p(1)$ is $(\alpha + n)/(\alpha + \beta + N)$; see, for example, DeGroot (1970, Section 9.2, Theorem 1). The resulting Bayes rule is

$$\begin{aligned}
 (20) \quad z(1, n) &= 0 \quad \text{for } (\alpha + n)/(\alpha + \beta + N) < p(0), \\
 z(1, n) &= \lambda \quad \text{for } (\alpha + n)/(\alpha + \beta + N) = p(0), \quad \text{where } 0 \leq \lambda \leq 1, \\
 z(1, n) &= 1 \quad \text{for } (\alpha + n)/(\alpha + \beta + N) > p(0).
 \end{aligned}$$

As (α, β) tend to zero, the Bayes rule approaches an empirical success rule. Moreover, the class of Bayes rules includes the data-invariant rules $z(1, \cdot) = 0$ and $z(1, \cdot) = 1$. The former occurs if the parameters (α, β) of the Beta prior satisfy $(\alpha + N)/(\alpha + \beta + N) < p(0)$, and the latter occurs if $\alpha/(\alpha + \beta + N) > p(0)$.

4.3. The Minimax-Regret Rule

The analysis of Section 4.2 shows that planners should restrict attention to monotone treatment rules, but does not recommend a particular selection among the class of monotone rules. Among the rules discussed above, the only ones that solve explicit and interpretable optimization problems are the Bayes rules and the maximin rule. However, Bayes rules rest on subjective priors and the maximin rule is data-invariant. Planners may be uncomfortable with the subjectivity of Bayes rules, and they may reasonably reject the maximin conclusion that the data are irrelevant.

The minimax-regret rule offers an appealing alternative to those discussed above. The minimax-regret rule solves a sensible, objective, optimization problem. Unlike the maximin rule, which also solves an objective optimization problem, the minimax-regret rule makes judicious use of sample data.

By (13), the regret of rule z in state of nature γ is

$$\begin{aligned}
 (21) \quad U^*(P_\gamma) - W(z, P_\gamma, N) &= \max[p(0), p_\gamma(1)] - \{p(0) + [p_\gamma(1) - p(0)] \cdot E_\gamma[z(1, n)]\} \\
 &= [p_\gamma(1) - p(0)] \cdot E_\gamma[z(0, n)] \quad \text{if } p_\gamma(1) \geq p(0), \\
 &= [p(0) - p_\gamma(1)] \cdot E_\gamma[z(1, n)] \quad \text{if } p(0) \geq p_\gamma(1).
 \end{aligned}$$

Thus, regret is the mean welfare loss when a member of the population is assigned the inferior treatment, multiplied by the expected fraction of the population assigned this treatment.

The minimax-regret rule does not have a closed-form expression, but it can be determined numerically. The maximum regret of treatment rule z is

$$(22) \quad R(z) \equiv \sup_{\gamma \in \Gamma} [p_\gamma(1) - p(0)] \cdot E_\gamma[z(0, n)] \cdot 1[p_\gamma(1) \geq p(0)] + [p(0) - p_\gamma(1)] \cdot E_\gamma[z(1, n)] \cdot 1[p(0) \geq p_\gamma(1)].$$

The expressions $[p_\gamma(1) - p(0)] \cdot E_\gamma[z(0, n)]$ and $[p(0) - p_\gamma(1)] \cdot E_\gamma[z(1, n)]$ are $(N+1)$ -order polynomials in $p_\gamma(1)$. The maximum regret of rule z is obtained by maximizing the first polynomial over $\{\gamma: p_\gamma(1) \geq p(0)\}$, the second over $\{\gamma: p(0) \geq p_\gamma(1)\}$, and selecting the larger of the two maxima. The minimax-regret rule is then obtained by minimizing $R(\cdot)$ over the collection of monotone treatment rules defined in (15), which are indexed by the threshold sample sizes and allocations (n_0, λ) , $n_0 = 0, \dots, N$, $\lambda \in [0, 1]$.

Table 1 reports the minimax-regret rule for specified values of $p(0)$ and N when all values of $p(1)$ are feasible; that is, when $[p_\gamma(1), \gamma \in \Gamma] = [0, 1]$. The top two panels display the value of (n_0, λ) for this rule. The third panel displays the value of minimax regret; that is, $R \equiv \min_{z \in Z} R(z)$. The bottom panel displays

the state of nature at which R is achieved; this value is denoted $p_R(1)$.

Consider first the entries for $N = 0$, indicating the absence of sample data. In this case, the only feasible treatment rules are $[z(0, 0), z(1, 0)] = (1 - \lambda, \lambda) \in [0, 1]$. For each value of λ , maximum regret is $\max\{[1 - p(0)](1 - \lambda), p(0)\lambda\}$. Hence, the minimax-regret rule is $[n_0 = 0, \lambda = 1 - p(0)]$ and the value of minimax regret is $R = p(0)[1 - p(0)]$. The rule allocates a positive fraction of the population to treatment 1 even if no experimental data are available; the magnitude of λ falls from 1 to 0 as $p(0)$ rises from 0 to 1.

The top panel of the table shows that the threshold n_0 of experimental successes for allocation of persons to treatment 1 increases with the sample size and with the success probability of treatment 0. The inequality $|n_0 - p(0)N| \leq 1$ holds everywhere in the table. Thus, the minimax-regret rule is well approximated by an empirical success rule.

The second panel shows that randomization when n_0 successes are observed is a generic feature of the minimax-regret rule. The third panel shows that the value of minimax regret decreases by roughly an order of magnitude as the sample size increases from 0 to 10. The specific pattern of decrease varies markedly with the value of $p(0)$. A curiosity is that R is a step function when $p(0) = 0.5$, with decreases occurring at odd values of N but not at even values.

The bottom panel shows the state of nature $p_R(1)$ which maximizes regret for the minimax-regret rule. This state of nature locates the maximum of an $(N+1)$ -order polynomial within the unit interval. Observe that $p_R(1)$ may be an extreme state of nature when $N \leq 3$, but it always is an interior value when $N \geq 4$. When $p(0) = 0.50$, the polynomial is symmetric and so has two global maxima that are equidistant from 0.50.

Savage on the Maximin and Minimax-Regret Criteria

The findings in Table 1 make plain that treatment choice using the minimax-regret rule differs fundamentally from treatment choice using the maximin rule. Savage (1951), whose review of Wald (1950) first explicitly distinguished between these criteria for decision making, argued strongly against application

of the minimax (here maximin) criterion, writing (p. 63):

Application of the minimax rule . . . is indeed ultra-pessimistic; no serious justification for it has ever been suggested, and it can lead to the absurd conclusion in some cases that no amount of relevant experimentation should deter the actor from behaving as though he were in complete ignorance.

Our finding that the maximin treatment rule is data-invariant illustrates this “absurd conclusion.” Savage emphasized that the minimax criterion is “ultra-pessimistic,” not the minimax-regret criterion. Our finding that the minimax-regret rule approximates the empirical success rule illustrates that the minimax-regret criterion is not particularly pessimistic.⁷

5. Using Covariate Information With Data From a Randomized Experiment

It was shown in Section 2 that a planner who knows the conditional response distributions $\{P[y(t)|x], t \in T\}$ should condition treatment choice on the observed covariates x . The appropriate use of covariate information is a more subtle matter when only sample data on treatment response are available. Researchers have long appreciated that it may not be judicious to condition predictions of outcomes on all observed covariates when the available sample is small, the reason being that estimates of best predictors tend to be less precise as the degree of conditioning increases.⁸ However, treatment choice and prediction of outcomes are different decision problems. So it is necessary to investigate the use of covariate information afresh from

⁷ Savage conjectured that Wald (1950), whose abstract study of minimax principles encompasses both criteria, had minimax regret in mind rather than minimax. This conjecture could not be verified because Wald died in 1950, shortly after Savage completed his review article.

