

Fractional Treatment Rules for Social Diversification of Indivisible Private Risks

Charles F. Manski

Department of Economics and Institute for Policy Research, Northwestern University

September 2005

Abstract

Should a social planner treat observationally identical persons identically? This paper shows that uniform treatment is not necessarily desirable when a planner has only partial knowledge of treatment response. Then there may be reason to implement a *fractional treatment rule*, with positive fractions of the observationally identical persons receiving different treatments. The planning problems studied here share some important features: treatment is individualistic, social welfare is a strictly increasing function of a population mean outcome, and outcomes depend on an unknown state of nature. They differ in the information that the planner has about the state of nature and in how he uses this information to make treatment choices. In particular, I compare treatment choice using Bayes rules and the minimax-regret criterion. Following the analysis, I put aside the literal notion of a planner who makes decisions on behalf of society and consider the feasibility of implementing fractional treatment rules in functioning democracies.

This research was supported in part by National Science Foundation grant SES-0314312. I am grateful to John Pepper, Jörg Stoye, and Alex Tetenov for comments. I have also benefitted from the opportunity to present this work in a seminar at the University of Wisconsin-Madison.

1. Introduction

Should a social planner treat observationally identical persons identically? The answer is positive when treatment is individualistic and a utilitarian planner knows the population distribution of treatment response (e.g., Manski, 2004, Section 2.1). However, uniform treatment is not necessarily desirable when a planner has only partial knowledge of treatment response. Then there may be reason to implement a *fractional treatment rule*, with positive fractions of the observationally identical persons receiving different treatments. Here is a simple, dramatic illustration.

Choosing Treatments for X-Pox: Suppose that a new viral disease called x-pox is sweeping the world. Medical researchers have proposed two mutually exclusive treatments, say a and b, which reflect alternative hypotheses, say H_a and H_b , about the nature of the virus. If H_a (H_b) is correct, all persons who receive treatment a (b) survive and all others die. It is not known which hypothesis is correct but there is consensus that H_a is correct with probability 0.3 and H_b with probability 0.7.

There are two singleton rules in this setting, one giving treatment a to the entire population and the other giving b. The latter rule, which implies that all of humanity survives with probability 0.7 and perishes with probability 0.3, is preferable to the former one, which reverses these probabilities. However, neither singleton rule is necessarily better than a fractional rule in which a fraction $\zeta \in (0, 1)$ of the population receives treatment b and the remaining $1 - \zeta$ receive treatment a. Under this rule, the fraction who survive is ζ with probability 0.7 and $1 - \zeta$ with probability 0.3. The fraction $\min(\zeta, 1 - \zeta)$ survives with certainty.

The optimal treatment rule depends on the social welfare function. Let $s(\zeta)$ denote the survival rate with allocation ζ . Let social welfare be $f[s(\zeta)]$, where $f(\cdot)$ is strictly increasing. Let the objective be to maximize expected social welfare $E\{f[s(\zeta)]\}$ over $\zeta \in [0, 1]$. If $f(\cdot)$ is linear, $E\{f[s(\zeta)]\} = 0.7\zeta + 0.3(1 - \zeta)$ and the optimal policy sets $\zeta = 1$. If $f(\cdot)$ is the log function, $E\{f[s(\zeta)]\} = 0.7\log(\zeta) + 0.3\log(1 - \zeta)$ and the optimal policy sets $\zeta = 0.7$. \square

The problem of treating x-pox illustrates how society can use a fractional treatment rule to diversify a risk that is privately indivisible. An individual cannot diversify; one receives treatment a or b and lives or dies. Yet society can diversify by having positive fractions of the population receive each treatment.

The x-pox problem also shows that society may not be able to rely on decentralized decision making to achieve the social optimum. If each person makes his own treatment choice, all will choose treatment b. These decentralized choices are socially optimal if welfare is linear in the survival rate but they are not optimal if welfare is the log of the survival rate or another concave function thereof.

The x-pox illustration is not an oddity. This paper shows that there are many planning problems in which partial knowledge of treatment response can make a fractional treatment rule desirable. The problems studied here share some important features: treatment is individualistic, social welfare is a strictly increasing function $f(\cdot)$ of a population mean outcome, and outcomes depend on an unknown state of nature. They differ in the information that the planner has about the state of nature and in how he uses this information to make treatment choices.¹

Section 2 formally sets out the common features of the problems to be studied. I introduce a special case that I repeatedly consider throughout the paper, this being choice between a status quo treatment and an innovation. I also discuss the planner's choice of a social welfare function.

Sections 3 through 5 examine three informational variations on the basic problem. Section 3 supposes that the planner is Bayesian. As in the x-pox illustration, the planner places a subjective probability distribution on the unknown state of nature and maximizes expected social welfare. The optimal treatment rule is generically a singleton if the function $f(\cdot)$ is linear but may be fractional if $f(\cdot)$ is nonlinear. Considering choice between a status quo treatment and an innovation, Proposition 1 gives a sufficient condition for the optimal rule to be fractional.

¹ The reasons for fractional treatment rules studied in this paper all stem from partial knowledge of treatment response. There may be other reasons for fractional rules as well. For example, aggregate resource constraints may make it infeasible to assign the same treatment to the entire population.

Section 4 supposes that the planner has no subjective or empirical information about the unknown state of nature and that he chooses a treatment rule minimizing maximum regret. In earlier work (Manski, 2005a, Chapter 2; 2005b), I have shown that the minimax-regret rule is fractional when there are two undominated treatments, the function $f(\cdot)$ is linear, and partial knowledge of treatment response is a consequence of a missing data problem. Here I generalize this finding substantially. Proposition 2 shows that the minimax-regret rule is generically fractional in choice problems with two undominated treatments, whatever the shape of $f(\cdot)$ and whatever the reason why knowledge of treatment response is incomplete.

Section 5 supposes that the planner has sample data on treatment outcomes and, following Wald (1950), evaluates a treatment rule by its frequentist expected performance across repeated samples. Manski (2005a, Chapter 3) and Manski and Tetenov (2005) have analyzed choice between a status quo treatment and an innovation when treatments have binary outcomes and when the data are the outcomes of an experiment assigning a random sample of subjects to the innovation. I summarize the findings here.

