

STATISTICAL TREATMENT RULES FOR HETEROGENEOUS POPULATIONS

BY CHARLES F. MANSKI¹

An important objective of empirical research on treatment response is to provide decision makers with information useful in choosing treatments. This paper studies minimax-regret treatment choice using the sample data generated by a classical randomized experiment. Consider a utilitarian social planner who must choose among the feasible *statistical treatment rules*, these being functions that map the sample data and observed covariates of population members into a treatment allocation. If the planner knew the population distribution of treatment response, the optimal treatment rule would maximize mean welfare conditional on all observed covariates. The appropriate use of covariate information is a more subtle matter when only sample data on treatment response are available. I consider the class of *conditional empirical success rules*; that is, rules assigning persons to treatments that yield the best experimental outcomes conditional on alternative subsets of the observed covariates. I derive a closed-form bound on the maximum regret of any such rule. Comparison of the bounds for rules that condition on smaller and larger subsets of the covariates yields *sufficient sample sizes* for productive use of covariate information. When the available sample size exceeds the sufficiency boundary, a planner can be certain that conditioning treatment choice on more covariates is preferable (in terms of minimax regret) to conditioning on fewer covariates.

KEYWORDS: Finite sample theory, large-deviations theory, minimax regret, randomized experiments, risk, social planner, statistical decision theory, treatment response.

1. INTRODUCTION

AN IMPORTANT OBJECTIVE of empirical research on treatment response is to provide decision makers with information useful in choosing treatments. This paper uses the Wald (1950) development of statistical decision theory and, in particular, the Savage (1951) minimax-regret criterion to study treatment choice using the sample data generated by a classical randomized experiment. Consider a utilitarian social planner who must choose among the feasible *statistical treatment rules*, these being functions that map the sample data and observed covariates of population members into a treatment allocation. If the planner knew the population distribution of treatment response, the optimal

¹The research presented here was supported in part by National Science Foundation Grants SES-0001436 and SES-0314312. An early version of this paper was circulated in May 1999 as National Bureau of Economic Research Technical Working Paper No. 242 and subsequent versions were circulated in October 2001 and May 2003. I am grateful for the opportunity to present aspects of this work in seminars at Duke University, Harvard-MIT, 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 benefited from the comments of Hidehiko Ichimura, Francesca Molinari, and anonymous referees. I am grateful to Jörg Stoye for able research assistance and for comments.

treatment rule would maximize mean welfare conditional on all observed covariates (Manski (2000)). However, the appropriate use of covariate information is a more subtle matter when only sample data on treatment response are available.²

Section 2 sets forth the planner's problem in generality and gives a simple illustration. Section 3 develops the original contribution of the paper. I consider the class of *conditional empirical success* rules; that is, rules assigning persons to treatments that yield the best experimental outcomes conditional on alternative subsets of the observed covariates. I use a large-deviations theorem of Hoeffding (1963) to obtain a closed-form bound on the maximum regret of any such rule. I show that comparison of the bounds for rules that condition on smaller and larger subsets of the covariates yields *sufficient sample sizes* for productive use of covariate information. When the available sample size exceeds the sufficiency boundary, a planner can be certain that conditioning treatment choice on more covariates is preferable (in terms of minimax regret) to conditioning on fewer covariates.

Stafford (1985, pp. 112–114) appears to have been the first to explicitly consider treatment choice as a utilitarian social planning problem. Maintaining the utilitarian paradigm, Manski (1996, 2000) showed how identification problems generate ambiguity about the identity of optimal treatment choices. Heckman, Smith, and Clements (1997) and Heckman and Smith (1998) have examined identification problems that arise when a social preference for some form of equity in outcomes motivates nonutilitarian treatment choice.

Some aspects of treatment choice using sample data have been studied from the Bayesian perspective by Rubin (1978) and Dehejia (2004), among others. There appears to be no precedent application of Wald's frequentist statistical decision theory in general, nor of the minimax-regret criterion in particular. Instead, frequentist analysis of treatment choice has been influenced by the theory of hypothesis testing; see the discussion in Section 3.5.

Treatment choice using sample data and observable covariates is conceptually similar to, but formally distinct from, prediction of outcomes using sample data and observable covariates. Statisticians studying prediction under square loss have used statistical decision theory to motivate *shrinkage estimators* (e.g., Copas (1983) and Lehmann (1983, Section 4.6)) and *variable selection procedures* (e.g., Kempthorne (1984) and Droge (1998)). In particular, Droge (1998) determined variable-selection procedures that minimize maximum regret within the family of orthogonal series estimators. The statistical literature on prediction with sample data is important, but its findings do not appear

²For example, the planner may be a physician who chooses medical treatments for his patients. The physician may observe each patient's medical history and the results of diagnostic tests. He may also read the findings of a randomized clinical trial showing how samples of patients with varying covariates respond to different treatments. The physician's problem is to use these data to make treatment choices.

to hold specific lessons for treatment choice. The reason is that prediction to minimize the expected value of square loss and treatment choice to maximize a utilitarian social welfare function are rather different mathematical problems.

Before commencing the analysis, I think it prudent to observe that the strength of this paper (or any other application of statistical decision theory) is also its vulnerability. The strength of the paper is that it takes an explicit stand on the problem of treatment choice with sample data and, in return, delivers specific conclusions about what constitutes a good treatment rule. The vulnerability is that reasonable persons may prefer other formulations of the problem and, hence, may conclude that the findings reported here do not meet their needs. Perspectives on what constitutes the problem of interest may vary in at least four respects.

First, a planner may or may not aim to maximize a utilitarian social welfare function. My assumption that the planner wants to maximize population mean welfare carries forward a long tradition in welfare economics and is analytically convenient. However, I would not assert that the utilitarian perspective is necessarily realistic in all settings. Analysis of treatment choice with sample data from nonutilitarian perspectives would be welcome, but is beyond the scope of this paper.

Second, a planner may or may not find Wald's frequentist statistical decision theory to be appealing. Decision theorists of the conditional Bayes school have long argued that statistical inference should be conditioned on the observed data and not rest on thought experiments that contemplate how a procedure would perform in repeated sampling; see, for example, Berger (1985, Chapter 1). Contrariwise, frequentists have observed with discomfort that Bayesian inference requires assertion of a subjective prior distribution on unknown states of nature. Some researchers have sought to blend aspects of Bayesian and frequentist thinking (e.g., Samaniego and Reneau (1994) and Chamberlain (2000)). The continuing doctrinal debates between and within the frequentist and Bayesian schools of thought will not be settled here.³

Third, a planner who does find Wald's theory appealing at a general level may or may not concur with its specific recommendation to measure the performance of statistical decision functions by *risk*; that is, by expected performance in repeated sampling. Risk is the appropriate measure if a planner is risk neutral and is analytically convenient as well. However, Wald's recommendation that decision functions be evaluated as procedures does not require use of risk. For example, a planner could reasonably measure performance by the median

³Going beyond conceptual criticism, Bayesians sometimes observe that the Wald approach can be computationally demanding. Indeed, computational difficulties may have been a reason why 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 relevant today. Moreover, Section 3 of this paper shows that it is possible to make analytical progress that simplifies computations.

value of welfare in repeated sampling. Or he could follow Vapnik (1982), who studied prediction under square loss from a frequentist perspective but measured performance by the probability that a decision function achieves a suitably satisfactory result in repeated sampling.