⁸ In particular, statisticians studying prediction under square loss have studied *shrinkage estimators* (e.g., Copas, 1983, and Lehmann, 1983, Section 4.6) and *variable selection procedures* (e.g., Kempthorne, 1984, and Droge, 1998).

the perspective of treatment choice. This is a large and difficult question, but I am able to report some progress.

I study here the use of covariate information with experimental data, from the minimax-regret perspective. In particular, I develop computable bounds on the maximum regret of alternative *conditional empirical success* rules. Conditional empirical success rules are empirical success rules that condition treatment choice on specified subsets of the available covariate information. I show that the bounds on maximum regret imply *sufficient sample sizes* for the beneficial use of covariate information. When the available sample size exceeds the sufficiency boundary, a planner can be certain that conditioning treatment choice on the observed covariates is preferable (in terms of minimax regret) to not conditioning.

Section 5.1 sets forth the maintained assumptions. Section 5.2 defines the class of conditional empirical success rules. Section 5.3 develops the bounds on maximum regret for experimental designs that use a stratified random sampling process to draw subjects. Section 5.4 extends the analysis to experimental designs that draw subjects by simple random sampling.⁹ Section 5.5 presents numerical findings that make evident the practical implications for treatment choice. Section 5.6 briefly discusses the use of conditional empirical success rules when the data are not from a classical randomized experiment.

5.1. The Setting

In this section, there are two treatments; thus, $T = \{0, 1\}$. There is a real-valued outcome with bounded range; without loss of generality, let $\inf(Y) = 0$ and $\sup(Y) = 1$. There is a finite covariate space with full support; thus, $|X| < \infty$ and $P(x = \xi) > 0$, all $\xi \in X$. The planner knows the population covariate

⁹ The stratified and simple random sampling designs studied here are polar cases, the former with the greatest feasible degree of stratification and the latter with the least. The proposition on simple random sampling developed in Section 5.4 can be modified to cover intermediate cases, in which the design stratifies on some but not all covariates.

distribution $P(x)$. Welfare is the outcome of a treatment; thus, $u[y_j(t), t, x_j] = y_j(t)$, $j \in J$, $t \in T$. Whereas the last condition was a normalization in Section 4, it now is a substantive assumption.

A randomized experiment is performed to learn about response to the two treatments, and the outcomes of all subjects are observed. I consider two alternative experimental designs, stratified random sampling and simple random sampling:

Stratified Random Sampling: The experimenter assigns to each treatment a specified number of subjects with each value of the covariates. Thus, for $(t, \xi) \in T \times X$, the experimenter draws $N_{t\xi}$ subjects at random from the population and assigns these subjects to treatment t . The set $N_{TX} \equiv [N_{t\xi}, (t, \xi) \in T \times X]$ of stratum sample sizes indexes the sampling process. For each $(t, \xi) \in T \times X$, let $N(t, \xi)$ be the sub-sample of subjects with covariates ξ who are assigned to treatment t . Then the sample data are the outcomes $Y_{TX} \equiv [y_j, j \in N(t, \xi); (t, \xi) \in T \times X]$. The feasible STRs are functions that map covariates and the data into a treatment allocation. Thus, for each value of x and Y_{TX} , rule z randomly allocates a fraction $z(1, x, Y_{TX})$ of persons with covariates x to treatment 1 and $z(0, x, Y_{TX})$ to treatment 0. The expected welfare of rule z is

$$(23) \quad W(z, P, N_{TX}) = \sum_{\xi \in X} P(x = \xi) \cdot \{E[y(0)|x = \xi] \cdot E[z(0, \xi, Y_{TX})] + E[y(1)|x = \xi] \cdot E[z(1, \xi, Y_{TX})]\}.$$

Simple Random Sampling: The experimenter draws N subjects at random from the population and randomly assigns them to treatments 0 and 1 with specified assignment probabilities, say q and $1 - q$. The pair (N, q) indexes the sampling process. The sample data are the stratum sample sizes N_{TX} and the outcomes Y_{TX} . An STR z allocates a fraction $z(0, x, N_{TX}, Y_{TX})$ of persons with covariates x to treatment 0 and $z(1, x, N_{TX}, Y_{TX})$ to treatment 1. The expected welfare of rule z is

$$(24) \quad W(z, P, N, q) = \sum_{\xi \in X} P(x = \xi) \cdot \{E[y(0)|x = \xi] \cdot E[z(0, \xi, N_{TX}, Y_{TX})] + E[y(1)|x = \xi] \cdot E[z(1, \xi, N_{TX}, Y_{TX})]\}.$$

Thus, stratified and simple random sampling differ in the status of the stratum sample sizes; N_{TX} is fixed by the experimenter in the former designs and is part of the sample data in the latter ones.

5.2. Conditional Empirical Success (CES) Rules

To the best of my knowledge, research to date in statistical decision theory does not shed light on decision problems in which expected welfare has form (23) or (24). No succinct characterization of the set of undominated treatment rules is available. Direct assessment of dominance by computation of expected welfare for all feasible treatment rules and states of nature is impractical.

Progress is possible if one restricts attention to a tractable class of treatment rules. Statisticians and econometricians studying efficiency in estimation have long made progress by restricting attention to tractable classes of estimators; for example, linear unbiased estimators or asymptotic normal estimators. Similarly, I make progress here by restricting attention to *conditional empirical success* rules.

In general, an empirical success rule emulates the optimal treatment rule by replacing unknown response distributions with sample analogs. When welfare is the outcome of a treatment, the optimal rule chooses treatments that solve the problems $\max_{t \in T} E[y(t)|x = \xi]$, $\xi \in X$. Hence, an empirical success rule replaces $E[y(t)|x = \xi]$, $(t, \xi) \in T \times X$ by corresponding sample-average outcomes and chooses treatments that maximize empirical success. Thus, $E[y(t)|x = \xi]$ is replaced by the average outcome $\bar{y}_{t\xi} \equiv (1/N_{t\xi}) \sum_{j \in N(t, \xi)} y_j$ and an empirical success rule solves the problems $\max_{t \in T} \bar{y}_{t\xi}$, $\xi \in X$.

Specification of an empirical success rule requires selection of the covariates on which treatment choice are conditioned. The idealized planning problem calls for conditioning on all observed covariates, but conditioning tends to diminish the statistical precision of the sample averages used to estimate population

means. Hence, it is reasonable to think that conditioning on some part of the observed covariates may be preferable when making treatment choices with sample data. This suggests comparison of empirical success rules that condition on alternative subsets of the available covariate information. Formally, let $v(\cdot): X \rightarrow V$ map the covariate space X into a space V . Then an empirical success rule can choose treatments that maximize empirical success conditional on a person's value of v rather than on his value of x .¹⁰

The proper way to measure empirical success conditional on v depends on the experimental design. Let $v \in V$. In a design with simple random sampling, the appropriate sample analog of $E[y(t)|v = v]$ is $\bar{y}_{tv} \equiv (1/N_{tv}) \sum_{j \in N(t, v)} y_j$, where $N(t, v)$ is the sub-sample of subjects with covariates v who are assigned to treatment t , and where $N_{tv} \equiv |N(t, v)|$. In a design with stratified random sampling, the appropriate sample analog is the design-weighted average $\bar{y}_{txv} \equiv \sum_{\xi \in X} \bar{y}_{t\xi} \cdot P(x = \xi | v = v)$.¹¹

To complete the definition of a CES rule requires a tie-breaking convention to determine treatment choice when both treatments have the same sample-average outcome. For simplicity, I suppose that the planner allocates all persons to treatment 0 in such cases.¹² Adopting this convention yields a singleton rule in which the expected fraction of persons assigned to treatment 1 is the probability, across repeated samples, that treatment 1 has greater empirical success than treatment 0. Hence, expected welfare is

¹⁰ Treatment choice using a CES rule is conceptually similar to, but technically distinct from, prediction of an outcome using an empirical predictor that conditions on a subset of the observed covariates. Research on variable-selection procedures has studied selection of covariates for purposes of prediction. Considering the classical normal linear regression model, Kempthorne (1984) finds that all variable-selection procedures are admissible. Considering a class of nonparametric regression problems, Droge (1998) determines variable-selection procedures that minimize maximum regret within the family of orthogonal series estimators. These and related findings on prediction are interesting, but they do not appear to hold lessons for the present problem of treatment choice.