The concluding Section 6 puts aside the literal notion of a planner who makes decisions on behalf of society and considers the feasibility of implementing fractional treatment rules in functioning democracies. A possible legal/ethical impediment is the normative principle calling for “equal treatment of equals.” A possible political advantage of fractional rules is that they convexify the collective problem of treatment choice—whereas discrete choice among singleton rules often is politically difficult, continuous choice among fractional rules may be less contentious.

Throughout this paper, I maintain the simplifying assumption that the members of the population are observationally identical. If persons have observable covariates, the planner can differentially treat persons with different covariates. Many planning problems are separable across covariance values if the function $f(\cdot)$ is linear; see, for example, Manski (2005a, Chapter 2). In such cases, the analysis of this paper can be applied independently to each subpopulation of observationally identical persons. However, planning problems typically are not separable if $f(\cdot)$ is nonlinear, and analysis becomes more complex.

I also maintain the simplifying assumption of a one-shot planning problem, which gives the planner

no opportunity to learn through experience. Situations that approximate one-shot problems, such as the x-pox illustration, do occur in practice. However, other planning problems play out over time, with the outcomes of early treatment decisions informing treatment choice later on. The possibility of learning introduces a further argument for fractional treatment rules. That is, observation of the outcomes experienced by persons who receive different treatments today may enable society to learn about treatment response and, hence, to improve treatment decisions tomorrow. The randomized clinical trials of medical research exemplify this reason for fractional rules.

2. Preliminaries

2.1. Basic Concepts

The basic concepts are as in my previous research on social choice with partial knowledge of treatment response (Manski , 2000, 2004, 2005a, 2005b), with one essential generalization. I have previously assumed that the planner wants to maximize a population mean outcome. I now suppose that the planner wants to maximize a strictly increasing function of this mean outcome. This generalization is immaterial if the planner knows the population distribution of treatment response, but it is consequential in settings with partial knowledge.

The planner's problem is to choose treatments from a finite set T of mutually exclusive and exhaustive treatments. Each member j of the population, denoted J , has a response function $y_j(\cdot): T \rightarrow Y$ mapping treatments $t \in T$ into responses $y_j(t) \in Y$. The socially relevant outcome from assigning treatment t to person j is $u_j(t) \equiv u[y_j(t), t]$, where $u(\cdot, \cdot): Y \times T \rightarrow R$. At the time of treatment choice, the planner knows the form of $u[y_j(t), t]$ but does not know the value of its first argument $y_j(t)$. For example, $u[y_j(t), t]$ may have the "benefit-cost" form $u[y_j(t), t] = y_j(t) - c_j(t)$, where $c(t)$ is the known real-valued cost of treatment t and

$y_j(t)$ is the unknown real-valued benefit of this treatment.

The members of J are observationally identical. A feasible treatment rule assigns all persons to one treatment or fractionally allocates persons across the treatments. Formally, let Z denote the unit simplex on \mathbb{R}^T ; that is, the space of functions that map T into the unit interval and that satisfy the adding-up conditions: $z \in Z \Rightarrow \sum_{t \in T} z(t) = 1$. The feasible treatment rules are the elements of Z . Rule z is singleton if, for some $t \in T$, it has the form $[z(t) = 1, z(t') = 0, \forall t' \neq t]$. Non-singleton rules are fractional.²

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 sense that J is uncountable and $P(j) = 0, j \in J$. For each treatment rule z , let $U(z, P) \equiv \sum_{t \in T} z(t) \cdot E[u(t)]$ denote the mean value of u realized under z . I assume that the planner wants to solve the *idealized problem*

$$(1) \quad \max_{z \in Z} f[U(z, P)],$$

where the function $f(\cdot)$ is strictly increasing in its argument. The solution is to assign the entire population to a treatment solving the problem

$$(2) \quad \max_{t \in T} E[u(t)].$$

The social welfare achieved by an optimal rule is

² Fractional treatment rules are randomized from the perspective of the population. The planner views the members of the population as observationally identical and, hence, must use some effectively random mechanism to assign specific persons to treatments. Thus, under rule z , each person j has probability $z(t)$ of receiving treatment t .

Fractional rules as defined here are not randomized in the decision theoretic sense of mixed strategies. Rule z assigns fixed fractions of the population to the treatments in T . A mixed strategy would have these fractions be random, with their realizations determined by an auxiliary randomizing device. Formally, a mixed strategy is not an element of Z but rather an element of Δ_Z , where Δ_Z is the space of all probability distributions on Z .

$$(3) \quad f^*(\mathbf{P}) \equiv f\{\max_{t \in T} E[u(t)]\}.$$

Observe that the solution to the idealized planning problem does not depend on the specific form of the function $f(\cdot)$, although the resulting social welfare does depend on $f(\cdot)$.

Researchers in public economics have studied many planning problems akin to (1), where members of the population differ in unobserved respects and the planner knows the population distribution of treatment response. In particular, the literature on optimal income taxation spawned by Mirrlees (1971) permits persons to differ in unobserved ability but assumes that the planner knows both the population distribution of ability and how ability affects labor supply. Similarly, studies of optimal income-maintenance programs (e.g., Besley and Coate, 1995), optimal disability benefits (e.g., Diamond and Sheshinski), and optimal criminal justice systems (e.g., Polinsky and Shavell, 2000) suppose that the planner does not observe certain individual attributes but does know both the population distribution of the unobserved attributes and how the attributes affect behavior.

In contrast, I study treatment choice when the planner has only partial knowledge of the quantities $E[u(t)]$, $t \in T$ and, hence, cannot solve the idealized problem. To formalize the idea of partial knowledge, let Γ denote the set of states of nature that the planner deems feasible. Then $(P_\gamma, \gamma \in \Gamma)$ are the feasible values for P and $\{E_\gamma[u(t)], t \in T\}$, $\gamma \in \Gamma$ are the feasible values for $\{E[u(t)], t \in T\}$, where $E_\gamma[u(t)] \equiv \int u(t) dP_\gamma$. The planner can solve the idealized problem if there exists a dominant treatment; that is, one solving all of the problems $\max_{t \in T} E_\gamma[u(t)]$, $\gamma \in \Gamma$. My concern is treatment choice when no dominant treatment exists.

