

**DIVERSIFIED POLICY CHOICE WITH PARTIAL KNOWLEDGE OF POLICY
EFFECTIVENESS**

Charles F. Manski
Department of Economics and Institute for Policy Research
Northwestern University

March 2009

This paper was prepared for presentation at a Plenary Session at the 2009 Annual Conference of the Society for Research on Educational Effectiveness. The writing was supported in part by NSF Grant SES-0549544.

1. Introduction

Why perform research on educational effectiveness? An important objective is to provide information useful in choosing education policy. This is not the only aim that researchers give for their work—some declare that they seek to infer causal effects or advance the science of learning. Nevertheless, informing policy certainly is and should be a central goal.

Consider a policy maker, henceforth termed a *planner*, who must choose treatments for a heterogeneous population. The population may, for example, be the students enrolled in a particular school or district, or perhaps all youth in the nation. The planner may correspondingly be a school principal, a district superintendent, or the federal government. The treatments may be alternative allocations of students to classrooms in a school, alternative curricula for students in a district, or alternative federal financial aid programs for college students in the nation.

Economists and other policy analysts have long asked how a planner should act. A standard exercise specifies a set of feasible treatment policies and a welfare function. The planner is presumed to know how persons respond to treatment. The goal is to characterize the optimal policy.

The practical relevance of the standard exercise is limited by the fact that available research typically yields only partial knowledge of treatment response. Hence, planners cannot determine optimal policies. Instead, they must choose treatments under ambiguity.

Treatment under ambiguity is common in education. Researchers have long sought to learn how class size, teacher attributes, curriculum, and grading procedures affect learning. Yet progress has been slow and public debates about the effectiveness of alternative pedagogical approaches are frequent. Policy choices are routinely made with no clear understanding of which policy is best.

Why has research provided only part of the knowledge needed to choose optimal policies?

How can planners cope with ambiguity, making reasonable policy choices with the knowledge available? How can we reduce ambiguity, enabling better policy choices? This paper addresses these questions.

Section 2 explains why partial knowledge of treatment response is pervasive. Statistical imprecision and identification problems both play roles, but identification problems usually are the dominant difficulty. Some identification problems are fundamental, while others stem from conventional research practices that can be improved.

Section 3 considers policy choice under ambiguity. I develop the broad theme that diversified treatment is appealing when the optimal policy is unknown. I show how the minimax-regret criterion, a venerable idea in decision theory, can be applied to choose a diversified policy.

Section 4 shows how *adaptive diversification* can reduce ambiguity over time. I explain how adaptive diversification differs from the conventional practice of randomized experimentation. Section 5 concludes.

Aiming to make this paper accessible to a broad audience, I give a mainly verbal exposition of my ideas here. These ideas have developed over the past decade in my research program on planning under ambiguity, which grew out of my earlier work on partial identification of treatment effects (Manski, 1990, 1995). I first connected identification problems with decisions under ambiguity in Manski (2000, 2002). The work has since advanced through Manski (2005, 2006, 2007a, 2007b, 2009a, 2009b). A rigorous yet only mildly technical exposition of the research program is given in my text *Identification for Prediction and Decision* (Manski, 2007b), particularly Chapter 11. The idea of adaptive diversification is introduced in a more recent technical paper Manski (2009b), which also presents many other extensions of my earlier analysis.

Although I shall mainly discuss educational research and policy in this paper, it should be evident that the ideas apply far more generally. Consider treatment of disease. The planner may be a healthcare agency choosing medical treatments for a population of persons who are susceptible to the disease. Medical research often yields only partial knowledge of response to alternative treatments. Hence, treatments often must be chosen when it is not clear which treatment is best.

Or consider provision of public assistance to unemployed workers. The population may be a group of workers who have lost their jobs in a plant closure and the planner may be a state department of employment security. One treatment may be conventional unemployment insurance, which provides cash assistance during unemployment but leaves job search to the worker. Other treatments may assist workers in job search or retrain them for positions in new occupations. Researchers have strived to assess the costs and benefits of alternative treatments, but no consensus has emerged. Hence, states are not sure how best to help displaced workers re-enter the labor market.

