CHOOSING TREATMENT POLICIES UNDER AMBIGUITY

Charles F. Manski

Department of Economics and Institute for Policy Research, Northwestern University

Forthcoming: *Annual Review of Economics*, Vol. 3.

Running Title: Treatment Under Ambiguity

contact information:

Prof. Charles F. Manski, Department of Economics, Northwestern University

2001 Sheridan Road, Evanston, IL 60208 USA

phone 847-491-8223, fax 847-491-7001

cfmanski@northwestern.edu

ABSTRACT

Economists studying choice with partial knowledge typically assume that the decision maker places a subjective distribution on unknown quantities and maximizes expected utility. Someone lacking a subjective distribution faces a problem of choice under ambiguity. This article reviews recent research on policy choice under ambiguity, when the task is to choose treatments for a population. Ambiguity arises when a planner has partial knowledge of treatment response and does not feel able to place a subjective distribution on the unknowns. I first discuss dominance and alternative criteria for choice among undominated policies. I then illustrate with choice of a vaccination policy by a planner who has partial knowledge of the effect of vaccination on illness. I next study a class of problems where a planner may want to cope with ambiguity by diversification, assigning observationally identical persons to different treatments. Lastly, I consider a setting where a planner should not diversify treatment.

Keywords: dominance, minimax regret, partial identification, planning, social choice, treatment response

CONTENTS

1. Introduction

When studying collective decision problems, economists have long asked how an optimizing social planner should behave. 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 is to characterize the optimal policy.

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

Research on optimal income taxation illustrates the problem. Stimulated by Mirrlees (1971), many theoretical studies have derived optimal tax schedules under the assumption that the planner knows how alternative tax schedules affect labor supply. However, knowledge of the actual responsiveness of labor supply to income taxes remains limited, despite the strenuous effort of empirical economists to shed light on the matter.

The informativeness of policy research is limited by both statistical imprecision and identification problems. Statistical imprecision arises when one attempts to make inference on a study population from sample data. Identification problems arise when one seeks to extrapolate from observable features of a study population to unobservable features or to other populations. Statistical imprecision contributes to lack of knowledge, but identification problems usually are the dominant difficulty in policy research. This review article focuses on partial knowledge stemming from identification problems.

A fundamental identification problem arises from the unobervability of counterfactual policy outcomes. At most one can observe the outcomes that have occurred under realized policies. The outcomes of unrealized policies are logically unobservable. Yet determination of an optimal policy requires prospective comparison of all feasible policies.

Practical problems of data collection enlarge the problem of identification of policy impacts. A planner may want to learn long-term policy outcomes, whereas empirical research may only measure short-term outcomes. Survey respondents may refuse to answer or may respond inaccurately to questions about their environments and outcomes. Experimental subjects may not comply with assigned treatments or may drop out of trials before their outcomes are measured.

These and other identification problems have long been central concerns of econometrics and of empirical research in economics. See Manski (2007) for a textbook exposition. Yet their implications for policy choice have remained largely unacknowledged in theoretical studies in public economics.

Social choice theory rarely makes any reference to uncertainty, never mind to specific inferential problems. The subject is not addressed in either the first or second edition New Palgrave articles on social choice (Sen, 1987; Bossert and Weymark, 2008). Research on mechanism design has studied planning in asymmetric-information settings, where heterogeneous agents possess private information about themselves. However, it is usually assumed that the planner knows the population distribution of unobserved agent characteristics and can optimize given this knowledge. For example, the planner of optimal income tax theory does not know the utility functions of individual agents but is assumed to know the population distribution of utility functions.

When economists have studied planning with partial knowledge, it has been standard to assert a subjective probability distribution over unknown decision-relevant quantities and propose choice of a policy that maximizes subjective expected welfare. For example, Nordhaus (2008) used this approach to express partial knowledge of parameter values in his assessment of global warming policy. Meltzer (2001) applied the expected utility criterion to medical decision making and Dehejia (2005) to evaluation of welfare programs.

Maximization of subjective expected welfare is reasonable when a planner has a credible basis for asserting a subjective distribution on unknown quantities. However, a subjective distribution is a form of knowledge, and a planner may not feel able to assert one. Then the planner faces a problem of choice under *ambiguity*.

Use of the term *ambiguity* to describe the absence of a subjective distribution appears to have originated in Ellsberg (1961). The term *uncertainty* was used in Keynes (1921) and Knight (1921). Some authors refer to ambiguity as *Knightian uncertainty*.

This article reviews my recent research on policy choice under ambiguity, which began with Manski (2000). The planning problems that I have studied share a relatively simple structure. The task is to choose treatments for a population whose members may vary in their response to treatment. The social welfare function sums the outcomes of the population members. Ambiguity arises when a planner has partial knowledge of treatment response and does not feel able to place a subjective distribution on the unknowns.

Here are three illustrations, among many that might be given.

Choosing Medical Treatments: Consider a health agency that must treat a population of persons who are susceptible to a disease. The relevant outcome is the health benefit of a treatment minus its cost, measured in comparable units. A utilitarian welfare function sums these net benefits across the population. The problem is that medical science yields only partial knowledge of treatment response. □

Choosing Sentences for Convicted Offenders: Consider a judge who must choose sentences for a population of convicted offenders. The relevant outcome may be recidivism by these offenders; that is, their future criminality. The problem is that criminologists have found it difficult to learn how sentencing affects recidivism. □

An Investor's Asset Allocation Decision: Consider an investor who must allocate 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. The analog of welfare is the aggregate return earned by the investor. At the time of the allocation decision, the investor may have only partial knowledge of investment returns. □

Whereas the first two illustrations concern policy choice, the third poses a classic problem of private decision making. Nevertheless, asset allocation shares the formal structure of the planning problems I have studied, and the findings apply to it.

This article is organized as follows. Section 2 sets forth the basic decision theoretic ideas on choice under ambiguity that have guided my study of policy choice when the welfare function is partially identified. I begin with the orthodox notion that a decision maker facing ambiguity should eliminate dominated actions from consideration. However, I depart from orthodoxy by arguing that maximization of subjective expected utility deserves no privileged status as a criterion for choosing an undominated action. I suggest that the maximin and the minimax-regret criteria merit serious consideration.

Abstract discussion of choice under ambiguity does not take one very far toward understanding the implications of dominance and alternative decision criteria. Study of particular policy problems is more illuminating. Sections 3 through 5 describe my analysis of three classes of treatment choice problems. Each section is of independent interest and may be read on its own. Juxtaposing the three sections also is rewarding, showing how the structure of a decision problem affects the choices that may be reasonable to contemplate.