¹¹ In a stratified design, $\bar{y}_{t\xi}$ is the sample analog of $E[y(t)|x = \xi]$, but \bar{y}_{tv} is not that of $E[y(t)|v = v]$. However, $E[y(t)|v = v] = \sum_{\xi \in X} E[y(t)|x = \xi] \cdot P(x = \xi | v = v)$. Hence \bar{y}_{txv} is the sample analog of $E[y(t)|v = v]$.

¹² Use of a tie-breaking rule that randomly allocates some persons to one treatment and the rest to the other treatment would not materially alter the analysis in Sections 5.3 and 5.4.

$$(25) \quad W(z_{Xv}, P, N_{TX}) = \sum_{v \in V} P(v = v) \cdot \{E[y(0)|v = v] \cdot P(\bar{y}_{0Xv} \geq \bar{y}_{1Xv}) + E[y(1)|v = v] \cdot P(\bar{y}_{1Xv} > \bar{y}_{0Xv})\}$$

in a design with stratified random sampling and

$$(26) \quad W(z_v, P, N, q) = \sum_{v \in V} P(v = v) \cdot \{E[y(0)|v = v] \cdot P(\bar{y}_{0v} \geq \bar{y}_{1v}) + E[y(1)|v = v] \cdot P(\bar{y}_{1v} > \bar{y}_{0v})\}$$

in a design with simple random sampling. Here, z_{Xv} denotes the CES rule conditioning on v in a design with stratified random sampling, and z_v denotes the corresponding rule in a design with simple random sampling.

Bounding the Expected Welfare and Maximum Regret of CES Rules

Direct analysis of the expected welfare of CES rules is arduous. The treatment selection probabilities are probabilities that one sample mean exceeds another. Such probabilities generically do not have closed-form expressions and are difficult to compute numerically. Hence, computation of the expected welfare of CES rules in all feasible states of nature is not practical.

Fortunately, it is possible to develop informative closed-form bounds on expected welfare. The upper bound is simple enough: The maximum population welfare attainable by any STR that conditions treatment choice on the covariates v is $\sum_{v \in V} P(v = v) \max\{E[y(0)|v = v], E[y(1)|v = v]\}$; hence, $W(z_{Xv}, P, N_{TX})$ and $W(z_v, P, N, q)$ cannot exceed this value. Section 5.3 develops the lower bound for experimental designs with stratified random sampling, and Section 5.4 does the same for designs with simple random sampling. In each case, the lower bound depends on sample size, the covariate distribution, and mean treatment response.

Sections 5.3 and 5.4 use the bounds on expected welfare to obtain bounds on the maximum regret of CES rules. Subtraction of the lower bound on expected welfare from the idealized optimum population welfare gives an upper bound on regret in any state of nature, and maximization of the upper bound on regret

over the feasible states of nature gives an upper bound on maximum regret. Similarly, subtraction of the upper bound on expected welfare from the idealized optimum population welfare gives a lower bound on regret, and maximization of the lower bound on regret over the feasible states of nature gives a lower bound on maximum regret. Sufficient sample sizes for the beneficial use of covariate information follow immediately from these bounds.

5.3. Bounding Expected Welfare and Maximum Regret in Designs with Stratified Random Sampling

Bounding Expected Welfare

The analysis of this section exploits an important large deviations theorem of Hoeffding (1963):

Large Deviations Theorem (Hoeffding, 1963, Theorem 2): Let w_1, w_2, \dots, w_K be independent real random variables, with bounds $a_k \leq w_k \leq b_k$, $k = 1, 2, \dots, K$. Let $\bar{w} \equiv (1/K) \sum_{k=1}^K w_k$ and $\mu \equiv E(\bar{w})$. Then, for $d > 0$, $\Pr(\bar{w} - \mu \geq d) \leq \exp[-2Kd^2 \{\sum_{k=1}^K (b_k - a_k)^2\}^{-1}]$. \square

The Hoeffding theorem is a powerful result. The only distributional assumption is that the random variables w_1, w_2, \dots, w_K are independent and have bounded supports. The upper bound on $\Pr(\bar{w} - \mu \geq d)$ does not depend on the specific distribution of \bar{w} , yet converges to zero at exponential rate as either K or d^2 grows.¹³

Proposition 2 uses the Hoeffding theorem to bound the expected welfare of a CES rule in a design with stratified sampling.

Proposition 2: Let subjects be drawn by stratified random sampling. Let $v(\cdot): X \rightarrow V$ and consider the CES

¹³ The price for derivation of such a simple large-deviations bound is that it need not be sharp. Hoeffding (1963), Theorem 1 gives tighter bounds on $\Pr(\bar{w} - \mu \geq d)$ that hold if w_1, w_2, \dots, w_K have the same range. However, these bounds are complicated and depend on some features of the distribution of \bar{w} .

rule z_{Xv} . For $v \in V$, define $M_v \equiv \max\{E[y(0)|v = v], E[y(1)|v = v]\}$ and $\delta_v \equiv |E[y(1)|v = v] - E[y(0)|v = v]|$. Then the expected welfare of rule z_{Xv} satisfies the inequality

$$(27) \quad \sum_{v \in V} P(v = v) M_v - D(z_{Xv}, P, N_{TX}) \leq W(z_{Xv}, P, N_{TX}) \leq \sum_{v \in V} P(v = v) M_v,$$

where

$$(28) \quad D(z_{Xv}, P, N_{TX}) \equiv \sum_{v \in V} P(v = v) \cdot \delta_v \cdot \exp[-2\delta_v^2 \{ \sum_{\xi \in X} P(x = \xi | v = v)^2 (N_{1\xi}^{-1} + N_{0\xi}^{-1}) \}^{-1}]. \quad \square$$

Proof. Let $v \in V$. I first write $\bar{y}_{1Xv} - \bar{y}_{0Xv}$ as the average of independent random variables and then apply the Hoeffding theorem. Let $N \equiv \sum_{\xi \in X} (N_{1\xi} + N_{0\xi})$ be the total sample size. Then

$$(29) \quad \begin{aligned} \bar{y}_{1Xv} - \bar{y}_{0Xv} &= \sum_{\xi \in X} P(x = \xi | v = v) (1/N_{1\xi}) \sum_{j \in N(1, \xi)} y_j - \sum_{\xi \in X} P(x = \xi | v = v) (1/N_{0\xi}) \sum_{j \in N(0, \xi)} y_j \\ &= (1/N) \{ \sum_{\xi \in X} \sum_{j \in N(1, \xi)} [y_j \cdot P(x = \xi | v = v) N/N_{1\xi}] + \sum_{\xi \in X} \sum_{j \in N(0, \xi)} [-y_j \cdot P(x = \xi | v = v) N/N_{0\xi}] \}. \end{aligned}$$

Thus, $\bar{y}_{1Xv} - \bar{y}_{0Xv}$ averages N independent random variables whose ranges are $[0, P(x = \xi | v = v) N/N_{1\xi}]$ and $[-P(x = \xi | v = v) N/N_{0\xi}, 0]$, $\xi \in X$.

