

DIVERSIFIED TREATMENT UNDER AMBIGUITY

Charles F. Manski

Department of Economics and Institute for Policy Research, Northwestern University

2009 Lawrence R. Klein Lecture
The International Economic Review

Abstract

This paper develops a broad theme about treatment under ambiguity through study of a particular decision criterion. The broad theme is that a planner may often want to cope with ambiguity by diversification, assigning observationally identical persons to different treatments. Study of the minimax-regret (MR) criterion substantiates the theme. The paper significantly extends my earlier analysis of one-period planning with individualistic treatment and a linear welfare function. I show that MR treatment allocations are fractional in a large class of planning problems with nonlinear welfare functions, interacting treatments, learning, and non-cooperative aspects. I also call attention to some problems of treatment under ambiguity in which the MR allocation is not fractional.

This research was supported in part by National Science Foundation grant SES-0549544. This paper revises and supercedes an earlier working paper "Fractional Treatment Rules for Social Diversification of Indivisible Private Risks," NBER Working Paper W11675, September 2005. I have benefitted from the opportunity to present this work at an invited session of the 2009 North American Winter Meeting of the Econometric Society and in seminars at Collegio Carlo Alberto, The Hebrew University of Jerusalem, Northwestern University, Princeton University, Universitat Pompeu Fabra, Université Paris 1, UCLA, University College London, and the University of Washington.

1. Introduction

When studying collective decision problems, economists have long asked how a planner should act. A standard exercise specifies a set of feasible policies and a welfare function. The planner is presumed to know the welfare achieved by each policy. The objective of the exercise is to characterize the optimal policy.

In practice planners often have only partial knowledge of the welfare achieved by alternative policies. Hence, they cannot determine optimal policies. This limits the relevance of the standard exercise to actual policy analysis.

In a research program that began in Manski (1990, 1995), I have studied how identification problems that are prevalent in empirical research generate ambiguity about the nature of optimal policies. Perhaps the most fundamental identification problem arises from the unobservability of counterfactual policy outcomes. At most one can observe the outcomes that occur under realized policies. The outcomes of unrealized policies are logically unobservable. Yet determination of an optimal policy requires comparison of all feasible policies.

My recent work considers how a planner might cope with ambiguity; see Manski (2000, 2004a, 2005a, 2006, 2007a, 2007b). The familiar Bayesian prescription is to assert a subjective probability distribution over the feasible states of nature and choose an action that maximizes subjective expected welfare. However, a subjective probability distribution is itself a form of knowledge, and a planner may have no credible basis for asserting one. I have studied problems of this type, with particular attention to application of the minimax-regret (MR) criterion.

Many of the planning problems that I have studied to date share a relatively simple structure. The planner must choose one of two treatments, say a and b, for each member of a population of observationally identical persons. The planner can treat population members differentially, assigning some persons to treatment a and the remainder to b. Treatment is individualistic, each person's outcome depending only on the treatment that he receives and not on the treatments of other persons. Outcomes take a bounded range

of values. The welfare function is linear, summing the outcomes of the population members. In this setting, the optimal policy assigns all persons to the treatment that yields the higher mean outcome. A planner faces ambiguity when he does not know which treatment is better.

Here are two illustrative applications among many that might be cited, the first being an instance of social planning and the second being one of private planning:

Choosing Medical Treatments for a Non-Infectious Disease: The planner is a public health agency which chooses treatments for a population of persons who are susceptible to a non-infectious disease. The relevant welfare outcome is the health benefit of a treatment minus its cost, measured in comparable units. A utilitarian welfare function sums net benefits across the population. The disease being non-infectious, each person's outcome depends only on his own treatment. Medical research yields only partial knowledge of treatment response, so the planner does not know which treatment is better.

An Investor's Asset Allocation Decision: The planner is an investor who allocates an endowment between two assets. The population members are dollars of endowment and the treatments are the two assets. The relevant outcome is the return on a dollar invested in an asset. Welfare is the aggregate return earned by the investor. At the time of the allocation decision, the investor has only partial knowledge of investment returns. Hence, he does not know what allocation maximizes profit.

While asset distribution has long been framed as a choice among alternative fractional allocations, medical treatment and other social planning problems have commonly been viewed as choices between two singleton allocations—either assign treatment a to all observationally identical persons or assign b to everyone. This specification of the choice set suffices when a planner knows the optimal treatment, but it may not when he faces ambiguity. Then a planner may reasonably select a fractional treatment allocation,

assigning some persons to treatment a and the remainder to treatment b. Fractional allocations enlarge the set of feasible policy choices by convexifying the singleton allocations. They cope with ambiguity through diversification.¹

The broad argument for diversification is that it enables a decision maker to balance two types of potential error. A Type A error occur when treatment a is chosen but is actually inferior to b, and a Type B error occurs when b is chosen but is inferior to a. The singleton allocation assigning everyone to treatment a entirely avoids type B errors but may yield Type A errors, and vice versa for singleton assignment to treatment b. Fractional allocations make both types of errors but reduce their potential magnitudes. Hence, it is conceivable that some fractional allocations may be preferred to the singleton ones.

A Bayesian planner with a linear welfare function generically chooses a singleton allocation. A necessary condition for Bayesian diversification is that the welfare function be strictly concave on at least part of its domain. In previous work (Manski 2005a; 2007a; 2007b, Chapter 11) I have shown that a planner with a linear welfare function who applies the minimax-regret criterion always diversifies when facing ambiguity. The MR criterion chooses an allocation that balances the potential welfare losses from Type A and Type B errors.

I build on this finding here, extending my earlier analysis to planning problems with nonlinear welfare functions (Section 3), interacting treatments (Section 4), learning (Section 5), and non-cooperative aspects (Section 6). I show that the minimax-regret treatment allocation is fractional in a broad class of planning problems under ambiguity. This contrasts with Bayesian and maximin planning, which diversify treatment choice in a much smaller class of settings. I also call attention to some problems of treatment under ambiguity in which the MR allocation is not necessarily fractional.

¹ It is important to distinguish differential treatment of persons who vary in observable respects from fractional allocations of persons who are observationally identical. It is well known that enabling treatment choice to vary systematically with observed covariates of population members can improve welfare if treatment response varies with these covariates; see, for example, Manski (2005a, Sec. 1.2). In contrast, fractional allocations randomly differentiate among persons who are observationally identical.

Section 2 reviews one-period planning with individualistic treatment and linear welfare. After introducing treatment choice under ambiguity, I compare Bayesian, maximin, and minimax-regret planning. I extend the analysis to settings where the planner observes covariates that distinguish among members of the population.

Section 3 considers a broader class of welfare functions than I have studied previously. I first permit welfare functions that monotonically transform the aggregate sum of outcomes. I show that the MR treatment allocation remains fractional when a planner faces ambiguity, the specific allocation depending on the transformation. I then address planning problems with non-additive cost of treatment, focusing on the polar cases of capacity constraints and fixed costs. I show that the MR allocation under ambiguity is fractional with all capacity constraints and with small fixed costs. However, it is singleton if fixed costs are sufficiently large.

I also consider planning with deontological welfare functions, which enable a planner to express a possible ethical objection to fractional treatment allocations. Fractional allocations violate one interpretation of the ethical principle calling for “equal treatment of equals.” They are consistent with this principle in the *ex ante* sense that all observationally identical people have the same probability of receiving a particular treatment, but they violate it in the *ex post* sense that observationally identical persons ultimately receive different treatments. The *ex post* sense of equal treatment expresses a deontological consideration that is absent from the consequentialist welfare functions usually assumed in economic analysis of planning. I formalize this idea and show that a special case is mathematically equivalent to planning with fixed costs.

Section 4 relaxes the assumption that treatment is individualistic and considers settings where treatments interact, each person’s outcome depending both on his own treatment and on the treatment allocation within the population. Problems of this type are more complex than ones with individualistic treatment. Ambiguity is more severe because determination of an optimal policy now requires knowledge not only of how outcomes vary with persons’ own treatments but also with the treatment allocation. I use

medical treatment of an infectious disease to illustrate.

Whereas Sections 2 through 4 suppose a one-period planning problem, Section 5 considers multi-period planning under ambiguity. I suppose that, in each of a sequence of periods, a planner chooses treatments for the current cohort of a population. This introduces the possibility of learning, with observation of treatment outcomes in earlier periods informing treatment choice in later periods.

I suggest use of the *adaptive minimax-regret (AMR)* criterion, which treats each cohort as well as possible in the static minimax-regret sense, using the information available at the time. The result is a fractional treatment allocation whenever the available knowledge does not suffice to determine which treatment is better. The criterion is adaptive because knowledge of treatment response accumulates over time, so successive cohorts may receive different fractional allocations.

Fractional allocations randomly assign observationally identical persons to different treatments. Hence, they automatically create randomized experiments, which are particularly informative for learning treatment response. I use medical treatment to illustrate. I explain how the AMR criterion differs from the current practice of randomized clinical trials in medicine and I discuss the drug approval process of the U. S. Food and Drug Administration.

Whereas Sections 2 through 5 are written from the perspective of a planner with the power to dictate policy, Section 6 considers situations in which treatment choice requires the agreement of two planners who may have different welfare functions and beliefs about the feasible states of nature. I show that if both planners face ambiguity and use the AMR criterion to compare allocations, then there exist fractional allocations that both prefer to the singleton ones. I use a union-management decision problem to illustrate. Section 7 concludes.

2. One-Period Planning with Individualistic Treatment and Linear Welfare

I review here the simple planning problem that forms the baseline for this paper. Section 2.1 sets out basic concepts and notation. Sections 2.2 and 2.3 consider Bayesian, maximin, and minimax-regret planning. Section 2.4 uses sentencing of convicted offenders to illustrate how treatment choice varies with the decision criterion used. Section 2.5 considers treatment choice when the planner observes covariates that distinguish among members of the population.

2.1. Basic Concepts and Notation

The Treatment-Choice Problem

There are two treatments, labeled a and b. The set of feasible treatments is $T \equiv \{a, b\}$. Each member j of a population denoted J has a response function $y_j(\cdot): T \rightarrow Y$ that maps treatments $t \in T$ into outcomes $y_j(t) \in Y$. The subscript j in $y_j(\cdot)$ indicates that treatment response may vary across the population. Let $u_j(t) \equiv u_j[y_j(t), t]$ denote the net contribution to welfare that occurs if person j receives treatment t and realizes outcome $y_j(t)$. For example, $u_j(t)$ may have the “benefit-cost” form $u_j(t) = y_{j1}(t) - y_{j2}(t)$, where $y_{j1}(t)$ is the benefit of treatment t and $y_{j2}(t)$ is its cost. Although treatment response may vary across the population, persons are observationally identical to the planner.

Let $P[y(\cdot)]$ denote the population distribution of treatment response. I suppose that the population is large in the formal sense of being atomless; that is, $P(j) = 0$ for all $j \in J$. This idealization implies that if the planner randomly assigns a positive fraction of the population to a treatment, the sub-population of persons who receive this treatment is infinite. This eliminates sampling variation as an issue when comparing alternative treatment allocations and analyzing treatment response.

The planner’s task is to allocate the population between the two treatments. A treatment allocation

is a $\delta \in [0, 1]$ that randomly assigns a fraction δ of the population to treatment b and the remaining $1 - \delta$ to treatment a.² I assume that the planner wants to choose a treatment allocation that maximizes mean welfare in the population. Let $\alpha \equiv E[u(a)]$ and $\beta \equiv E[u(b)]$ be the mean welfare that would result if all persons were to receive treatment a or b respectively. Welfare with allocation δ is

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

$W(\cdot)$ is a consequentialist welfare function that additively aggregates individual contributions to welfare. If the function $u(\cdot)$ expresses private preferences, then $W(\cdot)$ is the utilitarian welfare function often assumed in research on welfare economics.

The optimal treatment allocation is obvious if (α, β) are known. The planner should choose $\delta = 1$ if the average treatment effect $\beta - \alpha$ is positive and $\delta = 0$ if it is negative. The problem of interest is treatment choice when the sign of the average treatment effect is unknown.

To formalize the problem, let S index the feasible states of nature. Thus, the planner knows (or believes) that (α, β) lies in the set $[(\alpha_s, \beta_s), s \in S]$. This set is the identification region for (α, β) ; that is, the set of values of (α, β) that the planner concludes are feasible when he combines available empirical evidence with assumptions he finds credible to maintain. Manski (2007b) expositis many specific cases, showing the form that the identification region takes when observation of realized treatment outcomes in a study population is combined with various assumptions about treatment response and selection. The present analysis is applicable whenever $[(\alpha_s, \beta_s), s \in S]$ is bounded. I will denote the extreme feasible values of α

² Treatment is random from the perspective of population members, each person having a probability δ of receiving b and $1 - \delta$ of receiving a. However, a treatment allocation as defined here is not random in the decision theoretic sense of a mixed strategy. Allocation δ assigns fixed fractions δ and $1 - \delta$ of the population to treatments b and a respectively. A mixed strategy would make these fractions random, with their realizations determined by an auxiliary randomizing device. Formally, a mixed strategy is a probability distribution on $[0, 1]$.

and β as $\alpha_L \equiv \min_{s \in S} \alpha_s$, $\beta_L \equiv \min_{s \in S} \beta_s$, $\alpha_U \equiv \max_{s \in S} \alpha_s$, and $\beta_U \equiv \max_{s \in S} \beta_s$.