Fourth, a planner who does measure performance by risk may or may not want to use the minimax-regret criterion to choose a treatment rule. Section 2.3 discusses alternative criteria.

2. THE PLANNER'S PROBLEM

Section 2.1 poses the decision problem that, I presume, the planner wants to solve. This problem, which slightly generalizes one examined in Manski (2000), 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 problem* because planners usually do not have the knowledge of treatment response needed to implement the optimal rule. Sections 2.2 and 2.3 develop the problem of treatment choice with sample data.

2.1. *The Idealized Problem*

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 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 assign a treatment to each member of J . A fully specified treatment rule is a function $\tau(\cdot): J \rightarrow T$ assigning 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 covariates $x_j \in X$ for each person; thus, $x: J \rightarrow X$ is the random variable mapping persons into their covariates. 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.⁴

⁴Institutional constraints may restrict the feasible rules to a proper subset of these functions. For example, the planner may be precluded from using certain covariates (say race or gender)

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 is the set of *singleton* rules that assign all persons with the same observed covariates to one treatment; thus, $z(\cdot, \cdot)$ 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 $s \neq t$. Nonsingleton rules randomly allocate persons with covariates ξ across multiple treatments, with assignment shares $[z(t, \xi), t \in T]$. This definition of nonsingleton 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 has a utilitarian objective.

The planner wants to maximize 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 mean welfare that would be realized if the planner were to choose rule z is

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

Thus, the planner wants to solve the problem $\max_{z \in Z} U(z, P)$. The maximum is achieved by a singleton rule that allocates all persons with covariates ξ to a treatment solving the problem⁵

$$(2) \quad \max_{t \in T} E\{u[y(t), t, \xi] | x = \xi\}.$$

The population mean welfare achieved by an optimal rule is

$$(3) \quad U^*(P) \equiv \sum_{\xi \in X} P(x = \xi) \left\{ \max_{t \in T} E\{u[y(t), t, \xi] | x = \xi\} \right\}.$$

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

A planner facing the above decision problem can choose an optimal treatment rule if he knows the conditional response distributions $\{P[y(t) | x = \xi], (t, \xi) \in T \times X\}$. Full knowledge of these distributions is rare, but sample data

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.

⁵There is a unique optimal rule if problem (2) has a unique solution for every $\xi \in X$. There are multiple optimal rules if problem (2) 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.

on treatment response may be available. The question is: How may planners use sample data to make treatment choices?

To address this question, we first need to generalize the concept of a treatment rule. *Statistical treatment rules* 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.

Following Wald (1950), I evaluate statistical treatment rules as procedures applied as the sampling process is engaged repeatedly to draw independent data samples. I follow Wald in measuring the performance of a treatment rule by its risk; that is, its expected result across realizations of the sampling process. Risk will be called *expected welfare* here, where the objective is to maximize population mean welfare.

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

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

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, $E[z(t, \xi, \psi)] = Q[z(t, \xi, \psi) = 1]$.

ILLUSTRATION 1: Suppose that there are two treatments, a binary outcome, and no observed covariates; thus, $T = \{0, 1\}$, $Y = \{0, 1\}$, and $X = \emptyset$. 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, but does not know the response distribution $P[y(1)]$ of the innovation. Let welfare be the outcome of a treatment; thus, $u[y_j(t), t, x_j] = y_j(t), j \in J, t \in T$.

An experiment is performed to learn about response to the innovation, with N subjects randomly drawn from the population and assigned to treatment 1. A number n of these subjects realize $y = 1$ and the remaining $N - n$

realize $y = 0$. Assuming that all outcomes are observed, the sample size N indexes the sampling process and the number n of successes is a sufficient statistic for the data. The feasible statistical treatment rules are functions $z : T \times [0, \dots, N] \rightarrow [0, 1]$ that map the number of successes into a treatment allocation; 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. The expected welfare of rule z is

$$(5) \quad W(z, P, N) = p(0) \cdot E[z(0, n)] + p(1) \cdot E[z(1, n)] \\ = p(0) + [p(1) - p(0)] \cdot E[z(1, n)],$$

where $p(t) \equiv P[y(t) = 1]$, $t \in T$. The number of experimental successes is distributed binomial $\mathbf{B}[p(1), N]$, so

$$(6) \quad E[z(1, n)] = \sum_{i=0}^N z(1, i) \cdot b[n = i; p(1), N],$$

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

2.3. Implementable Criteria for Treatment Choice

Maximization of expected welfare yields the optimal treatment rule, but is not feasible without knowledge of P . To develop implementable criteria for treatment choice, let Γ index the set of feasible states of nature; thus, $[(P_\gamma, Q_\gamma), \gamma \in \Gamma]$ is the set of (P, Q) pairs that the planner deems possible. Then one might choose a rule that works well on average or one that works well uniformly over Γ .

Bayesian statistical decision theory suggests averaging over Γ . Thus, a Bayesian planner might place a subjective probability measure π on Γ and solve the problem $\sup_{z \in Z} \int W(z, P_\gamma, Q_\gamma) d\pi(\gamma)$.⁶ A non-Bayesian planner might choose a treatment rule that behaves uniformly well, in the sense of maximin or minimax-regret. The maximin criterion is $\sup_{z \in Z} \inf_{\gamma \in \Gamma} W(z, P_\gamma, Q_\gamma)$. The minimax-regret criterion is

$$(7) \quad \inf_{z \in Z} \sup_{\gamma \in \Gamma} U^*(P_\gamma) - W(z, P_\gamma, Q_\gamma),$$

⁶Solution of this problem, which views treatment rules as procedures, differs from the conditional Bayes prescription for decision making. The latter calls on the planner to form a posterior subjective distribution on Γ , conditional on the sample data, and to maximize the expected value of $U(z, P)$ with respect to this posterior distribution. Although the two problems differ, solution of the latter one at all points in the sample space yields the solution to the former one. See Berger (1985, Section 4.4.1).