Drawing on Manski (2010a), Section 3 considers choice of a vaccination policy when a health planner has partial knowledge of the external effect of vaccination on the illness rate of unvaccinated persons. Beyond its intrinsic interest, this work demonstrates how one may address a class of choice problems where a planner observes the outcome of a status-quo policy and feels able to partially extrapolate from the status quo to counterfactual policies. I first show how the planner can eliminate dominated vaccination rates and then how he can use the minimax or minimax-regret criterion to choose an undominated vaccination rate.

Section 4 describes some of the analysis of Manski (2009), which develops a broad theme about treatment under ambiguity through study of a particular decision criterion. The broad theme is that a planner may want to cope with ambiguity by diversification, assigning observationally identical persons to different treatments. Study of the minimax-regret criterion substantiates the theme. I show that this criterion always diversifies treatment when a planner must allocate the population to two treatments, and does not know which treatment is better. The adaptive minimax-regret criterion extends the analysis to dynamic settings, where the planner allocates a sequence of cohorts to treatment and can use outcomes observed from earlier treatment decisions to inform later decisions.

Sections 3 and 4 studied settings where each member of the population receives one of two treatments and where treatment response may vary across the population. Section 5 considers a scenario with a different structure, examined in Manski (2010b). Now the feasible treatments are a convex set, and treatment response is given by a common concave function that maps the treatment and the state of nature into an outcome. I show that the planner should not diversify treatment in this setting. Any fractional allocation is dominated by one that gives all members of the population the mean treatment.

Throughout this review article, I suppose that partial knowledge of treatment response stems purely from identification problems. A planner who observes only a sample of a study population rather than the population in its entirety must cope with statistical imprecision as well as with identification problems. The Wald (1950) development of statistical decision theory provides an appropriate framework for study of planning with sample data. Methodological research applying this framework includes Manski (2004; 2007, Chapter 12), Manski and Tetenov (2007), Hirano and Porter (2009), Stoye (2009, 2010), and Tetenov (2009). However, applications to practical problems of policy choice have yet to appear.

Readers with a macroeconomic orientation may ask how the work described here relates to the contemporaneous program of research on *robust* macroeconomic policy; see Barlevy (2010) for a review and references. The two research programs share a broad concern with policy choice under ambiguity, but they have differed in many important respects. Methodologically, the macroeconomic research has focused on the maximin criterion, whereas I have first studied dominance and then mainly applied the minimax-regret criterion to choose an undominated policy. Substantively, macroeconomists have studied problems requiring the planner to choose a single treatment for the entire population, whereas I have studied ones where the planner may choose a separate treatment for each member of the population.

Another important difference in the research programs concerns the maintained assumptions. Macroeconomists have usually assumed that the actual process driving the economy is a perturbation of some benchmark model. I have typically maintained relatively weak shape restrictions and distributional assumptions about treatment response.

Some readers may ask how the work described here relates to *sensitivity analysis*. See Weycker *et al.* (2005) for an example in the context of vaccination policy. Sensitivity analysis aims to determine optimal policy under a specified set of alternative assumptions. It does not provide a criterion for choice with partial knowledge, when one does not know which assumption is correct.

2. Choice under Ambiguity

This section reviews basic ideas about choice under ambiguity. I consider an agent—perhaps a firm, an individual, or a planner—who must choose an action yielding welfare that depends on an unknown state of nature. The agent has an objective function and beliefs, which I take as primitives. His problem is to choose an action without knowing the actual state of nature.

Formally, the agent faces choice set C and knows (or believes) that the actual state of nature lies in some set S. The objective function $w(\cdot, \cdot)$: $C \times S \rightarrow R^1$ maps actions and states into welfare. For example, $w(\cdot, \cdot)$ may be the profit function of a firm, the utility function of a consumer, or the welfare function of a planner. The agent wants to maximize $w(\cdot, r)$, where r is the actual state of nature, but he does not know r. He only knows that $r \in S$.

2.1. Dominance

How should the decision maker choose among the actions in C? The only prescription that I think warrants universal acceptance is *respect for weak dominance*. Action $c \in C$ is weakly dominated if there exists a $d \in C$ such that $w(d, s) \geq w(c, s)$ for all $s \in S$ and $w(d, s) > w(c, s)$ for some $s \in S$. Respect for weak dominance means than an agent should not choose a weakly dominated action. This prescription is uniquely compelling because weak dominance defines the circumstances in which an agent who wants to maximize $w(\cdot, r)$ knows that choice of one action improves on choice of another.

Let D denote the undominated subset of C. How should the decision maker choose among the elements of D? Let c and d be two undominated actions. Then either $[w(c, s) = w(d, s)$, all $s \in S]$ or there exist $s' \in S$ and $s'' \in S$ such that $w(c, s') > w(d, s')$ and $w(c, s'') < w(d, s'')$. In the former case, c and d are equally good choices and the decision maker is indifferent between them. In the latter case, the decision maker cannot order the two actions. Action c may yield a better or worse outcome than action d; the decision maker cannot say which. Thus, the normative question “How should the decision maker choose?” has no unambiguously correct answer.

2.2. Optimization of Known Transformations of the Welfare Function

Although there is no optimal choice among undominated actions, decision theorists have not wanted to abandon the idea of optimization. So they have proposed various ways of transforming the unknown objective function w(·, ·) into a function of actions alone, which can be maximized. In principle, one should maximize this function only over the undominated actions D. However, it often is difficult to determine which actions are undominated. Hence, it is common to perform the maximization over the full set C of feasible actions.

One idea is to average the elements of S and maximize the resulting function. This yields maximization of expected utility. Another is to seek an action that, in some well-defined sense, works uniformly well over all elements of S. This yields the maximin and minimax-regret criteria.

The Expected Utility Criterion

Many decision theorists suggest that a decision maker who knows only that the true state of nature lies in S should choose an action that maximizes some average of w(·, ·) over the elements of S. Let π be a specified probability distribution on S. For each feasible action c, let $\int w(c, s)d\pi$ be the mean value of w(c, s), calculated with respect to π. The expected utility criterion solves the optimization problem

$$(1) \qquad \max_{c \in C} \int w(c, s)d\pi.$$

In general, the solution to (1) depends on the distribution π placed on S. Decision theorists recommend that π should express the decision maker's personal beliefs about where r lies within S. Hence, π is called a *subjective* probability distribution.

Savage (1954) argued that a decision maker not only *might* but *should* place a subjective distribution on the states of nature and maximize expected utility. He famously proved a representation theorem showing that maximization of subjective expected utility is mathematically equivalent to behaving in accord with a set of consistency axioms on hypothetical choice behavior. Savage felt that adherence to the axioms constitutes the essence of rational choice with partial knowledge. Hence, he drew the normative conclusion that a decision maker should place a subjective distribution on unknowns and should maximize expected utility.