As yet another example, consider a judge choosing sentences for convicted offenders. Here the planner is the judge, the population to be treated is the group of offenders, and the feasible treatments are the legally permissible sentences. A judge may evaluate a sentence in part by its effect on recidivism. However, criminologists have not succeeded in learning much about how sentencing affects future criminality. Hence, judges often are unsure what sentence is best.

2. Reasons for Partial Knowledge of Treatment Response

The informativeness of research on treatment response is limited by 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 other features or to other populations.

Most researchers and policy makers have at least some familiarity with statistical imprecision. Yet identification problems, which are the dominant reason for partial knowledge of treatment response, remain less well understood. Section 2.1 describes some important sources of identification problems. Section 2.2 considers how researchers have sought to deal with these problems in practice.

2.1. Sources of Identification Problems

I will discuss multiple sources of identification problems. I first briefly discuss sources that are highly important yet easy to explain. I then give longer expositions of other sources whose significance has been less well appreciated by the research and policy communities.

Unobservability of Counterfactual Outcomes

An obvious but fundamental source of identification problems is that outcomes can be observed only for treatments that persons have received. The counterfactual outcomes that members

of a study population would have experienced under other treatments are logically unobservable. For example, we may be able to observe the outcomes that young children experience following receipt of a preschool treatment, but we cannot observe the outcomes that these children would have experienced had they not received the treatment. Yet optimal choice of a preschool policy requires comparison of the outcomes that children would experience with and without the treatment.

Measurement Problems

Measurement problems enlarge the gap between the information that planners would like to have and the evidence that empirical studies of treatment response provide. In observational studies, survey respondents may refuse to answer or may respond inaccurately to questions about the treatments that they have received and the outcomes that they have experienced. Experimental subjects may drop out of trials before their outcomes are measured.

A serious measurement problem often occurs when studies of treatment response have short durations. We often want to learn long-term outcomes of treatments, but short studies reveal only immediate outcomes. For example, when considering preschool policy, society may want to know how policy affects adult outcomes including college enrollment, work experience, and criminality. However, studies of short duration can only measure outcomes that are observable when children are still young. Credible extrapolation from such *surrogate outcomes* to the long-term outcomes of interest can be highly challenging.

Differing Composition of the Study and Treatment Populations

Researchers often study populations whose composition differs substantially from the population to be treated. Much research downplays the importance of correspondence between these populations. Donald Campbell argued that studies of treatment effects should be judged primarily by their *internal validity* and only secondarily by their *external validity* (e.g., Campbell and Stanley, 1963; Campbell, 1984). By internal validity, Campbell meant the credibility of findings within the study population, whatever it may be. By external validity, he meant the credibility of extrapolating findings from the study population to another population of interest.

Campbell's assertion has been used to argue for the universal primacy of experimental research over observational studies, whatever the study population may be. The reason given is that properly executed randomized experiments have high internal validity. However, from the perspective of policy choice, it makes no sense to value one type of validity above the other. What matters is the informativeness of a study for policy making in a population of interest.

Considering observational studies, Rosenbaum (1999) has recommended that such studies aim to approximate the conditions of laboratory experiments. Rosenbaum, like Campbell, downplays the importance of having the study population be similar to the population of interest, writing (page 259): "Studies of samples that are representative of populations may be quite useful in describing those populations, but may be ill-suited to inferences about treatment effects."

In accord with Campbell and Rosenbaum, many researchers using observational data to analyze treatment response focus on easy-to-study populations that differ fundamentally from the populations that planners must treat. One common practice, recommended by Campbell and his

followers, has been to perform *regression-discontinuity* analysis, which studies treatment response within a typically small sub-population of persons who may credibly be interpreted as having been randomized into treatment. Another has been to report the *effect of treatment on the treated*, where “the treated” are the members of a study population who actually received a specified treatment.