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

$$(8) \quad U^*(P_\gamma) \equiv \sum_{\xi \in X} P_\gamma(x = \xi) \left\{ \max_{t \in T} \int u[y(t), t, \xi] dP_\gamma[y(t)|x = \xi] \right\}.$$

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. Berger (1985) nicely exposit these decision criteria; see Section 1.5.2 and Chapter 5.

The maximin and minimax-regret criteria differently interpret the idea that the planner should choose a rule that behaves uniformly well across all states of nature. Whereas a maximin 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 state of nature. Savage (1951), whose review of Wald (1950) first explicitly distinguished between these criteria for decision making, argued 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.” Savage emphasized that the minimax-regret criterion is not similarly “ultra-pessimistic.” The illustration below shows how different the two criteria can be.

ILLUSTRATION 2: Continuing the illustration of Section 2.2, the only unknown determinant of expected welfare is $p(1)$. Hence, Γ is the set of feasible values of $p(1)$. Let $[p_\gamma(1), \gamma \in \Gamma] = (0, 1)$.

A Bayesian planner might assert a Beta prior, with parameters (α, β) . The posterior mean for $p(1)$ is $(\alpha + n)/(\alpha + \beta + N)$, so the resulting Bayes rule sets $z(1, n) = 0$ when $(\alpha + n)/(\alpha + \beta + N) < p(0)$ and $z(1, n) = 1$ when $(\alpha + n)/(\alpha + \beta + N) > p(0)$.

The maximin criterion always chooses treatment zero, regardless of the sample data. By (5), the infimum of expected welfare for rule z is $p(0) + \inf_{\gamma \in \Gamma} [p_\gamma(1) - p(0)] \cdot E_\gamma[z(1, n)]$, where $E_\gamma[z(1, n)]$ is the expression in (6) with $p_\gamma(1)$ replacing $p(1)$. The set $[p_\gamma(1), \gamma \in \Gamma]$ contains values smaller than $p(0)$. Hence, $\inf_{\gamma \in \Gamma} [p_\gamma(1) - p(0)] \cdot E_\gamma[z(1, n)] \leq 0$. The rule that always chooses treatment 0 makes $E_\gamma[z(1, n)] = 0, \forall \gamma \in \Gamma$. Hence, this is the maximin rule.

Now consider minimax-regret. The regret of rule z in state of nature γ is the mean welfare loss when a person is assigned the inferior treatment, multiplied by the expected fraction of the population assigned this treatment. Maximum

regret is

$$(9) \quad R(z) \equiv \sup_{\gamma \in \Gamma} [p_\gamma(1) - p(0)] \cdot E_\gamma[z(0, n)] \cdot \mathbb{1}[p_\gamma(1) \geq p(0)] \\ + [p(0) - p_\gamma(1)] \cdot E_\gamma[z(1, n)] \cdot \mathbb{1}[p(0) \geq p_\gamma(1)].$$

The problem $\min_{z \in Z} R(z)$ may be solved numerically to obtain the minimax-regret rule. Computations not reported here show that this rule is well approximated by the *empirical-success* rule which replaces $p(1)$ with the empirical success rate n/N ; the latter rule sets $z(1, n) = 0$ if $n/N < p(0)$ and $z(1, n) = 1$ when $n/N > p(0)$.

2.4. Discussion

Wald’s statistical decision theory is operational whenever expected welfare $W(\cdot, \cdot, \cdot)$ exists on its domain $Z \times [(P_\gamma, Q_\gamma), \gamma \in \Gamma]$ and a well-defined criterion is used to choose among the feasible treatment rules. Given these conditions, the theory enables comparison of all feasible statistical treatment rules. It addresses the problem of finite-sample statistical inference directly, without recourse to the large-sample approximations of asymptotic statistical theory.⁷ It applies whatever the sampling process may be and whatever information the planner may have about the population and the sampling process.

Wald’s theory also enables comparison of alternative sampling processes. Consider a two-period world, with 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 welfare. Then the expected welfare of (treatment rule, sampling process) pair (z, Q) is $W(z, P, Q) - C(Q)$. Section 3.5 illustrates how the minimax-regret criterion can be used to choose a sample design.

3. TREATMENT CHOICE USING DATA FROM A RANDOMIZED EXPERIMENT

Henceforth, I suppose that there are two treatments and that outcomes are real-valued with bounded range. The sample data are the realizations of a classical experiment randomly assigning subjects to these treatments. I consider designs that draw subjects at random within groups stratified by covariates and

⁷Indeed, the concept of regret provides a 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 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 $\lim_{N \rightarrow \infty} \{\sup_{\gamma \in \Gamma} U^*(P_\gamma) - W(z_N, P_\gamma, Q_N)\} = 0$.

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 which treatment to choose when the data are from an experiment with a finite sample of subjects. I address this question from the minimax-regret perspective, focusing on the class of conditional empirical success rules.

Section 3.1 sets forth the maintained assumptions. Section 3.2 defines conditional empirical success rules formally. Section 3.3 develops closed-form bounds on maximum regret for experimental designs that use a stratified random sampling process to draw subjects. These bounds yield sufficient sample sizes for productive use of covariate information. Section 3.4 extends the analysis to designs that draw subjects by simple random sampling. Section 3.5 presents numerical findings that make evident the practical implications for treatment choice. The concluding Section 3.6 briefly discusses the use of conditional empirical success rules when the data are not from a classical randomized experiment.

3.1. *The Setting*

Suppose that there are two treatments; thus, $T = \{0, 1\}$. There is a real 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 covariate distribution $P(x)$. Welfare is the outcome of a treatment; thus, $u[y_j(t), t, x_j] = y_j(t)$, $j \in J$, $t \in T$.⁸

A randomized experiment is performed to learn about treatment response, and the outcomes of all subjects are observed. I consider two experimental designs, stratified and simple random sampling.⁹

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 subpopulation with covariates ξ and assigns them 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 subsample of subjects with covariates ξ who are assigned to treatment t . Then the sample data are

⁸The analysis in this section extends to problems in which welfare has the general form $u[y(t), t, x]$, the function $u[\cdot, \cdot, x]$ having bounded range that may depend on x . This extension involves no conceptual difficulties but much complication of notation.

⁹These designs 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 3.4 can be modified to cover intermediate cases, in which the design stratifies on some but not all covariates.

the outcomes $Y_{TX} \equiv [y_j, j \in N(t, \xi); (t, \xi) \in T \times X]$. The feasible rules 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

$$(10) \quad W(z, P, N_{TX}) = \sum_{\xi \in X} P(x = \xi) \cdot \{E[y(0)|x = \xi] \cdot E[z(0, \xi, Y_{TX})] + E[y(1)|x = \xi] \cdot E[z(1, \xi, Y_{TX})]\}.$$

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} . Rule 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

$$(11) \quad W(z, P, N, q) = \sum_{\xi \in X} P(x = \xi) \cdot \{E[y(0)|x = \xi] \cdot E[z(0, \xi, N_{TX}, Y_{TX})] + E[y(1)|x = \xi] \cdot E[z(1, \xi, N_{TX}, Y_{TX})]\}.$$

3.2. Conditional Empirical Success (CES) Rules

Comparison of all feasible treatment rules is beyond the scope of this paper. Instead, I restrict attention to a tractable class of rules. Statisticians studying estimation have long made progress by restricting attention to tractable classes of estimators; for example, linear unbiased or asymptotic normal ones. Similarly, I make progress here by restricting attention to *conditional empirical success* rules.