Many economists and decision theorists find the Savage argument compelling. Hence, they see no need for normative analysis of choice under ambiguity. They may, however, still see a need for positive analysis, as there is ample empirical evidence that humans sometimes do not adhere to the Savage axioms in practice.

Other researchers do not agree that adherence to the Savage axioms constitutes the essence of rational choice. There are multiple schools of thought on the matter. A diverse body of modern axiomatic decision theory studies various systems of axioms on hypothetical choice behavior. Some of these axiom systems do not imply that a person will place a subjective distribution on unknowns, never mind maximize expected utility. See Binmore (2009) for a recent perspective. I personally have gone further and have argued that the basic concern of axiomatic decision theory, characterization of consistency in behavior across hypothetical choice scenarios, is not relevant to actual decision making (Manski, 2010c).

To find this review article of potential interest, a reader need not subscribe to a particular viewpoint on axiomatic decision theory. He or she need only find it reasonable to conceive of a decision maker who lacks a subjective distribution on the states of nature. Then this decision maker must somehow cope with ambiguity. I next discuss the two most prominent suggestions in the literature, the maximin and minimax-regret criteria.

The Maximin Criterion

The maximin criterion suggests that the decision maker choose an action that maximizes the minimum welfare attainable across the elements of S. For each feasible action c, consider the minimum feasible value of w(c, s); that is, $\min_{s \in S} w(c, s)$. A maximin rule chooses an action that solves the optimization problem

(2) $$\max_{c \in C} \ \min_{s \in S} \ w(c, s).$$

The maximin criterion has a clear normative foundation in *strictly competitive two-person games* (see Friedman, 1986, p. 30). In such a game, the decision maker chooses an action from C. Then a state from S is chosen by an opponent whose objective is to minimize the realized outcome. A decision maker who knows that he is a participant in a competitive game does not face ambiguity. He faces the problem of maximizing the known function specified in the maximin rule.

There is no compelling reason why a decision maker should or should not use a maximin rule when r is a fixed but unknown state. In this setting, the appeal of the maximin criterion is a personal rather than normative matter. Some decision makers may deem it essential to protect against worst-case scenarios, while others may not. Wald (1950), who studied the maximin criterion in depth, did not contend that a maximin rule is optimal, only that it is "reasonable." Considering the case in which the objective is to minimize rather than maximize w(·, r), he wrote (Wald, 1950, p. 18): "a minimax solution seems, in general, to be a reasonable solution of the decision problem."

The Minimax-Regret Criterion

The minimax-regret criterion, introduced in Savage (1951), has the decision maker choose an action that minimizes the maximum loss to welfare that results from not knowing the objective function. A minimax-regret choice solves the problem

$$(3) \qquad \min_{c \in C} \max_{s \in S} \; [\max_{d \in C} w(d, s) - w(c, s)].$$

Here $\max_{d \in C} w(d, s) - w(c, s)$ is the *regret* of action c in state of nature s; that is, the welfare loss associated with choice of c relative to an action that maximizes welfare in state s. The actual state is unknown, so one evaluates c by its maximum regret over all states and selects an action that minimizes maximum regret.

The maximin and minimax-regret criteria are sometimes confused with one another. Comparison of (2) and (3) shows that they are generally distinct. The two criteria coincide only in special cases. Suppose in particular that $\max_{d \in C} w(d, s)$ is constant for all $s \in S$. Then minimax-regret reduces to maximin.

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

Other Decision Criteria

The three criteria for decision making under ambiguity discussed above are particularly well-known, and I have found it revealing to compare their implications for behavior. However, these criteria are not the only ones that have received attention.

A decision maker who feels able to assert a partial subjective distribution on the states of nature could maximize minimum expected utility or minimize maximum expected regret. These criteria combine elements of expected utility theory with elements of research on choice under ambiguity. Statistical decision theorists refer to them as the Γ-maximin and Γ-minimax regret criteria (see Berger, 1985). The Γ-maximin criterion has drawn considerable attention from axiomatic decision theorists (Gilboa and Schmeidler, 1989).

3. Vaccination under Ambiguity

3.1.Background

The problem of choosing an optimal vaccination policy for a population susceptible to infectious disease has drawn considerable attention from epidemiologists and some from economists as well. Researchers studying optimal vaccination have typically assumed the planner knows how vaccination affects illness rates. See, for example, Ball and Lyne (2002), Becker and Starczak (1996), Brito, Sheshinski, and Intriligator (1991), Boulier, Datta, and Goldfarb (2007), Hill and Longini (2003), Patel, Longini, and Halloran (2005), and Scuffham and West (2002).

There are two reasons why a planner may have only partial knowledge of the effect of vaccination on illness. First, the planner may only partially know the *internal* effectiveness of vaccination in generating an immune response that prevents a vaccinated person from become ill or infectious. Second, he may only partially know the *external* effectiveness of vaccination in preventing transmission of disease to members of the population who are unvaccinated or unsuccessfully vaccinated.

The second issue is particularly problematic. A standard randomized clinical trial, which vaccinates an experimental group of individuals, enables evaluation of the internal effectiveness of vaccination. However, the trial does not reveal the external effect of applying different vaccination rates to the population. The outcome data only reveal the external effectiveness of the chosen vaccination rate. The outcomes with other vaccination rates remain counterfactual, yet choice of a vaccination policy requires comparison of alternative rates.

Attempting to cope with the absence of empirical evidence, researchers have used epidemiological models to forecast the outcomes that would occur with counterfactual vaccination policies. The articles on optimal vaccination cited earlier use a variety of such models. However, authors typically provide little information that would enable one to assess the accuracy of their assumptions about individual behavior, social interactions, and disease transmission.

Manski (2010a) studies choice of vaccination policy when a planner has partial knowledge of the external effectiveness of vaccination. I suppose that the planner's objective is to minimize the social cost of illness and vaccination. The consequences of alternative vaccination rates depend on the extent to which vaccination prevents illness. I suppose that the planner observes the illness rate of a study population whose vaccination rate has been chosen previously. He assumes that the illness rate of unvaccinated persons weakly decreases as the vaccination rate increases, but he does not know the magnitude of the preventive effect of vaccination.

In this setting, I first show how the planner can eliminate dominated vaccination rates and then how he can use the minimax or minimax-regret criterion to choose an undominated rate. Sections 3.2 through 3.4 summarize the analysis and findings. Section 3.5 discusses related planning problems.

3.2. Optimal Vaccination

As prelude to consideration of vaccination under ambiguity, I specify the optimization problem that the planner wants to solve and derive the solution in an illustrative case.