Illustration: Let $Y = [0, 1]$ and $u(\cdot) = y(\cdot)$. Let outcomes be observed in a study population J^* . Let $z_j = a$ if person $j \in J^*$ received treatment a and $z_j = b$ if j received b . Let $y_j \equiv y_j(z_j)$ be the outcome realized by j . Let $P^*[y(\cdot), z]$ be the distribution of treatment response and receipt in the study population. Thus, the planner observes $P^*(y, z)$.

Suppose the planner finds it credible to assume that the treatment and study populations have the same distribution of treatment response; that is, $P[y(\cdot)] = P^*[y(\cdot)]$. Manski (1990) showed that, in the absence of other assumptions, $[(\alpha_s, \beta_s), s \in S]$ is the rectangle $[\alpha_L, \alpha_U] \times [\beta_L, \beta_U]$, where

$$\begin{aligned} \alpha_L &= E^*(y|z = a)P^*(z = a), & \alpha_U &= E^*(y|z = a)P^*(z = a) + P^*(z = b), \\ \beta_L &= E^*(y|z = b)P^*(z = b), & \beta_U &= E^*(y|z = b)P^*(z = b) + P^*(z = a). \end{aligned} \quad \square$$

Partial knowledge is unproblematic for decision making if $(\alpha_s \geq \beta_s, s \in S)$ or if $(\alpha_s \leq \beta_s, s \in S)$; choosing $\delta = 0$ is optimal in the former case and $\delta = 1$ in the latter. The planner faces ambiguity if both treatments are undominated; that is, if $\alpha_s > \beta_s$ for some values of s and $\alpha_s < \beta_s$ for other values. I assume that the planner faces ambiguity.

Criteria for Choice Under Ambiguity

A planner facing ambiguity does not know the optimal treatment allocation. Yet he must somehow choose an allocation. How might he do so?

Decision theorists have proposed various ways of transforming the original optimization problem, which cannot be solved, into another one that can be solved. Bayesians recommend that a decision maker assert a subjective distribution on the feasible states of nature and choose an allocation that maximizes subjective mean welfare with respect to this distribution. The maximin and the minimax-regret criteria do

not use a subjective distribution. Instead they choose allocations that, in different senses, perform uniformly well over all states of nature.³

Bayesian decision theorists have often asserted pre-eminence for maximization of expected utility (welfare here), asserting not only that a decision maker *might* use this decision criterion but that he *should* do so. Reference is often made to representation theorems deriving the expected utility criterion from consistency axioms on hypothetical choice behavior, famously von Neumann and Morgenstern (1944) and Savage (1954). These and other contributions to axiomatic decision theory consider a decision maker who has formed a complete binary preference ordering over a specified class A of actions and, thus, who knows how he would behave if he were to face any choice set $D \subset A$. The theorems show that if the preference ordering adheres to certain consistency axioms, then the agent may be represented as maximizing expected utility. Thus, the theorems of axiomatic decision theory are interpretative rather than prescriptive.

Why then are the N-M and Savage theorems often considered to be prescriptive? Decision theorists often assert that an agent *should* form a complete binary preference ordering on the class A of actions and that preferences *should* adhere to the proposed axioms. If one accepts these assertions, the theorems imply that the agent *should* behave in a manner representable as maximization of expected utility. Thus, the theorems are prescriptive if one considers their consistency axioms to be compelling.⁴

³ These three approaches to decision making under ambiguity are particularly well-known, but they are not the only ones that have received attention. A decision maker who feels able to assert a subjective distribution on the states of nature need not maximize subjective mean welfare. He could instead maximize some quantile of the welfare distribution (see Manski, 1988). A decision maker who feels able to assert only a partial distribution on the states of nature could maximize minimum expected welfare or minimize maximum expected regret. These ideas have a long history in the literature on statistical decision theory, which refers to them as the Γ -maximin and Γ -minimax regret criteria (see Berger, 1985). The Γ -maximin approach has also drawn considerable attention from economists (e.g., Gilboa and Schmeidler, 1989; Hansen and Sargent, 2008).

⁴ A famous example is the Chernoff (1954) argument against the minimax regret criterion. Chernoff observed that this criterion can violate the consistency axiom called independence of irrelevant alternatives (IIA). The IIA axiom holds that if an agent is not willing to choose a given action from a hypothetical choice set, then he should not be willing to choose it from any larger hypothetical choice set; thus, for any $c \in D \subset E$, an agent who would not choose c from D should not choose c from E . Chernoff wrote (p. 426):

I will not use consistency axioms to argue for or against particular decision criteria. In Manski (2008), I have observed that a decision maker who wants to choose an optimal policy but lacks the knowledge to do so is not concerned with the consistency of his behavior across hypothetical choice sets. Rather, he wants to make a reasonable choice from the choice set that he actually faces. Hence, I reason that prescriptions for decision making should *respect actuality*. That is, they should promote welfare maximization in the choice problem the agent actually faces. Expected utility maximization respects actuality, but it has no special status from the actualist perspective.

2.2. Bayesian and Maximin Decision Making

Bayes Rules

A Bayesian planner places a subjective probability distribution π on the set S of feasible states of nature, computes the subjective mean value of welfare under each treatment allocation, and chooses an allocation that maximizes this subjective mean. Thus, the planner solves the optimization problem

$$(2) \quad \max_{\delta \in [0, 1]} E_{\pi}(\alpha) + [E_{\pi}(\beta) - E_{\pi}(\alpha)]\delta,$$

where $E_{\pi}(\alpha) = \int \alpha_s d\pi$ and $E_{\pi}(\beta) = \int \beta_s d\pi$ are the subjective means of α and β . The Bayes decision assigns everyone to treatment b if $E_{\pi}(\beta) > E_{\pi}(\alpha)$ and everyone to treatment a if $E_{\pi}(\alpha) > E_{\pi}(\beta)$. All allocations are

A third objection which the author considers very serious is the following. In some examples, the min max regret criterion may select a strategy d_3 among the available strategies d_1, d_2, d_3 , and d_4 . On the other hand, if for some reason d_4 is made unavailable, the min max regret criterion will select d_2 among d_1, d_2 , and d_3 . The author feels that for a reasonable criterion the presence of an undesirable strategy d_4 should not have an influence on the choice among the remaining strategies. This passage is the totality of Chernoff's argument. He introspected and concluded that any reasonable decision criterion should adhere to IIA, without explaining why he felt this way. He did not argue that minimax-regret decisions have adverse welfare consequences.

Bayes decisions if $E_{\pi}(\beta) = E_{\pi}(\alpha)$. Thus, a Bayesian planner behaves as would a planner who knows that the population means in (1) have the values in (2).

Although Bayesian planning is conceptually straightforward, it may not be straightforward to form a credible subjective distribution on the states of nature. The allocation chosen by a Bayesian planner depends on the subjective distribution used. Here, as always, the Bayesian paradigm is appealing only when a decision maker is able to form a subjective distribution that really expresses his beliefs.

The Maximin Criterion

To determine the maximin allocation, one first computes the minimum welfare attained by each allocation across all states of nature. One then chooses an allocation that maximizes this minimum welfare. Thus, the criterion is

$$(3) \quad \max_{\delta \in [0, 1]} \min_{s \in S} \alpha_s + (\beta_s - \alpha_s)\delta.$$

The solution has a simple form if (α_L, β_L) is a feasible value of (α, β) , as is so when the identification region is rectangular. Then the maximin allocation is $\delta = 0$ if $\alpha_L > \beta_L$, $\delta = 1$ if $\alpha_L < \beta_L$, and all $\delta \in [0, 1]$ if $\alpha_L = \beta_L$.

2.3. The Minimax-Regret Criterion

By definition, the regret of treatment allocation δ in state of nature s is the difference between the maximum achievable welfare and the welfare achieved with this allocation. The maximum welfare achievable in state of nature s is $\max(\alpha_s, \beta_s)$. Hence, allocation δ has regret $\max(\alpha_s, \beta_s) - [\alpha_s + (\beta_s - \alpha_s)\delta]$. The minimax-regret rule computes the maximum regret of each allocation over all states of nature and chooses an allocation to minimize maximum regret. Thus, the criterion is

$$(4) \quad \min_{\delta \in [0, 1]} \max_{s \in S} \max(\alpha_s, \beta_s) - [\alpha_s + (\beta_s - \alpha_s)\delta].$$

Let $S(a)$ and $S(b)$ be the subsets of S on which treatments a and b are superior. That is, let $S(a) \equiv \{s \in S: \alpha_s > \beta_s\}$ and $S(b) \equiv \{s \in S: \beta_s > \alpha_s\}$. Let $M(a) \equiv \max_{s \in S(a)} (\alpha_s - \beta_s)$ and $M(b) \equiv \max_{s \in S(b)} (\beta_s - \alpha_s)$ be maximum regret on $S(a)$ and $S(b)$ respectively. Define $M(a) = 0$ if $S(a)$ is empty and $M(b) = 0$ if $S(b)$ is empty. Manski (2007b, Complement 11A) proves that the MR criterion always makes a fractional treatment allocation when both treatments are undominated. The result is

$$(5) \quad \delta_{MR} = \frac{M(b)}{M(a) + M(b)}.$$

The proof is short and instructive, so I reproduce it here.

Proof: The maximum regret of rule δ on all of S is $\max [R(\delta, a), R(\delta, b)]$, where

$$(6a) \quad R(\delta, a) \equiv \max_{s \in S(a)} \alpha_s - [(1 - \delta)\alpha_s + \delta\beta_s] = \max_{s \in S(a)} \delta(\alpha_s - \beta_s) = \delta M(a),$$

$$(6b) \quad R(\delta, b) \equiv \max_{s \in S(b)} \beta_s - [(1 - \delta)\alpha_s + \delta\beta_s] = \max_{s \in S(b)} (1 - \delta)(\beta_s - \alpha_s) = (1 - \delta)M(b),$$

are maximum regret on $S(a)$ and $S(b)$. Both treatments are undominated, so $R(1, a) = M(a) > 0$ and $R(0, b) = M(b) > 0$. As δ increases from 0 to 1, $R(\cdot, a)$ increases linearly from 0 to $M(a)$ and $R(\cdot, b)$ decreases linearly from $M(b)$ to 0. Hence, the MR rule is the unique $\delta \in (0, 1)$ such that $R(\delta, a) = R(\delta, b)$. This yields (5). \square

The proof of (5) shows that the MR allocation balances the two types of potential error discussed in the Introduction. Recall that a Type A error occurs when treatment a is chosen but is actually inferior to

b, and a Type B error occurs when b is chosen but is inferior to a. For any allocation $\delta \in [0, 1]$, the quantities $R(\delta, b)$ and $R(\delta, a)$ give the potential welfare losses from Type A and B errors respectively. As δ increases from 0 to 1, the former potential loss decreases from $M(b)$ to 0 and the latter increases from 0 to $M(a)$. The MR criterion chooses δ to minimize the maximum potential loss, which occurs when $R(\delta, a) = R(\delta, b)$.⁵

When a planner uses allocation δ_{MR} , maximum regret is $M(a)M(b)/[M(a) + M(b)]$. It is interesting to compare this with the maximum regret that would result if the planner were only able to choose one of the singleton allocations. The solution would be $\delta = 0$ if $M(a) \geq M(b)$ and $\delta = 1$ if $M(a) \leq M(b)$. Maximum regret would be $\min[M(a), M(b)]$. Thus, permitting fractional allocations can reduce maximum regret to as little as one-half the value achievable with singleton allocations, this occurring when $M(a) = M(b)$.

Expressions $M(a)$ and $M(b)$ simplify when (α_L, β_U) and (α_U, β_L) are feasible values of (α, β) , as is so when the identification region is rectangular. Then $M(a) = \alpha_U - \beta_L$ and $M(b) = \beta_U - \alpha_L$. Hence,

$$(7) \quad \delta_{MR} = \frac{\beta_U - \alpha_L}{(\alpha_U - \beta_L) + (\beta_U - \alpha_L)}.$$

Result (7) simplifies further if either α or β is fully known. In particular, suppose that α is known. Then $\alpha_L = \alpha_U = \alpha$ and (7) becomes

$$(8) \quad \delta_{MR} = \frac{\beta_U - \alpha}{\beta_U - \beta_L}.$$

Full knowledge of α may be realistic if a is the status quo treatment and b is an innovation. Suppose,

⁵ Whereas the MR criterion balances the potential welfare losses from Type A and B errors, a planner may prefer to weigh these losses differently. Tetenov (2008) develops this idea, studying asymmetric MR criteria.

for example, that treatment a has been the standard therapy for a disease and treatment b is a proposed new therapy. Then one may be able to observe the outcomes experienced when earlier cohorts of patients were given treatment a, but no comparable data may be available for treatment b. Hence, the available empirical evidence may reveal α but not β .

The fractional character of the MR treatment allocation contrasts sharply with the generic singleton nature of the Bayesian allocation. Whereas the MR allocation minimizes maximum regret, a Bayesian allocation minimizes subjective expected regret.⁶ It is revealing to consider the special case where α is known. Bayesians sometime suggest that when a real quantity is known only to lie within some interval, a decision maker should assert a uniform distribution on the quantity and maximize expected welfare. Suppose that a planner places the uniform distribution $\mathbf{U}(\beta_L, \beta_U)$ on β and maximizes expected welfare. The subjective mean for β is $(\beta_L + \beta_U)/2$, so the Bayesian planner sets $\delta = 0$ if $(\beta_L + \beta_U)/2 < \alpha$ and $\delta = 1$ if $(\beta_L + \beta_U)/2 > \alpha$. In contrast, a MR planner sets $\delta_{MR} = (\beta_U - \alpha)/(\beta_U - \beta_L)$.