Suppose that $E[y(1)|v = v] < E[y(0)|v = v]$. Then $E(\bar{y}_{1Xv} - \bar{y}_{0Xv}) = -\delta_v$. Application of the Hoeffding theorem yields¹⁴

$$(30) \quad \begin{aligned} P(\bar{y}_{1Xv} > \bar{y}_{0Xv}) &= P(\bar{y}_{1Xv} - \bar{y}_{0Xv} + \delta_v > \delta_v) \\ &\leq \exp[-2N^2 \delta_v^2 \{ \sum_{\xi \in X} N_{1\xi} [P(x = \xi | v = v) \cdot N/N_{1\xi}]^2 + N_{0\xi} [P(x = \xi | v = v) \cdot N/N_{0\xi}]^2 \}^{-1}] \end{aligned}$$

¹⁴ Inequality (30) holds even when some stratum sample sizes are zero. If either $N_{0\xi}$ or $N_{1\xi}$ equals 0, the formalisms $0^{-1} = \infty$ and $\infty^{-1} = 0$ yield $P(x = \xi | v = v)^2 (N_{1\xi}^{-1} + N_{0\xi}^{-1})^{-1} = 0$.

$$= \exp[-2\delta_v^2 \{\sum_{\xi \in X} P(x = \xi | v = v)^2 (N_{1\xi}^{-1} + N_{0\xi}^{-1})\}^{-1}].$$

Similarly, if $E[y(1)|v = v] > E[y(0)|v = v]$, the Hoeffding theorem gives

$$(31) \quad P(\bar{y}_{0Xv} \geq \bar{y}_{1Xv}) \leq \exp[-2\delta_v^2 \{\sum_{\xi \in X} P(x = \xi | v = v)^2 (N_{1\xi}^{-1} + N_{0\xi}^{-1})\}^{-1}].$$

To obtain inequality (27), partition V into three regions:

$$V_a \equiv \{v \in V: E[y(1)|v = v] < E[y(0)|v = v]\}$$

$$V_b \equiv \{v \in V: E[y(1)|v = v] > E[y(0)|v = v]\}$$

$$V_c \equiv \{v \in V: E[y(1)|v = v] = E[y(0)|v = v]\}.$$

Then rewrite equation (25) as

$$\begin{aligned} (25') \quad W(z_{Xv}, P, N_{TX}) &= \sum_{v \in V_a} P(v = v) \{E[y(0)|v = v] \cdot P(\bar{y}_{0Xv} \geq \bar{y}_{1Xv}) + E[y(1)|v = v] \cdot P(\bar{y}_{1Xv} > \bar{y}_{0Xv})\} \\ &+ \sum_{v \in V_b} P(v = v) \{E[y(0)|v = v] \cdot P(\bar{y}_{0Xv} \geq \bar{y}_{1Xv}) + E[y(1)|v = v] \cdot P(\bar{y}_{1Xv} > \bar{y}_{0Xv})\} \\ &+ \sum_{v \in V_c} P(v = v) \{E[y(0)|v = v] \cdot P(\bar{y}_{0Xv} \geq \bar{y}_{1Xv}) + E[y(1)|v = v] \cdot P(\bar{y}_{1Xv} > \bar{y}_{0Xv})\}. \end{aligned}$$

The first term on the right side of (25') can be no smaller than the expression obtained by setting $P(\bar{y}_{1Xv} > \bar{y}_{0Xv})$ at the upper bound obtained in (30) and setting $P(\bar{y}_{0Xv} \geq \bar{y}_{1Xv})$ at its implied lower bound. The lower bound on the second term on the right side of (25') is obtained in the same manner, by applying (31). The third term is constant across all treatment allocations. Placing all three terms on the right side of (25') at their lower bounds yields the lower bound on $W(z_{Xv}, P, N_{TX})$ given in (27). The upper bound in (27) is obtained by analogous reasoning.

Q. E. D.

The bound on expected welfare obtained in Proposition 2 is a closed form function of the stratum sample sizes N_{TX} , the covariate distribution $P(x)$, and the mean treatment outcomes $\{E[y(0)|v = v], E[y(1)|v = v], v \in V\}$. The upper bound is the maximum population welfare achievable using covariates v . The lower bound differs from this ideal by the non-negative *finite-sample penalty* $D(z_{Xv}, P, N_{TX})$, which places an upper bound on the loss in welfare that results from estimating mean treatment outcomes rather than knowing them. The magnitude of the finite-sample penalty decreases with sample size, and it converges to zero at exponential rate if all elements of N_{TX} grow at the same rate.

Observe that, for each $v \in V$, the finite-sample penalty varies non-monotonically with δ_v . In particular, it varies as $\delta_v \cdot \exp(-C_{Nv} \delta_v^2)$, where $C_{Nv} \equiv 2 \{ \sum_{\xi \in X} P(x = \xi | v = v)^2 (N_{1\xi}^{-1} + N_{0\xi}^{-1}) \}^{-1}$. If $\delta_v = 0$, there is no finite-sample penalty because both treatments are equally good for persons with covariates v . As δ_v increases, the loss in welfare due to an error in treatment selection increases linearly, but the probability of making an error goes to zero at exponential rate. As a result, the penalty is maximized at $\delta_v = (2C_{Nv})^{-1/2}$. Inserting the worst-case values $[\delta_v = (2C_{Nv})^{-1/2}, v \in V]$ into equation (28) yields this uniform upper bound on the finite-sample penalty:

$$(32) \quad D(z_{Xv}, P, N_{TX}) \leq \frac{1}{2} e^{-1/2} \sum_{v \in V} P(v = v) \left\{ \sum_{\xi \in X} P(x = \xi | v = v)^2 (N_{1\xi}^{-1} + N_{0\xi}^{-1}) \right\}^{1/2}.$$

Bounding Maximum Regret

The bound on expected welfare in Proposition 2 immediately gives a bound on regret, namely

$$(33) \quad \sum_{\xi \in X} P(x = \xi) M_\xi - \sum_{v \in V} P(v = v) M_v \leq U^*(P) - W(z_{Xv}, P, N_{TX})$$

$$\leq \sum_{\xi \in X} P(x = \xi) M_\xi - \sum_{v \in V} P(v = v) M_v + D(z_{Xv}, P, N_{TX}).$$

The lower bound on regret is the idealized benefit of conditioning treatment choice on x rather than v . The upper bound is the idealized benefit of conditioning on x plus the finite sample penalty of the CES rule conditioning on v .

The above bound on regret is a function of the stratum sample sizes, the covariate distribution, and mean treatment response. The stratum sample sizes and the covariate distribution are known. By the law of iterated expectations, $E[y(t)|v] = \sum_{\xi \in X} E[y(t)|x = \xi] \cdot P(x = \xi|v)$. Hence, the only unknown quantities are $\{E[y(0)|x], E[y(1)|x]\}$. Maximizing the bound over the feasible values of $\{E[y(0)|x], E[y(1)|x]\}$ yields this bound on maximum regret:

$$(34) \quad R_{Lv} \leq R(z_{Xv}) \leq R_U(z_{Xv}),$$

where

$$\begin{aligned} R_{Lv} &\equiv \sup_{\gamma \in \Gamma} \sum_{\xi \in X} P(x = \xi) \max \{E_\gamma[y(0)|x = \xi], E_\gamma[y(1)|x = \xi]\} \\ &\quad - \sum_{v \in V} P(v = v) \max \{E_\gamma[y(0)|v = v], E_\gamma[y(1)|v = v]\}, \\ R_U(z_{Xv}) &\equiv \sup_{\gamma \in \Gamma} \sum_{\xi \in X} P(x = \xi) \max \{E_\gamma[y(0)|x = \xi], E_\gamma[y(1)|x = \xi]\} \\ &\quad - \sum_{v \in V} P(v = v) \max \{E_\gamma[y(0)|v = v], E_\gamma[y(1)|v = v]\} \\ &\quad + \sum_{v \in V} P(v = v) \cdot \delta_{v\gamma} \cdot \exp[-2\delta_{v\gamma}^2 \{ \sum_{\xi \in X} P(x = \xi|v = v)^2 (N_{1\xi}^{-1} + N_{0\xi}^{-1}) \}^{-1}], \end{aligned}$$

$\delta_{v\gamma} \equiv |E_\gamma[y(1)|v = v] - E_\gamma[y(0)|v = v]|$, and $E_\gamma[y(t)|v = v] = \sum_{\xi \in X} E_\gamma[y(t)|x = \xi] \cdot P(x = \xi|v = v)$. The lower

bound R_{L_v} applies to any STR that conditions treatment choice on covariates v , not only to the CES rule z_{Xv} .