For simplicity, I suppose here that the planner must choose the vaccination rate for a large population of observationally identical persons. Supposing that members of the population are observationally identical does not mean that persons actually are identical, only that the planner cannot distinguish them. If the planner observes health-relevant covariates for each person, he may want to choose vaccination rates that vary with these covariates. The present analysis extends easily to such settings if the external effect of vaccination occurs only within groups defined by observed covariates and not between groups. It also extends to settings where vaccination has imperfect, but known, internal effectiveness. See Manski (2010a), Sec. 4.

I assume that vaccination always prevents a vaccinated person from becoming ill. Let p(t) be the *external-response function*, giving the fraction of unvaccinated persons who become ill when the vaccination rate is t. Then the fraction of the population who become ill is p(t)(1 − t).

I suppose the planner wants to minimize a social cost function with two additive components. These are the harm caused by illness and the cost of vaccination. Let a > 0 denote the mean social harm caused by illness and let c > 0 denote the mean social cost per vaccination, measured in commensurate units. The social cost of vaccination rate t is

(4) $$K(t) = ap(t)(1 - t) + ct.$$

The first term on the right-hand side gives the aggregate cost of illness, and the second gives the aggregate cost of vaccination. This simple social cost function expresses the core tension of vaccination policy: a higher vaccination rate is more effective in preventing illness but is more costly.

The planner wants to solve the problem $\min_{t \in [0, 1]} K(t)$. The optimization problem is invariant to the scale of $K(\cdot)$. Hence, without loss of generality, I let a = 1 and interpret c as the ratio of the mean social cost of vaccination to the mean social cost of illness.

In the notation of Section 2, the set S of states of nature is the set of external-response functions that the planner deems feasible. The choice set is C = [0, 1]. Welfare function w is the negative of social cost function K.

The planner's problem is solvable if the external-response function is known. Suppose it is known to be linear, with $p(t) = \rho(1 - t)$ and $0 < \rho \leq 1$. Thus, the illness rate of unvaccinated persons is ρ if no one is vaccinated and decreases linearly to zero as the vaccination rate rises. Then the optimal vaccination rate is

$$(5) \qquad t^* = \underset{t \in [0, 1]}{\operatorname{argmin}}\ \rho(1 - t)^2 + ct.$$

The quadratic first term of the social cost function is minimized at t = 1, and the linear second term at t = 0. The optimal vaccination rate must resolve this tension. The optimal rate is

$$(6) \qquad \begin{aligned} t^* &= 0 && \text{if } 2\rho < c. \\ &= 1 - c/(2\rho) && \text{if } 2\rho \geq c. \end{aligned}$$

Observe that for no value of parameters (c, ρ) is it optimal to vaccinate the entire population. It is, however, optimal to vaccinate no one if 2ρ < c.

3.3. Partial Knowledge of External Effectiveness

To demonstrate how a planner with partial knowledge of external effectiveness may choose a vaccination rate, I consider decision making in a particular informational setting. I suppose that the planner observes the vaccination and illness rates of a study population, whose vaccination rate has been chosen previously to be some value less than one. I have the planner maintain two assumptions. First, he assumes that the study population and the treatment population have the same external-response function. Second, he assumes that the illness rate of unvaccinated persons weakly decreases as the vaccination rate increases. However, he makes no assumption about the magnitude of the external effect of vaccination and does not place a subjective distribution on this magnitude.

Let $r < 1$ denote the observed vaccination rate in the study population and $q(1 - r)$ denote the observed realized illness rate. The two maintained assumptions are

Assumption 1 (Study Population): The planner observes r and $q(1 - r)$. He knows that $q = p(r)$.

Assumption 2(Vaccination Weakly Prevents Illness): The planner knows that $p(t)$ is weakly decreasing in t.

Taken together, these assumptions imply that

(7) $t \le r \;\Rightarrow\; p(t) \ge q,$

$t \ge r \;\Rightarrow\; p(t) \le q.$

This knowledge of the external response function is much weaker than the traditional assumption that the planner knows the function. Moreover, Assumptions 1 and 2 often are credible. It often is possible to observe the vaccination and illness rates of a study population, and credible to assume that the study and treatment populations have similar if not identical external-response functions. It usually is credible to assume that vaccination weakly prevents illness. Assumption 2 is a specific instance of the general idea of *monotone treatment response* developed in Manski (1997).

Given the empirical evidence and assumptions, I show that a candidate vaccination rate t is strictly dominated if any of these conditions hold:

(a) Let $c < q$. Then t is strictly dominated if $t < r$.

(b) Let $c > q$. Then t is strictly dominated if $t > r + q(1 - r)/c$.

(c) Let $c > 1$. Then t is strictly dominated if $(1 - q)/(c - q) < t \leq r$ or if $t > \max(r, 1/c)$.

It might have been thought that Assumptions 1 and 2 are too weak to yield interesting dominance findings. However, the proposition shows that these assumptions have considerable power. The broad finding is that small (large) values of t are dominated when the vaccination cost c is sufficiently small (large). Parts (a) through (c) give the specifics.

With dominated vaccination rates eliminated from consideration, the planner must still choose among the undominated rates. I derive the minimax and minimax-regret rates.

The minimax criterion selects the vaccination rate that minimizes maximum social cost over all feasible external-response functions. The minimax rate turns out to be

$$
\begin{aligned}
(8) \quad t^m &= 0 && \text{if} \quad c > 1 \text{ and } 1 \leq q(1 - r) + cr, \\
&= r && \text{if} \quad c > 1 \text{ and } 1 \geq q(1 - r) + cr \\
& && \text{or if} \quad q < c < 1, \\
&= \text{all } t \in [0, 1] && \text{if } c = q \text{ and } q = 1, \\
&= \text{all } t \in [r, 1] && \text{if } c = q \text{ and } q < 1, \\
&= 1 && \text{if} \quad c < q.
\end{aligned}
$$

Thus, the minimax rate generically takes one of the three values (0, r, 1), the only exception being when $c = q$, which has multiple maximin rates. All else equal, the minimax rate weakly decreases with the vaccination cost c. It weakly increases with the realized illness rate q if $c < 1$ and decreases with q otherwise.

The regret of vaccination rate t measures the difference between the social cost delivered by rate t and that delivered by the best possible rate. The minimax-regret criterion selects the vaccination rate that minimizes maximum regret across all feasible external response functions. I show the following:

(a) Let $c \leq q$. Then the minimax-regret vaccination rate is

(9a) $$t^{mr} = (q + cr)/(q + c).$$

(b) Let $c > q$. Then the minimax-regret vaccination rate is

(9b) $$t^{mr} = \underset{t \in [0, 1]}{\operatorname{argmin}}\ 1[t < r]\cdot\{\max[(1 - q)(1 - t), (1 - t) + c(t - r), (c - q)t]\} + 1[t \geq r]\cdot\{\max[q(1 - t), c(t - r), (c - q)t]\}.$$

Thus, as the cost c of vaccination increases from 0 to q, the minimax-regret vaccination rate decreases continuously from 1 to $(1 + r)/2$. In contrast, the minimax rate equals 1 whenever $c \leq q$. When $c > q$, the solution to the minimax-regret problem generally does not have an explicit form of simplicity comparable to the minimax problem. However, the abstract finding in (9b) simplifies in the polar case r = 0, where no one was vaccinated in the study population. Then $t^{mr} = q/(q + c)$.

