

TREATMENT CHOICE WHEN TREATMENT RESPONSE IS PARTIALLY IDENTIFIED

Charles F. Manski
Northwestern University

An important objective of studies of treatment response is to provide decision makers with information useful in choosing treatments. Often the decision maker is a *planner* who must choose treatments for a heterogeneous population. The planner may want to choose treatments whose outcomes maximize the welfare of this population.

Examples:

- (a) a physician choosing medical treatments for a population of patients.
- (b) a judge choosing sentences for a population of convicted offenders.

Studies of treatment response are useful to planners to the extent that they reveal how outcomes vary with treatments and observable covariates.

Identification problems and the need for statistical inference from finite samples limit the information that studies can provide.

How might planners with partial knowledge of treatment response make treatment choices?

A Class of Planning Problems

A planner must choose treatments for a heterogeneous population. Each member of the population has a response function that maps treatments into an outcome of interest. The planner observes some covariates for each member of the population. The observed covariates determine the set of *treatment rules* that are feasible to implement. These are functions that map the observed covariates into a treatment allocation.

The planner wants to choose a treatment rule that maximizes population mean welfare, but does not have the knowledge of treatment response needed to implement the optimal rule.

The planner observes a study population in which treatments have already been selected and outcomes realized. The problem is to use this empirical evidence and credible assumptions to choose treatments.

Practices That Limit the Usefulness of Research on Treatment Response

Hypothesis Testing

Research on treatment response has been strongly influenced by the theory of hypothesis testing, especially by the idea of testing the null hypothesis of *zero average treatment effect*.

(a) This null hypothesis is prominent in experimental design, where researchers use norms for statistical power to choose sample sizes.

(b) Research findings may be unreported or deemed “insignificant” if they do not meet test-based criteria for statistical precision.

(c) Hypothesis testing has been particularly influential in medical research using randomized clinical trials. Testing the hypothesis of zero average treatment effect is institutionalized in the U. S. Food and Drug Administration drug approval process, which calls for comparison of a new treatment with a placebo. Approval of the new treatment normally requires rejection of the null hypothesis of zero average treatment effect in two independent clinical trials.

Hypothesis testing is remote from treatment choice:

(a) The classical practice of handling the null and alternative hypotheses asymmetrically, fixing the probability of a type I error and seeking to minimize the probability of a type II error, makes no sense from the perspective of treatment choice.

(b) Error probabilities only measure the chance of choosing a sub-optimal rule; they do not measure the damage resulting from a sub-optimal choice.

The Study Population and the Treatment Population

Much research downplays the importance of correspondence between the study population J and the treatment population J*.

Campbell argued that studies of treatment effects should be judged primarily by their *internal validity* and only secondarily by their *external validity*. Internal validity means the credibility of findings within study population J, whatever J may be. External validity is the credibility of extrapolating findings from J to another population, say J*.

Rosenbaum (1999) downplays the importance of having the study population be similar to the population of interest, writing:

“Studies of samples that are representative of populations may be quite useful in describing those populations, but may be ill-suited to inferences about treatment effects.”

Many economists evaluating social programs focus on easy-to-study populations that differ sharply from the populations that planners must treat.

(a) A common practice has been to report the “effect of treatment on the treated,” where “the treated” are the members of a study population who actually received a specified treatment.

(b) To cope with noncompliance in randomized experiments, Angrist, Imbens and Rubin (1996) recommend that treatment effects be reported for the sub-population of “compliers,” these being persons who would comply with their designated experimental treatments whatever they might be.

These practices yield findings useful to a planner only if treatment response among “the treated” or “compliers” can credibly be assumed to be the same as in the treatment population.

Reporting Observable Variation in Treatment Response

To inform treatment choice, research on treatment response should aim to learn how treatment response varies with the covariates that planners observe. The key to success is to determine which persons should receive which treatments. For example,

(a) Judges may be able to lower recidivism among criminal offenders by sentencing some to prison and others to probation.

(b) Social workers may be able to increase the earnings of welfare recipients by placing some in job training and others in basic skills classes.

Yet the prevalent research practice has been to report treatment response in the population as a whole or within broad sub-populations, rather than conditional on the covariates that planners observe.