When (α_L, β_L) is a feasible value of (α, β) , there is a similarly sharp contrast between the fractional character of the MR allocation and the singleton nature of the maximin allocation. However, the maximin allocation may be fractional when (α_L, β_L) is not feasible. Indeed, the maximin and MR criteria are equivalent to one another when $\max(\alpha_s, \beta_s)$ is constant across states of nature.⁷

⁶ Bayesian decision making is usually described as maximization of expected welfare, but it is mathematically equivalent to minimization of expected regret. Consider an abstract setting where a decision maker faces choice set C , the feasible states of nature are S , and the objective is to maximize a welfare function $w(\cdot, \cdot): C \times S \rightarrow \mathbb{R}^1$ mapping actions $a \in C$ and states $s \in S$ into welfare. Let the decision maker assert a subjective distribution π on S .

The usual description of the Bayesian criterion is $\max_{a \in C} E_\pi[w(a, s)]$. The expected regret of action a is $E_\pi[\max_{c \in C} w(c, s) - w(a, s)] = E_\pi[\max_{c \in C} w(c, s)] - E_\pi[w(a, s)]$. The first term on the right-hand side does not vary with action a . Hence, minimization of expected regret is equivalent to maximization of expected welfare.

⁷ This is an instance of a general result. Consider again an abstract setting where a decision maker faces choice set C , the feasible states of nature are S , and the objective is to maximize a welfare function $w(\cdot, \cdot): C \times S \rightarrow \mathbb{R}^1$ mapping actions $a \in C$ and states $s \in S$ into welfare. Then the maximin decision criterion is $\max_{a \in C} \min_{s \in S} w(a, s)$ and the MR criterion is $\min_{a \in C} \max_{s \in S} [\max_{c \in C} w(c, s) - w(a, s)]$. Suppose that $\max_{c \in C} w(c, s)$ is constant for all $s \in S$. Then MR reduces to the maximin criterion.

I caution the reader that the MR allocation is not always fractional when a planner allocates the population among more than two treatments. Considering a setting with three treatments, Manski (2005b) gave an example in which all three treatments are undominated yet there exists a singleton MR allocation. Stoye (2007a) has studied a class of planning problems with multiple treatments and has found that the MR allocations are subtle to characterize. They often are fractional, but he gives an example in which there exists a unique singleton allocation. Throughout this paper I restrict attention to planning problems with two treatments.

2.4. Choosing Sentences for Convicted Juvenile Offenders

To illustrate Bayesian, maximin, and minimax-regret planning, consider the problem of choosing sentences for a population of convicted offenders. I apply findings reported in Manski and Nagin (1998), who studied the sentencing and recidivism of male youth in the state of Utah who were convicted of offenses before they reached age 16.

In this illustration, the planner is the state of Utah and the population are males under age 16 who are convicted of an offence. Treatment a is the status quo policy, this being a decentralized system where judges have discretion to choose between residential confinement and a sentence that does not involve confinement. Treatment b is an innovation mandating confinement for all convicted offenders. I take the outcome of interest to be a binary measure of recidivism. Specifically, $y(t) = 1$ if an offender who receives treatment t is not convicted of another crime in the two-year period following sentencing, and $y(t) = 0$ if the offender is convicted of a subsequent crime. Let $u(t) = y(t)$. Then $\alpha = P[y(a) = 1]$ and $\beta = P[y(b) = 1]$.

Analyzing data on outcomes under the status quo policy, Manski and Nagin (1998) find that $\alpha = 0.61$. The data do not fully identify β . In the absence of knowledge of how judges choose sentences or how juveniles respond to their sentences, the data reveal only that $\beta \in [0.03, 0.92]$. Thus, the innovation may be

much better or worse than the status quo. Manski and Nagin (1998) argue that this “worst-case” bound on β is germane to policy making because criminologists have found it difficult to learn how sentencing affects recidivism. Researchers have long debated the counterfactual outcomes that offenders would experience if they were to receive other sentences.

Consider policy choice when the state of Utah knows that $\alpha = 0.61$ and $\beta \in [0.03, 0.92]$. If the state applies the Bayesian paradigm, it fully adopts the innovation of mandatory confinement if $E_{\pi}(\beta) > 0.61$ and leaves the status quo of judicial discretion in place if $E_{\pi}(\beta) < 0.61$. If the state applies the maximin criterion, it leaves the status quo in place because $\beta_L = 0.03 < 0.61$. If the state applies the MR criterion, it randomly sentences to confinement $(\beta_U - \alpha)/(\beta_U - \beta_L) = (0.92 - 0.61)/(0.92 - 0.03) = 0.35$ of the offenders and leaves judicial discretion in place for the remaining fraction 0.65.

2.5. Planning with Observable Covariates

Sections 2.1 through 2.4 considered treatment of a population whose members are observationally identical to the planner. In practice, persons may have observable covariates and a planner may be able to differentially treat persons with different covariates. A simple approach to decision making is to separate persons by their covariates and apply the findings of Sections 2.2 and 2.3 to each group. This works when the relevant objective function is separable in the covariates but not otherwise. I explain here.

Suppose that the planner observes covariates $x_j \in X$ for each person j , where the covariate space X is finite and $P(x = \xi) > 0$, $\xi \in X$. Then the planner can distinguish among persons with different covariates, choosing a vector δ_{ξ} , $\xi \in X$ of treatment allocations that may vary with ξ . Welfare with this vector of allocations is

$$(9) \quad W(\delta_{\xi}, \xi \in X) = \sum_{\xi \in X} \alpha_{\xi} + (\beta_{\xi} - \alpha_{\xi})\delta_{\xi},$$

where $\alpha_\xi \equiv E[u(a)|x = \xi]$ and $\beta_\xi \equiv E[u(b)|x = \xi]$. The welfare function is separable in ξ . Hence, an allocation vector is optimal if, for each $\xi \in X$, it sets $\delta_\xi = 1$ when $\beta_\xi > \alpha_\xi$ and $\delta_\xi = 0$ when $\beta_\xi < \alpha_\xi$.

Now consider treatment choice under ambiguity. A Bayesian planner solves the problem

$$(10) \quad \max_{\delta_\xi \in [0, 1], \xi \in X} \sum_{\xi \in X} E_\pi(\alpha_\xi) + [E_\pi(\beta_\xi) - E_\pi(\alpha_\xi)]\delta_\xi,$$

where $E_\pi(\alpha_\xi) = \int \alpha_{\xi s} d\pi$ and $E_\pi(\beta_\xi) = \int \beta_{\xi s} d\pi$ are the subjective means of α_ξ and β_ξ . The objective function is separable in ξ . Hence, the Bayes decision assigns all persons with covariates ξ to treatment b if $E_\pi(\beta_\xi) > E_\pi(\alpha_\xi)$ and all such persons to treatment a if $E_\pi(\alpha_\xi) > E_\pi(\beta_\xi)$.

Maximin planning is more subtle because the objective function may not be separable in ξ . The optimization problem is

$$(11) \quad \max_{\delta_\xi \in [0, 1], \xi \in X} \min_{s \in S} \sum_{\xi \in X} \alpha_{\xi s} + (\beta_{\xi s} - \alpha_{\xi s})\delta_\xi.$$

This objective function generally is nonseparable, and determination of the maximin allocation requires joint choice of $\delta_\xi, \xi \in X$. However, the objective function is separable if $(\alpha_{\xi L}, \beta_{\xi L}), \xi \in X$ is a feasible value of $(\alpha_\xi, \beta_\xi), \xi \in X$. Then (11) reduces to

$$(12) \quad \max_{\delta_\xi \in [0, 1], \xi \in X} \sum_{\xi \in X} \alpha_{\xi L} + (\beta_{\xi L} - \alpha_{\xi L})\delta_\xi.$$

In this setting, the maximin decision assigns all persons with covariates ξ to treatment b if $\beta_{\xi L} > \alpha_{\xi L}$ and all such persons to treatment a if $\alpha_{\xi L} > \beta_{\xi L}$.

Minimax-regret planning is similarly subtle. The optimization problem is

$$(13) \quad \min_{\delta_\xi \in [0, 1], \xi \in X} \max_{s \in S} \sum_{\xi \in X} \max(\alpha_{\xi s}, \beta_{\xi s}) - [\alpha_{\xi s} + (\beta_{\xi s} - \alpha_{\xi s})\delta_\xi].$$

This objective function is generally nonseparable, and determination of the MR allocation requires joint choice of $\delta_\xi, \xi \in X$. However, the objective function is separable if $(\alpha_{\xi L}, \beta_{\xi U}), \xi \in X$ and $(\alpha_{\xi U}, \beta_{\xi L}), \xi \in X$ are feasible values of $(\alpha_\xi, \beta_\xi), \xi \in X$. Then (13) reduces to

$$(14) \quad \min_{\delta_\xi \in [0, 1], \xi \in X} \sum_{\xi \in X} \max [(\alpha_{\xi U} - \beta_{\xi L})\delta_\xi, (\beta_{\xi U} - \alpha_{\xi L})(1 - \delta_\xi)].$$

Hence, the MR allocation for persons with covariates ξ is

$$(15) \quad \delta_{\xi MR} = \frac{\beta_{\xi U} - \alpha_{\xi L}}{(\alpha_{\xi U} - \beta_{\xi L}) + (\beta_{\xi U} - \alpha_{\xi L})}.$$

An important open problem is to characterize the maximin and MR allocations when the set S is such that the objective function is nonseparable. Nonseparability generically occurs when the planner has information that relates the values of (α_ξ, β_ξ) across $\xi \in X$. For example, X could be an ordered set and the planner may know that β_ξ is monotone in ξ .

3. Nonlinear Welfare

This section studies planning when the welfare function is nonlinear in various ways. Section 3.1 considers monotonic transformations of the aggregate sum of outcomes. Section 3.2 addresses planning problems with non-additive cost of treatment, specifically capacity constraints or fixed costs. Section 3.3 studies planning with deontological welfare functions.

3.1. Monotone Transformations of the Welfare Function

Consider monotone transformations of welfare function (1) of the form

$$(16) \quad W(\delta) = f[\alpha + (\beta - \alpha)\delta],$$

where $f(\cdot)$ is strictly increasing in its argument. The shape of $f(\cdot)$ is immaterial to treatment choice when one treatment is superior in all states of nature. Whatever monotone function $f(\cdot)$ may be, $\delta = 0$ is optimal if $(\alpha_s \geq \beta_s, s \in S)$ and $\delta = 1$ if $(\alpha_s \leq \beta_s, s \in S)$. However, shape may matter when a planner faces ambiguity.

In Bayesian planning, the shape of $f(\cdot)$ expresses the planner's risk preferences, with linear $f(\cdot)$ implying indifference between mean-preserving spreads of a gamble and concave $f(\cdot)$ implying a preference for gambles with smaller spreads. However, the Bayesian definition of risk preferences does not carry over to other decision criteria. Hence, I do not associate the shape of $f(\cdot)$ with risk preferences here. Indeed, I show below that how $f(\cdot)$ matters, if at all, depends on the decision criterion that the planner uses.

Bayes Rules

A Bayesian planner with welfare function (16) solves the optimization problem

$$(17) \quad \max_{\delta \in [0, 1]} \int f[\alpha_s + (\beta_s - \alpha_s)\delta] d\pi.$$

The solution is generically singleton if $f(\cdot)$ is convex, but it may be fractional if $f(\cdot)$ has concave segments. Manski and Tetenov (2007, Proposition 5) consider the special case where the planner knows α and is uncertain only about β . We show that the Bayes allocation is $\delta = 0$ if $f(\cdot)$ is concave and $E_\pi(\beta) < \alpha$. It is fractional if $f(\cdot)$ is continuously differentiable, $E_\pi(\beta) > \alpha$, and $\int f(\beta) d\pi < f(\alpha)$. The allocation may be fractional

or the singleton $\delta = 1$ if $\int f(\beta)d\pi \geq f(\alpha)$.

The Maximin Criterion

The maximin problem

$$(18) \quad \max_{\delta \in [0, 1]} \min_{s \in S} f[\alpha_s + (\beta_s - \alpha_s)\delta]$$

has the same solution for all strictly increasing $f(\cdot)$. Thus, the shape of $f(\cdot)$ does not affect the allocation.

The Minimax-Regret Criterion

The shape of $f(\cdot)$ does affect the solution to the minimax-regret problem

$$(19) \quad \min_{\delta \in [0, 1]} \max_{s \in S} \max [f(\alpha_s), f(\beta_s)] - f[(1 - \delta)\alpha_s + \delta\beta_s].$$

Nevertheless, the central qualitative finding of Section 2.3 continues to hold with almost complete generality.

I show here that the MR allocation is fractional whenever $f(\cdot)$ is continuous.