Sections 3 through 5 differ in the information the planner possesses beyond knowledge of Γ . In Section 3, he places a subjective probability distribution on Γ . In Section 4, he has no information beyond knowledge of Γ . In Section 5, he observes the treatment outcomes realized by a sample of the population.

2.2. Choice Between a Status Quo Treatment and an Innovation

The basic concepts set forth above make plain the nature of the planner's problem but are too abstract to permit an entirely transparent analysis. Hence, I repeatedly find it useful to consider a simple special case. This is the problem of choice between a status quo treatment and an innovation. It is sometimes realistic to suppose that historical experience reveals the population mean outcome of the status quo treatment but not that of the innovation. Thus, the states of nature index the possible mean outcomes for the innovation.

Formally, there are two treatments, $t = a$ denoting the status quo and $t = b$ the innovation. The mean outcomes if the entire population were to receive one treatment are $\alpha \equiv E[u(a)]$ and $\beta = E[u(b)]$ respectively. The planner knows α and knows that β lies in some set $(\beta_\gamma, \gamma \in \Gamma)$ of possible mean outcomes.

Consider a rule that assigns a fraction ζ of the population to treatment b and the remaining $1 - \zeta$ to treatment a . The mean outcome under this rule is

$$(4) \quad \alpha(1 - \zeta) + \beta\zeta = \alpha + (\beta - \alpha)\zeta.$$

Social welfare is $f[\alpha + (\beta - \alpha)\zeta]$.

The optimal treatment rule is obvious if (α, β) are known. The planner should choose $\zeta = 1$ if $\beta > \alpha$ and $\zeta = 0$ if $\beta < \alpha$; all values of ζ yield the same welfare if $\beta = \alpha$. The problem of interest is treatment choice when α is known but it is only known that $\beta \in (\beta_\gamma, \gamma \in \Gamma)$.

Problems of choice between a status quo treatment and an innovation occur often in practice. In the medical arena, the status quo may be a "standard" treatment for a health condition and the innovation may be a new treatment proposed by medical researchers. Historical experience administering the standard treatment to populations of patients may have made its properties well understood. In contrast, the safety and efficacy of the proposed treatment may be uncertain, the available information deriving from biomedical theorizing, animal studies, and small clinical trials with human subjects. In such circumstances, public

agencies such as the Food and Drug Administration (FDA) in the United States decide whether to approve new treatments for administration to the population at large. That done, medical practitioners collectively decide whether the status quo treatment and the innovation should co-exist, each finding some application, or whether the innovation should replace the status quo, taking its place as the standard treatment.

2.3. Specifying the Social Welfare Function

Consideration of planning problems where the objective is to maximize a population mean outcome have been commonplace in public economics, from Bentham's utilitarianism onward. Yet there appears to be no precedent for taking the social welfare function to be a nonlinear increasing function of a population mean outcome, as I do here through the function $f(\cdot)$. The reason is that previous research on social planning has not had to specify the shape of $f(\cdot)$. As noted in Section 2.1, researchers have assumed full knowledge of the distribution of treatment response. This makes the shape of $f(\cdot)$ immaterial. Hence, economists have found it convenient to suppose that $f(\cdot)$ is linear.

I cannot give a general prescription for specification of $f(\cdot)$. Choice of a social welfare function must depend on the context. If I were the planner in the x-pox illustration, I likely would specify $f(\cdot)$ to be a concave function that tends to $-\infty$ as the survival rate tends to zero. Perhaps the log function would be suitable, but I might want to use a function with more curvature, such as $f(s) = -s^{-2}$ or $f(s) = -\exp(1/s)$. I am not sure which particular function I would use, but the analysis of this paper makes clear that the choice is consequential. Whether a planner is Bayesian or applies the minimax-regret criterion, the shape of $f(\cdot)$ affects treatment choice.

It is tempting to view the shape of $f(\cdot)$ as expressing social risk preferences, with linear $f(\cdot)$ expressing risk neutrality and concave $f(\cdot)$ expressing risk aversion. This perspective has a clear interpretation in Bayesian planning, where linear $f(\cdot)$ implies indifference between mean-preserving spreads of a gamble and concave $f(\cdot)$ implies a preference for gambles with smaller spreads. Thinking of $f(\cdot)$ as social risk preferences

is less well-grounded in minimax-regret planning; which does not use the probabilistic apparatus of expected utility theory. Hence, I do not associate the shape of $f(\cdot)$ with risk preferences in this paper.

3. Bayesian Planning

A Bayesian planner places a subjective probability distribution, say Π , on the states of nature and solves the problem

$$(5) \quad \max_{z \in Z} \int f[U(z, P_\gamma)] d\Pi.$$

In this way, a Bayes rule depends on the planner's beliefs and the shape of $f(\cdot)$.

If $f(\cdot)$ is linear, problem (5) reduces to

$$(6) \quad \max_{z \in Z} \sum_{t \in T} z(t) \cdot E_\Pi[u(t)],$$

where $E_\Pi[u(t)] \equiv \int E_\gamma[u(t)] d\Pi$ is the subjective expected value of $E[u(t)]$. Problem (6) is solved by assigning the entire population to a treatment solving the problem

$$(7) \quad \max_{t \in T} E_\Pi[u(t)].$$

Thus, a Bayesian planner with linear $f(\cdot)$ generically chooses a singleton rule. Fractional rules are optimal only when problem (7) has multiple solutions, in which case any mixture of these solutions is optimal.

If $f(\cdot)$ is nonlinear, the specific forms of Π and f determine whether the solution to (5) is a singleton or fractional rule. It is difficult to reach conclusions in abstraction, so I focus on the choice between a status

quo treatment and an innovation. In this case, problem (5) becomes

$$(8) \quad \max_{\zeta \in [0, 1]} \int f[\alpha + (\beta_\gamma - \alpha)\zeta] d\Pi.$$

Proposition 1 shows that, given some regularity and substantive conditions, one can conclude whether the solution to (8) is a singleton or fractional rule.

Proposition 1: Consider problem (8). Let Π be non-degenerate. Let $\beta_\Pi \equiv \int \beta_\gamma d\Pi$ denote the subjective mean of β .