The upper bound $R_U(z_{Xv})$ is specific to this CES rule.

Examination of (34) shows that the tractability of computing the bound on maximum regret does not vary with sample size. However, tractability varies inversely with the cardinality of the covariate space X . A state of nature is a value for the $(|X| \times 2)$ -dimensional vector of mean treatment responses $\{E_\gamma[y(0)|x = \xi], E_\gamma[y(1)|x = \xi], \xi \in X\}$. Hence, solution of the extremum problems defining $R_L(z_{Xv})$ and $R_U(z_{Xv})$ becomes increasingly difficult as $|X|$ grows.

The CES Rule Conditioning on All Observed Covariates

The bound on maximum regret simplifies considerably in the case of rule z_{Xx} , which conditions treatment choice on all observed covariates. Then $V = X$ and $v(x) = x$, so (34) reduces to

$$(35) \quad 0 \leq R(z_{Xx}) \leq \sup_{\gamma \in \Gamma} \sum_{\xi \in X} P(x = \xi) \cdot \delta_{\xi\gamma} \cdot \exp[-2\delta_{\xi\gamma}^2 (N_{1\xi}^{-1} + N_{0\xi}^{-1})^{-1}].$$

The upper bound in (35) is the supremum of the finite-sample penalty across all feasible states of nature.

Suppose that all distributions of treatment response are feasible. Then the derivation of the uniform upper bound on the finite-sample penalty given in (32) shows that (35) further reduces to¹⁵

¹⁵ The upper bound in (36) can be improved when all stratum sizes are positive but some are extremely small; specifically, when some strata have one or two observations. When all stratum sample sizes are positive, the regret of rule z_{Xx} cannot exceed the value 1/4 in any state of nature. The upper bound in (36) can exceed 1/4 when some strata have one or two observations. In such cases, the value 1/4 is a better upper bound than the one given in (36).

To prove that 1/4 is an upper bound on the regret of rule z_{Xx} , observe that the regret of this rule in state of nature γ is

$$\begin{aligned} U^*(P_\gamma) - W(z_{Xx}, P_\gamma, N_{TX}) &= \sum_{\xi \in X} P(x = \xi) \cdot \delta_{\xi\gamma} \cdot \{P_\gamma(\bar{y}_{0X\xi} \geq \bar{y}_{1X\xi}) \cdot 1\{E_\gamma[y(1)|x = \xi] > E_\gamma[y(0)|x = \xi]\} \\ &\quad + P_\gamma(\bar{y}_{1X\xi} > \bar{y}_{0X\xi}) \cdot 1\{E_\gamma[y(0)|x = \xi] > E_\gamma[y(1)|x = \xi]\} \}. \end{aligned}$$

$$(36) \quad 0 \leq R(z_{xx}) \leq \frac{1}{2} e^{-1/2} \sum_{\xi \in X} P(x = \xi) (N_{1\xi}^{-1} + N_{0\xi}^{-1})^{1/2}.$$

Asymptotic Implications

Inequalities (35) and (36) imply that rule z_{xx} is uniformly consistent, and yield lower bounds on its pointwise and uniform rates of convergence. Inequality (35) shows that if all elements of N_{TX} grow at the same rate, the regret of rule z_{xx} converges to zero at least at exponential rate in every state of nature. Inequality (36) shows that if all elements of N_{TX} grow at the same rate, the maximum regret of rule z_{xx} converges to zero at least with the square root of sample size.

Inequality (36) also implies that the minimax-regret treatment rule itself is uniformly consistent.

Fix $\xi \in X$ and suppose that $E_\gamma[y(1)|x = \xi] > E_\gamma[y(0)|x = \xi]$. Then

$$\begin{aligned} & \delta_{\xi\gamma} \cdot \{P_\gamma(\bar{y}_{0X\xi} \geq \bar{y}_{1X\xi}) \cdot 1\{E_\gamma[y(1)|x = \xi] > E_\gamma[y(0)|x = \xi]\} + P_\gamma(\bar{y}_{1X\xi} > \bar{y}_{0X\xi}) \cdot 1\{E_\gamma[y(0)|x = \xi] > E_\gamma[y(1)|x = \xi]\}\} \\ &= \delta_{\xi\gamma} \cdot P_\gamma(\bar{y}_{0X\xi} \geq \bar{y}_{1X\xi}) = E_\gamma(\bar{y}_{1X\xi} - \bar{y}_{0X\xi}) \cdot P_\gamma(\bar{y}_{0X\xi} \geq \bar{y}_{1X\xi}). \end{aligned}$$

The second equality holds because $\bar{y}_{1X\xi}$ and $\bar{y}_{0X\xi}$ are unbiased estimates of $E_\gamma[y(1)|x = \xi]$ and $E_\gamma[y(0)|x = \xi]$. Now let $d \equiv \bar{y}_{1X\xi} - \bar{y}_{0X\xi}$ and observe that

$$\begin{aligned} E_\gamma(\bar{y}_{1X\xi} - \bar{y}_{0X\xi}) \cdot P_\gamma(\bar{y}_{0X\xi} \geq \bar{y}_{1X\xi}) &= E_\gamma(d) \cdot P_\gamma(d \leq 0) = [E_\gamma(d|d > 0) \cdot P_\gamma(d > 0) + E_\gamma(d|d \leq 0) \cdot P_\gamma(d \leq 0)] \cdot P_\gamma(d \leq 0) \\ &\leq P_\gamma(d > 0) \cdot P_\gamma(d \leq 0) \leq 1/4. \end{aligned}$$

The same inequality holds if $E_\gamma[y(1)|x = \xi] < E_\gamma[y(0)|x = \xi]$. Hence, regret is bounded above by 1/4. I am grateful to Jörg Stoye for this argument.

Although the form of the minimax-regret rule is unknown, this rule necessarily has maximum regret no larger than the upper bound in (36). Hence, minimax regret converges to zero at a rate no lower than the square root of sample size.

Sufficient Sample Sizes for the Beneficial Use of Covariate Information

Sufficient sample sizes for the beneficial use of covariate information follow immediately from the bounds on maximum regret. Suppose that all states of nature are feasible. Let $v(\cdot): X \rightarrow V$ be a specified many-to-one mapping of x into a covariate v . The maximum regret of any STR conditioning treatment choice on v must exceed that of CES rule z_{xx} if $R_{Lv} > R_U(z_{xx})$, where R_{Lv} is given in (34) and $R_U(z_{xx})$ in (36). Thus, treatment choice conditioning on x is necessarily preferable (in term of minimax regret) to conditioning on v if $R_{Lv} > \frac{1}{2} e^{-\frac{1}{2}} \sum_{\xi \in X} P(x = \xi) (N_{1\xi}^{-1} + N_{0\xi}^{-1})^{\frac{1}{2}}$. The right side of this inequality is decreasing in each component of the vector of stratum sample sizes N_{Tx} . Hence, sufficient sample sizes for conditioning on x to be preferable to conditioning on v are solutions to the problem

$$(37) \quad \min N_{Tx}: R_{Lv} > \frac{1}{2} e^{-\frac{1}{2}} \sum_{\xi \in X} P(x = \xi) (N_{1\xi}^{-1} + N_{0\xi}^{-1})^{\frac{1}{2}}.$$

Several aspects of this derivation warrant comment. First, N_{Tx} being a vector, problem (37) generically has multiple solutions; hence, I refer to sufficient sample sizes (plural). Second, when N_{Tx} exceeds a sufficiency boundary, the maximum regret of CES rule z_{xx} is smaller than that of *all* STRs conditioning treatment choice on v , not just smaller than that of rule z_{xv} . Third, the sufficiency boundaries provide a sufficient condition for superiority of rule z_{xx} to rules that condition on v , not a necessary condition. There may not exist any rule conditioning on v whose maximum regret attains the lower bound R_{Lv} , and the maximum regret of rule z_{xx} may be less than the upper bound $R_U(z_{xx})$. Fourth, the present findings do not show that rule z_{xx} is the best STR conditioning treatment choice on x . There may exist a non-CES rule that

is superior to z_{Xx} .