Proof: Recall that $S(a) \equiv \{s \in S: \alpha_s > \beta_s\}$ and $S(b) \equiv \{s \in S: \beta_s > \alpha_s\}$. Let

$$(20a) \quad R(\delta, a) \equiv \max_{s \in S(a)} f(\alpha_s) - f[(1 - \delta)\alpha_s + \delta\beta_s],$$

$$(20b) \quad R(\delta, b) \equiv \max_{s \in S(b)} f(\beta_s) - f[(1 - \delta)\alpha_s + \delta\beta_s],$$

be the maximum regret of allocation δ on $S(a)$ and $S(b)$ respectively. The maximum regret of δ on all of S

is $\max [R(\delta, a), R(\delta, b)]$. As δ increases from 0 to 1, $R(\cdot, a)$ strictly increases from 0 to $R(1, a) > 0$ and $R(\cdot, b)$ strictly decreases from $R(0, b) > 0$ to 0. Continuity of $f(\cdot)$ and boundedness of $[(\alpha_s, \beta_s), s \in S]$ imply that $R(\cdot, a)$ and $R(\cdot, b)$ are continuous functions of δ . Hence, there exists a unique $\delta \in (0, 1)$ such that $R(\delta, a) = R(\delta, b)$. This is the MR allocation. \square

This proof, as did the proof to result (5), shows that the MR allocation balances the welfare losses from errors of Type A and B. Again, the quantities $R(\delta, b)$ and $R(\delta, a)$ give the potential welfare losses from Type A and B errors respectively. As δ increases from 0 to 1, the former potential loss decreases from $R(0, b)$ to 0 and the latter increases from 0 to $R(1, a)$. The MR criterion chooses δ to minimize the maximum potential loss, which occurs when $R(\delta, a) = R(\delta, b)$.

Logarithmic Welfare

Section 2.3 showed that the minimax-regret allocation has the simple form (7) when $f(\cdot)$ is linear and $\{\alpha_L, \beta_U\}, (\alpha_U, \beta_L)\}$ are feasible values of (α, β) . The MR allocation typically must be determined numerically when $f(\cdot)$ is nonlinear. However, a simple form emerges when $f(\cdot)$ is the log function and $\{\alpha_L, \beta_U\}, (\alpha_U, \beta_L)\}$ are feasible values of (α, β) . Then

$$(21a) \quad R(\delta, a) = \max_{s \in S(a)} \log\{\alpha_s / [(1 - \delta)\alpha_s + \delta\beta_s]\} = \log\{\alpha_U / [(1 - \delta)\alpha_U + \delta\beta_L]\},$$

$$(21b) \quad R(\delta, b) = \max_{s \in S(b)} \log\{\beta_s / [(1 - \delta)\alpha_s + \delta\beta_s]\} = \log\{\beta_U / [(1 - \delta)\alpha_L + \delta\beta_U]\}.$$

Hence, the MR allocation solves the equation

$$(22) \quad \alpha_U / [(1 - \delta)\alpha_U + \delta\beta_L] = \beta_U / [(1 - \delta)\alpha_L + \delta\beta_U].$$

The solution is

$$(23) \quad \delta_{MR} = \frac{\alpha_U(\beta_U - \alpha_L)}{\alpha_U(\beta_U - \alpha_L) + \beta_U(\alpha_U - \beta_L)} .$$

Comparison of (7) and (23) shows that the MR allocations under linear and logarithmic welfare coincide when $\beta_U = \alpha_U$, but they otherwise generally differ from one another. In particular, the two allocations differ when the planner knows α and has partial knowledge of β . Then (23) reduces to

$$(24) \quad \delta_{MR} = \frac{\alpha(\beta_U - \alpha)}{\alpha(\beta_U - \alpha) + \beta_U(\alpha - \beta_L)} .$$

By assumption $\beta_U > \alpha$. Hence, the fraction of the population allocated to treatment b when welfare is logarithmic is smaller than when welfare is linear. For example, in the sentencing illustration of Section 2.4, the MR allocation with logarithmic welfare is 0.26 rather than the 0.35 found with linear welfare.

Allocation of an Endowment Between a Safe and a Risky Asset

To illustrate the above findings, consider an investor with concave welfare function who must allocate an endowment between a safe and a risky asset. Let the safe asset be treatment a, with known return α . Let the risky asset be b, whose return is known to lie in the interval $[\beta_L, \beta_U]$, where $\beta_L < \alpha < \beta_U$. Thus, the safe and risky assets are analogous to a status quo treatment and an innovation.

As discussed above, a Bayesian investor allocates fully to the safe asset if $E_\pi(\beta) < \alpha$, diversifies his portfolio if $E_\pi(\beta) > \alpha$ and $\int f(\beta)d\pi < f(\alpha)$, and may either diversify or allocate fully to the risky asset if $\int f(\beta)d\pi \geq f(\alpha)$. A maximin investor allocates fully to the safe asset. A minimax-regret investor always diversifies,

the specific fractional allocation depending on the shape of $f(\cdot)$. See Brock and Manski (2008) for further analysis of this investment decision in the context of competitive lending.

3.2. Non-Additive Cost of Treatment

Treatment may be costly. The foregoing analysis covers settings where the aggregate cost of a treatment allocation is the sum of individual treatment costs. This was alluded to in Section 2.1, where I observed that $u_j(t)$ may have the benefit-cost form $u_j(t) = y_{j1}(t) - y_{j2}(t)$, where $y_{j1}(t)$ is the benefit when person j receives treatment t and $y_{j2}(t)$ is the cost. There are many ways in which cost might be non-additive. This section considers the polar cases of capacity constraints and fixed costs.

Capacity Constraints

I have thus far assumed that all treatment allocations $\delta \in [0, 1]$ are feasible. Capacity constraints may place an upper bound on the fraction of the population who receive each treatment. A capacity constraint is a cost that equals zero when the fraction of persons who receive a treatment is below the upper bound and infinity thereafter.

Let the maximum fractions of the population who may receive treatments a and b be $\eta(a)$ and $\eta(b)$ respectively. Then the feasible allocations are $\delta \in [1 - \eta(a), \eta(b)]$. A constrained Bayesian, maximin, or minimax-regret allocation solves the relevant extremum problem over the feasible δ . I focus on the MR allocation when welfare has form (16) and $f(\cdot)$ is continuous.

Let δ_{MR} denote the unconstrained MR allocation and δ_{CMR} the constrained MR allocation. Section 2.3 showed that maximum regret at any allocation δ equals $R(\delta, b)$ for $\delta \leq \delta_{MR}$ and $R(\delta, a)$ for $\delta \geq \delta_{MR}$, where $R(\cdot, b)$ is strictly decreasing in δ and $R(\cdot, a)$ is strictly increasing. It follows that the constrained MR allocation is the feasible allocation closest to the unconstrained MR allocation. That is,

$$(25) \quad \delta_{\text{CMR}} = 1 - \eta(a) \text{ if } \delta_{\text{MR}} < 1 - \eta(a),$$

$$\delta_{\text{MR}} \quad \text{if } \delta_{\text{MR}} \in [1 - \eta(a), \eta(b)],$$

$$\eta(b) \quad \text{if } \delta_{\text{MR}} > \eta(b).$$

Fixed Costs

A fixed cost is a cost component that equals zero when no one receives a treatment and takes a constant positive value when any positive fraction of the population receives the treatment. Fixed costs make singleton allocations advantageous relative to fractional ones. Suppose that treatments a and b have non-negative fixed costs $C(a)$ and $C(b)$ respectively. Then allocations $\delta = 0$ and $\delta = 1$ have fixed costs $C(a)$ and $C(b)$, but any $\delta \in (0, 1)$ bears the larger fixed cost $C(a) + C(b)$. I show here that the MR allocation is fractional if the fixed costs are small but is singleton if they are sufficiently large.

For simplicity, I suppose that the welfare function is linear and that the fixed costs have known values that do not vary with the state of nature. Thus, the welfare function is

$$(26) \quad W(\delta) = \alpha + (\beta - \alpha)\delta - C(a) \cdot 1[\delta < 1] - C(b) \cdot 1[\delta > 0].$$

Allocation $\delta = 0$ is optimal if $\alpha - C(a) \geq \beta - C(b)$ and $\delta = 1$ is optimal if $\alpha - C(a) \leq \beta - C(b)$.

The problem of interest is treatment choice under ambiguity. Let $S(a)$ and $S(b)$ be the subsets of S on which treatments a and b are superior. That is, $S(a) = \{s \in S: \alpha_s - C(a) > \beta_s - C(b)\}$ and $S(b) = \{s \in S: \beta_s - C(b) > \alpha_s - C(a)\}$. The planner faces ambiguity if both $S(a)$ and $S(b)$ are non-empty.

As earlier, let $M(a) \equiv \max_{s \in S(a)} (\alpha_s - \beta_s)$ and $M(b) \equiv \max_{s \in S(b)} (\beta_s - \alpha_s)$. Recall from (5) that the MR allocation in the absence of fixed costs is $\delta_{\text{MR}} = M(b)/[M(a) + M(b)]$. Let δ_{FMR} denote the MR allocation in the presence of fixed costs. The result is

(27)

$$\delta_{\text{FMR}} = 0 \quad \text{if} \quad M(b) + C(a) - C(b) \leq \min \{M(a) - C(a) + C(b), \delta_{\text{MR}}M(a) + C(a)[1 - \delta_{\text{MR}}] + C(b)\delta_{\text{MR}}\},$$

$$\delta_{\text{FMR}} = \delta_{\text{MR}} + \frac{C(a) - C(b)}{M(a) + M(b)}$$

$$\text{if} \quad \delta_{\text{MR}}M(a) + C(a)[1 - \delta_{\text{MR}}] + C(b)\delta_{\text{MR}} \leq \min \{M(a) - C(a) + C(b), M(b) + C(a) - C(b)\},$$

$$\delta_{\text{FMR}} = 1 \quad \text{if} \quad M(a) - C(a) + C(b) \leq \min \{M(b) + C(a) - C(b), \delta_{\text{MR}}M(a) + C(a)[1 - \delta_{\text{MR}}] + C(b)\delta_{\text{MR}}\}.$$

Proof: For any $\delta \in [0, 1]$, the maximum regret of δ on all of S is $\max [R(\delta, a), R(\delta, b)]$, where

$$\begin{aligned} (28a) \quad R(\delta, a) &\equiv \max_{s \in S(a)} \alpha_s - C(a) - \{\alpha_s + \delta(\beta_s - \alpha_s) - C(a) \cdot 1[\delta < 1] - C(b) \cdot 1[\delta > 0]\} \\ &= \delta M(a) - C(a) \cdot 1[\delta = 1] + C(b) \cdot 1[\delta > 0] \end{aligned}$$

$$\begin{aligned} (28b) \quad R(\delta, b) &\equiv \max_{s \in S(b)} \beta_s - C(b) - \{\alpha_s + \delta(\beta_s - \alpha_s) - C(a) \cdot 1[\delta < 1] - C(b) \cdot 1[\delta > 0]\} \\ &= (1 - \delta)M(b) + C(a) \cdot 1[\delta < 1] - C(b) \cdot 1[\delta = 0]. \end{aligned}$$

are maximum regret on $S(a)$ and $S(b)$.

Application of (28) at $\delta = 0$ and $\delta = 1$ gives the maximum regret values

$$\max [R(0, a), R(0, b)] = M(b) + C(a) - C(b),$$

$$\max [R(1, a), R(1, b)] = M(a) - C(a) + C(b).$$

Application of (28) at $\delta \in (0, 1)$ gives

$$\max [R(\delta, a), R(\delta, b)] = \max [\delta M(a) + C(b), (1 - \delta)M(b) + C(a)].$$

The minimum of maximum regret over $\delta \in (0, 1)$ solves the equation

$$\delta M(a) + C(b) = (1 - \delta)M(b) + C(a).$$

Hence, the minimax regret allocation on $\delta \in (0, 1)$ is

$$\frac{M(b) + C(a) - C(b)}{M(a) + M(b)} = \delta_{MR} + \frac{C(a) - C(b)}{M(a) + M(b)}$$

and the minimax regret value on $\delta \in (0, 1)$ is $\delta_{MR}M(a) + C(a)[1 - \delta_{MR}] + C(b)\delta_{MR}$. The final step is to minimize maximum regret over $\delta \in (0, 1)$, $\delta = 0$, and $\delta = 1$. This yields (27). \square

Result (27) simplifies if use of each treatment incurs equal fixed cost. Let $C \equiv C(a) = C(b)$. Then (27) reduces to

$$(29) \quad \begin{aligned} \delta_{FMR} &= 0 && \text{if } M(b) \leq \min \{M(a), \delta_{MR}M(a) + C\}, \\ \delta_{FMR} &= \delta_{MR} && \text{if } \delta_{MR}M(a) + C \leq \min \{M(a), M(b)\}, \\ \delta_{FMR} &= 1 && \text{if } M(a) \leq \min \{M(b), \delta_{MR}M(a) + C\}. \end{aligned}$$

Thus, a common fixed cost smaller than $\min \{M(a), M(b)\} - \delta_{MR}M(a)$ has no effect on the minimax-regret allocation. However, a larger fixed cost makes the allocation singleton.

3.3. Deontological Welfare Functions

I have thus far maintained the traditional consequentialist assumption of welfare economics. That is, policy choices matter only for the outcomes they yield. This section considers welfare functions that embrace deontological considerations. Deontological ethics supposes that actions may have intrinsic value, apart from their consequences. I first give a brief abstract introduction and then focus on the specific idea of “equal treatment of equals.”