From the perspective of treatment choice, analysis of treatment response in an easy-to-study population is sensible if treatment response is homogeneous. Then planners can be confident that research findings can be extrapolated to the populations they must treat. In human populations, however, homogeneity of treatment response may be the exception rather than the rule. Whether the context be educational or medical or social, it is reasonable to think that persons vary in their response to treatment. To the degree that treatment response is heterogeneous, a planner must take care when extrapolating research findings from a study population to a treatment population, as optimal treatments in the two may differ.

Differing Treatment Selection in the Study and Treatment Populations

A common policy choice takes a set of treatments as given and contemplates changing the process of treatment selection relative to that observed in a study population. I first discuss the well-known *selection problem* and then describe at greater length the less familiar *mixing problem*.

The Selection Problem: Most researchers recognize the selection problem in observational studies. Suppose that treatments were selected in a decentralized manner in the study population. The problem is to predict response to a policy mandating that all persons receive a specified treatment.

Then one must predict the counterfactual outcomes that would have been experienced by members of the study population who did not receive the specified treatment, had this treatment been mandated.

Many researchers recognize that a version of the selection problem arises in randomized experiments with partial compliance. When some experimental subjects do not comply with their assigned treatments, the experiment is equivalent to an observational study in which treatment assignment serves as an instrumental variable.

The Mixing Problem: I have found that few researchers recognize the mixing problem that arises when data from classical randomized experiments are used to predict the outcomes of policies that offer but do not mandate treatments. The mixing problem is the converse of the selection problem. The selection problem arises when treatment selection in the study population is decentralized, but treatment would be mandated under the policy of interest. Conversely, the mixing problem arises when treatment selection in the study population is mandated (as in a classical randomized experiment), but treatment would be decentralized under the policy of interest.

Consider, for example, choice among alternative preschool policies. Traditional randomized experiments with preschool interventions have sought to learn the outcomes that occur when members of treatment and control groups respectively enroll and do not enroll in a proposed program. Observation of the treatment/control group reveals outcomes if the program is mandated/unavailable. Suppose that the policy under consideration would make the program available, but not mandated. The experiment does not reveal outcomes under this policy. Outcomes depend on program participation and on the joint distribution of treatment response, quantities that

the experiment does not reveal.

I introduced and studied the mixing problem in Manski (1997). To illustrate findings on identification, I took the subjects in the Perry Preschool Project to be the study population. The high school graduation rate of the treatment group was 0.67 and that of the control group was 0.49. This evidence has been interpreted as revealing what graduation rates would occur if the preschool treatment were mandated or unavailable. However, the evidence does not reveal what graduation rate would occur under a policy that makes the treatment available but does not mandate it.

One might think that the graduation rate would necessarily fall between 0.49 and 0.67, with the exact rate depending on program participation. This conclusion is correct if one knows that receiving the preschool treatment can never make a child worse off. It is also correct if one knows that parents would randomly decide whether to enroll their children in preschool. Suppose, however, that one finds neither assumption credible. My formal analysis shows that, if one does not know how parents would behave or how children would respond to treatment, one can conclude only that the graduation rate would be between 0.16 and 1.

2.2. Dealing with Identification Problems

A traditional response of researchers to identification problems has been to make assumptions that, when combined with available data, point-identify quantities of interest. In the literature on analysis of treatment response, this is evident in the prevalence of articles reporting point estimates of average treatment effects.

The trouble with this is that powerful incentives often influence researchers to maintain

assumptions far stronger than they can persuasively defend, in order to draw strong conclusions. The scientific community rewards those who produce unambiguous findings. The public rewards those who offer unequivocal policy recommendations.

Research on treatment response has rested on a variety of strong assumptions that have much identifying power but little substantive foundation. Some prominent assumptions include these:

* Treatment selection in observational studies is random conditional on specified covariates. This assumption underlies propensity score and matching methods.

