TREATMENT CHOICE WITH PARTIAL KNOWLEDGE OF TREATMENT RESPONSE

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 provide.

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

I use elements of statistical decision theory to address this question.

Bibliography in Progress

Monographs

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

Identification for Prediction and Decision, Chaps. 11 and 12, in process.

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, 2004, pp. 1221-1246.

"Search Profiling with Partial Knowledge of Deterrence," *Economic Journal*, forthcoming.

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

"Admissible Treatment Rules for a Risk-Averse Planner with Experimental Data on an Innovation," with A. Tetenov, *Journal of Statistical Planning and Inference*, forthcoming.

Related Work

Brock, W. (2006), "Profiling Problems with Partially Identified Structure," *Economic Journal*, forthcoming.

Hirano, K. and J. Porter (2005), "Asymptotics for Statistical Treatment Rules."

Stoye, J. (2004), "Minimax Regret Treatment Choice with Finite Samples."

Stoye, J. (2005), "Minimax Regret Treatment Choice with Incomplete Data and Many Treatments."

Schlag, K. (2006), "ELEVEN–Tests Need for a Recommendation"

SEARCH PROFILING WITH PARTIAL KNOWLEDGE OF DETERRENCE

C. Manski, (*Economic Journal*, forthcoming)

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, where 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

Research on social planning recognizes that a planner may want to treat observationally different persons differently. This is the essence of profiling.

However, research on social planning usually presumes that a planner should treat observationally identical persons identically.

Uniform treatment is appropriate when a utilitarian planner knows the population distribution of treatment response.

It may not be desirable with partial knowledge of treatment response.

Identification problems or issues of statistical inference can make fractional rules desirable.

Implementation of a fractional rule enables society to diversify a risk that is privately indivisible. An individual cannot diversify—a person receives either treatment a or b. Society can diversify by having positive fractions of the population receive each treatment.

The Ethics of Fractional Rules

A possible ethical objection to fractional rules is that they violate the normative principle calling for "equal treatment of equals."

Fractional rules are consistent with this principle in the *ex ante* sense that observationally identical people have the same probability of receiving a particular treatment. Fractional rules violate the principle in the *ex post* sense that observationally identical persons ultimately receive different treatments.

Societies sometimes implement the ex ante sense of "equal treatment" in the design of major policies. Examples include random drug testing, calls for jury service, and the American Green Card and Vietnam draft lotteries.

Experiments with randomized assignment of treatments provide equal treatment in the ex ante sense. Indeed, the prevailing standard of medical ethics permits randomized clinical trials only when partial knowledge of treatment response prevents a determination that one treatment is superior to another.

I have studied several simple planning problems in which partial knowledge of treatment response can makes a fractional treatment rule desirable. These problems share certain features:

(a) treatment is individualistic

(b) social welfare is an increasing function of a population mean outcome

(c) outcomes depend on an unknown state of nature.

(d) members of the population are observationally identical

(e) a one-period planning horizon.

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.

I will discuss

– Choosing Treatments for X-Pox

– Choice between a status quo treatment and an innovation

Data = large randomized experiment with partial compliance

Data = small classical randomized experiment

– Minimax-regret choice between two undominated treatments

A general result

Planning with missing outcomes.

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, $t = a$ and $t = b$, which reflect alternative hypotheses, say H_a and H_b, about the nature of the virus. If H_t is correct, all persons who receive treatment t survive and all others die. It is known that one of the two hypotheses is correct, but it is not known which one; thus, there are two states of natures, $\gamma = H_a$ and $\gamma = H_b$. Suppose that the objective is to maximize the survival rate of the population

There are two singleton rules in this setting, one giving treatment a to the entire population and the other giving b. Each rule provides equal treatment of equals in the ex post sense. Each also equalizes realized outcomes. The entire population either survives or dies.

Consider the rule in which a fraction $\delta \in [0, 1]$ of the population receives treatment b and the remaining $1 - \delta$ receives treatment a. Under this rule, the fraction who survive is

$$\delta \cdot 1[\gamma = H_b] + (1 - \delta) \cdot 1[\gamma = H_a].$$

The maximin and the minimax-regret rule both set $\delta = ½$. These rules treat everyone equally ex ante, each person having a 50 percent chance of receiving each treatment. They do not treat people equally ex post. Nor do they equalize outcomes. Half the population lives and half dies.

Choice Between a Status Quo Treatment and an Innovation

$T = \{a, b\}$. $t = a$ is the *status quo* and $t = b$ is the *innovation*.

The outcomes $y(t)$ are binary. Let $\alpha \equiv P[y(a) = 1]$ and $\beta \equiv P[y(b) = 1]$.