In Section 3.2, the fixed costs $C(a)$ and $C(b)$ made the treatment allocation affect welfare directly,

regardless of the resulting outcomes. Although I then described $C(a)$ and $C(b)$ in ordinary economic language as fixed costs, welfare function (26) can be interpreted as expressing the deontological idea that any use of treatment a or b is normatively bad per se, with $C(a)$ and $C(b)$ expressing the respective welfare losses. More generally, a planner might use a welfare function of the form

$$(30) \quad W(\delta) = f[\alpha + (\beta - \alpha)\delta + g(\delta)],$$

where $g(\cdot)$ is a planner-specified function of δ .

Welfare function (30) expresses consequentialism through its $\alpha + (\beta - \alpha)\delta$ component and deontological ethics through its $g(\delta)$ component. A planner using such a welfare function trades off consequentialist and deontological considerations, choosing a deontologically inferior allocation if it yields sufficiently superior outcomes, and vice versa. Permitting trade offs among the attributes of an action is almost universally accepted by economists. However, philosophical discussions generally take the position that deontological considerations should supercede consequentialist ones. This suggests a lexicographic decision process in which one first restricts attention to actions that are deontologically acceptable and only then considers the consequences of these actions.

Equal Treatment of Equals

When considering fractional treatment allocations, a particularly salient deontological idea is the normative principle calling for equal treatment of equals. Fractional allocations are consistent with this principle in the sense that observationally identical persons have equal probabilities of receiving particular treatments. They are inconsistent with the principle in the sense that observationally identical persons do not actually receive the same treatment. Thus, equal treatment holds ex ante but not ex post.

If one is concerned only with the ex ante sense of equal treatment, then all values of δ are

deontologically equivalent. In terms of welfare function (30), the function $g(\cdot)$ is constant. If one is concerned with the ex post sense of equal treatment, singleton allocations are advantageous relative to fractional ones. In terms of (30), $g(0) = g(1) > g(\delta)$ for $\delta \in (0, 1)$. The equal fixed-cost case considered at the end of Section 3.2 has the form $g(0) = g(1) = -C$ and $g(\delta) = -2C$ for $\delta \in (0, 1)$. Thus, placing value C on the deontological consideration of equal ex post treatment does not affect the minimax-regret allocation if $C < \min \{M(a), M(b)\} - \delta_{MR}M(a)$. However, it makes the minimax-regret allocation singleton if C is larger.

A dramatic illustration of the difference between the ex ante and ex post senses of equal treatment occurs in this hypothetical problem of treatment choice considered in Manski (2007b, Section 11.7).

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_i 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 nature, $s = H_a$ and $s = H_b$. Let welfare be the survival rate of the population. If a fraction δ of the population receives treatment b and the remaining $1 - \delta$ receives treatment a , the fraction who survive is $(1 - \delta) \cdot 1[s = H_a] + \delta \cdot 1[s = H_b]$.

The singleton allocations $\delta = 0$ and $\delta = 1$ provide equal treatment in both the ex ante and ex post senses. These allocations also equalize realized outcomes—the entire population either survives or dies. The minimax-regret allocation is $\delta = 1/2$. Everyone is treated equally ex ante, each person having a 50 percent chance of receiving each treatment, but not ex post. Nor are outcomes equalized—half the population lives and half dies.⁸ \square

⁸ The maximin allocation also is $\delta = 1/2$, as the maximin and MR criteria are equivalent in this illustration. In contrast, a Bayesian planner would allocate the entire population to the treatment with the higher subjective probability of success.

Democratic societies ordinarily adhere to the ex post sense of equal treatment. Americans who have the same income, deductions, and exemptions are required to pay the same federal income tax. The Equal Protection clause in the 14th Amendment to the U. S. Constitution is held to mean that all persons in a jurisdiction are subject to the same laws, not that all persons have the same chance of being subject to different laws.

Nevertheless, some important policies adhere to the ex ante sense of equal treatment but explicitly violate the ex post sense. American examples include random tax audits, drug testing and airport screening, random calls for jury service, and the Green Card and Vietnam draft lotteries. These policies have not been prompted by the desire to cope with and reduce ambiguity that motivates treatment diversification. Yet they do indicate some willingness of society to accept ex post unequal treatment.

Reduction of ambiguity is the explicit objective of randomized clinical trials in medicine and other randomized social experiments. Combining ex ante equal treatment with ex post unequal treatment is precisely what makes randomized experiments useful in learning about treatment response. Modern medical ethics permits randomized trials only under conditions of *clinical equipoise*; that is, when partial knowledge of treatment response prevents a determination that one treatment is superior to another. Clinical equipoise is essentially a synonym for ambiguity.

4. Interacting Treatments

This section relaxes the assumption that treatment is individualistic and considers settings where treatments interact, each person's outcome depending both on his own treatment and on the treatment allocation within the population. A good illustration is medical treatment of an infectious disease.

Let there be two treatments: (a) therapy after infection and (b) vaccination before infection. Suppose

that vaccination of a person always prevents his infection and that the infection rate of unvaccinated persons decreases with the fraction of the population who are vaccinated. However, existing epidemiological science yields only partial knowledge of the relationship between the fraction vaccinated and the infection rate of unvaccinated persons. Then determination of an optimal treatment rule may not be possible.

Section 4.1 sets up general concepts and notation. Section 4.2 considers treatment of an infectious disease. See Manski (2009a) for further discussion.

4.1. Concepts and Notation

It is easy to extend the concepts and notation of Section 2 and 3 to planning problems where outcomes depend on a person's own treatment and on the population treatment allocation. This only requires expansion of the domains of the response function $y_j(\cdot)$ and welfare function $u_j(\cdot)$.

Suppose now that person j has a response function $y_j(\cdot, \cdot): T \times [0, 1] \rightarrow Y$ that maps own treatments t and population allocations δ into outcomes $y_j(t, \delta) \in Y$. Let $u_j(t, \delta) \equiv u_j[y_j(t, \delta), t, \delta]$ denote the net contribution to welfare. Let $\alpha(\delta) \equiv E[u(a, \delta)]$ and $\beta(\delta) \equiv E[u(b, \delta)]$ be mean welfare when a randomly drawn person receives treatment a or b and the population allocation is δ . Then the welfare function is

$$(31) \quad W(\delta) = \alpha(\delta) + [\beta(\delta) - \alpha(\delta)]\delta.$$

This extension of equation (1) remains a consequentialist welfare function that aggregates individual contributions to welfare. However, in contrast to (1), $W(\delta)$ now is not linear in δ . Whereas the optimal allocation under (1) is generically singleton, the optimal allocation under (31) may be singleton or fractional, depending on how $\alpha(\cdot)$ and $\beta(\cdot)$ vary with δ .

Consider planning under ambiguity. Whereas the set S of states of nature previously indexed the

feasible pairs of values (α, β) , it now indexes the feasible pairs of functions $[\alpha(\cdot), \beta(\cdot)]$. For $s \in S$, define

$$(32) \quad W_s(\delta) \equiv \alpha_s(\delta) + [\beta_s(\delta) - \alpha_s(\delta)]\delta.$$

Partial knowledge is unproblematic for decision making if there exists a value of δ such that $W_s(\delta) \geq W_s(\delta')$ for all $\delta' \in [0, 1]$ and all $s \in S$. If so, this allocation is optimal. The planner faces ambiguity if there exists no δ that is superior in all states of nature.

As previously, a planner may cope with ambiguity by computing a Bayes decision or by applying the maximin or minimax-regret criterion. The Bayes decision with subjective distribution π solves the problem

$$(33) \quad \max_{\delta \in [0, 1]} E_\pi[\alpha(\delta)] + \{E_\pi[\beta(\delta)] - E_\pi[\alpha(\delta)]\}\delta,$$

where $E_\pi[\alpha(\delta)] = \int \alpha_s(\delta)d\pi$ and $E_\pi[\beta(\delta)] = \int \beta_s(\delta)d\pi$ are the subjective means of $\alpha(\delta)$ and $\beta(\delta)$. Thus, a Bayesian planner behaves as does a planner who knows that the population means appearing in (32) have the values in (33).

The maximin criterion is

$$(34) \quad \max_{\delta \in [0, 1]} \min_{s \in S} \alpha_s(\delta) + [\beta_s(\delta) - \alpha_s(\delta)]\delta.$$

For each $\delta \in [0, 1]$, let $\alpha_L(\delta) \equiv \inf_{s \in S} \alpha_s(\delta)$ and $\beta_L(\delta) \equiv \inf_{s \in S} \beta_s(\delta)$. Criterion (34) simplifies if $[\alpha_L(\delta), \beta_L(\delta)]$ is a feasible value of $[\alpha(\delta), \beta(\delta)]$ for each $\delta \in [0, 1]$. Then the maximin allocation solves the problem

$$(35) \quad \max_{\delta \in [0, 1]} \alpha_L(\delta) + [\beta_L(\delta) - \alpha_L(\delta)]\delta.$$

Thus, the maximin planner behaves as would a planner who knows that the population means in (32) have the values in (35).

The minimax-regret criterion is

$$(36) \quad \min_{\delta \in [0, 1]} \max_{s \in S} \left\{ \max_{d \in [0, 1]} \{ \alpha_s(d) + [\beta_s(d) - \alpha_s(d)]d \} - \{ \alpha_s(\delta) + [\beta_s(\delta) - \alpha_s(\delta)]\delta \} \right\}.$$

Problem (36) does not simplify in general, so I leave it in this form for now. Section 4.2 analyzes a tractable special case.

Severity of Ambiguity

Ambiguity with interacting treatments is more severe than with individualistic treatment. Optimization previously required knowing only whether the number α is larger or smaller than the number β . It now requires knowledge of the functions $\alpha(\cdot)$ and $\beta(\cdot)$.

Consider in particular the identification problem stemming from the unobservability of counterfactual outcomes. Suppose that one has data on the treatments received and outcomes experienced in a study population. When treatments are individualistic, each member of the study population has only one counterfactual outcome. If person j received treatment a , then we observe outcome $y_j(a)$ but not $y_j(b)$, and vice versa. When treatments interact, each member of the study population has many counterfactual outcomes. If person j received treatment a and the population allocation was δ , then we observe $y_j(a, \delta)$ but we do not observe $[y_j(a, d), d \neq \delta]$ or $\{y_j(b, d), d \in [0, 1]\}$. Thus, data on realized treatments and outcomes yield only a small part of the knowledge needed to solve the optimization problem.

4.2. Choosing Treatments for an Infectious Disease

To illustrate the abstract ideas developed above, I consider treatment of an infectious disease. I assume that vaccination of a person always prevents his infection and that the infection rate of unvaccinated persons decreases with the fraction of the population who are vaccinated. The illustration supposes a single-period planning problem. Multi-period planning would require attention to learning and also to the possibility that past treatment allocations may affect mutation of the pathogen.

Formally, let treatment a be therapy after infection and let treatment b be vaccination. Let the planner know that $\alpha(\delta)$ weakly increases with δ and that $\beta(\delta)$ does not vary with δ . Then the optimal value of δ solves the problem

$$(37) \quad \max_{\delta \in [0, 1]} \alpha(\delta)(1 - \delta) + \beta\delta,$$

where β denotes the constant value of $\beta(\cdot)$.

The optimization problem is trivial if the vaccine has no side effects and no cost. Then $\beta \geq \alpha(1)$ and it is optimal to vaccinate everyone. The optimal allocation is less obvious when $\beta < \alpha(1)$. To obtain a concrete sense of the planning problem, I first solve (37) when β has a known value and $\alpha(\cdot)$ is a known linear function. I then consider planning under ambiguity.

Linear Treatment Interactions

Suppose it is known that $\alpha(\cdot)$ has the linear form $\alpha(\delta) = \alpha(0)(1 - \delta) + \alpha(1)\delta$. Then (37) has the quadratic form

$$(38) \quad \max_{\delta \in [0, 1]} \alpha(0)(1 - \delta)^2 + \alpha(1)\delta(1 - \delta) + \beta\delta.$$

Suppose that the planner knows the values of $[\alpha(0), \alpha(1), \beta]$. Then he can determine the optimal treatment allocation.

The first-order-condition derived from (38) is

$$(39) \quad -2\alpha(0) + 2\alpha(0)\delta + \alpha(1) - 2\alpha(1)\delta + \beta = 0.$$

The second derivative is $2[\alpha(0) - \alpha(1)]$. By assumption, $\alpha(\cdot)$ is weakly increasing. Hence, the second derivative is non-positive.

Suppose that $\alpha(1) = \alpha(0)$. Then treatment is individualistic. The optimal allocation is $\delta = 1$ if $\beta \geq \alpha$ and $\delta = 0$ if $\beta \leq \alpha$, where α is the constant value of $\alpha(\cdot)$.

Suppose that $\alpha(1) > \alpha(0)$. Then vaccination has a positive external effect. The first-order-condition gives the maximum of the quadratic welfare function at

$$(40) \quad \delta^* = \frac{1}{2} + \frac{1}{2}[\beta - \alpha(0)]/[\alpha(1) - \alpha(0)].$$

The optimal allocation is δ^* if $\delta^* \in [0, 1]$. It is optimal to vaccinate no one if $\delta^* < 0$ and everyone if $\delta^* > 1$.

Observe that the fraction vaccinated is greater than one-half when $\beta > \alpha(0)$ and equals one when $\beta \geq \alpha(1)$. The fraction vaccinated is positive but less than one-half when $2\alpha(0) - \alpha(1) < \beta < \alpha(0)$. It is optimal to vaccinate no one when $\beta \leq 2\alpha(0) - \alpha(1)$.