- (a) Let $f(\cdot)$ be strictly concave. Then the Bayes rule is unique. The solution is $\zeta = 0$ if $\beta_\Pi \leq \alpha$.
- (b) Let $f(\cdot)$ be continuously differentiable. Let $f(\cdot)$ and Π be sufficiently regular that

$$\partial \{ \int f[\alpha + (\beta_\gamma - \alpha)\zeta] d\Pi \} / \partial \zeta = \int \{ \partial f[\alpha + (\beta_\gamma - \alpha)\zeta] / \partial \zeta \} d\Pi$$

in a neighborhood of $\zeta = 0$. Then all solutions satisfy $\zeta > 0$ if $\beta_\Pi > \alpha$. All solutions satisfy $\zeta \in (0, 1)$ if $\beta_\Pi > \alpha$ and $\int f(\beta_\gamma) d\Pi < f(\alpha)$. □

Proof:

(a) Strict concavity of $f(\cdot)$ implies that the integral $\int f[\alpha + (\beta_\gamma - \alpha)\zeta] d\Pi$ is strictly concave in ζ . Hence, problem (8) has a unique solution. If $\zeta = 0$, then $\int f[\alpha + (\beta_\gamma - \alpha)\zeta] d\Pi = f(\alpha)$. For each $\zeta > 0$, $f[\alpha + (\beta_\gamma - \alpha)\zeta]$ is strictly concave as a function of β_γ . Hence, $\int f[\alpha + (\beta_\gamma - \alpha)\zeta] d\Pi < f[\alpha + (\beta_\Pi - \alpha)\zeta]$. Hence, $\beta_\Pi \leq \alpha \Rightarrow \int f[\alpha + (\beta_\gamma - \alpha)\zeta] d\Pi < f(\alpha)$.

(b) $\{ \partial f[\alpha + (\beta_\gamma - \alpha)\zeta] / \partial \zeta \}_{\zeta=0} = (\beta_\gamma - \alpha) \cdot [df(x)/dx]_{x=\alpha}$. It follows that $\{ \partial \{ \int f[\alpha + (\beta_\gamma - \alpha)\zeta] d\Pi \} / \partial \zeta \}_{\zeta=0} = (\beta_\Pi - \alpha) \cdot [df(x)/dx]_{x=\alpha}$. By assumption, $\beta_\Pi > \alpha$ and $[df(x)/dx]_{x=\alpha} > 0$. Hence, $\int f[\alpha + (\beta_\gamma - \alpha)\zeta] d\Pi$ strictly

increases with ζ in a neighborhood of $\zeta = 0$, implying that solutions to (8) are positive. If $\int f(\beta_\gamma) d\Pi < f(\alpha)$, then $\zeta = 1$ does not solve problem (8). Hence, solutions are fractional.

Q. E. D.

Observe that the concavity and differentiability restrictions placed on $f(\cdot)$ are used in different parts of the proposition. The proof of part (a) only uses the assumption that $f(\cdot)$ is strictly concave. The proof of part (b) only uses the assumption that $f(\cdot)$ is continuously differentiable and the stated regularity condition.

Also observe that, as in the x-pox problem, society may not be able to rely on decentralized decision making to achieve the social optimum. Suppose that persons makes their own treatment choices and that all are Bayesian, with Π the consensus subjective distribution on Γ . Then all choose treatment a when $\beta_\Pi < \alpha$ and all choose b when $\beta_\Pi > \alpha$. Decentralized decision making may yield a fractional treatment allocation only when $\beta_\Pi = \alpha$, in which case persons are indifferent between the two treatments.

A Simple Special Case

A simple special case is amenable to complete analysis and sheds further light on the Bayes rule. Let $f(\cdot)$ be the log function, let Γ contain the two elements $\{0, 1\}$, with $\beta_0 < \alpha < \beta_1$, and let $\pi \equiv \Pi(\gamma = 0)$. Then

$$(9) \quad \int f[\alpha + (\beta_\gamma - \alpha)\zeta] d\Pi = \pi \cdot \log[\alpha + (\beta_0 - \alpha)\zeta] + (1 - \pi) \cdot \log[\alpha + (\beta_1 - \alpha)\zeta].$$

The derivative of the left hand side with respect to ζ is

$$(10) \quad \partial \{ \int f[\alpha + (\beta_\gamma - \alpha)\zeta] d\Pi \} / \partial \zeta = \pi(\beta_0 - \alpha) / [\alpha + (\beta_0 - \alpha)\zeta] + (1 - \pi)(\beta_1 - \alpha) / [\alpha + (\beta_1 - \alpha)\zeta].$$

Inspection of the sign of this derivative shows when the Bayes rule is singleton or fractional.

Multiplication of the right hand side of (10) by the two positive terms in the denominators yields

$$\pi(\beta_0 - \alpha)[\alpha + (\beta_1 - \alpha)\zeta] + (1 - \pi)(\beta_1 - \alpha)[\alpha + (\beta_0 - \alpha)\zeta] = (\beta_0 - \alpha)(\beta_1 - \alpha)\zeta + \alpha(\beta_{\Pi} - \alpha),$$

where $\beta_{\Pi} = \pi\beta_0 + (1 - \pi)\beta_1$. Observe that $(\beta_0 - \alpha)(\beta_1 - \alpha) < 0$. If $\beta_{\Pi} \leq \alpha$, the sign of the derivative is negative at all $\zeta \in (0, 1]$; hence, the Bayes rule sets $\zeta = 0$ as shown in Proposition 1a. If $\beta_{\Pi} > \alpha$, the sign of the derivative is positive in a neighborhood of $\zeta = 0$; hence, the Bayes rule sets $\zeta > 0$ as shown in Proposition 1b. The derivative decreases with ζ and equals zero when

$$(11) \quad \zeta = \alpha(\beta_{\Pi} - \alpha) / |(\beta_0 - \alpha)(\beta_1 - \alpha)|.$$

If the right hand side of (11) is less than one, then the Bayes rule is fractional with this value of ζ . Otherwise, the Bayes rule sets $\zeta = 1$.

Bayesian Planning with Sample Data

Suppose that the planner has sample data on treatment outcomes. If he is Bayesian, the above analysis applies as stated, with Π being the planner's subjective distribution given the sample data.