The planner knows α. He knows that $\beta \in B$, with $\alpha \in \text{int}(B)$. B are the states of nature.

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

$$\alpha(1 - \delta) + \beta\delta \;=\; \alpha + (\beta - \alpha)\delta.$$

Social welfare is $f[\alpha + (\beta - \alpha)\delta]$, where $f(\cdot)$ is increasing.

The planner should choose $\delta = 1$ if $\beta > \alpha$ and $\delta = 0$ if $\beta < \alpha$. The problem is treatment choice when α is known but it is only known that $\beta \in B$.

Treatment Using Data from an Experiment with Partial Compliance

A 'large' randomized experiment is performed on a study population. Subjects assigned to the innovation can refuse to comply and choose the status quo treatment instead. Those assigned to the status quo cannot cross over to receive the innovation.

The empirical evidence point-identifies α but only partially identifies β. Its identification region is the interval $B = [\beta_L, \beta_H]$, where

$$\beta_L \equiv P(y = 1 | \zeta = b, z = b)P(z = b | \zeta = b),$$
$$\beta_U \equiv P(y = 1 | \zeta = b, z = b)P(z = b | \zeta = b) + P(z \neq b | \zeta = b).$$

Here ζ is the treatment assigned and z is the treatment received.

A planner must choose treatments for a new population known to be identical to the study population in its distribution of treatment response. The planner's objective is to maximize the rate of treatment success $\alpha + (\beta - \alpha)\delta$.

I consider Bayes rules, the maximin criterion, and the minimax-regret criterion.

Bayes Rules

A Bayesian planner places a subjective probability distribution π on the interval $[\beta_L, \beta_U]$, computes the subjective mean value of social welfare, and chooses a treatment allocation that maximizes this subjective mean. Thus, the planner solves the optimization problem

$$\max_{\delta \in [0, 1]} \quad \alpha + [E_\pi(\beta) - \alpha]\delta,$$

where $E_\pi(\beta) = \int \beta d\pi$ is the subjective mean of β. The Bayes decision assigns everyone to the innovation if $E_\pi(\beta) > \alpha$ and everyone to the status quo if $\alpha > E_\pi(\beta)$. All treatment allocations are Bayes decisions if $E_\pi(\beta) = \alpha$.

The Maximin Criterion

A maximin planner acts as if β equals its smallest feasible value, β_L. Thus, the planner solves the optimization problem

$$\max_{\delta \in [0, 1]} \quad \alpha + (\beta_L - \alpha)\delta.$$

The maximin rule sets $\delta = 0$, which assigns everyone to the status quo.

The Minimax-Regret Criterion

Suppose that the planner chooses allocation δ and that $\beta = k$, where $k \in [\beta_L, \beta_U]$. Then regret is

$$\max(\alpha, k) - [\alpha + (k - \alpha)\delta] = (\alpha - k)\delta \cdot 1[k < \alpha] + (k - \alpha)(1 - \delta) \cdot 1[k > \alpha].$$

Maximum regret across all the feasible values of k is

$$\max [(\alpha - \beta_L)\delta, (\beta_U - \alpha)(1 - \delta)].$$

A minimax-regret rule solves the optimization problem

$$\min_{\delta \in [0, 1]} \max [(\alpha - \beta_L)\delta, (\beta_U - \alpha)(1 - \delta)].$$

The solution is obtained by choosing δ to solve the equation

$$(\alpha - \beta_L)\delta = (\beta_U - \alpha)(1 - \delta).$$

The minimax-regret treatment allocation is the fraction

$$\delta_{MR} = (\beta_U - \alpha)/(\beta_U - \beta_L).$$

δ_{MR} decreases linearly from 1 to 0 as α increases from β_L to β_U.

Illustration: The Illinois UI Experiment

The status quo is conventional UI and the innovation is UI with a wage subsidy. Let y(t) = 1 if an unemployed person is rehired within 11 weeks and y(t) = 0 otherwise. Dubin and Rivers (1993) report that

$$\alpha = 0.35, \qquad P(y = 1 \mid \zeta = b, z = b) = 0.38, \qquad P(z = b \mid \zeta = b) = 0.68.$$

Hence, $\beta_L = 0.26$ and $\beta_U = 0.58$.

Let the objective be to maximize the fraction of unemployed persons who are rehired within 11 weeks.

A Bayes rule assigns everyone to UI with the wage subsidy if $E_\pi(\beta) > 0.35$ and everyone to conventional UI if $E_\pi(\beta) < 0.35$.

The maximin rule assigns everyone to conventional UI.

The minimax-regret rule assigns 72 percent of all unemployed persons to UI with the wage subsidy and 38 percent to conventional UI.