Treatment Under Ambiguity

For simplicity, I will suppose that the value of β is known, perhaps from performance of a randomized clinical trial with the vaccine. Then the only ambiguity is partial knowledge of the function $\alpha(\cdot)$

that describes how the outcomes of unvaccinated persons vary with the population vaccination rate.

The Bayes decision with subjective distribution π solves the problem

$$(41) \quad \max_{\delta \in [0, 1]} E_{\pi}[\alpha(\delta)](1 - \delta) + \beta\delta.$$

The maximin criterion is

$$(42) \quad \max_{\delta \in [0, 1]} \alpha_L(\delta)(1 - \delta) + \beta\delta.$$

Thus, a Bayesian or maximin planner behaves as would one who knows that $\alpha(\cdot) = E_{\pi}[\alpha(\cdot)]$ or that $\alpha(\cdot) = \alpha_L(\cdot)$ respectively.

The minimax-regret criterion is

$$(43) \quad \min_{\delta \in [0, 1]} \max_{s \in S} \left\{ \max_{d \in [0, 1]} [\alpha_s(d)(1 - d) + \beta d] - [\alpha_s(\delta)(1 - \delta) + \beta\delta] \right\}.$$

Although this problem is complex in abstraction, it turns out to have a remarkably simple solution in an important special case. Suppose the planner knows that $\alpha(\cdot)$ is bounded and weakly increasing, the known bounds being real numbers L and U such that $L \leq \alpha(0) \leq \alpha(1) \leq U$. When this is the only information available about $\alpha(\cdot)$, the MR allocation is (see the Appendix for the proof)

$$(44) \quad \delta_{MR} = 1 \quad \text{if } \beta > U,$$

$$= \frac{U - L}{(U - L) + (U - \beta)} \quad \text{otherwise.}$$

Observe that δ_{MR} is always positive, no matter how small β may be. For fixed (L, U) , the fraction vaccinated increases continuously from 0 to $\frac{1}{2}$ as β increases from $-\infty$ to L and then increases further to 1 when $\beta = U$. The behavior of the minimax-regret allocation as a function of β differs sharply from that of the maximin allocation, which is zero if $\beta \leq L$ and one if $\beta \geq U$.

The minimax-regret allocation computed here differs from the minimax-regret allocation computed under the assumption that treatment is individualistic. Equation (7) gave the allocation under that assumption, with $\alpha_U = U$, $\alpha_L = L$, and $\beta_U = \beta_L = \beta$. Thus, the MR allocation assuming individualistic treatment equals 0 if $\beta \leq L$, $(\beta - L)/(U - L)$ if $L < \beta < U$, and 1 if $\beta \geq U$.

5. Multi-Period Planning Problems

Whereas Sections 2 through 4 assumed a one-period planning problem, I now consider multi-period planning under ambiguity. I suppose that, in each period, a planner must choose treatments for the current cohort of a population. The planner wants to maximize the welfare of each cohort. As in Sections 2 and 3, treatment is individualistic.

The essential new feature of multi-period problems is that learning is possible, with observation of the outcomes experienced by earlier cohorts informing treatment choice for later cohorts. Fractional treatment allocations are advantageous for learning because they generate randomized experiments yielding outcome data on both treatments. Sampling variation is not an issue when cohorts are large, so all fractional allocations yield the same information. Hence, the choice among fractional allocations may be based on other grounds.⁹

⁹ Different fractional allocations are not equivalent from the perspective of learning when treatments interact as in Section 4. Then observation of outcomes under allocation δ is only informative about $\alpha(\delta)$ and $\beta(\delta)$.

I suggest use of the *adaptive minimax-regret (AMR)* criterion. In each period, the AMR criterion applies the static minimax-regret criterion, using the information available at the time. In the absence of large fixed costs or deontological considerations, the result is a fractional allocation whenever both treatments are undominated. The AMR criterion is adaptive because successive cohorts may receive different allocations as knowledge of treatment response accumulates over time.

Section 5.1 formalizes the AMR criterion. Section 5.2 gives a numerical illustration showing how a centralized health planning system could use the criterion to choose treatments for a non-infectious disease. Section 5.3 discusses how the AMR criterion differs from the current medical practice of randomized clinical trials. Section 5.4 suggests a modification to the drug approval process of the Food and Drug Administration.

5.1. The Adaptive Minimax-Regret Criterion

To frame the multi-period planning problem we need to extend the concepts and notation used earlier. Let $n = 0, 1, \dots, N$ denote the periods in which treatment allocations must be chosen. Let $P_n[y(\cdot)]$ denote the distribution of treatment response across cohort n . I assume that all cohorts are large and have the same distribution of treatment response. Thus, $P_n[y(\cdot)] = P[y(\cdot)]$ for all n , where $P[y(\cdot)]$ is a time-invariant distribution.

The assumption of a time-invariant outcome distribution enables learning. Observation of the outcomes experienced by earlier cohorts yields information about $P[y(\cdot)]$ that can inform treatment choice for later cohorts. Formally, learning occurs when observation of outcomes enables the planner to shrink the set of feasible states of nature.

In each period, the set of feasible treatments is $T = \{a, b\}$. The planner's problem is to allocate each cohort between the two treatments. A treatment allocation is a vector $\delta \equiv (\delta_n, n = 0, \dots, N)$ that randomly assigns a fraction δ_n of cohort n to treatment b and the remaining $1 - \delta_n$ to treatment a . The optimal

allocation in each period is $\delta_n = 1$ if $\beta \geq \alpha$ and $\delta_n = 0$ if $\beta \leq \alpha$.

Let S_n index the feasible states of nature in period n . The planner chooses an allocation δ_n with knowledge of $(\delta_{n'}, n' < n)$ and $(S_{n'}, n' \leq n)$, but without knowledge of the information $(S_{n'}, n' > n)$ that he will possess later on. It is conceptually subtle and computationally daunting to approach choice of δ_n in a forward-looking manner, considering all logically possible future sequences of information sets and choices. It is much simpler to proceed myopically, choosing δ_n as if n is the sole period of a static choice problem.

The AMR criterion provides an appealing myopic decision rule. The criterion in period n is

$$(45) \quad \min_{\delta_n \in [0, 1]} \max_{s \in S_n} [f(\alpha_s), f(\beta_s)] - f[(1 - \delta_n)\alpha_s + \delta_n\beta_s].$$

When $f(\cdot)$ is linear, the AMR allocation follows immediately from (5). Let $S_n(a) \equiv \{s \in S_n : \alpha_s > \beta_s\}$ and $S_n(b) \equiv \{s \in S_n : \beta_s > \alpha_s\}$. Let $M_n(a) \equiv \max_{s \in S_n(a)} (\alpha_s - \beta_s)$ and $M_n(b) \equiv \max_{s \in S_n(b)} (\beta_s - \alpha_s)$. Then

$$(46) \quad \delta_{n\text{AMR}} = \frac{M_n(b)}{M_n(a) + M_n(b)}.$$

All of the other findings in Sections 2 and 3 extend in the same way.

The AMR criterion has practical and normative appeal. The practical appeal is its simplicity. The static minimax-regret allocation has a particularly transparent form when welfare is linear. The AMR allocation (46) inherits this transparency.

The normative appeal is that the AMR criterion treats each cohort as well as possible, in the minimax-regret sense, given the available knowledge. It does not ask the members of one cohort to sacrifice its own welfare for the benefit of future cohorts. Nevertheless, the AMR criterion is informationally beneficial to future cohorts in the broad class of settings where it yields a fractional treatment allocation

under ambiguity. Unless large fixed costs or deontological considerations make the AMR allocation singleton, application of the criterion maximizes cross-cohort learning about treatment response.

5.2 Application of the AMR rule in Centralized Health Care Systems

Close approximations to the AMR rule could be implemented in centralized health care systems where government agencies directly assign treatments. Examples include the Veterans Health Administration (VA) in the United States and the National Health Service (NHS) in England. Implementation could also occur in employer-based and other private systems where health maintenance organizations (HMOs) directly provide medical care. I say that the VA, NHS, and HMOs could only implement approximate versions of the AMR rule because these planners do not have fixed patient populations as assumed here. Patients who are unhappy with the care provided by the VA, NHS, or HMOs can opt-out and seek medical care elsewhere. However, strong financial disincentives typically limit opting-out to the relatively wealthy.

I present a hypothetical numerical illustration concerning treatment of a life-threatening disease. The planner faces ambiguity because the outcome of interest unfolds over multiple periods following receipt of treatment. As empirical evidence accumulates, the AMR treatment allocation changes accordingly.

In the illustration, a is a status quo treatment whose outcome distribution is known from historical experience, and b is an innovation with initially unknown outcome distribution. The adaptive minimax-regret rule applies the MR rule to each successive cohort, using the knowledge of β available at the time. Thus, the AMR decision at each n is

$$(47) \quad \delta_{\text{AMR}(n)} = \begin{cases} 0 & \text{if } \beta_{U_n} < \alpha, \\ (\beta_{U_n} - \alpha) / (\beta_{U_n} - \beta_{L_n}) & \text{if } \beta_{L_n} \leq \alpha \leq \beta_{U_n}, \end{cases}$$

$$= 1 \quad \text{if } \beta_{L_n} > \alpha.$$

The minimax value of regret for cohort n is $\max[0, (\beta_{U_n} - \alpha)(\alpha - \beta_{L_n})/(\beta_{U_n} - \beta_{L_n})]$ and mean welfare is $\alpha + (\beta - \alpha)\delta_{AMR(n)}$. Minimax regret is computed using the information available at the time of treatment. Mean welfare depends on the value of β , which is not known.

Treating a Life-Threatening Disease

Consider treatment of a life-threatening non-infectious disease. I take the outcome of interest to be the number of years that a patient survives within a specified time horizon following treatment. For this illustration, the horizon is five years and $y(t)$ is the number of years that a patient lives during the five years following receipt of treatment t . Thus, $y(t)$ has the time-additive form

$$(48) \quad y_j(t) = \sum_{k=1}^K y_{jk}(t),$$

where $y_{jk}(t) = 1$ if patient j is alive k years after treatment, $y_{jk}(t) = 0$ otherwise, and $K = 5$.

If patient j receives treatment b , outcome $y_j(b)$ gradually becomes observable as time passes. At the time of treatment, $y_j(b)$ can take any of the values $[0, 1, 2, 3, 4, 5]$. A year later, one can observe whether j is still alive and hence can determine whether $y_j(b) = 0$ or $y_j(b) \geq 1$. And so on until year five, when the outcome is fully observable.

Table 1 presents hypothetical data on annual death rates following treatment by the status quo and the innovation. The entries show that 20 (10) percent of the patients who receive the status quo (innovation) die within the first year after treatment. In each of the later years, the death rates are 5 and 2 percent respectively. Overall, the mean numbers of years lived after treatment are $\alpha = 3.5$ and $\beta = 4.3$. The former value is known at the outset from historical experience. The latter gradually becomes observable.

Table 1: Treating a Life-Threatening Disease						
cohort or year (n or k)	death rate in k th year after treatment		bound on β for cohort n	AMR allocation for cohort n	minimax value of regret for cohort n	mean life span for cohort n
	Status Quo	Innovation				
0			[0, 5]	0.30	1.05	3.74
1	0.20	0.10	[0.90, 4.50]	0.28	0.72	3.72
2	0.05	0.02	[1.78, 4.42]	0.35	0.60	3.78
3	0.05	0.02	[2.64, 4.36]	0.50	0.43	3.90
4	0.05	0.02	[3.48, 4.32]	0.98	0.02	4.28
5	0.05	0.02	[4.30, 4.30]	1	0	4.30

Assume that the planner measures welfare by a patient's length of life; thus, $u(t) = y(t)$. Also assume that the planner has no initial knowledge of β . That is, he does not know whether the innovation will be disastrous, with all patients dying in the first year following treatment, or entirely successful, with all patients living five years or more. Then the initial bound on β is $[\beta_{L0}, \beta_{U0}] = [0, 5]$. Hence, the initial AMR treatment allocation is $\delta_0 = 0.30$.

In year 1 the planner observes that, of the patients in cohort 0 assigned to the innovation, 10 percent died in the first year following treatment. This enables him to deduce that $P[y(b) \geq 1] = 0.90$. The planner uses this information to tighten the bound on β to $[\beta_{L1}, \beta_{U1}] = [0.90, 4.50]$. It follows that $\delta_1 = 0.28$. In each subsequent year the planner observes another annual death rate, tightens the bound on β , and computes the treatment allocation accordingly. The result is that $\delta_2 = 0.35$, $\delta_3 = 0.50$, and $\delta_4 = 0.98$. In year 5 he knows that the innovation is better than the status quo, and so sets $\delta_5 = 1$.

5.3. The AMR Criterion and the Practice of Randomized Clinical Trials

The illustration of Section 5.2 exemplifies a host of settings in which a medical planner must choose between a well-understood status quo treatment and an innovation whose properties are only partially known. When facing situations of this kind, it has been common to commission randomized clinical trials (RCTs) to learn about the innovation. The fractional allocations produced by the AMR criterion are randomized experiments, so it is natural to ask how application of the AMR criterion differs from the current practice of RCTs. There are many major differences, described below.

Fraction of the Population Receiving the Innovation: The AMR treatment allocation δ_{AMR} can take any value in the interval $[0, 1]$. In contrast, the sample receiving the innovation in current RCTs is typically a very small fraction of the relevant population, with sample size determined by conventional calculations of statistical power. For example, in trials conducted to obtain U. S. Food and Drug Administration (FDA) approval of new drugs, the sample receiving the innovation typically comprises two to three thousand persons, whereas the relevant patient population may contain hundreds of thousands or millions of persons. Thus, the value of δ in an RCT is generally less than 0.01 and often less than 0.001.