Statistical decision theorists have studied some interesting Bayesian planning problems with sample data. In particular, work on the design and analysis of randomized clinical trials has examined two-period planning problems in which some members of a population are randomly assigned to treatments in the first period and the experimental findings are used to inform treatment choice in the second period. See, for example, Canner (1970), Cheng, Su, and Berry (2003), and Dehejia (2005).

These contributions and (as far as I am aware) other studies of Bayesian planning have assumed without comment that $f(\cdot)$ is linear. Hence, research on Bayesian planning has taken for granted that Bayes rules are singleton. Proposition 1 shows that this is not necessarily the case when $f(\cdot)$ is nonlinear.

4. Minimax-Regret Planning

The minimax-regret criterion for treatment choice uses no information beyond the planner's knowledge that the actual state of nature lies in Γ . The criterion is

$$(12) \quad \inf_{z \in Z} \sup_{\gamma \in \Gamma} f^*(P_\gamma) - f[U(z, P_\gamma)].$$

Here, $f^*(P_\gamma)$ is the optimal social welfare achievable if it were known that $P = P_\gamma$; that is, by (3),

$$(13) \quad f^*(P_\gamma) \equiv f\{\max_{t \in T} E_\gamma[u(t)]\}.$$

The quantity $f^*(P_\gamma) - f[U(z, P_\gamma)]$ is the *regret* of rule z in state of nature γ ; that is, the loss in social welfare from not knowing the actual state of nature.³

Manski (2005a, Chapter 2; 2005b) derived the minimax-regret rule when there are two treatments and a planner with linear $f(\cdot)$ observes the actual treatment outcomes of a study population in which the process of treatment selection is unknown. This empirical evidence yields only partial knowledge of treatment response because outcomes under counterfactual treatments are not observable. I found that the minimax-regret criterion yields a singleton rule only when one treatment weakly dominates the other, yielding weakly higher social welfare in all states of nature and strictly higher welfare in some state of nature. The minimax-regret rule is fractional when both treatments are weakly undominated, each being better in some state of nature.

I show here that the minimax-regret rule is generically fractional in problems with two undominated

³ The maximin rule is another criterion that uses no information beyond knowledge of Γ . It calls for solution of the problem $\sup_{z \in Z} \inf_{\gamma \in \Gamma} f[U(z, P_\gamma)]$. This rule is deservedly viewed as ultra-conservative. Consider choice between a status quo treatment and innovation. If there is any feasible state of nature in which the innovation is inferior to the status quo, the maximin rule assigns the entire population to the latter.

treatments. The function $f(\cdot)$ need not be linear and missing outcome data need not be the reason why the planner has partial knowledge of treatment response. Let $T = \{a, b\}$. Then the minimax-regret problem is

$$(14) \quad \inf_{\zeta \in [0, 1]} \sup_{\gamma \in \Gamma} \max \{f\{E_{\gamma}[u(a)]\}, f\{E_{\gamma}[u(b)]\}\} - f\{(1 - \zeta)E_{\gamma}[u(a)] + \zeta E_{\gamma}[u(b)]\}.$$

Proposition 2 shows that solutions to (14) are fractional if a modest regularity condition holds.

Proposition 2: Let $\Gamma(a) \equiv \{\gamma \in \Gamma: E_{\gamma}[u(a)] \geq E_{\gamma}[u(b)]\}$ and $\Gamma(b) \equiv \{\gamma \in \Gamma: E_{\gamma}[u(b)] \geq E_{\gamma}[u(a)]\}$. Let

$$R(\zeta; a) \equiv \sup_{\gamma \in \Gamma(a)} f\{E_{\gamma}[u(a)]\} - f\{(1 - \zeta)E_{\gamma}[u(a)] + \zeta E_{\gamma}[u(b)]\}$$

$$R(\zeta; b) \equiv \sup_{\gamma \in \Gamma(b)} f\{E_{\gamma}[u(b)]\} - f\{(1 - \zeta)E_{\gamma}[u(a)] + \zeta E_{\gamma}[u(b)]\}$$

be the maximum regret of rule ζ on $\Gamma(a)$ and $\Gamma(b)$ respectively. Suppose that $R(\cdot; a)$ and $R(\cdot; b)$ are continuous on $[0, 1]$. Also suppose that both treatments are weakly undominated; that is, there exist states of nature γ and γ' such that $E_{\gamma}[u(a)] > E_{\gamma}[u(b)]$ and $E_{\gamma'}[u(a)] < E_{\gamma'}[u(b)]$. Then problem (14) has a unique solution, which lies in the open interval $(0, 1)$. \square

Proof: The maximum regret of rule ζ on all of Γ is $\max [R(\zeta; a), R(\zeta; b)]$. As ζ increases from 0 to 1, $R(\cdot; a)$ continuously increases from 0 to $R(1; a)$ and $R(\cdot; b)$ continuously decreases from $R(0; b)$ to 0. The fact that both treatments are weakly undominated implies that $\{R(1; a) > 0, R(0; b) > 0\}$ and, moreover, that $R(\cdot; a)$ and $R(\cdot; b)$ are strictly monotone functions of ζ . Hence, the minimax-regret rule is the unique $\zeta \in (0, 1)$ such that $R(\zeta; a) = R(\zeta; b)$.

Q. E. D.

The basic idea underlying Proposition 2 is that singleton and fractional treatment rules have different extremum properties across states of nature. Each singleton rule is the best rule in some states of nature and the worst in the others. In contrast, fractional rules yield intermediate social welfare in all states of nature. These facts suggest, but do not quite prove, that the minimax-regret rule must be fractional rather than singleton. To complete the proof requires some regularity in the behavior of $R(\cdot; a)$ and $R(\cdot; b)$; the assumption that these functions are continuous suffices.