The above discussion compares conditioning on x with conditioning on v when all states of nature are feasible. Sufficient sample sizes for other comparisons may be generated in the same manner. Let $v(\cdot): X \rightarrow V$ and $w(\cdot): X \rightarrow W$ be distinct mappings of X into covariates v and w respectively. The maximum regret of all STRs conditioning treatment choice on v must exceed that of CES rule z_{Xw} if $R_{Lv} > R_U(z_{Xw})$. The upper bound $R_U(z_{Xw})$, defined in (34), is decreasing in each component of N_{TX} . Hence, sufficient sample sizes for conditioning on w to be preferable to conditioning on v are solutions to the problem

$$(38) \quad \min N_{TX}: R_{Lv} > R_U(z_{Xw}).$$

Problem (38) generically has solutions whenever $v(\cdot)$ produces a coarser partition of X than does $w(\cdot)$.

5.4. Bounding Expected Welfare and Maximum Regret in Designs with Simple Random Sampling

In experimental designs with simple random sampling, the stratum sample sizes N_{TX} are random rather than fixed. However, the reasoning used to prove Proposition 2 continues to be applicable conditional on any realization of N_{TX} . Hence, Proposition 2 provides the “inner loop” for analysis of simple random sampling. Here is the result.

Proposition 3: Let subjects be drawn by simple random sampling. Let $v(\cdot): X \rightarrow V$ and consider the CES rule z_v . For $v \in V$, let B_{Nv} denote the Binomial distribution $\mathbf{B}[P(v = v), N]$. For $n = 0, \dots, N$, let B_{nv} denote the Binomial distribution $\mathbf{B}(q, n)$.¹⁶ Then the expected welfare of rule z_v satisfies the inequality

¹⁶ If $n = 0$, define this distribution to be degenerate with all mass on the value zero.

$$(39) \quad \sum_{v \in V} P(v = v) M_v - D(z_v, P, n, q) \leq W(z_v, N, P, q) \leq \sum_{v \in V} P(v = v) M_v,$$

where

$$(40) \quad D(z_v, p, N, q) \equiv \sum_{v \in V} P(v = v) \cdot \delta_v \sum_{n=0}^N \sum_{m=0}^n B_{Nv}(n) \cdot B_{nq}(m) \cdot \exp\{-2\delta_v^2[(n - m)^{-1} + m^{-1}]^{-1}\}. \quad \square$$

Proof. It suffices to prove the result for rule z_x , which conditions on all observed covariates. Rule z_x is algebraically the same as rule z_{xx} , considered in Section 5.3. Rules z_{xx} and z_x differ only in that the stratum sample sizes N_{TX} are fixed with stratified sampling and are random with simple random sampling. The bounds on treatment-selection probabilities obtained in the proof of Proposition 2 hold under simple random sampling, conditional on the realization of N_{TX} . For $\xi \in X$, these bounds are

$$(41a) \quad P(\bar{y}_{1\xi} > \bar{y}_{0\xi} | N_{TX}) \leq \exp[-2\delta_\xi^2 (N_{1\xi}^{-1} + N_{0\xi}^{-1})^{-1}]$$

if $E[y(1) | x = \xi] < E[y(0) | x = \xi]$ and

$$(41b) \quad P(\bar{y}_{0\xi} \geq \bar{y}_{1\xi} | N_{TX}) \leq \exp[-2\delta_\xi^2 (N_{1\xi}^{-1} + N_{0\xi}^{-1})^{-1}]$$

if $E[y(1) | x = \xi] > E[y(0) | x = \xi]$.

The random variable $N_{0\xi} + N_{1\xi}$ is distributed $B_{N\xi}$. Conditional on the event $\{N_{0\xi} + N_{1\xi} = n\}$, $N_{0\xi}$ is distributed B_{nq} . Hence the unconditional treatment-selection probabilities satisfy the inequalities

$$(42a) \quad P(\bar{y}_{1\xi} > \bar{y}_{0\xi}) \leq \sum_{n=0}^N \sum_{m=0}^n B_{N\xi}(n) B_{nq}(m) \exp\{-2\delta_\xi^2[(n - m)^{-1} + m^{-1}]^{-1}\}$$

if $E[y(1)|x = \xi] < E[y(0)|x = \xi]$ and

$$(42b) \quad P(\bar{y}_{0\xi} \geq \bar{y}_{1\xi}) \leq \sum_{n=0}^N \sum_{m=0}^n B_{N\xi}(n) B_{nq}(m) \exp\{-2\delta_\xi^2[(n-m)^{-1} + m^{-1}]^{-1}\}$$

if $E[y(1)|x = \xi] > E[y(0)|x = \xi]$. The remainder of the proof is the same as that of Proposition 2.

Now let $v(\cdot): X \rightarrow V$ be any specified function and consider rule z_v . The same argument as above holds if one applies Proposition 2 to a sampling process that stratifies on v rather than on x .

Q. E. D.

The bound on expected welfare obtained in Proposition 3 is a closed form function of the sample size N , the treatment assignment probability q , and the mean treatment outcomes $\{E[y(0)|v = v], E[y(1)|v = v], v \in V\}$. As in Proposition 2, the upper bound is the maximum population welfare achievable using covariates v , and the lower bound differs from this ideal by a finite-sample penalty, here $D(z_v, N, P, q)$.

The bound on expected welfare in Proposition 3 gives this bound on regret:

$$(43) \quad \sum_{\xi \in X} P(x = \xi) M_\xi - \sum_{v \in V} P(v = v) M_v \leq U^*(P) - W(z_v, P, N, q) \\ \leq \sum_{\xi \in X} P(x = \xi) M_\xi - \sum_{v \in V} P(v = v) M_v + D(z_v, P, N, q).$$

Maximizing over the feasible values of $\{E[y(0)|x], E[y(1)|x]\}$ yields this bound on maximum regret:

$$(44) \quad R_{Lv} \leq R(z_v) \leq R_{Uv}(z_v),$$

where

$$\begin{aligned}
R_U(z_v) \equiv & \sup_{\gamma \in \Gamma} \sum_{\xi \in X} P(x = \xi) \max \{E_\gamma[y(0)|x = \xi], E_\gamma[y(1)|x = \xi]\} \\
& - \sum_{v \in V} P(v = v) \max \{E_\gamma[y(0)|v = v], E_\gamma[y(1)|v = v]\} \\
& + \sum_{v \in V} P(v = v) \cdot \delta_{v\gamma} \sum_{n=0}^N \sum_{m=0}^n B_{Nv}(n) \cdot B_{nq}(m) \cdot \exp \{-2\delta_{v\gamma}^2 [(n - m)^{-1} + m^{-1}]^{-1}\}.
\end{aligned}$$

A sufficient sample size for the beneficial use of covariate information follows immediately from the bounds on maximum regret. Let $v(\cdot): X \rightarrow V$ and $w(\cdot): X \rightarrow W$ be distinct mappings of X into covariates v and w respectively. The maximum regret of any STR conditioning treatment choice on v must exceed that of CES rule z_{Xw} if $R_{Lv} > R_U(z_w)$. Hence, the sufficient sample size for conditioning on covariates w to be preferable to conditioning on v solves the problem

$$(45) \quad \min N: R_{Lv} > R_U(z_w).$$

5.5. Numerical Findings for Binary Covariates

Computation of the bounds on maximum regret is particularly simple and revealing when the covariate x is a binary random variable; thus, let $X = \{a, b\}$. Then there are only two CES rules under any experimental design; one rule conditions treatment choice on x and the other does not. A state of nature is a quadruple $\{E[y(t)|x]; t = 0, 1; x = a, b\}$. The present analysis supposes that all states of natures are feasible.