* Treatment selection and response in observational studies agree with a specified econometric model. This assumption underlies structural modeling approaches.

* Noncompliance in experiments is the same as it would be in actual interventions. This assumption underlies intention-to-treat analysis.

Researchers often claim that, even if the above assumptions are not strictly correct, they are adequate approximations. However, they rarely present any evidence or argument to buttress the claim.

Research findings based on untenable assumptions are not useful to a planner facing a policy-choice problem. The objective of a planner is to maximize *actual* social welfare, not the social welfare that would prevail if untenable assumptions were to hold. Researchers and policy makers should keep in mind what I have called (Manski, 2003):

The Law of Decreasing Credibility: The credibility of inference decreases with the strength of the assumptions maintained.

3. Policy Choice under Ambiguity

Henceforth I consider the situation of a planner who finds that credible assumptions do not suffice to determine an optimal policy. This planner must somehow cope with ambiguity. Section 3.1 develops the broad theme that diversified treatment is appealing when the optimal policy is unknown. Section 3.2 explains how the minimax-regret criterion can be used to choose a diversified policy.

3.1. Diversified Treatment Policies

Financial diversification has long been a familiar idea in discussions of portfolio allocation, where an investor allocates wealth across a set of investments. A portfolio is said to be diversified if the investor allocates positive fractions of wealth to different investments, rather than all to one investment.

The broad rationale for financial diversification is that it enables someone who is uncertain about the returns to investments to balance different potential errors. Suppose that there are two investments, say stocks and bonds. One type of error occurs if the investor allocates a dollar to stocks, but it turns out that bonds yield a higher return. Another type occurs if he allocates a dollar

to bonds, but it happens that stocks yield a higher return. Allocating all wealth to stocks entirely avoids errors of the second kind but may yield massive errors of the first kind, and vice versa for allocation of all wealth to bonds. Diversified allocations necessarily make errors, but limit the frequency with which they can occur. Hence, an investor may reasonably prefer a diversified allocation to committing all wealth to one investment.

Treatment diversification is a less familiar idea that resembles financial diversification. To introduce the idea, I describe a dire, hopefully hypothetical, scenario first considered in Manski (2007b, Chapter 11).

Treating X-Pox

Suppose that a new disease called x-pox is sweeping a community. It is impossible to avoid infection. If untreated, infected persons always die. Thus, the entire population will die in the absence of effective treatment.

Suppose that medical researchers propose two treatments. The researchers know that one treatment is effective, but they do not know which one. They know that administering both treatments in combination is fatal. Thus, a person will survive if and only if she is administered the effective treatment alone. There is no time to experiment to learn which treatment is effective. Everyone must be treated right away.

Suppose that a public health agency must decide how to treat the community. The agency can select one treatment and administer it to everyone. Then the entire population will either live or die. Or it can diversify, giving one treatment to some fraction of the community and the other

treatment to the remaining fraction. Then the survival rate will be one of the two chosen fractions. If half the population receives each treatment, the survival rate is certain to be fifty percent.

From the perspective of the health agency, the problem of treating x-pox has the same structure as portfolio allocation. The agency must allocate members of the community across the two treatments, which we may label A and B. The agency commits a Type A error if it administers treatment A to a person, but it turns out that B is the effective treatment. Similarly, a Type B error occurs if a person receives treatment B, but treatment A is effective. Giving everyone treatment A entirely avoids type B errors but may yield massive Type A errors, and vice versa for giving everyone treatment B. Diversified allocations necessarily make errors but reduce their potential frequency. If half the population receive each treatment, errors occur half the time.

Observe that an individual cannot diversify her own treatment for x-pox. Each person receives one treatment and either lives or dies. Yet the community can diversify by having positive fractions of the population receive each treatment. Thus, private diversification of treatment for x-pox is impossible, but communal diversification is possible.