An empirical success rule emulates the optimal treatment rule by replacing unknown response distributions with sample analogs. When welfare is the outcome of 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 $\bar{y}_{t\xi} \equiv (1/N_{t\xi}) \sum_{j \in N(t, \xi)} y_j$ and an empirical success rule solves the problems $\max_{t \in T} \bar{y}_{t\xi}$, $\xi \in X$.

The above rule conditions treatment choice on all observed covariates, as is optimal in the idealized problem. However, conditioning tends to diminish the

statistical precision of the sample averages used to estimate population means. Hence, 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 observed covariates. Formally, let $v(\cdot) : X \rightarrow V$ map the covariate space X into a space V . Then a conditional empirical success rule chooses treatments that maximize empirical success conditional on a person's value of v .

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 $E[y(t)|v = \nu]$ is $\bar{y}_{t\nu} \equiv (1/N_{t\nu}) \sum_{j \in N(t,\nu)} y_j$, where $N(t, \nu)$ is the subsample of subjects with covariates ν who are assigned to treatment t , and where $N_{t\nu} \equiv |N(t, \nu)|$. In a design with stratified random sampling, the appropriate sample analog is the design-weighted average $\bar{y}_{tX\nu} \equiv \sum_{\xi \in X} \bar{y}_{t\xi} \cdot P(x = \xi|v = \nu)$.¹⁰

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.¹¹ 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.

Let z_{Xv} denote the rule conditioning on v in a design with stratified random sampling, and let z_v denote the corresponding rule in a design with simple random sampling. Then expected welfare is

$$(12) \quad W(z_{Xv}, P, N_{TX}) = \sum_{\nu \in V} P(v = \nu) \cdot \{ E[y(0)|v = \nu] \cdot P(\bar{y}_{0X\nu} \geq \bar{y}_{1X\nu}) + E[y(1)|v = \nu] \cdot P(\bar{y}_{1X\nu} > \bar{y}_{0X\nu}) \}$$

in a design with stratified random sampling and

$$(13) \quad W(z_v, P, N, q) = \sum_{\nu \in V} P(v = \nu) \cdot \{ E[y(0)|v = \nu] \cdot P(\bar{y}_{0\nu} \geq \bar{y}_{1\nu}) + E[y(1)|v = \nu] \cdot P(\bar{y}_{1\nu} > \bar{y}_{0\nu}) \}$$

in a design with simple random sampling.

¹⁰In a stratified design, $\bar{y}_{t\xi}$ is the sample analog of $E[y(t)|x = \xi]$, but $\bar{y}_{t\nu}$ is not that of $E[y(t)|v = \nu]$. However, $E[y(t)|v = \nu] = \sum_{\xi \in X} E[y(t)|x = \xi] \cdot P(x = \xi|v = \nu)$. Hence $\bar{y}_{tX\nu}$ is the sample analog of $E[y(t)|v = \nu]$.

¹¹Use of a tie-breaking convention that randomly allocates some persons to one treatment and the rest to the other would not materially alter the analysis in Sections 3.3 and 3.4.

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

Direct analysis of the expected welfare of CES rules is arduous. The treatment-selection probabilities are probabilities that one sample average exceeds another. Such probabilities generically do not have closed-form expressions and are difficult to compute numerically.

Fortunately, it is possible to develop useful closed-form bounds on expected welfare. The upper bound is simple. The highest population welfare attainable by any rule conditioning treatment choice on covariates v is $\sum_{v \in V} P(v = \nu) \max\{E[y(0)|v = \nu], E[y(1)|v = \nu]\}$; hence, $W(z_{X\nu}, P, N_{TX})$ and $W(z_\nu, P, N, q)$ cannot exceed this value. This section develops the lower bound for designs with stratified random sampling, and Section 3.4 does the same for ones with simple random sampling. In each case, I use the bounds on expected welfare to obtain bounds on maximum regret.

The analysis exploits this 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 $d > 0$, $\Pr(\bar{w} - \mu \geq d) \leq \exp[-2K^2 d^2 \{\sum_{k=1}^K (b_k - a_k)^2\}^{-1}]$.

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.¹²

Bounding Expected Welfare: Proposition 1 uses the Hoeffding theorem to bound the expected welfare of a CES rule in a design with stratified sampling.

PROPOSITION 1: *Let subjects be drawn by stratified random sampling. Let $v(\cdot) : X \rightarrow V$ and consider rule $z_{X\nu}$. For $\nu \in V$, define $M_\nu \equiv \max\{E[y(0)|v = \nu], E[y(1)|v = \nu]\}$, and $\delta_\nu \equiv |E[y(1)|v = \nu] - E[y(0)|v = \nu]|$. Then the expected welfare of rule $z_{X\nu}$ satisfies the inequality*

$$(14) \quad \sum_{\nu \in V} P(v = \nu) M_\nu - D(z_{X\nu}, P, N_{TX}) \leq W(z_{X\nu}, P, N_{TX}) \leq \sum_{\nu \in V} P(v = \nu) M_\nu,$$

¹²The 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 nuisance parameters.

where

$$(15) \quad D(z_{X\nu}, P, N_{TX}) \equiv \sum_{\nu \in V} P(\nu = \nu) \cdot \delta_\nu \exp \left[-2\delta_\nu^2 \left\{ \sum_{\xi \in X} P(x = \xi | \nu = \nu)^2 (N_{1\xi}^{-1} + N_{0\xi}^{-1}) \right\}^{-1} \right].$$

PROOF: Let $\nu \in V$. Recall the design-weighted sample average outcomes $\bar{y}_{1X\nu}$ defined in Section 3.2. I first write $\bar{y}_{1X\nu} - \bar{y}_{0X\nu}$ 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

$$(16) \quad \begin{aligned} \bar{y}_{1X\nu} - \bar{y}_{0X\nu} &= \sum_{\xi \in X} P(x = \xi | \nu = \nu) (1/N_{1\xi}) \sum_{j \in N(1, \xi)} y_j \\ &\quad - \sum_{\xi \in X} P(x = \xi | \nu = \nu) (1/N_{0\xi}) \sum_{j \in N(0, \xi)} y_j \\ &= (1/N) \left\{ \sum_{\xi \in X} \sum_{j \in N(1, \xi)} [y_j \cdot P(x = \xi | \nu = \nu) N/N_{1\xi}] \right. \\ &\quad \left. + \sum_{\xi \in X} \sum_{j \in N(0, \xi)} [-y_j \cdot P(x = \xi | \nu = \nu) N/N_{0\xi}] \right\}. \end{aligned}$$