The lower bound on maximum regret for any STR that does not condition treatment choice on x is

(46) $R_{L\phi} =$

$$\sup_{\gamma \in \Gamma} P(x = a) \max \{E_{\gamma}[y(0)|x = a], E_{\gamma}[y(1)|x = a]\} + P(x = b) \max \{E_{\gamma}[y(0)|x = b], E_{\gamma}[y(1)|x = b]\} \\ - \max \{P(x = a)E_{\gamma}[y(0)|x = a] + P(x = b) E_{\gamma}[y(0)|x = b], P(x = a)E_{\gamma}[y(1)|x = a] + P(x = b) E_{\gamma}[y(1)|x = b]\},$$

where $\{E_{\gamma}[y(t)|x]; t = 0, 1; x = a, b\}$ can take values in the unit hypercube $[0, 1]^4$. The unique solution of problem (46) is $\min\{P(x = a), P(x = b)\}$.¹⁷ Thus, $R_{L\phi} = \min\{P(x = a), P(x = b)\}$.

The upper bounds on maximum regret must be computed numerically. I first consider designs with simple random sampling and then ones with stratified sampling.

Simple Random Sampling

Under simple random sampling, the CES rules are z_x and z_{ϕ} . Table 2 computes the upper bound on maximum regret of each rule in designs with equal treatment assignment probabilities (i. e., $q = 0.5$) and sample sizes ranging from 1 to 200. Three covariate distributions are considered, with $P(x = a) = 0.05, 0.25$, or 0.5 .

Conditioning treatment choice on x is definitely better (in terms of maximum regret) than not conditioning whenever the upper bound on the maximum regret of rule z_x is smaller than the lower bound on the maximum regret of rule z_{ϕ} . Examination of Table 2 shows that the sufficient sample size for beneficial use of covariate information lies between $N = 100$ and $N = 200$ when $P(x = a) = 0.05$, is $N = 15$ when $P(x = a) = 0.25$, and is $N = 6$ when $P(x = a) = 0.50$.

These numerical findings indicate that prevailing practices in the use of covariate information in

¹⁷ Without loss of generality, let $P(x = a) \leq P(x = b)$ and $E[y(0)|x = a] \leq E[y(1)|x = a]$. Then it can be shown that a state of nature which solves problem (46) is $E_{\gamma}[y(0)|x = a] = E_{\gamma}[y(1)|x = b] = 0$, $E_{\gamma}[y(1)|x = a] = 1$, $E_{\gamma}[y(0)|x = b] = P(x = a)/P(x = b)$.

treatment choice are too conservative. It is commonly thought that treatment choice should be conditioned on covariates only if treatment response varies in a “statistically significant” manner across covariate values. Statistical significance is conventionally taken to mean rejection of the null hypothesis that the average treatment effect is invariant with respect to x ; that is, rejection of the hypothesis

$$(47) \quad E[y(1)|x = a] - E[y(0)|x = a] = E[y(1)|x = b] - E[y(0)|x = b].$$

This hypothesis is rarely rejected in small samples, so use of covariate information in treatment choice is commonly viewed as imprudent. The numerical findings presented here indicate that conditioning on covariates is warranted in samples far smaller than those required to show statistical significance.

Numerical findings aside, it should be clear that testing hypothesis (47) is remote in principle from the problem of treatment choice. A planner needs to assess the performance of alternative treatment rules, whether measured by maximum regret or by some other criterion. Hypothesis tests simply do not address the planner’s problem.

Quasi-Optimal Stratified Designs

A stratified random sampling design is indexed by the stratum sample sizes N_{TX} . The usual rationale for stratification is to improve statistical precision relative to simple random sampling. Hence, a natural way to study treatment choice under stratified sampling is to fix the overall sample size N and, for each feasible STR, determine the value of N_{TX} that minimizes maximum regret subject to the constraint $\sum_{(t, x)} N_{tx} = N$.

The analysis of Section 5.3 does not identify optimal stratified designs, but it does identify *quasi-optimal* designs that minimize the upper bound on maximum regret.¹⁸ Table 3 computes the quasi-optimal

¹⁸ I call these “quasi” optimal designs for two reasons. One is that I restrict attention to CES rules; it is unknown whether these designs are optimal for STRs not in the CES family. The other is that the criterion considered here is minimization of an upper bound on maximum regret, not maximum regret itself.

designs for the two feasible CES rules, z_{xx} and $z_{x\phi}$, when the overall sample size ranges from 1 to 52 and when $P(x = a) = 0.05$ or 0.25 .¹⁹ The covariate distribution with $P(x = a) = 0.5$ is omitted from the table because the quasi-optimal design in this case generically is the equal-shares allocation $\{N_{tx} = N/4; t = 0, 1; x = a, b\}$.

The first two rows of each panel of Table 3 show the quasi-optimal stratification, and the third row gives the upper bound on maximum regret. A generic finding is that, for each value of x , equal numbers of subjects should be assigned to each treatment; thus, $N_{0a} = N_{1a}$ and $N_{0b} = N_{1b}$. However, the fact that stratum sample sizes are integers implies that this treatment-balance condition is not strictly implementable when the overall sample size is an odd number. The entries (1, 2) for rule z_{xx} when $N = 3$ show that $(N_{0b} = 1, N_{1b} = 2)$ and $(N_{0b} = 2, N_{1b} = 1)$ are both quasi-optimal in this case of an odd value of N . The blank entries that sometimes appear when $N \leq 3$ indicate that all allocations yield the same trivial upper bound on maximum regret, namely 1.

Table 3 shows that the quasi-optimal stratification generically draws more subjects with covariate value $x = b$ than with value $x = a$. This is unsurprising given that the two covariate distributions considered in the table have $P(x = a) < P(x = b)$. The best designs for rule $z_{x\phi}$ are approximately self-weighting; that is, $N_{ta}/N_{tb} \approx P(x = a)/P(x = b)$. This too is unsurprising given that rule $z_{x\phi}$ does not condition treatment choice on the covariate. More interesting is the fact that the best designs for rule z_{xx} generically over-sample subjects with covariates value $x = a$; that is, $N_{ta}/N_{tb} \geq P(x = a)/P(x = b)$. The fact that sample sizes are integers makes it difficult to draw conclusions about the degree of over-sampling when $P(x = a) = 0.05$. However, it appears that when $P(x = a) = 0.25$, the sampling ratio $N_{ta}/N_{tb} \approx 1/2$.

Comparison of Tables 2 and 3 shows how quasi-optimal stratified random sampling improves treatment choice relative to simple random sampling. When N is very small, rule z_{xx} sometimes substantially

¹⁹ In each panel of the table and for each value of N , the quasi-optimal design is determined by computing the upper bound on maximum regret for each of the finitely many feasible values of N_{TX} and selection of the stratification that minimizes this upper bound.

outperforms its simple random sampling counterpart z_x . The advantage of stratification declines as N grows but remains non-negligible throughout the range of sample sizes considered here. The tables indicate that stratification is beneficial only when one conditions treatment choice on x ; rule $z_{x\phi}$ does not outperform its simple random sampling counterpart z_ϕ .

5.6. Using CES Rules in Non-classical Settings

To conclude this section, I point out that CES rules may be used to choose treatments when the data are not from a classical randomized experiment. Rules z_{xv} and z_v are well-defined when some experimental subjects do not comply with their assigned treatments, and they remain well-defined when the data are observational rather than experimental. CES rules may also be used when some data are missing. A common practice in the empirical analysis of treatment response has been to ignore cases with missing data and to compute empirical success for the cases with complete data. Formally, this means application of rule z_{xv} or z_v with the sub-samples $N(t, v)$ now defined to be the cases with complete data.

The expressions for expected welfare given in equations (25) and (36) remain valid in these non-classical settings. The lower bound R_{Lv} on maximum regret also remains valid, as this bound applies to all STRs conditioning treatment choice on v . However, the upper bounds on maximum regret proved in Propositions 2 and 3 do not necessarily hold.