Group Subject to Randomization: Under the AMR criterion, the persons receiving the innovation are randomly drawn from the full patient population. In contrast, present clinical trials randomly draw subjects from pools of persons who volunteer to participate. Hence, a trial at most reveals the distribution of treatment response within the sub-population of volunteers, not within the full patient population.

Measurement of Outcomes: Under the AMR criterion, one observes the health outcomes of real interest as they unfold over time and one uses these data to inform subsequent treatment decisions. In contrast, RCTs

performed to obtain FDA approval of new drugs typically have durations of only two to three years. For example, a three-year trial on the disease described in Table 1 would only reveal that $\beta \in [2.64, 4.36]$. Attempting to learn from trials of short duration, researchers often measure surrogate outcomes rather than health outcomes of real interest. For example, treatments for heart disease may be evaluated using data on patient cholesterol levels and blood pressure rather than heart attacks and life span. Extrapolation from surrogate outcomes to outcomes of interest can be difficult; see Fleming and Demets (1996), Psaty *et al.* (1999), and Sculpher and Claxton (2005).¹⁰

Blinding of Treatment Assignment: When the AMR criterion is applied, assigned treatments are known to patients and their physicians. In contrast, blinded treatment assignment has been the norm in clinical trials of new drugs. Hence, a trial at most reveals the distribution of response in a setting where patients and physicians are uncertain what treatment has been assigned. It does not reveal the distribution of response in a real clinical setting where patients and physicians would know the assigned treatment.

Use of Empirical Evidence in Decision Making: Choosing a treatment allocation to minimize maximum regret is remote from the way that the findings of RCTs are now used in decision making. The conventional approach is to perform a statistical hypothesis test, the null hypothesis being that the innovation is no better than the status quo treatment and the alternative being that it is better. If the null hypothesis is not rejected, the status quo treatment continues in force and no one subsequently receives the innovation. If the null is rejected, the innovation replaces the status quo as the treatment of choice. A decision mechanism of this type

¹⁰ Measurement of surrogate outcomes rather than outcomes of welfare interest is also common in research assessing non-medical treatments. Consider, for example, evaluation of educational interventions in early childhood. The outcomes of interest may be years of schooling completed by adulthood and job performance in adulthood. Not wanting to wait for these outcomes to unfold over time, researchers have often used performance in the early grades of school to judge the success of innovations. Here, as in medical trials, extrapolation from surrogate outcomes to outcomes of interest can be difficult.

is institutionalized in the FDA drug approval process; see Fisher and Moyé (1999).

Adaptive Clinical Trials

The AMR criterion shares a broad familial relationship with the idea of *adaptive clinical trials*, but differs in important respects. Adaptive trials sequentially draw subjects into traditional clinical trials and use a frequentist or Bayesian statistical criterion to make the allocation of new subjects across treatments a function of the outcomes observed to date for subjects drawn earlier. The objective, as stated in Tamura *et al.* (1994, p. 768), is to “use the observed response data to adapt the allocation probabilities, so that more patients will hopefully receive the better treatment.”

The AMR criterion shares with adaptive trials the broad objective of using observed treatment responses to inform subsequent treatment choices. However, these ideas differ in two ways. First, the AMR criterion proposes fractional allocation of the entire patient population, not a sample of volunteers. Second, the AMR criterion as developed here is intended to cope with ambiguity rather than the statistical imprecision that motivates adaptive trials. Indeed, the large-population assumption maintained in this paper renders statistical imprecision a negligible concern.

The idealization of a large population approximates well the actual environment for treatment of widespread conditions such as diabetes, heart disease, and various cancers. However, statistical imprecision in empirical findings on treatment response may be a serious cause of errors when the patient population is small. In a one-period planning problem with no identification problems, a planner with clinical trial data can apply findings on finite-sample minimax-regret treatment choice developed in Manski (2004a, 2005a), Manski and Tetenov (2007), Stoye (2009), and Schlag (2007). Manski (2007a) and Stoye (2007b) consider one-period planning problems with certain identification problems.

Multi-period planning problems are more complex. In multi-period settings, there is a tension at each point in time between achievement of two desirable objectives. One is to choose a treatment allocation that

performs as well as possible for the current cohort, given the information currently available. The other is to produce as much new information on treatment response as possible, in order to improve the treatment of future cohorts. These objectives do not conflict in large patient populations because even a small fractional allocation to the innovation produces a large treatment group. They may conflict in small populations.

Cheng, Su, and Berry (2003) study the tension between the two objectives in multi-period planning problems with no identification problems. They approach the problem from the Bayesian perspective. As this paper is written, the interaction of identification problems with statistical imprecision in multi-period planning problems is entirely an open question.

5.4. Adaptive Partial Drug Approval

At several points in Section 5.3 I referred to the FDA drug approval process. The FDA presently makes binary decisions on New Drug Applications submitted by pharmaceutical firms, either approving or disapproving the drug after reviewing the findings of clinical trials. The FDA is a planner, but one with more limited power than the centralized health care systems of Section 5.2.

Let a denote the status quo treatment for a disease and let b denote a new drug treatment subject to FDA approval. FDA disapproval of the new drug sets $\delta = 0$. However, approval does not fix a value for δ . Rather, approval opens a decentralized decision process in which use of the new drug is determined by the market interaction of the pharmaceutical firm, health insurance agencies, physicians, and patients.

Formal analysis of drug approval as a problem of planning under ambiguity is well beyond the scope of this paper. In addition to the fact that approval only permits use of a drug rather than requiring it, the FDA must be cognizant that its approval process may affect the pharmaceutical innovation process that generates new drugs. Nevertheless, I think that our theme of adaptive diversification under ambiguity can usefully be applied to the FDA. I explain here. See also Manski (2009b).

Recent public discourse on drug approval has been dominated by debate over the length of the FDA process. Pharmaceutical firms eager for returns on investments and patient advocates wanting fast access to new drugs argue for shortening the length of the process. Health researchers and patient advocates concerned that approval decisions are made with inadequate knowledge of treatment response argue that trials of longer duration should be performed on subjects who more closely resemble patient populations.

Attention has focused so heavily on the length of the approval process because, as matters stand, the permitted use of a new drug has a sharp discontinuity at the date of the FDA approval decision. Beforehand, a typically tiny fraction of the patient population receives the new drug in clinical trials. Afterwards, use of the drug is unconstrained if approval is granted and zero if approval is not granted. Thus, the date of the approval decision is the central feature of the present process.

Our consideration of the AMR criterion suggests that framing FDA action as a binary decision between full approval and complete disapproval needlessly constrains the set of policy options that warrant consideration. It may be beneficial to empower the FDA to institute an *adaptive partial approval* process, where the extent of the permitted use of a new drug would vary smoothly as empirical evidence accumulates. The stronger the evidence on health outcomes of interest, the more that use of a new drug would be permitted.

An adaptive approval process would eliminate the present discontinuity between the pre-approval and post-approval periods. Instead, the FDA would monitor the available evidence and, over time, adjust the degree of approval. Whereas the AMR criterion fixes the fraction $\delta_{\text{AMR}(n)}$ of the population who receive treatment b in period n , the FDA would set an upper bound on this fraction, leaving it to the market to determine actual usage of the new drug subject to the upper bound.¹¹

¹¹ To learn as much as possible about treatment response, it would be best to randomly allocate the patient population between the status quo treatment and the new drug. But random allocation would be a non-starter politically. Partial approval with market determination of usage of the new drug is intended as a pragmatic second-best approach.

I anticipate that with a partial-approval process in place, pharmaceutical firms would perform longer clinical trials than at present and conduct them on populations that more closely resemble actual patient populations. As discussed in Section 5.3, trials currently provide only weak evidence on treatment response and binary approval decisions are made under much ambiguity. In a partial approval regime, firms would have an incentive to provide stronger evidence in order to secure a higher degree of FDA approval.

6. Two-Planner Games

Sections 2 through 5 studied decision making by a planner who can dictate the treatment allocation. Planners possessing close to unilateral decision power exist in some important settings. For example, centralized health care systems more or less choose medical treatments for their patients. I touched on planning with weaker power only in the discussion of the FDA at the end of Section 5.

This section considers situations in which treatment choice requires the agreement of two planners who may have different welfare functions and beliefs about the feasible states of nature. Section 6.1 shows that if both planners face ambiguity and use the AMR criterion to compare allocations, then there exist fractional allocations that both prefer to the singleton allocations. Section 6.2 uses a union-management decision problem to illustrate.

6.1. Noncooperative Application of the AMR Criterion

Reconsider the multi-period planning problem of Section 5. Rather than a single planner, there now are two planners, denoted $m = 1$ and $m = 2$. Let $u_m(t) \equiv u_m[y(t), t]$ denote the net contribution to the welfare function of planner m when a person receives treatment t and realizes outcome $y(t)$. Let $\alpha_m \equiv E[u_m(a)]$ and

$\beta_m \equiv E[u_m(b)]$. Let S_{mn} index the feasible states of nature in period n , as perceived by planner m .

Suppose that both planners use the AMR criterion to compare treatment allocations. Let δ_{mnAMR} be the allocation that planner m would choose in period n if he were able to dictate policy choice. Without loss of generality, suppose that $\delta_{1nAMR} \leq \delta_{2nAMR}$.

The analysis of Section 5 shows that, for each planner m , the maximum static regret of an allocation δ_n strictly decreases on the interval $[0, \delta_{mnAMR}]$ and increases on the interval $[\delta_{mnAMR}, 1]$. Hence, both planners prefer δ_{1nAMR} to all $\delta_n < \delta_{1nAMR}$, both prefer δ_{2nAMR} to all $\delta_n > \delta_{2nAMR}$, and preferences differ for $\delta_n \in [\delta_{1nAMR}, \delta_{2nAMR}]$. Thus, the set of pareto efficient allocations is $[\delta_{1nAMR}, \delta_{2nAMR}]$.

Suppose that a is a status quo treatment and b is an innovation. Consider a decision mechanism that calls on each planner to announce his preferred allocation and then, giving deference to the status quo, selects the smaller of the two reported values. This mechanism makes it incentive compatible for each planner to reveal his preferred allocation truthfully, regardless of what the other planner announces. Thus, δ_{1nAMR} is the chosen allocation. This result is pareto efficient. Alternatively, one could give deference to the innovation and select the larger of the two reported values. This mechanism also makes truthtelling incentive compatible and yields the pareto efficient allocation δ_{2nAMR} .

When there is conflict between the preferences of two planners, modern societies usually defer to the status quo rather than to the innovation. This is especially evident in the American legal system. A longstanding tenet of the legal system is that the plaintiff in a civil proceeding bears the *burden of proof* to show that an action by the defendant (the status quo) is improper.¹²

¹² This tenet was recently applied by the U.S. Supreme Court to choice between a status quo treatment and an innovation. The Individuals with Disabilities Education Act is a federal statute requiring that public schools provide to each disabled child “an individualized education program.” The language of the statute does not specify who bears the burden of proof when parents believe that a school has not properly complied with the statute. In the case *Shaffer v. Weast* (Supreme Court of the United States, 2005), the parents of a disabled child challenged the adequacy of the educational services provided by his school (the status quo policy) and proposed an alternative (the innovation). The Court ruled that the parents have the of showing the status quo to be inadequate, writing “We hold that the burden lies, as it typically does, on the party seeking relief.”

6.2. Teacher Evaluation in New York City

To illustrate the two-planner decision problem, I consider an educational setting where the problem is to choose between a status quo policy for teacher evaluation and an innovation. The two planners are a school district and a teacher's union. The status quo is the traditional system basing evaluation on scrutiny of teacher preparation and observation of classroom lesson delivery. The innovation bases teacher evaluation on student performance in standardized tests. The contract between the school district and the union requires that any departure from the status quo be approved by both planners.

A potential instance of this teacher evaluation problem was described in a recent article in the *New York Times* (Medina, 2008):

“New York City has embarked on an ambitious experiment, yet to be announced, in which some 2,500 teachers are being measured on how much their students improve on annual standardized tests. While officials say it is too early to determine how they will use the data, which is already being collected, they say it could eventually be used to help make decisions on teacher tenure or as a significant element in performance evaluations and bonuses. . . . Randi Weingarten, the union president, said she had grave reservations about the project, and would fight if the city tried to use the information for tenure or formal evaluations or even publicized it. She and the city disagree over whether such moves would be allowed under the contract.”

Thus, New York City is acting unilaterally to collect data that could potentially be used to evaluate teachers. The contemplated change from the status quo differs from a fractional allocation as defined in this paper because the participating schools were not randomly drawn from the population of New York City schools. This difference aside, the allocation that the City has in mind is fractional with δ about equal to 0.10.

New York City appears to see itself as a planner with unilateral power to implement the innovation.

However, the teacher's union asserts that any departure from the status quo policy requires their agreement. The *Times* reporter writes that an attempted unilateral decision by the City "would undoubtedly open up a legal battle with the teacher's union."

Suppose that implementation of a new policy requires agreement by the City and the union. As presently framed, the decision problem is a static noncooperative choice between $\delta = 0$ and $\delta = 0.10$. The analysis of Section 5 suggests that it would be better to frame it as a noncooperative choice of $\delta_n \in [0, 1]$ in a sequence of periods n , with observed outcomes in earlier periods informing treatment choice in later ones.