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

Suppose that $E[y(1) | \nu = \nu] < E[y(0) | \nu = \nu]$. Then $E(\bar{y}_{1X\nu} - \bar{y}_{0X\nu}) = -\delta_\nu$. Application of the Hoeffding theorem yields¹³

$$(17) \quad \begin{aligned} P(\bar{y}_{1X\nu} > \bar{y}_{0X\nu}) &= P(\bar{y}_{1X\nu} - \bar{y}_{0X\nu} + \delta_\nu > \delta_\nu) \\ &\leq \exp \left[-2N^2\delta_\nu^2 \left\{ \sum_{\xi \in X} N_{1\xi} [P(x = \xi | \nu = \nu) \cdot N/N_{1\xi}]^2 \right. \right. \\ &\quad \left. \left. + N_{0\xi} [P(x = \xi | \nu = \nu) \cdot N/N_{0\xi}]^2 \right\}^{-1} \right] \\ &= \exp \left[-2\delta_\nu^2 \left\{ \sum_{\xi \in X} P(x = \xi | \nu = \nu)^2 (N_{1\xi}^{-1} + N_{0\xi}^{-1}) \right\}^{-1} \right]. \end{aligned}$$

¹³Inequality (17) 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 $\{P(x = \xi | \nu = \nu)^2 (N_{1\xi}^{-1} + N_{0\xi}^{-1})\}^{-1} = 0$.

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

$$(18) \quad P(\bar{y}_{0X\nu} \geq \bar{y}_{1X\nu}) \leq \exp \left[-2\delta_\nu^2 \left\{ \sum_{\xi \in X} P(x = \xi|v = \nu)^2 (N_{1\xi}^{-1} + N_{0\xi}^{-1}) \right\}^{-1} \right].$$

To obtain inequality (14), partition V into three regions:

$$\begin{aligned} V_a &\equiv \{v \in V : E[y(1)|v = \nu] < E[y(0)|v = \nu]\}, \\ V_b &\equiv \{v \in V : E[y(1)|v = \nu] > E[y(0)|v = \nu]\}, \\ V_c &\equiv \{v \in V : E[y(1)|v = \nu] = E[y(0)|v = \nu]\}. \end{aligned}$$

Then rewrite equation (12) as

$$\begin{aligned} (12') \quad W(z_{X\nu}, P, N_{TX}) &= \sum_{v \in V_a} P(v = \nu) \{ E[y(0)|v = \nu] \cdot P(\bar{y}_{0X\nu} \geq \bar{y}_{1X\nu}) \\ &\quad + E[y(1)|v = \nu] \cdot P(\bar{y}_{1X\nu} > \bar{y}_{0X\nu}) \} \\ &\quad + \sum_{v \in V_b} P(v = \nu) \{ E[y(0)|v = \nu] \cdot P(\bar{y}_{0X\nu} \geq \bar{y}_{1X\nu}) \\ &\quad + E[y(1)|v = \nu] \cdot P(\bar{y}_{1X\nu} > \bar{y}_{0X\nu}) \} \\ &\quad + \sum_{v \in V_c} P(v = \nu) \{ E[y(0)|v = \nu] \cdot P(\bar{y}_{0X\nu} \geq \bar{y}_{1X\nu}) \\ &\quad + E[y(1)|v = \nu] \cdot P(\bar{y}_{1X\nu} > \bar{y}_{0X\nu}) \}. \end{aligned}$$

The first term on the right-hand side of (12') can be no smaller than the expression obtained by setting $P(\bar{y}_{1X\nu} > \bar{y}_{0X\nu})$ at the upper bound obtained in (17) and setting $P(\bar{y}_{0X\nu} \geq \bar{y}_{1X\nu})$ at its implied lower bound. The lower bound on the second term on the right-hand side of (12') is similarly obtained, applying (18). The third term is constant across all treatment allocations. Placing all three terms on the right-hand side of (12') at their lower bounds yields the lower bound on $W(z_{X\nu}, P, N_{TX})$ given in (14). The upper bound in (14) is similarly obtained. Q.E.D.

The bound on expected welfare obtained in Proposition 1 is a closed-form function of the stratum sample sizes N_{TX} , the covariate distribution $P(x)$, and the mean treatment outcomes $\{E[y(0)|v = \nu], E[y(1)|v = \nu], \nu \in V\}$. The upper bound is the maximum welfare achievable using covariates v . The lower bound differs from this ideal by the nonnegative *finite-sample penalty* $D(z_{X\nu}, P, N_{TX})$, 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.

For each $\nu \in V$, the finite-sample penalty varies nonmonotonically with δ_ν . In particular, it varies as $\delta_\nu \cdot \exp(-C_{N\nu}\delta_\nu^2)$, where $C_{N\nu} \equiv 2\{\sum_{\xi \in X} P(x = \xi | v = \nu)^2(N_{1\xi}^{-1} + N_{0\xi}^{-1})\}^{-1}$. If $\delta_\nu = 0$, there is no finite-sample penalty because both treatments are equally good for persons with covariates ν . As δ_ν 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 $\delta_\nu = (2C_{N\nu})^{-1/2}$. Inserting the worst-case values $[\delta_\nu = (2C_{N\nu})^{-1/2}, \nu \in V]$ into (15) yields this uniform upper bound on the finite-sample penalty:

$$(19) \quad D(z_{X\nu}, P, N_{TX}) \leq \frac{1}{2}e^{-1/2} \sum_{\nu \in V} P(v = \nu) \left\{ \sum_{\xi \in X} P(x = \xi | v = \nu)^2 (N_{1\xi}^{-1} + N_{0\xi}^{-1}) \right\}^{1/2}.$$

Bounding Maximum Regret: Proposition 1 immediately yields a bound on regret, this being

$$(20) \quad \sum_{\xi \in X} P(x = \xi)M_\xi - \sum_{\nu \in V} P(v = \nu)M_\nu \leq U^*(P) - W(z_{X\nu}, P, N_{TX}) \leq \sum_{\xi \in X} P(x = \xi)M_\xi - \sum_{\nu \in V} P(v = \nu)M_\nu + D(z_{X\nu}, P, N_{TX}).$$

The lower bound 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 rule conditioning on v .