3.4. Numerical Examples

Numerical examples are useful to illustrate the above findings. Here are three, each modifying the preceding example in some respect.

First consider a scenario where the mean cost of vaccination (relative to illness) is c = 0.05. The planner observes a study population with no vaccination (r = 0) and with illness rate q = 1/5. In this setting, a planner who believes the external-response function is linear would conclude that $\rho = 1/5$ and would choose the vaccination rate $t^* = 7/8$. A planner who only knows the function to be weakly decreasing would not be able to conclude that any vaccination rates are dominated, because c < q and r = 0. The minimax vaccination rate is $t^m = 1$ and the minimax-regret rate is $t^{mr} = 4/5$.

Next revise the scenario by supposing that the planner observes a study population where r = 1/2 and q = 1/10. Continue to assume that c = 0.05. A planner who believes the external-response function is linear would still conclude that $\rho = 1/5$ and choose $t^* = 7/8$. A planner who only knows the function to be weakly decreasing can determine that any vaccination rate smaller than 1/2 is strictly dominated. The minimax vaccination rate remains $t^m = 1$ and the minimax-regret rate is $t^{mr} = 5/6$.

Now revise the scenario again by supposing that vaccination is more costly relative to illness, say c = 0.25. Continue to assume that r = 1/2 and q = 1/10. In this case, a planner who believes the external-response function is linear would choose $t^* = 3/8$. A planner who only knows the function to be weakly decreasing can conclude that any vaccination rate larger than 7/10 is strictly dominated. The minimax and the minimax-regret vaccination rates are both 1/2.

3.5. Related Planning Problems

The scenario considered above is realistic enough to demonstrate key ideas about vaccination under ambiguity, but is idealized enough to yield simple findings. Looking beyond vaccination, the analysis demonstrates how one may address a class of choice problems where a planner observes the outcome of a status-quo policy and feels able to partially extrapolate to counterfactual policies. Manski (2006) gives another demonstration. There I studied the criminal-justice problem of choosing a rate of search for evidence of crime, when a planner has partial knowledge of the deterrent effect of search on the rate of crime commission. I considered a planner who wants to minimize the social cost of crime, search, and punishment. The planner observes the crime rate under a status-quo search rate and assumes that the crime rate falls as the search rate rises. The formal structure of this planning problem is similar to that of the vaccination problem, the substantive difference between the two notwithstanding.

4. Diversified Treatment Choice

4.1. Background

In the Introduction, I cited allocation of an endowment between two assets as a planning problem. When an investor is unsure which asset will yield the higher return, it is common to recommend that he should hold a diversified portfolio. That is, he should allocate a positive fraction of the endowment to each asset. Traditional formal arguments for diversification assume that the investor maximizes subjective expected utility and is risk averse.

Diversification may also be appealing when a social planner must treat a population of persons and does not know the optimal treatment. Let there be two feasible treatments, labeled a and b. The broad argument for diversification is that it enables the planner to balance two types of potential error. 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. The singleton allocation assigning the entire population 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.

A formal argument for diversified policy choice may be made by supposing that the planner maximizes subjective expected welfare and is risk averse in the traditional sense of the term. That is, the planner aggregates population outcomes, takes a monotone-concave transformation of the aggregate outcome, and maximizes the expected value of this quantity. But what about scenarios where the planner lacks a subjective distribution over treatment outcomes? Manski (2007, Chapter

11; 2009) considers the minimax-regret (MR) criterion and shows that it always yields a diversified treatment allocation when the planner faces ambiguity. The MR criterion chooses an allocation that balances the potential welfare losses from Type A and Type B errors. This allocation turns out always to be fractional when the better treatment is not known.

In this section, I summarize the most basic parts of my analysis. I focus on a simple setting where treatment is known to be individualistic, welfare is a linear function of individual outcomes, and members of the population are observationally identical. This setting eliminates several possible reasons for differential treatment of a population.

Individualistic treatment means that the outcome experienced by a person may depend only on the treatment he receives, not on the treatments of other members of the population. This eliminates the external effectiveness that was essential to analysis of vaccination in Section 3. Welfare being a linear function of individual outcomes means that a planner who maximizes expected welfare is risk-neutral. This eliminates the traditional argument for diversification based on risk aversion.

Considering a population of observationally identical people eliminates the possibility of profiling; that is, systematic differentiation among persons who vary in observable respects. 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 (2007, Chapter 11). In contrast, diversification randomly differentiates among persons who are observationally identical.

Section 4.2 summarizes the basic analysis, which considers a one-period planning problems. Section 4.3 discusses the ethical issue of "equal treatment of equals" as it arises with diversified treatment. Section 4.4 extends the basic analysis to multi-period planning problems, where the planner may use observation of treatment outcomes in earlier periods to inform treatment choice in later periods. This yields a recommendation for *adaptive diversification*.

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

4.2.1. Concepts and Notation

Let there be 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 $y_j(\cdot)\colon T \rightarrow Y$ 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 $u_j(t) \equiv u_j[y(t), t]$ 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 eliminates sampling variation as an issue when considering diversified treatment choice.

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. 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 mean welfare if all persons were to receive treatment a or b respectively. Welfare with allocation δ is

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

$W(\cdot)$ is a consequentialist welfare function that additively aggregates individual contributions to welfare.

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 and the planner feels unable to place a subjective distribution on it.

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. The present analysis is applicable when $[(\alpha_s, \beta_s), s \in S]$ is bounded. Let the extreme feasible values of α and β be $\alpha_L \equiv \min_{s \in S} \alpha_s$, $\beta_L \equiv \min_{s \in S} \beta_s$, $\alpha_U \equiv \max_{s \in S} \alpha_s$, and $\beta_U \equiv \max_{s \in S} \beta_s$. Manski (2007) exposits many specific cases, showing the form that the region takes when observation of realized treatment outcomes in a study population is combined with various assumptions about treatment response and selection.

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.

Then all $\delta \in [0, 1]$ are undominated. I henceforth assume that the planner faces ambiguity.

4.2.2. The Expected Utility and the Maximin Criteria

There is no uniquely correct way to choose an undominated allocation. This section briefly discusses the choices made by a planner who uses the expected utility or the maximin criterion. The next section develops minimax-regret treatment choice.