What should the agency do? It could give everyone the same treatment and hope to make the right choice, recognizing the possibility that the outcome may be calamitous. Or it could give half the population each treatment, ensuring that half the community lives and half dies. One can reasonably argue for either alternative.

The argument for treatment diversification strengthens if everyone need not be treated at the same time. Suppose instead that infection occurs in two waves, with some members of the community becoming ill right away and the remainder later on. Then diversified treatment of the early cases is tantamount to performance of a randomized experiment. Observation of the outcomes

of early treatment reveals which treatment is effective, enabling the agency to treat all later cases effectively. Administering each treatment to half of the early cases ensures the survival of fifty percent of these persons and all those who become ill later.

Treating x-pox exemplifies a large class of medical, educational, and other policy decisions that call for a planner to choose treatments under ambiguity. The general rationales for treatment diversification are the same as in the x-pox illustration. First, diversification limits the frequency with which treatment errors can occur. Second, it performs a randomized experiment that yields evidence on treatment response. Thus, diversification copes with ambiguity in the short run and reduces ambiguity in the long run.

Equal Treatment of Equals

Treatment diversification as illustrated in the x-pox example calls for assigning different treatments to persons who are observationally identical. This differs from profiling, where one differentially treats persons who differ in observable characteristics. Differential treatment of persons with different attributes is often uncontroversial when there is reason to think that treatment response varies with the observed attributes. However, it is commonly thought that observationally identical persons should receive the same treatment. Why so?

A possible ethical objection to treatment diversification is that it violates the normative principle calling for “equal treatment of equals.” Diversification is consistent with this principle in the *ex ante* sense that all observationally identical people have the same probability of receiving a particular treatment. It violates the principle in the *ex post* sense that observationally identical

persons ultimately receive different treatments. Thus, equal treatment holds ex ante but not ex post.

The x-pox scenario dramatically illustrates the difference between the ex ante and ex post senses of equal treatment. Administering treatment A to the entire population provides equal treatment in the ex post sense. Moreover, it equalizes realized outcomes, as the entire population either survives or dies. Administering each treatment to half the population treats everyone equally ex ante, each person having a 50 percent chance of receiving each treatment. However, it does not treat people equally ex post. Nor do it equalize outcomes, as half the population lives and half dies.

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.

Further support for ex post unequal treatment is found in American federalism, which permits many laws and administrative practices to vary across the states. The American Progressive movement has long appreciated that federalism enables the states to experiment with new policy ideas, the lessons eventually benefitting the entire nation. A century ago, Theodore Roosevelt, in his introduction to McCarthy (1912), wrote this about the Progressive leader Robert La Follette: “Thanks to the movement for genuinely democratic popular government which Senator La Follette led to overwhelming victory in Wisconsin, that state has become literally a laboratory for wise experimental legislation aiming to secure the social and political betterment of the people as a whole.” Twenty years later, Justice Louis Brandeis, in his dissent to the 1932 Supreme Court case *New York State Ice Co. v. Liebmann* (285 U.S. 311), added what has become a famous remark on this theme: “It is one of the happy incidents of the federal system that a single courageous State may, if its citizens choose, serve as a laboratory; and try novel social and economic experiments without risk to the rest of the country.”

3.2. The Minimax-Regret Criterion

I now move beyond broad discussion of treatment diversification and consider how the idea can be implemented. For ease of exposition, I focus on choice between a status quo treatment and an innovation. It is often reasonable to suppose that the effectiveness of a status quo treatment is known from experience, but that less is known about the effectiveness of a newly proposed

treatment. Thus, I shall suppose that the locus of ambiguity is lack of knowledge of response to the innovation.

Let the status quo treatment be called A and the innovation be called B. Suppose that a planner must allocate a population between these treatments. When making choices of this type, it has been common to place the burden of proof on the innovation. That is, it has been common to administer the status quo treatment unless there is considerable evidence that the innovation is better. This practice places more weight on Type B errors (choosing the innovation when the status quo is better) than on Type A errors (choosing the status quo when the innovation is better). However, there is no intrinsic reason why a planner should weigh the two types of error asymmetrically.