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 in bound (20) are $\{E[y(0)|x], E[y(1)|x]\}$. Maximizing the lower and upper bounds over the feasible values of $\{E[y(0)|x], E[y(1)|x]\}$ yields this bound on maximum regret:

$$(21) \quad R_{Lv} \leq R(z_{Xv}) \leq R_U(z_{Xv}),$$

where

$$R_{Lv} \equiv \sup_{\gamma \in \Gamma} \sum_{\xi \in X} P(x = \xi) \max\{E_\gamma[y(0)|x = \xi], E_\gamma[y(1)|x = \xi]\} - \sum_{\nu \in V} P(v = \nu) \max\{E_\gamma[y(0)|v = \nu], E_\gamma[y(1)|v = \nu]\},$$

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

$\delta_{v\gamma} \equiv |E_\gamma[y(1)|v = v] - E_\gamma[y(0)|v = v]|$, 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_{Lv} applies to any rule that conditions treatment choice on covariates v , not only to z_{Xv} . The upper bound $R_U(z_{Xv})$ is specific to this rule.

The bound on maximum regret simplifies in the case of rule z_{Xx} , which conditions treatment choice on all observed covariates. Then $V = X$ and $v(x) = x$, so (21) reduces to

$$(22) \quad 0 \leq R(z_{Xx}) \leq \sup_{\gamma \in \Gamma} \sum_{\xi \in X} P(x = \xi) \cdot \delta_{\xi\gamma} \cdot \exp[-2\delta_{\xi\gamma}^2 (N_{1\xi}^{-1} + N_{0\xi}^{-1})^{-1}].$$

The upper bound in (22) 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 (19) shows that (22) further reduces to¹⁴

$$(23) \quad 0 \leq R(z_{Xx}) \leq \frac{1}{2} e^{-1/2} \sum_{\xi \in X} P(x = \xi) (N_{1\xi}^{-1} + N_{0\xi}^{-1})^{1/2}.$$

Sufficient Sample Sizes for Productive Use of Covariate Information: Sufficient sample sizes for productive use of covariate information follow immediately from the bounds on maximum regret. Let $v(\cdot) : X \rightarrow V$ be a many-to-one mapping of x into a covariate v . The maximum regret of any rule conditioning treatment choice on v must exceed that of rule z_{Xx} if $R_{Lv} > R_U(z_{Xx})$, where

¹⁴Inequalities (22) and (23) show that rule z_{Xx} is uniformly consistent, and they yield lower bounds on its pointwise and uniform rates of convergence. Inequality (22) shows that if all elements of N_{Tx} grow at the same rate, regret converges to zero at least at exponential rate in every state of nature. Inequality (23) shows that if all elements of N_{Tx} grow at the same rate, maximum regret converges to zero at least with the square root of sample size. Inequality (23) also implies that the minimax-regret rule is uniformly consistent. Although the form of this rule is unknown, it necessarily has maximum regret no larger than the upper bound in (23). Hence, minimax regret converges to zero at a rate no lower than the square root of sample size.

R_{L_v} is given in (21) and $R_U(z_{X_x})$ in (22). The quantity $R_U(z_{X_x})$ is decreasing in each component of the vector N_{TX} of stratum sample sizes. Hence, sufficient sample sizes for conditioning on x to be preferable to conditioning on v are solutions to the problem

$$(24) \quad \min N_{TX} : R_{L_v} > \sup_{\gamma \in I} \sum_{\xi \in X} P(x = \xi) \cdot \delta_{\xi\gamma} \cdot \exp[-2\delta_{\xi\gamma}^2 (N_{1\xi}^{-1} + N_{0\xi}^{-1})^{-1}].$$

Several aspects of this derivation warrant comment. First, N_{TX} being a vector, problem (24) generically has multiple solutions; hence, I refer to sufficient sample sizes (plural). Second, when N_{TX} exceeds a sufficiency boundary, the maximum regret of rule z_{X_x} is smaller than that of *all* rules conditioning treatment choice on v , not just smaller than that of rule z_{X_v} . Third, the sufficiency boundaries provide a sufficient condition for superiority of rule z_{X_x} 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_{L_v} , and the maximum regret of rule z_{X_x} may be less than $R_U(z_{X_x})$. Fourth, the present analysis does not show that rule z_{X_x} is the best rule conditioning treatment choice on x ; there may exist a non-CES rule that is superior to z_{X_x} .

The above discussion compares conditioning on x with conditioning on v . Sufficient sample sizes for other covariate 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. Then conditioning on w is necessarily preferable to conditioning on v if N_{TX} is such that $R_{L_v} > R_U(z_{X_w})$, where $R_U(z_{X_w})$ is defined in (21).

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

In designs with simple random sampling, the stratum sample sizes N_{TX} are random rather than fixed. However, the reasoning used to prove Proposition 1 continues to apply conditional on any realization of N_{TX} . Hence, Proposition 1 provides the “inner loop” for analysis of simple random sampling. Here is the result.

PROPOSITION 2: *Let subjects be drawn by simple random sampling. Let $v(\cdot) : X \rightarrow V$ and consider rule z_v . For $v \in V$, let B_{N_v} denote the Binomial distribution $\mathbf{B}[P(v = v), N]$. For $n = 0, \dots, N$, let B_{n_q} denote the Binomial distribution $\mathbf{B}(q, n)$.¹⁵ Then the expected welfare of rule z_v satisfies the inequality*

$$(25) \quad \sum_{v \in V} P(v = v)M_v - D(z_v, P, N, q) \leq W(z_v, N, P, q) \leq \sum_{v \in V} P(v = v)M_v,$$

¹⁵If $n = 0$, define this distribution to be degenerate with all mass on the value zero.

where

$$(26) \quad D(z_v, p, N, q) \equiv \sum_{v \in V} P(v = \nu) \cdot \delta_\nu \sum_{n=0}^N \sum_{m=0}^n B_{N\nu}(n) \cdot B_{nq}(m) \cdot \exp\{-2\delta_\nu^2[(n - m)^{-1} + m^{-1}]^{-1}\}.$$

PROOF: Consider rule z_x , which conditions on all observed covariates. Rule z_x is algebraically the same as rule z_{Xx} , considered in Section 3.3. These rules 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 1 hold under simple random sampling, conditional on the realization of N_{TX} . For $\xi \in X$, these bounds are

$$(27a) \quad P(\bar{y}_{1\xi} > \bar{y}_{0\xi} | N_{TX}) \leq \exp[-2\delta_\xi^2(N_{1\xi}^{-1} + N_{0\xi}^{-1})^{-1}]$$

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

$$(27b) \quad P(\bar{y}_{0\xi} \geq \bar{y}_{1\xi} | N_{TX}) \leq \exp[-2\delta_\xi^2(N_{1\xi}^{-1} + N_{0\xi}^{-1})^{-1}]$$

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

$$(28a) \quad P(\bar{y}_{1\xi} > \bar{y}_{0\xi}) \leq \sum_{n=0}^N \sum_{m=0}^n B_{N\xi}(n) B_{nq}(m) \exp\{-2\delta_\xi^2[(n - m)^{-1} + m^{-1}]^{-1}\}$$

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

$$(28b) \quad P(\bar{y}_{0\xi} \geq \bar{y}_{1\xi}) \leq \sum_{n=0}^N \sum_{m=0}^n B_{N\xi}(n) B_{nq}(m) \exp\{-2\delta_\xi^2[(n - m)^{-1} + m^{-1}]^{-1}\}$$

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

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 1 to a sampling process that stratifies on v rather than on x . Q.E.D.