A Simple Special Case

Consideration of choice between a status quo treatment and an innovation sheds further light on the minimax-regret rule. Let Γ contain the two elements $\{0, 1\}$, with $\beta_0 < \alpha < \beta_1$. Then the minimax-regret problem is

$$(15) \quad \inf_{\zeta \in [0, 1]} \max \{f(\alpha) - f\{(1 - \zeta)\alpha + \zeta\beta_0\}, f(\beta_1) - f\{(1 - \zeta)\alpha + \zeta\beta_1\}\}.$$

Let $f(\cdot)$ be continuous. As ζ rises from 0 to 1, the first term inside the brackets increases strictly and continuously from 0 to $f(\alpha) - f(\beta_0)$ and the second term similarly decreases from $f(\beta_1) - f(\alpha)$ to 0. Hence, the minimax-regret rule is the unique $\zeta \in (0, 1)$ that equalizes the two terms, solving the equation

$$(16) \quad f(\alpha) - f\{(1 - \zeta)\alpha + \zeta\beta_0\} = f(\beta_1) - f\{(1 - \zeta)\alpha + \zeta\beta_1\}.$$

The solution depends on the shape of $f(\cdot)$. If $f(\cdot)$ is linear, the minimax-regret rule sets

$$(17) \quad \zeta = (\beta_1 - \alpha)/(\beta_1 - \beta_0).$$

If $f(\cdot)$ is the log function, the minimax-regret rule sets

$$(18) \quad \zeta = \alpha(\beta_1 - \alpha) / [\alpha(\beta_1 - \alpha) + \beta_1(\alpha - \beta_0)].$$

Comparison of (17) and (18) shows that a planner with linear $f(\cdot)$ allocates a larger fraction of the population to treatment b than does one using the log transformation. This illustrates that, although Proposition 2 shows the minimax-regret rule is generically fractional, the specific fractional allocation varies with $f(\cdot)$.

Multiple Treatments

Does a result as broad and simple as Proposition 2 hold when there are more than two treatments? The answer appears to be negative. The proof of the proposition makes critical use of the fact that the functions $R(\cdot; a)$ and $R(\cdot; b)$ cross exactly once as ζ increases from 0 to 1. This single-crossing property does not hold when there are three or more treatments.

Although minimax-regret rules often are fractional, there exist non-trivial planning problems with multiple treatments which have singleton minimax-regret rules. Here is an example. Let ζ^* denote the unique minimax-regret rule determined in Proposition 2. Now introduce a third treatment c, so the treatment set is $\{a, b, c\}$ rather than $\{a, b\}$. Then the minimax-regret problem changes from (14) to

$$(19) \quad \inf_{z \in Z} \sup_{\gamma \in \Gamma} \max \{f\{E_\gamma[u(a)]\}, f\{E_\gamma[u(b)]\}, f\{E_\gamma[u(c)]\}\} \\ - f\{z(a)E_\gamma[u(a)] + z(b)E_\gamma[u(b)] + z(c)E_\gamma[u(c)]\}.$$

Suppose that, in every state of nature, treatment c has the same mean outcome as rule ζ^* ; thus,

$$(20) \quad E_\gamma[u(c)] = (1 - \zeta^*)E_\gamma[u(a)] + \zeta^*E_\gamma[u(b)], \quad \text{all } \gamma \in \Gamma.$$

Then problem (19) reduces to

$$(21) \quad \inf_{z \in Z} \sup_{\gamma \in \Gamma} \max \{ f\{E_{\gamma}[u(a)]\}, f\{E_{\gamma}[u(b)]\} \} \\ - f\{[z(a) + (1 - \zeta^*)z(c)]E_{\gamma}[u(a)] + [z(b) + \zeta^*z(c)]E_{\gamma}[u(b)]\}.$$

Problems (14) and (21) have the same structure. Hence, there are multiple minimax-regret rules, each satisfying the condition $\zeta^* = z(b) + \zeta^*z(c)$. Many fractional rules satisfy this condition and the singleton rule $[z(a) = 0, z(b) = 0, z(c) = 1]$ does as well.

This example obviously has a special structure. In every state of nature, the welfare of treatment c is intermediate between that of treatments a and b . Hence, c is not weakly dominated and, moreover, there exists no state of nature in which it is the unique best treatment. Nevertheless, the example does demonstrate that the behavior of minimax-regret rules can be more complex when there are multiple treatments than when there are two treatments.

5. Treatment Choice with Sample Data

Suppose that the planner has sample data on treatment outcomes. As observed earlier, the analysis of Section 3 continues to apply as stated if the planner is Bayesian, with Π being the planner's subjective distribution given the sample data. Hence, it is only necessary here to consider how a non-Bayesian planner might use sample data.

Wald (1950) introduced a frequentist version of statistical decision theory. Wald's idea is easy to explain, but working out the implications for treatment choice can be rather difficult. For this reason, I focus on a simple special case analyzed in Manski (2005b, Chapter 3) and Manski and Tetenov (2005). This

concerns choice between a status quo treatment and an innovation when treatments have binary outcomes and the data are the outcomes of an experiment assigning a random sample of subjects to the innovation. Section 5.1 gives the general idea and Section 5.2 summarizes the findings for the special case. Section 5.3 briefly discusses research on other treatment-choice problems with sample data.

5.1. Generalities

Wald proposed that treatment rules be evaluated by their expected performance as the sampling process is engaged repeatedly to draw independent data samples. Let Q denote the sampling process and let Ψ denote the sample space; that is, Ψ is the set of data samples that may be drawn under Q . Let Z now denote the space of functions that map $T \times \Psi$ into the unit interval and that satisfy the adding-up conditions: $z \in Z \Rightarrow \sum_{t \in T} z(t, \psi) = 1, \forall \psi \in \Psi$. Then each $z \in Z$ defines a *statistical treatment rule*.

Repeated engagement of the sampling process to draw independent samples makes social welfare a random variable. In state of nature γ , the *expected welfare* yielded by z in repeated samples is ⁴

$$(22) \quad W(z, P_\gamma, Q_\gamma) \equiv \int f\left\{ \sum_{t \in T} z(t, \psi) \cdot E_\gamma[u(t)] \right\} dQ_\gamma(\psi).$$

Rule z is said to be *admissible* if there exists no other rule that yields at least the expected welfare of z in all states of nature and greater expected welfare in some state of nature; rule z is *inadmissible* otherwise. There is considerable consensus among statistical decision theorists that inadmissible rules should not be chosen. There is no similar consensus on how to choose among the admissible rules, but one possibility is the finite-sample minimax-regret criterion

⁴The expressions *statistical treatment rule* and *expected welfare* are from Manski (2004). In Wald's terminology, the former is a *statistical decision function* and the latter is the negative of *risk*.

$$(23) \quad \inf_{z \in Z} \sup_{\gamma \in \Gamma} f^*(P_\gamma) - W(z, P_\gamma, Q_\gamma).$$

This replaces the social welfare function $f[U(z, P_\gamma)]$ of the earlier minimax-regret criterion (12) with its expected value $W(z, P_\gamma, Q_\gamma)$ across repeated samples.