A planner using the expected utility criterion places a subjective distribution π on set S, 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

$$(11)\qquad \max_{\delta \in [0, 1]} \quad E_\pi(\alpha) + [E_\pi(\beta) - E_\pi(\alpha)]\delta,$$

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 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 maximize expected utility if $E_\pi(\beta) = E_\pi(\alpha)$. Thus, the planner behaves as would one who knows that the population means in (10) have the values in (11).

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

$$(12)\qquad \max_{\delta \in [0, 1]} \quad \min_{s \in S} \quad \alpha_s + (\beta_s - \alpha_s)\delta.$$

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 $\alpha_L = \beta_L$.

4.2.3. The Minimax-Regret Criterion

By definition, the regret of 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 s is $\max(\alpha_s, \beta_s)$. Hence, 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

$$(13) \quad \min_{\delta \in [0, 1]} \ \max_{s \in S} \ \max(\alpha_s, \beta_s) - [\alpha_s + (\beta_s - \alpha_s)\delta].$$

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. Manski (2007, Complement 11A) proves that the MR criterion always makes a fractional treatment allocation when both treatments are undominated. The result is

$$(14) \quad \delta_{MR} = \frac{M(b)}{M(a) + M(b)}.$$

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

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

$$(15a) \quad R(\delta, a) \equiv \max_{s \in S(a)} \alpha_s - [(1-\delta)\alpha_s + \delta\beta_s] = \max_{s \in S(a)} \delta(\alpha_s - \beta_s) = \delta M(a),$$

$$(15b) \quad R(\delta, b) \equiv \max_{s \in S(b)} \beta_s - [(1-\delta)\alpha_s + \delta\beta_s] = \max_{s \in S(b)} (1-\delta)(\beta_s - \alpha_s) = (1-\delta)M(b),$$

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(·, a) increases linearly from 0 to M(a) and R(·, b) decreases linearly from M(b) to 0. Hence, the MR rule is the unique δ ∈ (0, 1) such that R(δ, a) = R(δ, b). This yields (14). □

The proof of (14) shows that the MR allocation balances the potential losses from the two types of error. 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 δ ∈ [0, 1], the quantities R(δ, b) and R(δ, 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(δ, a) = R(δ, b).

When a planner uses allocation δ_{MR}, maximum regret is $\delta_{MR}M(A) = M(a)M(b)/[M(a) + M(b)]$. 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,

$$(16) \quad \delta_{MR} = \frac{\beta_U - \alpha_L}{(\alpha_U - \beta_L) + (\beta_U - \alpha_L)}.$$

Result (16) simplifies further if either α or β is fully known. In particular, suppose that α is known. Then $\alpha_L = \alpha_U = \alpha$ and (16) becomes

$$(17) \quad \delta_{MR} = \frac{\beta_U - \alpha}{\beta_U - \beta_L}.$$

Full knowledge of α may be realistic if a is the status quo treatment and b is an innovation. Suppose, 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 expected-utility allocation. 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 utility. 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 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 $\delta_{MR} = (\beta_U - \alpha)/(\beta_U - \beta_L)$.

I caution the reader that analysis of the minimax-regret criterion when there are more than two feasible treatments is less transparent than with two treatments. Stoye (2007) has studied a class of planning problems with multiple qualitatively different 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.

4.2.4. Illustration: Choosing Sentences for Convicted Juvenile Offenders

To illustrate planning using the expected utility, maximin, and minimax-regret criteria, 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, this being judicial discretion to sentence an offender to residential confinement or to order 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 and suppose that the planner's objective is to minimize 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 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 partially 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 maximizes expected utility, 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 minimax-regret criterion, it applies (17). Thus, it subjects a randomly chosen fraction $(\beta_U - \alpha)/(\beta_U - \beta_L) = (0.92 - 0.61)/(0.92 - 0.03) = 0.35$ of offenders to mandatory confinement, leaving judicial discretion in place for the remaining fraction 0.65. The maximum regret of the MR allocation is $(0.35)(0.61 - 0.03) = 0.20$.

4.3. Diversification and "Equal Treatment of Equals"

The analysis in Section 4.2 maintained the traditional consequentialist assumption of welfare economics. That is, policy choices matter only for the outcomes they yield. In contrast, deontological ethics supposes that actions may have intrinsic value, apart from their consequences.

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.

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 (2007, 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_t is correct, all persons who receive treatment t survive and all others die. It is known that one of the two hypotheses is correct, but it is not known which one; thus, there are two states of natures, $s = H_a$ 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 δ = 0 and δ = 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. A planner applying the expected utility criterion makes a singleton allocation, allocating the entire population to the treatment with the higher subjective probability of success.

The maximin and minimax-regret allocations are both δ = ½. 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. □

Democratic societies ordinarily adhere to the ex post sense of equal treatment. However, 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 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.

Manski (2009) considers planning with deontological welfare functions, which enable a planner to formally express ethical objections to fractional treatment allocations for their violation of ex post equal treatment of equals. I characterize concern with ex post equal treatment as a fixed cost incurred when the planner chooses a fractional allocation. Posing a welfare function that generalizes (10) by incorporating this fixed cost, I show that the MR allocation remains the fraction given in (14) when the fixed cost is small, but is singleton if the fixed cost is sufficiently large.

4.4. Adaptive Diversification

I now consider a multi-period setting where, in each period, a planner must choose treatments for the current cohort of a population. An 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. Diversification of treatment is advantageous for learning because it generates randomized experiments yielding outcome data on both treatments.

Multi-period planning with learning raises subtle normative issues concerning inter-cohort equity. Earlier cohorts are inherently disadvantaged relative to later ones, as observation of earlier cohorts provides information useful to later ones but not vice versa. The inequity may be exacerbated if a planner purposefully treats earlier cohorts in a manner intended to enhance learning, for the benefit of later cohorts.

Supposing that utility is transferable across cohorts, utilitarian welfare economics suggests compensation of earlier cohorts for the positive externalities that observation of their outcomes provides to later cohorts. I will suppose that compensation is not possible and, hence, will not go down this path.

I instead discuss the *adaptive minimax-regret* (*AMR*) criterion. Each period, the AMR criterion diversifies treatment by applying the minimax-regret criterion using the information available at the time. The result is a fractional allocation whenever neither treatment is dominated. The AMR criterion is adaptive because successive cohorts may receive different allocations as knowledge of treatment response accumulates over time.

I do not claim optimality properties for the AMR criterion, but I feel that it has normative appeal. The criterion treats each cohort as well as possible, in the MR sense, given the knowledge available at the time. It does not ask the members of one cohort to sacrifice for the benefit of future cohorts. Nevertheless, the diversification of treatment performed for the benefit of the current cohort enables learning about treatment response. When the population is large, all fractional allocations generate the same learning. Hence, future cohorts benefit as much as possible from observation of the outcomes of earlier ones.