The bound on expected welfare obtained in Proposition 2 is a closed form function of the sample size N , the treatment assignment probability q , and the mean treatment outcomes $\{E[y(0)|v = \nu], E[y(1)|v = \nu], \nu \in V\}$. As in Propo-

sition 1, 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 $D(z_v, N, P, q)$.

Proposition 2 yields this bound on regret:

$$\begin{aligned}
 (29) \quad & \sum_{\xi \in X} P(x = \xi)M_\xi - \sum_{v \in V} P(v = \nu)M_\nu \\
 & \leq U^*(P) - W(z_v, P, N, q) \\
 & \leq \sum_{\xi \in X} P(x = \xi)M_\xi - \sum_{v \in V} P(v = \nu)M_\nu + D(z_v, P, N, q).
 \end{aligned}$$

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

$$(30) \quad R_{L,v} \leq R(z_v) \leq R_U(z_v),$$

where $R_{L,v}$ was defined in (21) and

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

Sufficient sample sizes for productive use of covariate information follow from (30). Let $v(\cdot) : X \rightarrow V$ and $w(\cdot) : X \rightarrow W$ be distinct mappings of X into covariates v and w respectively. Then conditioning on w is necessarily preferable to conditioning on v if N is such that $R_{L,v} > R_U(z_w)$.

3.5. Numerical Findings for Binary Covariates

Computation of the bounds on maximum regret is simple and revealing when 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 $\{E[y(t)|x = \xi]; t = 0, 1; \xi = a, b\}$. The present computations suppose that all states of natures are feasible.

The upper bounds on maximum regret can be computed numerically. The lower bound for the rule that does not condition treatment choice on x is

$$\begin{aligned}
 (31) \quad R_{L\phi} = & \sup_{\gamma \in \Gamma} P(x = a) \max\{E_\gamma[y(0)|x = a], E_\gamma[y(1)|x = a]\} \\
 & + P(x = b) \max\{E_\gamma[y(0)|x = b], E_\gamma[y(1)|x = b]\} \\
 & - \max\{P(x = a)E_\gamma[y(0)|x = a] + P(x = b)E_\gamma[y(0)|x = b], \\
 & \quad P(x = a)E_\gamma[y(1)|x = a] + P(x = b)E_\gamma[y(1)|x = b]\},
 \end{aligned}$$

where $\{E_\gamma[y(t)|x = \xi]; t = 0, 1; \xi = a, b\}$ can take values in the unit hypercube $[0, 1]^4$. The right-hand side of (31) is $\min\{P(x = a), P(x = b)\}$.¹⁶ Thus, $R_{L\phi} = \min\{P(x = a), P(x = b)\}$.

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 I computes the upper bound on maximum regret of each rule in designs with equal treatment assignment probabilities ($q = .5$) and sample sizes ranging from 1 to 200. Three covariate distributions are considered, with $P(x = a) = .05, .25, \text{ or } .5$. The table shows that the sufficient sample size for productive use of covariate information lies between $N = 100$ and $N = 200$ when $P(x = a) = .05$, is $N = 15$ when $P(x = a) = .25$, and is $N = 6$ when $P(x = a) = .50$.

These numerical findings suggest that prevailing practices in the use of covariate information in 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 mean response is the same across values of x . This hypothesis is rarely rejected in small samples, so use of covariate information in treatment choice is commonly viewed as imprudent. The findings in Table I suggest that conditioning treatment choice on covariates is warranted in samples far smaller than those required to show statistically significant differences in treatment response across covariate values.

Numerical findings aside, testing hypotheses 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 do not address the planner’s problem.

¹⁶Without 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 (31) is $E_\gamma[y(0)|x = a] = E_\gamma[y(1)|x = b] = 0, E_\gamma[y(1)|x = a] = 1, E_\gamma[y(0)|x = b] = P(x = a)/P(x = b)$.

TABLE I
 UPPER BOUND ON MAXIMUM REGRET OF CES RULES UNDER SIMPLE
 RANDOM SAMPLING ($q = .5$)

N	Rule $z_x, P(x = a)$			Rule $z_\phi, P(x = a)$		
	.05	.25	.50	.05	.25	.50
1	1	1	1	1	1	1
2	.729	.862	.921	.700	.695	.684
3	.510	.712	.812	.506	.644	.5
4	.404	.582	.699	.410	.598	.5
5	.352	.503	.592	.358	.558	.5
6	.318	.456	.496	.323	.523	.5
7	.293	.414	.418	.299	.499	.5
8	.275	.382	.372	.278	.479	.5
9	.258	.355	.340	.264	.465	.5
10	.243	.331	.313	.253	.452	.5
11	.235	.309	.295	.242	.440	.5
12	.227	.289	.278	.232	.432	.5
13	.219	.270	.262	.222	.425	.5
14	.212	.256	.252	.214	.418	.5
15	.204	.243	.242	.210	.412	.5
16	.197	.233	.233	.205	.406	.5
17	.191	.223	.224	.201	.400	.5
18	.184	.215	.215	.197	.395	.5
19	.180	.207	.207	.193	.390	.5
20	.177	.200	.202	.189	.387	.5
24	.166	.177	.184	.175	.376	.5
28	.155	.164	.169	.162	.367	.5
32	.145	.153	.154	.153	.358	.5
36	.135	.144	.143	.149	.349	.5
40	.126	.135	.137	.146	.342	.5
44	.117	.128	.132	.142	.339	.5
48	.109	.121	.127	.138	.337	.5
52	.104	.115	.122	.135	.334	.5
100	.076	.083	.078	.103	.308	.5
200	.047	.055	.061	.087	.289	.5

Quasi-Optimal Stratified Designs: A stratified 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. A natural way to study treatment choice with stratified data is to fix the overall sample size N and, for each feasible rule, determine the value of N_{TX} that minimizes maximum regret subject to the constraint $\sum_{(t,x)} N_{tx} = N$.