A reason may be pre-occupation with “statistical significance.” Conditioning on covariates reduces the precision of estimates of treatment effects, so findings may be “insignificant” by conventional criteria of hypothesis testing.

Researchers wanting to inform treatment choice should not use statistical insignificance as a reason to refrain from studying observable variation in treatment response. A planner must be concerned with the quantitative variation of outcomes with treatments and covariates. Conventional hypothesis tests do not address this question.

Untenable Assumptions

Powerful incentives influence researchers studying treatment response to maintain assumptions far stronger than they can persuasively defend, in order to draw strong conclusions.

(a) The scientific community rewards those who produce unambiguous findings.

(b) The public rewards those who offer simple analyses leading to unequivocal policy recommendations.

Research findings based on untenable assumptions are not much use to a planner. The objective of a planner should be to maximize *actual* social welfare, not the social welfare that would prevail if untenable assumptions were to hold.

STUDIES OF TREATMENT CHOICE WITH PARTIAL KNOWLEDGE OF TREATMENT RESPONSE

Charles F. Manski

Book

Social Choice with Partial Knowledge of Treatment Response, Princeton University Press, 2005.

Articles

“Identification problems and decisions under ambiguity: Empirical analysis of treatment response and normative analysis of treatment choice,” *Journal of Econometrics*, Vol. 95, 2000, pp. 415-442.

“Treatment Choice Under Ambiguity Induced by Inferential Problems,” *Journal of Statistical Planning and Inference*, Vol. 105, 2002, pp. 67-82.

“Social Learning from Private Experiences: The Dynamics of the Selection Problem,” *Review of Economic Studies*, Vol. 71, 2004, pp. 443-458.

“Social Learning and the Adoption of Innovations,” in L. Blume and S. Durlauf (editors), *The Economy as an Evolving Complex System III*, Oxford University Press, 2005.

“Statistical Treatment Rules for Heterogeneous Populations,” *Econometrica*, Vol. 72, pp. 1221-1246.

“Admissible Treatment Rules for a Risk-Averse Planner with Experimental Data on an Innovation,” with A. Tetenov, September 2005.

“Minimax-Regret Treatment Choice with Missing Outcome Data,” *Journal of Econometrics*, forthcoming.

“Fractional Treatment Rules for Social Diversification of Indivisible Private Risks,” September 2005.

“Search Profiling with Partial Knowledge of Deterrence,” November 2005.

SEARCH PROFILING WITH PARTIAL KNOWLEDGE OF DETERRENCE

C. Manski, November 2005

Normative research in public economics has generally assumed that the relevant social planner knows how policy affects population behavior. Economists studying optimal income taxation assume that the planner knows how the tax schedule affects labor supply (Mirrlees, 1971). Those studying optimal criminal justice systems assume that the planner knows how policing and sanctions affect offense rates (Polinsky and Shavell, 2000).

Planners may not possess the knowledge that economists assume them to have. Hence, there is reason to consider policy formation when a planner has only partial knowledge of policy impacts.

I consider the choice of a search profiling policy, wherein decisions to search for evidence of crime may vary with observable covariates of the persons at risk of being searched. Recent research on profiling has sought to define and detect racial discrimination (Knowles, Persico, and Todd, 2003). My concern is to understand how a social planner might reasonably choose a profiling policy..

I suppose that the objective is to minimize the utilitarian social cost of crime and search. Search is costly per se, and search that reveals a crime entails costs for punishment of offenders. Search is beneficial to the extent that it deters or prevents crime. Deterrence is expressed through the *offense function*, which describes how the offense rate of persons with given covariates varies with the search rate applied to these persons. Prevention occurs when search prevents an offense from causing social harm.

I examine the planning problem when the planner has only partial knowledge of the offense function and, hence, is unable to determine what policy is optimal. In particular, I suppose that the planner observes the offense rates of a study population whose search rule has previously been chosen. He knows that the study population and the population of interest have the same offense function. He also knows that search weakly deters crime; that is, the offense rate weakly decreases as the search rate increases. However, the planner does not know the magnitude of the deterrent effect of search. (This is the *monotone-treatment-response* setting of Manski, 1997).

I first show how the planner can eliminate dominated search rules, which are inferior whatever the actual offense function may be. Broadly speaking, low (high) search rates are dominated when the cost of search is low (high). I then show how the planner can use the minimax or minimax-regret criterion to choose an undominated search rule.