Section 4.4.1 formalizes the AMR criterion. Section 4.4.2 gives a numerical illustration showing how a centralized health planning system could use the criterion to choose treatments for a non-infectious disease. Section 4.4.3 discusses how the AMR criterion differs from the current practice of randomized experiments.

4.4.1. The Adaptive Minimax-Regret Criterion

To frame the multi-period planning problem we need to extend the concepts and notation used earlier. Let $n = 0, 1, \ldots, N$ 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 $\delta \equiv (\delta_n, n = 0, \ldots, N)$ 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 period n is $\delta_n = 1$ if $\beta \geq \alpha$ and $\delta_n = 0$ if $\beta \leq \alpha$. The planner faces ambiguity in period n if he does not know whether α is larger than β.

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

$$(18) \qquad \min_{\delta_n \in [0, 1]} \ \max_{s \in S_n} \ \max(\alpha_s, \beta_s) - [(1 - \delta_n)\alpha_s + \delta_n\beta_s].$$

The AMR allocation follows immediately from (14). 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 $M_n(a) \equiv \max_{s \in Sn(a)} (\alpha_s - \beta_s)$ and $M_n(b) \equiv \max_{s \in Sn(b)} (\beta_s - \alpha_s)$. Then

$$(19) \quad \delta_{nAMR} = \frac{M_n(b)}{M_n(a) + M_n(b)}.$$

4.4.2. Application to Centralized Health Care Systems

Here is a 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. Early on, the planner does not yet know how successful treatment will be. As time passes, empirical evidence accumulates and 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 (17) to each successive cohort, using the knowledge of β available at the time. Thus, the AMR decision at each n is

$$
\begin{aligned}
(20) \quad \delta_{AMR(n)} &= 0 && \text{if } \beta_{Un} < \alpha, \\
&= (\beta_{Un} - \alpha)/(\beta_{Un} - \beta_{Ln}) && \text{if } \beta_{Ln} \le \alpha \le \beta_{Un}, \\
&= 1 && \text{if } \beta_{Ln} > \alpha.
\end{aligned}
$$

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. Let the horizon be five years and let y(t) denote 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

$$(21) \qquad y_j(t) = \sum_{k=1}^{K} y_{jk}(t),$$

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.

[place Table 1 here]

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 $\delta_5 = 1$.

4.4.3. The AMR Criterion and the Practice of Randomized Experiments

The above illustration exemplifies a host of settings in which a 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 perform randomized experiments 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 randomized experiments. There are many major differences. I describe three here.

Fraction of the Population Receiving the Innovation

The AMR treatment allocation δ_{nAMR} can take any value in the interval [0, 1]. In contrast, the group receiving the innovation in current experiments is typically a very small fraction of the relevant population. 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 a drug trial is generally less than 0.01 and often less than 0.001.

Group Subject to Randomization

Under the AMR criterion as illustrated above, 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 as illustrated above, 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, the trials performed to obtain FDA approval of new drugs typically have durations of only two to three years. A three-year trial on the disease described in Table 1would only reveal that $\beta \in [2.64, 4.36]$.

Attempting to learn from trials of short duration, medical researchers often measure surrogate outcomes rather than 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. Medical researchers have cautioned that extrapolation from surrogate outcomes to outcomes of interest can be difficult; see Fleming and Demets (1996) and Psaty *et al.* (1999). Nevertheless, the practice has persisted.

Extrapolation from surrogate outcomes is similarly problematic in non-medical contexts. For example, preschool interventions are often evaluated using test performance in the early grades of school. However, the outcomes of real interest measure the long-term development of children into adults.

5. Treatment with a Convex Choice Set and a Common Concave Outcome Function

The analysis summarized in Section 4 provides a formal foundation for diversified treatment choice when the optimal treatment is not known. Nevertheless, I should caution the reader that diversified treatment is not always desirable. I have previously observed that fixed treatment costs may make singleton allocations preferable. In this section, I call attention to a class of problems of policy choice under ambiguity where all diversified allocations are dominated.

Sections 3 and 4 considered settings where each member of the population receives one of two treatments: vaccination or no vaccination in Section 3, treatment a or b in Section 4. Manski (2010b) studies a one-period planning problem with a different structure. Now the feasible treatments are a convex set and treatment response is given by a common concave function that maps the treatment and the state of nature into an outcome. The planner should not diversify treatment in this setting. Any fractional allocation is dominated by one that gives all members of the population the mean treatment.

The finding rests on a simple application of Jensen's inequality. Again let S list the feasible states of nature. Let X be the convex set of feasible treatments. Let $f(\cdot, \cdot): X \times S \rightarrow R$ be a known function that maps treatments and states into the real line. Suppose that, for each $s \in S$, $f(\cdot, s)$ is concave on X. Whereas I earlier permitted treatment response to vary across the population, the present analysis assumes that treatment response is homogenous. This is manifest in the notation, as $f(\cdot, \cdot)$ is the same for all members of the population.

Suppose that a planner can assign each agent any feasible treatment. Let $(x_j, j \in J)$ be a treatment allocation with mean $\mu_x \equiv \int x_j dP(j)$. As earlier, let the welfare of an allocation add up outcomes across the population. Thus, the welfare of allocation $(x_j, j \in J)$ in state s is $\int f(x_j, s)dP(j)$.

Jensen's inequality gives

$$(22) \quad f(\mu_x, s) \;\geq\; \int f(x_j, s)dP(j), \quad \text{all } s \in S.$$

Result (22) shows that, in each state, welfare when the planner assigns every agent treatment μ_x is at least as large as welfare with allocation $(x_j, j \in J)$. Thus, diversified treatment of the population is dominated by assigning the associated mean treatment to all persons.

Application to Medical Treatment

Medical treatment with partial knowledge of treatment response illustrates when diversification is and is not a reasonable strategy. Consider first an organ disease with two alternative treatments. One is surgery to repair the organ and the other is replacement of the organ with a transplant. Convex combinations of these treatments are not feasible—one can only repair or replace. In a setting of this sort, diversification warrants consideration when it is not known which treatment is better. Some fraction of patients would have the organ repaired and the remaining fraction would receive transplants. The minimax-regret criterion provides a coherent method to choose the fractions.

Now consider exercise as a treatment intended to increase life span. Here convex combinations of treatments are feasible—one can exercise in low, high, or intermediate intensities. Suppose that the objective function is concave and homogeneous across the relevant patient population, with diminishing marginal returns to higher intensity of exercise. Then a planner should not vary intensity across patients. Any diversified treatment strategy is dominated by one in which all patients exercise at the mean of the diversified intensities.