Suppose instead that the planner gives equal weight to Type A and Type B errors and, consequently, wants to balance their potential welfare effects. Decision theorists have proposed various ways to formalize this idea. One prominent formalization, espoused in Bayesian decision theory, calls on the planner to place a subjective probability distribution on whatever aspects of treatment response are unknown and then maximize subjective expected welfare. This approach to decision making may or may not yield a diversified treatment allocation, depending on the specific form of the subjective distribution and the welfare function.

I will instead discuss the *minimax-regret* (MR) criterion, first proposed by Savage (1951). The MR criterion only calls on the planner to place lower and upper bounds on the welfare achievable by a treatment. It is very simple and always yields a diversified treatment allocation in the setting to be discussed here.

By definition, the regret of a treatment allocation is the loss that one suffers by choosing this allocation rather than the best allocation. One obviously would like to choose the best allocation,

in which case one would suffer no regret. However, a decision maker facing ambiguity does not know the best allocation. The minimax-regret criterion selects an allocation that minimizes the maximum regret that could potentially materialize.

Applying the MR criterion to choice between a status quo treatment and an innovation shows that it is best to have a positive fraction of the population receive each treatment. The fraction receiving the innovation balances its upside potential against its downside risk. The specific fraction of the population who should receive the innovation depends on what the planner knows about treatment response.

I sketch the derivation of the MR allocation and give the result here. More technically inclined readers who seek a formal exposition may read Manski (2007b, Chapter 11).

The Minimax-Regret Allocation

Consider treatment choice immediately after the innovation is discovered. Let $W(A)$ and $W(B)$ respectively denote the welfare that would be achieved by administering the status quo treatment or the innovation to all persons. Treatment A is better if $W(A)$ is larger than $W(B)$; treatment B is better otherwise.

Given that A is the status quo treatment, it is reasonable to suppose that the planner knows $W(A)$ from past experience. Treatment B is new, so the planner does not know $W(B)$. Suppose it only knows that $W(B)$ lie between certain lower and upper bounds, say $L(B)$ and $U(B)$. Here $L(B)$ expresses a worst case analysis of the welfare achieved by the innovation and $U(B)$ a best-case analysis. Suppose that the superior treatment is ambiguous. This is so if $L(B)$ is smaller than $W(A)$,

which in turn is smaller than $U(B)$.

It may be that treatment B is superior to A. If so, maximum regret equals $U(B) - W(A)$ times the fraction of the population who receive treatment A. Or treatment A may be superior to B, in which case maximum regret equals $W(A) - L(B)$ times the fraction who receive treatment B. Given this, it is easy to show that overall maximum regret is minimized by administering treatment B to this fraction of the population:

$$\frac{U(B) - W(A)}{U(B) - L(B)}.$$

The form of the minimax-regret allocation is remarkably simple and intuitive. Given a fixed value for $W(A)$, the fraction of the population who receive the innovation increases with $L(B)$ and $U(B)$. The fraction is positive when $U(B)$ is larger than $W(A)$ and is less than one when $L(B)$ is smaller than $W(A)$. Thus, the MR criterion yields a diversified treatment allocation whenever the better treatment is ambiguous.

4. Adaptive Diversification

4.1. The Adaptive Minimax-Regret Criterion

The argument for diversification, whether through use of the MR criterion or some other decision method, strengthens if a planner looks ahead to treatment of future cohorts of persons. Administration of an innovation to a randomly chosen group of persons can reveal how effective the

new treatment is.

The *adaptive minimax-regret (AMR)* criterion provides an appealing way to cope with ambiguity and reduce it over time. Each period, AMR diversifies treatment by applying the minimax-regret criterion using the information available at the time. Thus, it treats each cohort as well as possible, in the MR sense, given the available knowledge. It does not ask the members of one cohort to sacrifice for the benefit of future cohorts. Yet over time, it enables learning about treatment response.