FRACTIONAL TREATMENT RULES FOR SOCIAL DIVERSIFICATION OF INDIVISIBLE PRIVATE RISKS

C. Manski, September 2005

Should a social planner treat observationally identical persons identically?

yes, when treatment is individualistic and a utilitarian planner knows the distribution of treatment response.

not necessarily, when a planner has 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.

Illustration: Choosing Treatments for X-Pox

Suppose that a new viral disease called x-pox is sweeping the world. Researchers have proposed two mutually exclusive treatments, a and b, which reflect alternative hypotheses, 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, giving everyone a or b. These rules imply that all of humanity survives with probability 0.3 (0.7) and perishes with probability 0.7 (0.3). 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. Then 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 rule depends on the social welfare function. Let $s(\zeta)$ denote the survival rate with allocation ζ . Let the social welfare for this survival rate be $f[s(\zeta)]$, where $f(\cdot)$ is strictly increasing. The objective is to solve the problem

$$\max_{\zeta \in [0, 1]} E\{f[s(\zeta)]\}.$$

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, then $E\{f[s(\zeta)]\} = 0.7\log(\zeta) + 0.3\log(1 - \zeta)$ and the optimal policy sets $\zeta = 0.7$.

The x-pox problem 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. Society can diversify by having positive fractions of the population receive each treatment.

There are many planning problems in which partial knowledge of treatment response can make a fractional treatment rule desirable. I study several such problems which share some important features: treatment is individualistic, social welfare is an 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.

Bayesian Planning: The planner places a subjective distribution on the unknown state of nature and maximizes subjective expected social welfare. The generic finding is that the optimal treatment rule is a singleton if society is risk neutral and is fractional if society is sufficiently risk averse.

Minimax-Regret Planning: The planner has no subjective or empirical information about the unknown state of nature and chooses a treatment rule minimizing maximum regret. I have previously analyzed this problem when there are two treatments, $f(\cdot)$ is linear, and partial knowledge of treatment response is a consequence of a missing data problem. I found that the minimax-regret rule is fractional whenever both treatments are undominated. The present analysis extends this finding to general problems of choice between two treatments.

Simplifying Assumptions: (1) the members of the population are observationally identical and (2) a one-period planning problem.

Basic Concepts

The basic concepts are as in my earlier work, 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.

The problem is to choose treatments from a finite set T of treatments. Each member j of population J has a response function $y_j(\cdot): T \rightarrow Y$ mapping treatments $t \in T$ into responses $y_j(t) \in Y$. The contribution to welfare from assigning treatment t to person j is $u_j(t) \equiv u[y_j(t), t]$. At the time of treatment choice, the planner knows the form of $u[y_j(t), t]$ but not the value of its first argument $y_j(t)$.

The members of J are observationally identical. A feasible treatment rule assigns all persons to one treatment or fractionally allocates persons across the treatments. Let Z denote the unit simplex on R^T . 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.

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. For each 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 . The planner wants to solve the problem

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

$f(\cdot)$ is strictly increasing.

A solution to (1) assigns the entire population to a treatment solving the problem

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

The social welfare achieved by the optimal rule is

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

Determination of the optimal rule requires knowledge of $E[u(t)]$, $t \in T$.

Let Γ denote the feasible states of nature. 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 $\{E[u(t)], t \in T\}$. The planner can solve the idealized problem if there exists a dominant treatment; that is, one solving all of the problems $\max_{t \in T} E_\gamma[u(t)], \gamma \in \Gamma$. Our concern is with treatment choice when no dominant treatment exists.

Choice Between a Status Quo Treatment and an Innovation

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 planner should choose $\zeta = 1$ if $\beta > \alpha$ and $\zeta = 0$ if $\beta < \alpha$. The problem of interest is choice when α is known but it is only known that $\beta \in (\beta_\gamma, \gamma \in \Gamma)$.