5.2. Treatment Choice With Binary Outcomes Using Experimental Data on an Innovation

Consider the problem of choice between a status quo treatment and an innovation when treatments have binary outcomes. Let the sample data be the outcomes of an experiment assigning a random sample of N subjects to the innovation. Of these subjects, a number n experience outcome $y(b) = 1$ and the remaining $N - n$ experience $y(b) = 0$. Let the outcomes of all subjects be observed. Treatment response distributions are point-identified in this setting; partial knowledge of treatment response occurs only because the sample size N is finite.

Let $\alpha = P[y(a) = 1]$ and $\beta = P[y(b) = 1]$ be the population success rates that would occur if all persons were to receive treatment a or b respectively. 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 the functions $\zeta(\cdot): [0, \dots, N] \rightarrow [0, 1]$ that map the number of experimental successes into a treatment allocation. Thus, for each value of n , rule ζ allocates a fraction $\zeta(n)$ of the population to treatment b and the remaining $1 - \zeta(n)$ to treatment a .

Let $p(n; \beta) \equiv N!/[n!(N - n)!] \beta^n (1 - \beta)^{N-n}$ denote the Binomial probability of n successes in N trials.

In state of nature γ , the expected welfare yielded by rule $\zeta(\cdot)$ across repeated samples is

$$(24) \quad W(\zeta, \beta_\gamma, N) \equiv \sum_{n=0}^N p(n; \beta_\gamma) \cdot f[\alpha + (\beta_\gamma - \alpha) \cdot \zeta(n)].$$

It turns out that the admissibility of fractional and singleton rules depends on the shape of $f(\cdot)$.

Linear $f(\cdot)$

Manski (2005a, Chapter 3) considers the case in which $f(\cdot)$ is linear. Then the expected welfare of rule ζ in state of nature γ is

$$(25) \quad W(\zeta, \beta_\gamma, N) = \alpha + (\beta_\gamma - \alpha) \cdot E_\gamma[\zeta(n)],$$

where $E_\gamma[\zeta(n)] = \sum_n p(n; \beta_\gamma) \cdot \zeta(n)$. In this setting, a theorem of Karlin and Rubin (1956) shows that the admissible treatment rules are ones which assign all persons to the status quo treatment if the number of experimental successes is below a specified threshold, all to the innovation if the success rate is above the threshold, and permits a fractional allocation only if the number of experimental successes exactly equals the threshold. Karlin and Rubin called these rules *monotone*, but I will refer to them as *KR-monotone*. Formally, a KR-monotone treatment rule has the structure

$$(26) \quad \begin{aligned} \zeta(n) &= 0 & \text{for } n < k, \\ \zeta(n) &= \lambda & \text{for } n = k, \\ \zeta(n) &= 1 & \text{for } n > k, \end{aligned}$$

where k is any integer such that $0 \leq k \leq N$ and λ is any real number such that $0 \leq \lambda \leq 1$.

Manski (2005, Proposition 3.1) applies Karlin and Rubin (1956, Theorem 4) to show that the admissible and KR-monotone rules coincide if $(\beta_\gamma, \gamma \in \Gamma)$ contains values that are smaller and larger than α but excludes the extreme values 0 and 1. This is a striking result. It says that if $f(\cdot)$ is linear, statistical treatment rules are inadmissible if they are fractional at more than one value of n .

A particular consequence is that the finite-sample minimax-regret rule, which is necessarily admissible, is KR-monotone. In contrast, Proposition 2 showed that the minimax-regret rule not using sample data is generically fractional. This difference in findings is intriguing but is open to two interpretations. One

might interpret the difference as showing that, from the minimax-regret perspective, even small randomized experiments qualitatively alter treatment choice relative to the case with no data at all. Or one might interpret it as showing that the finite-sample minimax-regret criterion is not so obvious an extension of the basic minimax-regret idea as it might seem at first inspection.

Concave $f(\cdot)$

The Karlin and Rubin theorems do not apply when $f(\cdot)$ is nonlinear, but a weaker notion of monotonicity remains relevant if $f(\cdot)$ is concave. Define a *fractional monotone* rule to be one in which the fraction of the population assigned to the innovation weakly increases with the experimental success rate. Manski and Tetenov (2005) show that the class of fractional monotone rules is essentially complete if $f(\cdot)$ is concave. That is, given any rule which is not fractional monotone, there exists a fractional monotone rule which performs at least as well in all feasible states of nature.

Investigating particular treatment rules, we find that Bayes rules and the minimax-regret rule depend on the curvature of $f(\cdot)$. These rules are KR-monotone if the curvature is sufficiently weak, but they deliver interior fractional treatment allocations if the curvature is sufficiently strong. If $f(\cdot)$ has sufficient curvature and B contains positive values arbitrarily close to 0, the minimax-regret rule never assigns the entire population to treatment b .

5.3. Other Treatment-Choice Problems

Outside of the setting of Section 5.2, there have been only a few non-Bayesian studies of treatment choice with sample data, and these have all assumed that $f(\cdot)$ is linear.

The Karlin-Rubin result cited above does not require that treatments have binary outcomes, but it does require that outcome distributions have the monotone-likelihood-ratio (MLR) property. Bernoulli distributions have the MLR property, so this requirement is unrestrictive when outcomes are binary.

However, distributions with the MLR property are special when outcomes can take more than two values.