<u>Admissible Treatment Rules for a Risk-averse Planner with Experimental Data on an Innovation</u>

C. Manski and A. Tetenov (*JSPI*, 2007)

Let $B = (0, 1)$. A classical randomized experiment with N subjects is performed on the innovation. The feasible treatment rules are functions $z(\cdot)$ that map the number of experimental successes into a treatment allocation.

Rule z is *admissible* if there exists no other rule z' such that $W(z; \beta) \leq W(z'; \beta)$ for all $\beta \in B$ and $W(z; \beta) < W(z'; \beta)$ for some $\beta \in B$.

A class of treatment rules is *essentially complete* if, given any rule outside this class, there exists a member of the class that performs at least as well in all states of nature.

A class of rules is *complete* if, given any rule outside this class, there exists a member of the class that performs at least as well in all states of nature and better in some state of nature.

A class of rules is *minimal complete* if the class is complete and all of its members are admissible.

Let the objective be to maximize the population rate of treatment success.

Karlin and Rubin (1956) show that the admissible rules assign all members of the population to the status quo if the number of experimental successes is below a specified threshold and all to the innovation if the number of successes is above the threshold.

An interior fractional allocation is admissible only when the number of experimental successes exactly equals the threshold.

Karlin and Rubin called this class of treatment rules *monotone*, but we say *KR-monotone*.

Let the objective be maximization of a concave-monotone function f(·) of the success rate. Thus, the planner is "risk-averse."

We show that the admissible rules depend on the curvature of f(·). With sufficient curvature, admissible treatment rules need not be KR-monotone and some KR-monotone rules are inadmissible.

Define a *fractional monotone rule* to be one where the fraction of persons assigned to the innovation weakly increases with the number of experimental successes. We show that the class of fractional monotone rules is complete if f(·) is concave and strictly monotone.

Define an *M-step monotone rule* to be a fractional monotone rule with an interior fractional treatment assignment for no more than M consecutive values of the number of experimental successes. We show that the M-step monotone rules are a complete class if f(·) is differentiable and has sufficiently weak curvature.

We show that Bayes rules and the minimax-regret rule depend on the curvature of the welfare function. They are KR-monotone if the curvature is sufficiently weak, but give interior fractional treatment allocations if the curvature is sufficiently strong.

Minimax-Regret Choice with Two Undominated Treatments

General Finding

T = {a, b}. The objective is to maximize the mean outcome. Both treatments are undominated on the set Γ of states of nature..

The minimax-regret criterion is

$$\min_{\delta \in [0, 1]} \max_{\gamma \in \Gamma} \max \{E_\gamma[y(a)], E_\gamma[y(b)]\} - \{(1 - \delta)E_\gamma[y(a)] + \delta E_\gamma[y(b)]\}.$$

Proposition: There is a unique fractional minimax-regret rule whenever a continuity condition holds.

Continuity Condition: Let Γ(a) and Γ(b) be the subsets of Γ on which treatments a and b are superior. That is, let $\Gamma(a) \equiv \{\gamma \in \Gamma: E_\gamma[y(a)] \geq E_\gamma[y(b)]\}$ and $\Gamma(b) \equiv \{\gamma \in \Gamma: E_\gamma[y(b)] \geq E_\gamma[y(a)]\}$. Let

$$R(\delta, a) \equiv \sup_{\gamma \in \Gamma(a)} E_\gamma[y(a)] - \{(1 - \delta)E_\gamma[y(a)] + \delta E_\gamma[y(b)]\}$$

$$R(\delta, b) \equiv \sup_{\gamma \in \Gamma(b)} E_\gamma[y(b)] - \{(1 - \delta)E_\gamma[y(a)] + \delta E_\gamma[y(b)]\}$$

be the maximum regret of rule δ on Γ(a) and Γ(b) respectively. Suppose that R(·, a) and R(·, b) are continuous on [0, 1].

Proof: The maximum regret of rule δ on all of Γ is max $[R(\delta, a), R(\delta, b)]$. As δ 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 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 δ. Hence, the minimax-regret rule is the unique $\delta \in (0, 1)$ such that $R(\delta, a) = R(\delta, b)$. □

Treatment Choice with Missing Outcome Data
(Manski, *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

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

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

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

Minimax-Regret Planning and the Selection Problem

Let outcomes take values in the interval [0, 1[. Let utility be the outcome of treatment.

Consider an observational study with no assumptions on the process of treatment selection, Then

$$\delta_{MR} = [1 - E(y|z = a)]P(z = a) + E(y|z = b)P(z = b).$$

Drug Approval at the FDA

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.