Bayesian Planning

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

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

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

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

If $f(\cdot)$ is concave, the specific forms of Π and $f(\cdot)$ determine whether the Bayesian rule is a singleton or fractional. Consider 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: Consider problem (8). Let $f(\cdot)$ be strictly concave and continuously differentiable. Let Π be non-degenerate. Let $f(\cdot)$ and Π be sufficiently regular that

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

in a neighborhood of $\zeta = 0$. Let $\beta_\Pi \equiv \int \beta_\gamma d\Pi$ denote the subjective mean of β . Then

(a) If $\beta_\Pi \leq \alpha$, the unique solution is $\zeta = 0$.

(b) If $\beta_\Pi > \alpha$, all solutions satisfy $\zeta > 0$.

(c) If $\beta_\Pi > \alpha$ and $\int f(\beta_\gamma) d\Pi < f(\alpha)$, all solutions satisfy $\zeta \in (0, 1)$. □

A Simple Special Case: Let $f(\cdot)$ be the log function, let Γ contain the two elements $\{0, 1\}$, with $\beta_0 < \alpha < \beta_1$. Let $\pi \equiv \Pi(\gamma = 0)$ and $\beta_\Pi = \pi\beta_0 + (1 - \pi)\beta_1$. Then the Bayes rule sets

$$\zeta = \max \left\{ 0, \min \left\{ \alpha(\beta_\Pi - \alpha) / |(\beta_0 - \alpha)(\beta_1 - \alpha)|, 1 \right\} \right\}.$$

Minimax-Regret Planning

The minimax-regret criterion uses only the information that the state of nature lies in Γ . The criterion is

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

$f^*(P_\gamma)$ is the optimal social welfare achievable if it were known that $P = P_\gamma$:

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

$f^*(P_\gamma) - f[U(z, P_\gamma)]$ 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.

Problems with Two Treatments

The minimax-regret rule is generically fractional in problems with two treatments. 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: 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 $(0, 1)$. □

Reasoning: Singleton and fractional rules have different extremum properties across states of nature. Each singleton rule is best in some states of nature and worst in the others. In contrast, fractional rules yield intermediate social welfare in all states of nature.

Choice Between a Status-Quo Treatment and an Innovation

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

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

If $f(\cdot)$ is linear,

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

If $f(\cdot)$ is the log function,

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

Treatment Choice with Missing Outcome Data
(*SCPKTR*, 2005, and *JoE* forthcoming)

Let J_t denote the sub-population of persons whose outcome $y(t)$ is observable.
By the Law of Total Probability,

$$P[y(t)] = P[y(t) | J_t] \cdot P(J_t) + P[y(t) | \text{not } J_t] \cdot P(\text{not } J_t).$$

$P(J_t)$ and $P[y(t) | J_t]$ can be learned empirically, but $P[y(t) | \text{not } J_t]$ cannot.

Proposition: Let $T = \{a, b\}$. Let $\{P[y(t) | J_t], P(J_t); t \in T\}$ be known. Let $u_{0t} \equiv \inf_{y \in Y} u(y, t)$ and $u_{1t} \equiv \sup_{y \in Y} u(y, t)$ be finite. Let $e_t \equiv E[u(t) | J_t]$ and $p_t \equiv P(J_t)$. Then the minimax-regret rule is

$$z^*(b) = 1 \text{ if } (e_a - u_{1a})p_a + (u_{0b} - e_b)p_b + (u_{1a} - u_{0b}) < 0,$$

$$= 0 \text{ if } (e_b - u_{1b})p_b + (u_{0a} - e_a)p_a + (u_{1b} - u_{0a}) < 0,$$

$$= \frac{(e_b - u_{1b})p_b + (u_{0a} - e_a)p_a + (u_{1b} - u_{0a})}{(u_{0b} - u_{1b})p_b + (u_{0a} - u_{1a})p_a + (u_{1b} - u_{0b}) + (u_{1a} - u_{0a})} \quad \text{otherwise.}$$

Implementing Fractional Treatment Rules

The notion of a planner who acts on behalf of society is a useful fiction. Nevertheless, we must ultimately consider the feasibility of implementing fractional treatment rules in democratic societies.

A possible legal/ethical objection to fractional rules is that they violate the normative principle calling for “equal treatment of equals.” Fractional rules do 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.

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. The existence of a consensus minimax-regret rule requires only that the population agree on the feasible states of nature.

A specific aspect of American public policy that seems well-suited for implementation of fractional rules 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 demonstration of “substantial evidence of effect.” 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.