The analysis of Section 3.3 does not yield optimal stratified designs, but it does yield *quasi-optimal* designs that minimize the upper bound on maximum

regret.¹⁷ Table II shows the quasi-optimal designs for the two feasible rules z_{X_x} and z_{X_ϕ} , when N ranges from 1 to 52 and when $P(x = a) = .05$ or $.25$. The covariate distribution with $P(x = a) = .5$ is omitted from the table because the quasi-optimal design in this case generically is the equal-shares allocation $\{N_{t\xi} = N/4; t = 0, 1; \xi = a, b\}$.

The first two columns of each panel of Table II show the quasi-optimal stratification, and the third column 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, stratum sample sizes being integers, this condition is not strictly implementable when N is odd. The entries (1, 2) for rule z_{X_x} 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 II shows that the quasi-optimal stratification generically draws more subjects with covariate value $x = b$ than with $x = a$. 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_ϕ} 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_ϕ} does not condition treatment choice on the covariate. Perhaps more interesting is the fact that the best designs for rule z_{X_x} 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 $P(x = a) = .05$. However, it appears that when $P(x = a) = .25$, the sampling ratio $N_{ta}/N_{tb} \approx 1/2$.

Comparison of Tables I and II shows how quasi-optimal stratified random sampling improves treatment choice relative to simple random sampling. When N is very small, rule z_{X_x} sometimes substantially outperforms its simple random sampling counterpart z_x . The advantage of stratification declines as N grows but remains nonnegligible throughout the range of sample sizes considered here. The tables indicate that stratification is productive only when one conditions treatment choice on x ; rule z_{X_ϕ} does not outperform its simple random sampling counterpart z_ϕ .

3.6. Using CES Rules in Nonclassical Settings

To conclude, I point out that CES rules may be used to choose treatments when the data are not from a classical randomized experiment. Rules z_{X_v} and z_v are well defined when some experimental subjects do not comply

¹⁷I call these "quasi" optimal designs for two reasons. One is that I restrict attention to CES rules; these designs may not be optimal for other rules. The other is that the criterion considered here is minimization of the upper bound on maximum regret, not maximum regret itself.

TABLE II
QUASI-OPTIMAL STRATIFIED RANDOM SAMPLING DESIGNS

N	N_{0a}, N_{1a}	N_{0b}, N_{1b}	$R_U(z_{Xx})$	N_{0a}, N_{1a}	N_{0b}, N_{1b}	$R_U(z_{X\phi})$
	Rule $z_{Xx}, P(x = a) = .05$			Rule $z_{X\phi}, P(x = a) = .05$		
1			1			1
2	0	1	.457			1
3	0	(1, 2)	.403			1
4	0	2	.338	1	1	.457
8	0	4	.250	1	3	.284
12	1	5	.203	1	5	.232
16	1	7	.173	1	7	.204
20	2	8	.154	1	9	.187
24	2	10	.143	1	11	.172
28	2	12	.133	1	13	.159
32	2	14	.124	1	15	.152
36	2	16	.115	1	17	.148
40	2	18	.108	1	19	.145
44	2	20	.101	1	21	.141
48	2	22	.094	1	23	.137
52	3	23	.088	1	25	.134
	Rule $z_{Xx}, P(x = a) = .25$			Rule $z_{X\phi}, P(x = a) = .25$		
1			1			1
2	0	1	.572			1
3	0	(1, 2)	.529			1
4	1	1	.423	1	1	.585
8	1	3	.293	1	3	.464
12	2	4	.234	2	4	.427
16	3	5	.205	2	6	.400
20	3	7	.182	3	7	.385
24	4	8	.162	3	9	.374
28	5	9	.153	4	10	.365
32	6	10	.144	4	12	.356
36	6	12	.137	5	13	.348
40	7	13	.129	5	15	.342
44	8	14	.122	6	16	.339
48	8	16	.116	6	18	.336
52	8	18	.110	7	19	.333

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 $z_{X\nu}$ or z_ν with the subsamples $N(t, \nu)$ now defined to be the cases with complete data.

The expressions for expected welfare given in (12) and (13) remain valid in these nonclassical settings. The lower bound R_{L_v} on maximum regret also remains valid, as this bound applies to all sampling processes and all rules conditioning treatment choice on v . However, the upper bounds on maximum regret proved in Propositions 1 and 2 need not hold. The crux of the problem is the step in the proof of Proposition 1 stating: "Suppose that $E[y(1)|v = \nu] < E[y(0)|v = \nu]$. Then $E(\bar{y}_{1X\nu} - \bar{y}_{0X\nu}) = -\delta_\nu$." This step rested on the assumption that the data are from a classical randomized experiment; in particular, on the fact that $E(\bar{y}_{tXx}) = E[y(t)|x]$. This equality holds in some nonclassical settings, but not in others.

Dept. of Economics, Northwestern University, 2001 Sheridan Rd., Evanston, IL 60208-2600, U.S.A.; cfmanski@northwestern.edu.

Manuscript received October, 2001; final revision received November, 2003.

REFERENCES

- BERGER, J. (1985): *Statistical Decision Theory and Bayesian Analysis*. New York: Springer-Verlag.
- CHAMBERLAIN, G. (2000): "Econometrics and Decision Theory," *Journal of Econometrics*, 95, 255–283.
- COPAS, J. (1983): "Regression, Prediction, and Shrinkage," *Journal of the Royal Statistical Society B*, 45, 311–354.
- DEHEJIA, R. (2004): "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.
- HECKMAN, J., AND J. SMITH (1998): "Evaluating the Welfare State," in *Econometrics and Economic Theory in the 20th Century*, ed. by S. Strøm. Cambridge: Cambridge University Press, 241–318.
- 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.
- 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: John Wiley & Sons.
- MANSKI, C. (1996): "Learning About Treatment Effects from Experiments with Random Assignment of Treatments," *Journal of Human Resources*, 31, 707–733.
- (2000): "Identification Problems and Decisions Under Ambiguity: Empirical Analysis of Treatment Response and Normative Analysis of Treatment Choice," *Journal of Econometrics*, 95, 415–442.
- RUBIN, D. (1978): "Bayesian Inference for Causal Effects: The Role of Randomization," *The Annals of Statistics*, 6, 34–58.
- SAMANIEGO, F., AND D. RENEAU (1994): "Toward a Reconciliation of the Bayesian and Frequentist Approaches to Point Estimation," *Journal of the American Statistical Association*, 89, 947–957.
- 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 *Social Experimentation*, ed. by J. Hausman and D. Wise. Chicago: University of Chicago Press.
- VAPNIK, V. (1982): *Estimation of Dependencies Based on Empirical Data*. New York: Springer-Verlag.
- WALD, A. (1950): *Statistical Decision Functions*. New York: John Wiley & Sons.