The crux of the problem is the step in the proof of Proposition 2 stating: “Suppose that $E[y(1)|v = v] < E[y(0)|v = v]$. Then $E(\bar{y}_{1Xv} - \bar{y}_{0Xv}) = -\delta_v$.” This step rested on the assumption that the data are from a classical randomized experiment; in particular, on the fact that $E(\bar{y}_{tXx}) = E[y(t)|x]$. In non-classical settings, this step in the proof may or may not hold. A planner may apply Proposition 2 to evaluate CES rules in non-classical settings if he is willing to assume that $E(\bar{y}_{tXx}) = E[y(t)|x]$, but should not do so otherwise.

6. Conclusion

The Wald (1950) development of statistical decision theory gives a general prescription for treatment choice with sample data: use expected welfare to evaluate the performance of alternative STRs, eliminate dominated rules, then choose among the undominated rules. This prescription applies in both parametric and nonparametric settings, enables comparison of all feasible statistical treatment rules and, moreover, enables comparison of alternative sampling processes. I find it enormously appealing that Wald's theory directly addresses the problem of finite-sample statistical inference, without recourse to the large-sample approximations of asymptotic statistical theory. The minimax regret criterion, first proposed explicitly in Savage (1951), provides an attractive decision theoretic expression of the idea of efficient statistical inference.

The original methodological work of this paper has been to study treatment choice using experimental data, from the minimax-regret perspective. The analysis in Section 4 of evaluation of an innovation yielded a concise description of the set of undominated treatment rules and tractable computation of the minimax-regret rule. Section 5 yielded bounds on the maximum regret of alternative CES rules and consequent sufficient sample sizes for conditioning treatment choice on observed covariates. These findings begin, but only begin, to exploit the potential usefulness for treatment choice of Wald's general prescription and of the minimax-regret criterion in particular.

References

- Berger, J. (1985), *Statistical Decision Theory and Bayesian Analysis*, New York: Springer-Verlag.
- Copas, J. (1983), "Regression, Prediction, and Shrinkage," *Journal of the Royal Statistical Society B*, 45, 311–354.
- DeGroot, M. (1970), *Optimal Statistical Decisions*, New York: McGraw-Hill.
- Dehejia, R. (2003), "Program Evaluation as a Decision Problem," *Journal of Econometrics*, forthcoming.
- Droge, B. (1998), "Minimax Regret Analysis of Orthogonal Series Regression Estimation: Selection Versus Shrinkage," *Biometrika*, 85, 631–643.
- Fisher, L. and L. Moyé (1999), "Carvedilol and the Food and Drug Administration Approval Process: An Introduction," *Controlled Clinical Trials*, 20, 1–15.
- Heckman, J., J. Smith, and N. Clements (1997) "Making the Most out of Programme Evaluations and Social Experiments: Accounting for Heterogeneity in Programme Impacts," *Review of Economic Studies*, 64, 487–535.
- Hoeffding, W. (1963), "Probability Inequalities for Sums of Bounded Random Variables," *Journal of the American Statistical Association*, 58, 13–30.
- Karlin, S. and H. Rubin (1956), "The Theory of Decision Procedures for Distributions with Monotone Likelihood Ratio," *Annals of Mathematical Statistics*, 27, 272–299.
- Kempthorne, P. (1984), "Admissible Variable-Selection Procedures When Fitting Regression Models by Least Squares for Prediction," *Biometrika*, 71, 593–597.
- Lehmann, E. (1983), *Theory of Point Estimation*, New York: Wiley.
- Manski, C. (2000), "Identification Problems and Decisions Under Ambiguity: Empirical Analysis of Treatment Response and Normative Analysis of Treatment Choice," *Journal of Econometrics*, 95, 415–442.
- Manski, C. (2002), "Treatment Choice Under Ambiguity Induced by Inferential Problems," *Journal of Statistical Planning and Inference*, 105, 67–82.
- Manski, C. (2003), *Partial Identification of Probability Distributions*, New York: Springer-Verlag.
- Rubin, D. (1978), "Bayesian Inference for Causal Effects: The Role of Randomization," *The Annals of Statistics*, 6, 34–58.
- Savage, L. (1951), "The Theory of Statistical Decision," *Journal of the American Statistical Association*, 46, 55–67.

Stafford, F. (1985), "Income-Maintenance Policy and Work Effort: Learning from Experiments and Labor-Market Studies." in J. Hausman and D. Wise. eds. *Social Experimentation*. Chicago: University of Chicago Press.

Wald, A. (1950), *Statistical Decision Functions*, New York: Wiley.

TABLE 1: MINIMAX-REGRET TREATMENT RULES

n_0: threshold sample size	N = 0	N = 1	N = 2	N = 3	N = 4	N = 5	N = 6	N = 7	N = 8	N = 9	N = 10
$p(0) = 0.10$	0	0	0	0	0	0	0	0	0	1	1
$p(0) = 0.25$	0	0	0	1	1	1	1	2	2	2	2
$p(0) = 0.50$	0	1	1	2	2	3	3	4	4	5	5
$p(0) = 0.75$	0	1	2	2	3	4	5	5	6	7	8
$p(0) = 0.90$	0	1	2	3	4	5	6	7	8	8	9

λ: threshold allocation	N = 0	N = 1	N = 2	N = 3	N = 4	N = 5	N = 6	N = 7	N = 8	N = 9	N = 10
$p(0) = 0.10$	0.90	0.67	0.52	0.41	0.32	0.26	0.18	0.09	0.00	0.89	0.78
$p(0) = 0.25$	0.75	0.36	0.17	0.93	0.67	0.42	0.18	0.93	0.67	0.43	0.18
$p(0) = 0.50$	0.50	1.00	0.50	1.00	0.50	1.00	0.50	1.00	0.50	1.00	0.50
$p(0) = 0.75$	0.25	0.64	0.83	0.07	0.33	0.58	0.82	0.07	0.33	0.57	0.82
$p(0) = 0.90$	0.10	0.33	0.48	0.59	0.68	0.74	0.82	0.91	1.00	0.11	0.22

R: minimax regret value	N = 0	N = 1	N = 2	N = 3	N = 4	N = 5	N = 6	N = 7	N = 8	N = 9	N = 10
$p(0) = 0.10$	0.09	0.067	0.052	0.041	0.033	0.027	0.022	0.019	0.017	0.017	0.017
$p(0) = 0.25$	0.19	0.090	0.052	0.039	0.038	0.035	0.030	0.027	0.027	0.025	0.023
$p(0) = 0.50$	0.25	0.063	0.063	0.044	0.044	0.035	0.035	0.030	0.030	0.027	0.027
$p(0) = 0.75$	0.19	0.090	0.052	0.039	0.038	0.035	0.030	0.027	0.027	0.025	0.023
$p(0) = 0.90$	0.09	0.067	0.052	0.041	0.033	0.027	0.022	0.019	0.017	0.017	0.016

$p_R(1)$: minimax regret state	N = 0	N = 1	N = 2	N = 3	N = 4	N = 5	N = 6	N = 7	N = 8	N = 9	N = 10
$p(0) = 0.10$	0, 1	0	0	0	0.28	0.01	0.23	0.21	0.05	0.05	0.20
$p(0) = 0.25$	0, 1	0	0.50	0.45	0.45	0.42	0.40	0.15	0.38	0.15	0.36
$p(0) = 0.50$	0, 1	0.25 0.75	0.25 0.75	0.32 0.68	0.32 0.68	0.35 0.65	0.35 0.65	0.37 0.63	0.37 0.63	0.38 0.62	0.38 0.62
$p(0) = 0.75$	0, 1	1	0.50	0.55	0.55	0.58	0.60	0.85	0.62	0.85	0.64
$p(0) = 0.90$	0, 1	1	1	1	0.72	0.99	0.77	0.79	0.95	0.95	0.80