Manski (2004), Stoye (2005), and Hirano and Porter (2005) study minimax-regret treatment choice when there are two treatments with unknown outcome distributions and a randomized experiment provides data on the outcomes of both treatments. Manski (2004) assumes that outcomes have bounded support and derives an upper bound on the finite-sample maximum regret of *empirical-success rules* that use sample average outcomes to estimate population mean outcomes. Stoye (2005) assumes that outcomes are binary and shows that a particular empirical-success rule is finite-sample minimax-regret. He also provides a partial analysis of cases in which outcomes have bounded support. Hirano and Porter (2005) use the Le Cam asymptotic version of Wald's statistical decision theory to study a class of problems in which the researcher may have some a priori knowledge of the treatment outcome distributions. They find that empirical success rules are asymptotic minimax-regret, in the Le Cam sense, when the researcher has no a priori knowledge of the outcome distributions.

Whereas the above research assumes that the available data are from a randomized experiment, Manski (2005a, Chapter 4; 2005b) considers treatment choice using observational data when treatments were selected in an unknown way. In this case, outcome distributions are partially identified. I show that a treatment rule estimating certain point-identified population means by their sample analogs is finite-sample minimax-regret.

6. Implementing Fractional Treatment Rules

The treatment-choice problems examined in this paper, while hardly exhaustive, give ample reason for a planner to consider fractional rules. The standard practice in research on social planning has been to assume that the choice set only contains singleton rules. This is inconsequential in the idealized problem of Section 2.1 and in similar settings where the planner knows the population distribution of treatment response.

However, partial knowledge of treatment response may make fractional rules preferable to singleton ones. As shown in Sections 3 through 5, the specifics depend on the planner's information and on the shape of $f(\cdot)$.

The notion of a planner who acts on behalf of society is a useful fiction that greatly simplifies analysis of collective decision problems. Nevertheless, if the findings of this paper are to have practical value in functioning democracies, we must ultimately consider the feasibility of implementing fractional treatment rules in such societies.

A possible ethical objection to fractional rules is that they violate the normative principle calling for "equal treatment of equals." Fractional rules violate this principle if one interprets "equal treatment" to require that observationally identical persons receive the same element of the choice set T . However, fractional rules are consistent with the principle if it is enough for observationally identical people to have equal probabilities of receiving particular treatments.

Democracies occasionally accept the probabilistic interpretation in the design of major policies, such as random drug testing, calls for jury service, and the American Green Card and Vietnam draft lotteries. They routinely accept it when they permit the conduct of randomized clinical trials. Indeed, the prevailing standard of medical ethics permits randomized trials only when partial knowledge of treatment response prevents a determination that one treatment is superior to another. The open question is whether democracies would accept broader application of the idea that, when knowledge of treatment response is incomplete, "equal treatment of equals" can ethically be probabilistic rather than deterministic.

A political argument in support of fractional rules is that they convexify collective problems of treatment choice. Political processes often find it difficult to make discrete choices among alternative singleton rules, each of which may be favored by a different segment of the population. Fractional rules open opportunities for compromise, transforming a problem of discrete choice into one of continuous choice among alternative fractional allocations.

Political considerations may also favor application of the minimax-regret criterion rather than a Bayes rule. Determination of a Bayes rule requires specification of a subjective distribution on the states of nature,

but different segments of the population may disagree in their beliefs. In contrast, the existence of a consensus minimax-regret rule requires only that the population agree on the feasible states of nature.

It may be that, in some circumstances, enabling the population to trade their assigned treatments would increase public acceptance of fractional rules. During the American Civil War, draftees were permitted to hire substitutes to fulfill their military service. In recent times, various countries have established emissions trading systems for control of pollution, where firms are permitted to trade government-determined pollution endowments. Future research should explore the implications of treatment trading for social welfare.

To conclude, I suggest a specific aspect of American public policy that seems well-suited for implementation of fractional rules. This is the drug approval process of the Food and Drug Administration. The present process essentially makes a binary choice between unconstrained approval and total disapproval of a new drug. With only these two options on the table, the FDA sets a high bar for approval, requiring that the manufacturers of pharmaceuticals demonstrate “substantial evidence of effect” for their products (see Gould, 2002). It may be preferable to implement a fractional approval process setting a knowledge-dependent ceiling on the production and marketing of new drugs—the stronger the evidence of effect, the higher the ceiling.

References

- Besley, T. and S. Coate, "The Design of Income Maintenance Programmes," *Review of Economic Studies*, 62, 187-221.
- Canner, P. (1970), "Selecting One of Two Treatments When the Responses Are Dichotomous," *Journal of the American Statistical Association*, 65, 293-306.
- Cheng, Y., F. Su, and D. Berry (2003), "Choosing Sample Size for a Clinical Trial Using Decision Analysis," *Biometrika*, 90, 923-936.
- Dehejia, R. (2005), "Program Evaluation as a Decision Problem," *Journal of Econometrics*, 125, 141-173.
- Diamond, P. and E. Sheshinski (1995), "Economic Aspects of Optimal Disability Benefits," *Journal of Public Economics*, 57, 1-23.
- Gould, A. (2002), "Substantial Evidence of Effect," *Journal of Biopharmaceutical Statistics*, 12, 53-77.
- Hirano, K. and R. Porter (2005), "Asymptotics for Statistical Treatment Rules," Department of Economics, University of Arizona.
- Karlin, S. and H. Rubin (1956), "The Theory of Decision Procedures for Distributions with Monotone Likelihood Ratio," *Annals of Mathematical Statistics*, 27, 272-299.
- 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. (2004), "Statistical Treatment Rules for Heterogeneous Populations," *Econometrica*, 72, 1221-1246.
- Manski, C. (2005a), *Social Choice with Partial Knowledge of Treatment Response*, Princeton: Princeton University Press, forthcoming.
- Manski, C. (2005b), "Minimax-Regret Treatment Choice with Missing Outcome Data," *Journal of Econometrics*, forthcoming.
- Manski, C. and A. Tetenov (2005), "Admissible Treatment Rules for a Risk-Averse Planner with Experimental Data on an Innovation," Department of Economics, Northwestern University.
- Mirrlees, J. (1971), "An Exploration in the Theory of Optimal Income Taxation," *Review of Economic Studies*, 38, 175-208.
- Polinsky, M. and S. Shavell (2000), "The Economic Theory of Public Enforcement of Law," *Journal of Economic Literature*, 38, 45-76.
- Stoye, J. (2005), "Minimax-Regret Treatment Choice with Finite Sample," Department of Economics, New York University.
- Wald, A. (1950), *Statistical Decision Functions*, New York: Wiley.