6. Conclusion

Optimal policy choice under ambiguity is not achievable, but reasonable choices based on coherent decision-theoretic principles are achievable. Planners should not seek to hide ambiguity behind untenable assumptions. They should face up to ambiguity when decisions must be made and seek to reduce it over time. This paper has described some principles for policy choice under ambiguity and has applied them to various settings.

An important general lesson is to first study dominance in order to eliminate clearly bad policies, and then study particular decision criteria to choose an undominated policy. Although all policies were undominated in the setting of Section 4, substantial subsets of policies were found to be dominated in the settings of Sections 3 and 5. Thus, study of dominance can pay off.

Another important general lesson is that there is no objectively correct way to choose an undominated policy. A planner might maximize subjective expected welfare, maximize minimum welfare, minimize maximum regret, or use another decision criterion. Research of the type described in this article cannot prescribe a "best" decision criterion. However, it can inform policy choice by characterizing the properties of various criteria in specific settings.

ACKNOWLEDGEMENTS

This research was supported in part by National Science Foundation grant SES-0911181. I am grateful to Gadi Barlevy and Tim Bresnahan for comments.

LITERATURE CITED

Ball F, Lyne O. 2002. Optimal Vaccination Policies for Stochastic Epidemics Among a Population of Households. *Mathematical Biosciences*. 177&178: 333–354

Barlevy G. 2010. Robustness and Macroeconomic Policy. *Annual Review of Economics*. In press

Becker N, Starczak D. 1997. Optimal Vaccination Strategies for a Community of Households. *Mathematical Biosciences*. 139: 117-132

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

Binmore K. 2009. *Rational Decisions*. Princeton: Princeton University Press

Brito D, Sheshinski E, Intriligator M. 1991. Externalities and Compulsory Vaccinations. *Journal of Public Economics*. 45: 69-90

Bossert W,Weymark J. 2008. Social Choice (New Developments). *The New Palgrave Dictionary of Economics, Second Edition*. London: Palgrave McMillan

Boulier B, Datta T, Goldfarb R. 2007. Vaccination Externalities. *The B.E. Journal of Economic Analysis & Policy*. 7: issue 1, article 23. http://www.bepress.com/bejeap/vol7/iss1/art23

Campbell D. 1984. Can We Be Scientific in Applied Social Science?. *Evaluation Studies Review Annual*. 9: 26-48

Campbell D, Stanley R. 1963. *Experimental and Quasi-Experimental Designs for Research*. Chicago: Rand McNally

Chernoff H. 1954. Rational Selection of Decision Functions. *Econometrica*. 22: 422-443

Dehehia, R. 2004. Program Evaluation as a Decision Problem. *Journal of Econometrics*. 125: 141-173

Ellsberg D. 1961. Risk, Ambiguity, and the Savage Axioms. *Quarterly Journal of Economics*. 75: 643-669

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

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

Friedman, J. 1986. *Game Theory with Applications to Economics*. Oxford: Oxford University Press.

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

Hill A, Longini I. 2003. The Critical Vaccination Fraction for Heterogeneous Epidemic Models. *Mathematical Biosciences*. 181: 85–106

Hirano K, Porter J. 2009. Asymptotics for Statistical Treatment Rules. *Econometrica*. 77: 1683-1701

Hodges J, Lehmann E. 1950. Some Problems in Minimax Point Estimation. *Annals of Mathemtical Statistics*. 21: 182-197

Keynes J. 1921. *A Treatise on Probability*. London: MacMillan

Knight F. 1921. *Risk, Uncertainty, and Profit*. Boston: Houghton-Mifflin

Manski C. 1997. Monotone Treatment Response. *Econometrica*. 65: 1311-1334

Manski C. 2000. Identification Problems and Decisions under Ambiguity: Empirical Analysis of Treatment Response and Normative Analysis of Treatment Choice. *Journal of Econometrics*. 95: 415–442

Manski C. 2004. Statistical Treatment Rules for Heterogeneous Populations. *Econometrica*. 72: 1221-1246

Manski C. 2006. Search Profiling with Partial Knowledge of Deterrence. *Economic Journal*. 116: F385-F401

Manski C. 2007. *Identification for Prediction and Decision*. Cambridge, MA: Harvard University Press

Manski C. 2009. Diversified Treatment under Ambiguity. *International Economic Review*. 50: 1013-1041

Manski C. 2010a. Vaccination with Partial Knowledge of External Effectiveness. *Proceedings of the National Academy of Sciences*. 107: 3953-3960

Manski C. 2010b. When Consensus Choice Dominates Individualism: Jensen's Inequality and Collective Decisions under Uncertainty. *Quantitative Economics*. In press

Manski C. 2010c. Actualist Rationality. *Theory and Decision*. In press

Manski C, Nagin D. 1998. Bounding Disagreements about Treatment Effects: a Case Study of Sentencing and Recidivism. *Sociological Methodology.* 28: 99–137

Manski C, Tetenov A. 2007. Admissible Treatment Rules for a Risk-averse Planner with Experimental Data on an Innovation. *Journal of Statistical Planning and Inference*. 137: 1998-2010

Meinert C. 1986. *Clinical Trials: Design, Conduct, and Analysis*. Oxford: Oxford University Press

Meltzer D. 2001. Addressing Uncertainty in Medical Cost-Effectiveness: Implications of Expected Utility Maximization for Methods to Perform Sensitivity Analysis and the Use of Cost-Effectiveness Analysis to Set Priorities for Medical Research. *Journal of Health Economics.* 20: 109-129

Mirrlees J. 1971. An Exploration in the Theory of Optimal Income Taxation. *Review of Economic Studies*. 38: 175-208

Nordhaus W. 2008. *A Question of Balance: Weighing the Options on Global Warming Policy*. New Haven: Yale University Press

Patel R, Longini I, Halloran E. 2005. Finding Optimal Vaccination Strategies for Pandemic Influenza Using Genetic Algorithms. *Journal of Theoretical Biology*. 234: 201-212

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

Scuffham P, West P. 2002. Economic Evaluation of Strategies for the Control and Management of Influenza in Europe. *Vaccine*. 20: 2562–2578

Sen A. 1987. Social Choice. *The New Palgrave: A Dictionary of Economics*. London: Palgrave McMillan

Stoye J. 2007. Minimax Regret Treatment Choice with Incomplete Data and Many Treatments. *Econometric Theory*. 23: 190-199

Stoye J. 2009. Minimax Regret Treatment Choice with Finite Samples. *Journal of Econometrics*. 151: 70-81

Tetenov A. 2009. Statistical Treatment Choice Based on Asymmetric Minimax Regret Criteria.

Collegio Carlo Alberto, Turin, Italy

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

Wald A. 1950. *Statistical Decision Functions*. New York: Wiley

Weycker D et al. 2005. Population-wide Benefits of Routine Vaccination of Children Against Influenza. *Vaccine*. 23: 1284-1293

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