The fact that the City currently contemplates a fractional allocation suggests that it views itself as facing a problem of policy choice under ambiguity. The union's perception is not apparent, because it currently has no way to voice its preference except to state its opposition to unilateral decision making by the City. The analysis of Section 6.1 suggests that it would be better to have the City and the union each announce their preferred allocation and then select the smaller of the announced allocations.

7. Conclusion

Diversification is widely accepted as a desirable feature of private planning with incomplete information, enshrined in the aphorism "Don't put all your eggs in one basket." Yet discussions of social planning commonly presume that all observationally identical persons should receive the same treatment. There is no intrinsic reason why a social planner should limit consideration to singleton allocations.

Considering alternative approaches to planning under ambiguity, I have focused on the minimax-regret criterion, which has received remarkably little attention in the long period since it was proposed by Savage (1951). Considering one-period planning problems with individualistic treatment and a linear welfare function, Manski (2007b, Chapter 11) showed that the MR allocation always is fractional under ambiguity

and has a simple explicit form. This paper substantially extends my earlier analysis. I have shown here that the MR criterion diversifies treatment in a large class of problems, including ones with nonlinear welfare functions, interacting treatments, learning, and non-cooperative aspects. Considering multi-period settings, I have emphasized that adaptive diversification of treatments enables a planner to cope with ambiguity in the short run and reduce it in the long run. Thus, this analysis considerably strengthens the formal foundations for the widespread belief that diversification is desirable when making decisions with partial knowledge.

Nevertheless, the reader should be aware that diversified treatment is not always desirable or feasible. In some circumstances, fixed costs of treatment may make singleton allocations preferable or society may insist on equal ex post treatment of equals. Moreover, humanity faces collective threats ranging from strategic nuclear war to global warming to asteroid strikes. As long as we inhabit just one world, these are undiversifiable risks.

Whereas this paper has studied treatment allocation as a planning problem, economists often seek to predict the allocation that would result from decentralized decision making, with members of the population choosing their own treatments. An individual cannot diversify in the setting considered here, as treatments a and b are indivisible. However, decentralization may yield a diversified allocation if individuals use heterogeneous objectives, information, or decision criteria to make their private choices. Then some persons may choose treatment a, while others choose b. It is not possible to make general predictions about the nature of decentralized allocations, but study of specific settings can be productive. I have made some progress in Manski (2004b). See also the discussion in Manski (2007, Sec. 11.8).

Appendix: Proof of Result (44)

Recall the minimax-regret criterion given in (43):

$$\min_{\delta \in [0, 1]} \max_{s \in S} \left\{ \max_{d \in [0, 1]} [\alpha_s(d)(1-d) + \beta d] - [\alpha_s(\delta)(1-\delta) + \beta\delta] \right\}.$$

In this derivation, S indexes the set of weakly increasing functions $\alpha(\cdot)$ such that $L \leq \alpha(0) \leq \alpha(1) \leq U$. I first prove the result when $\beta \leq L$ and then when $L \leq \beta \leq U$. I omit the case $\beta > U$ as it is obvious that vaccinating everyone is best in this circumstance.

$\beta \leq L$: For each value of δ ,

$$\begin{aligned} & \max_{s \in S} \left\{ \max_{d \in [0, 1]} [\alpha_s(d)(1-d) + \beta d] - [\alpha_s(\delta)(1-\delta) + \beta\delta] \right\} \\ &= \max_{\gamma \in [L, U]} \max_{s \in S: \alpha_s(\delta) = \gamma} \max_{d \in [0, 1]} [\alpha_s(d)(1-d) + \beta d] - [\gamma(1-\delta) + \beta\delta] \\ &= \max_{\gamma \in [L, U]} \max_{d \in [0, 1]} \max_{s \in S: \alpha_s(\delta) = \gamma} [\alpha_s(d)(1-d) + \beta d] - [\gamma(1-\delta) + \beta\delta] \\ &= \max_{\gamma \in [L, U]} \max_{\substack{d \in [0, \delta] \\ d \in (\delta, 1]}} \max_{s \in S: \alpha_s(\delta) = \gamma} [\alpha_s(d)(1-d) + \beta d] - [\gamma(1-\delta) + \beta\delta]. \end{aligned}$$

Observe that

$$\begin{aligned} & \max_{\gamma \in [L, U]} \max_{d \in [0, \delta]} \max_{s \in S: \alpha_s(\delta) = \gamma} [\alpha_s(d)(1-d) + \beta d] - [\gamma(1-\delta) + \beta\delta] = \max_{\gamma \in [L, U]} (\gamma - \beta)\delta = (U - \beta)\delta. \\ & \max_{\gamma \in [L, U]} \sup_{d \in (\delta, 1]} \max_{s \in S: \alpha_s(\delta) = \gamma} [\alpha_s(d)(1-d) + \beta d] - [\gamma(1-\delta) + \beta\delta] = \max_{\gamma \in [L, U]} (U - \gamma)(1 - \delta) = (U - L)(1 - \delta). \end{aligned}$$

Hence, problem (19) reduces to $\min_{\delta \in [0, 1]} [(U - \beta)\delta, (U - L)(1 - \delta)]$. As δ increases from 0 to 1, $(U - \beta)\delta$ increases from 0 to $U - \beta$ and $(U - L)(1 - \delta)$ decreases from $U - L$ to 0. Hence, the minimum over δ solves the equation $(U - L)(1 - \delta) = (U - \beta)\delta$. This gives result (44).

$L < \beta \leq U$: For each value of δ ,

$$\begin{aligned} & \max_{s \in S} \left\{ \max_{d \in [0, 1]} [\alpha_s(d)(1 - d) + \beta d] - [\alpha_s(\delta)(1 - \delta) + \beta\delta] \right\} \\ &= \max_{\gamma \in [L, U]} \max_{s \in S: \alpha_s(\delta) = \gamma} \max_{d \in [0, 1]} [\alpha_s(d)(1 - d) + \beta d] - [\gamma(1 - \delta) + \beta\delta] \\ &= \max_{\gamma \in [L, U]} \max_{d \in [0, 1]} \max_{s \in S: \alpha_s(\delta) = \gamma} [\alpha_s(d)(1 - d) + \beta d] - [\gamma(1 - \delta) + \beta\delta] \\ &= \max_{\substack{\gamma \in [L, \beta] \\ \gamma \in [\beta, U]}} \max_{\substack{d \in [0, \delta] \\ d \in (\delta, 1]}} \max_{s \in S: \alpha_s(\delta) = \gamma} [\alpha_s(d)(1 - d) + \beta d] - [\gamma(1 - \delta) + \beta\delta]. \end{aligned}$$

Observe that

$$\max_{\gamma \in [L, \beta]} \max_{d \in [0, \delta]} \max_{s \in S: \alpha_s(\delta) = \gamma} [\alpha_s(d)(1 - d) + \beta d] - [\gamma(1 - \delta) + \beta\delta] = 0.$$

$$\max_{\gamma \in [L, \beta]} \sup_{d \in (\delta, 1]} \max_{s \in S: \alpha_s(\delta) = \gamma} [\alpha_s(d)(1 - d) + \beta d] - [\gamma(1 - \delta) + \beta\delta] = \max_{\gamma \in [L, \beta]} (U - \gamma)(1 - \delta) = (U - L)(1 - \delta).$$

$$\max_{\gamma \in [\beta, U]} \max_{d \in [0, \delta]} \max_{s \in S: \alpha_s(\delta) = \gamma} [\alpha_s(d)(1 - d) + \beta d] - [\gamma(1 - \delta) + \beta\delta] = \max_{\gamma \in [\beta, U]} (\gamma - \beta)\delta = (U - \beta)\delta.$$

$$\max_{\gamma \in [\beta, U]} \sup_{d \in (\delta, 1]} \max_{s \in S: \alpha_s(\delta) = \gamma} [\alpha_s(d)(1 - d) + \beta d] - [\gamma(1 - \delta) + \beta\delta] = \max_{\gamma \in [\beta, U]} (U - \gamma)(1 - \delta) = (U - \beta)(1 - \delta).$$

Hence, the minimum over δ again solves the equation $(U - L)(1 - \delta) = (U - \beta)\delta$. This again gives (44).

References

Berger, J. (1985), *Statistical Decision Theory and Bayesian Analysis*, New York: Springer-Verlag.

Brock, W. and C. Manski (2008), "Competitive Lending with Partial Knowledge of Loan Repayment," Department of Economics, Northwestern University.

Cheng, Y., F. Su, and D. Berry (2003), "Choosing Sample Size for a Clinical Trial Using Decision Analysis," *Biometrika*, 90, 923-936.

Chernoff, H. (1954), "Rational Selection of Decision Functions," *Econometrica*, 22, 422-443.

Fisher, L. and L. Moyé (1999), "Carvedilol and the Food and Drug Administration Approval Process: An Introduction," *Controlled Clinical Trials*, 20, 1-15.

Fleming, T. and D. Demets (1996), "Surrogate End Points in Clinical Trials: Are We Being Misled?" *Annals of Internal Medicine*, 125, 605-613.

Gilboa, I. and D. Schmeidler (1989), "Maxmin Expected Utility with Non-Unique Prior," *Journal of Mathematical Economics*, 18, 141-153.

Hansen, L. and T. Sargent (2008), *Robustness*, Princeton: Princeton University Press.

Manski, C. (1988), "Ordinal Utility Models of Decision Making Under Uncertainty," *Theory and Decision*, 25, 79-104.

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. (2004a), "Statistical Treatment Rules for Heterogeneous Populations," *Econometrica*, 72, 1221-1246.

Manski, C. (2004b), "Social Learning from Private Experiences: The Dynamics of the Selection Problem," *Review of Economic Studies*, 71, 443-458.

Manski, C. (2005a), *Social Choice with Partial Knowledge of Treatment Response*, Princeton: Princeton University Press.

Manski, C. (2005b), "Fractional Treatment Rules for Social Diversification of Indivisible Private Risks," National Bureau of Economic Research Working Paper W11675.

Manski, C. (2006), "Search Profiling with Partial Knowledge of Deterrence." *Economic Journal*, 116, F385-F401.

Manski, C. (2007a), "Minimax-Regret Treatment Choice with Missing Outcome Data," *Journal of Econometrics*, 139, 105-115.

- Manski, C. (2007b), *Identification for Prediction and Decision*, Cambridge, Mass.: Harvard University Press.
- Manski, C. (2008b), "Actualist Rationality," Department of Economics, Northwestern University.
- Manski, C. (2009a), "Social Planning with Partial Knowledge of Social Interactions," *Journal of Institutional and Theoretical Economics*, 165, 146-158.
- Manski, C. (2009b), "Adaptive Partial Drug Approval: A Health Policy Proposal," *The Economists' Voice*, 6, article 9.
- Manski, C. and D. Nagin (1998), "Bounding Disagreements about Treatment Effects: a Case Study of Sentencing and Recidivism," *Sociological Methodology*, 28, 99–137.
- Manski, C. and A. Tetenov (2007), "Admissible Treatment Rules for a Risk-Averse Planner with Experimental Data on an Innovation," *Journal of Statistical Planning and Inference*, 137, 1998-2010.
- Medina, J. (2008), "New York Measuring Teachers by Test Score," *The New York Times*, January 21, 2008.
- Psaty, B. et al. (1999), "Surrogate End Points, Health Outcomes, and the Drug-Approval Process for the Treatment of Risk Factors for Cardiovascular Disease," *Journal of the American Medical Association*, 282, 786-790.
- Savage, L. (1951), "The Theory of Statistical Decision," *Journal of the American Statistical Association*, 46, 55–67.
- Savage, L. (1954), *The Foundations of Statistics*, New York: Wiley.
- Sculpher, M. and K. Claxton (2005), "Establishing the Cost-Effectiveness of New Pharmaceuticals under Conditions of Uncertainty—When Is There Sufficient Evidence?" *Value in Health*, 8, 433-446.
- Schlag, K. (2006), "Eleven," Technical Report, European University Institute.
- Stoye, J. (2007a), "Minimax Regret Treatment Choice with Incomplete Data and Many Treatments," *Econometric Theory* 23, 190-199.
- Stoye, J. (2007b), "Minimax-Regret Treatment Choice with Finite Samples and Missing Outcome Data," Proceedings of the 5th International Symposium on Imprecise Probability, Prague, Czech Republic.
- Stoye, J. (2009), "Minimax-Regret Treatment Choice with Finite Samples," *Journal of Econometrics*, forthcoming.
- Supreme Court of the United States (2005), "No. 04698, Brian Schaffer, a Minor, by His Parents and next Friends, Jocelyn and Martin Schaffer, et Al., Petitioners V. Jerry Weast, Superintendent, Montgomery County Public Schools, et al.," 546 U. S. __ 2005.
- Tamura, R., D. Faries, J. Andersen, and J. Heiligenstein (1994), "A Case Study of an Adaptive Clinical Trial in the Treatment of Out-Patients with Depressive Disorder," *Journal of the American Statistical Association*, 89, 768-776.

Tetenov, A. (2008), "Statistical Treatment Choice Based on Asymmetric Minimax Regret Criteria," Collegio Carlo Alberto, Turin, Italy.

von Neumann, J. and O. Morgenstern (1944), *Theory of Games and Economic Behavior*, Princeton: Princeton University Press.