Choice Between a Status Quo Middle-School Curriculum and an Innovation with a High-School Preparatory Program

Here is an illustration. Let treatment A be the status quo middle-school curriculum in a district. Let B be the status quo curriculum combined with a high-school preparatory program the summer before high-school entry. Let the outcome of interest be the number of grades of high school that a student successfully completes in the four years following completion of middle school. Let the welfare of a policy be measured by mean grade completion across all students.

Table 1 shows the AMR treatment allocation each year in a scenario where grade completion rates under the status quo treatment in years one through four are (0.8, 0.7, 0.6, 0.6) respectively, while those under the innovation are (0.9, 0.8, 0.7, 0.7). Suppose that the planner knows from experience that mean grade completion under the status quo curriculum is 2.7 years. However, the planner has no initial knowledge of grade completion under the innovation—mean grade completion could be anything from 0 to 4 years. Then the initial AMR treatment allocation assigns the fraction $(4 - 2.7)/(4 - 0) = 0.325$ of all students to the innovation.

A year later, the planner observes that 0.9 of the students receiving the innovation successfully complete their first year of high school. Then he can conclude that mean grade completion under the innovation is at least 0.9 and at most 3.9 years. Hence, the updated AMR allocation to the innovation is $(3.9 - 2.7)/(3.9 - 0.9) = 0.4$. The planner learns more from observing outcomes in years two through four, and further updates the AMR allocation to 0.5, 0.7, and 1 respectively.

Table 1: AMR Treatment with a High-School Preparatory Program					
cohort or year (n or k)	grade completion rate in k th year after treatment		bound on W(b), cohort n	AMR allocation, cohort n	mean grade completion, cohort n
	Status Quo	Innovation			
0			[0, 4]	0.325	2.83
1	0.80	0.90	[0.9, 3.9]	0.4	2.86
2	0.70	0.80	[1.7, 3.7]	0.5	2.90
3	0.60	0.70	[2.4, 3.4]	0.7	2.98
4	0.60	0.70	[3.1, 3.1]	1	3.10

4.2. 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 commission 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 current experimental practice. There are several major differences, described below.

Fraction of the Population Receiving the Innovation: As shown in the illustration, the AMR treatment allocation can take any value between zero and one. In contrast, the sample 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. The most famous of all education experiments, the Perry Preschool Project, had only about sixty children in its treatment group.

Group Subject to Randomization: Under the AMR criterion, the persons receiving the innovation are randomly drawn from the full population of policy interest. In contrast, present experiments typically draw subjects from sub-populations. For example, medical experiments draw subjects from pools of persons who volunteer to participate and educational ones draw subjects in particular schools or communities.

Measurement of Outcomes: Under the AMR criterion, one observes the outcomes of real interest as they unfold over time and one uses these data to inform subsequent treatment decisions. In contrast, present experiments often have short durations. Hence, researchers often measure surrogate outcomes rather than outcomes of 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. Preschool interventions may be evaluated using test performance in the early grades of school rather than long- term achievement.

Use of Empirical Evidence in Decision Making: Choosing a treatment allocation to minimize maximum regret differs sharply from the way that the findings of experiments 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. This decision procedure is institutionalized in the FDA drug-approval process.

The philosophy underlying hypothesis testing is remote from treatment choice. The classical practice of handling the null and alternative hypotheses asymmetrically, fixing the probability of a type I error and seeking to minimize the probability of a type II error, makes no sense from the perspective of treatment choice. Moreover, error probabilities at most measure the chance of choosing a sub-optimal rule; they do not measure the damage resulting from a sub-optimal choice. See Manski (2007b, Chapter 12) for further discussion.

5. Conclusion

Optimal policy choices with partial knowledge of policy effectiveness generally are not achievable. However, reasonable choices based on coherent decision-theoretic principles are achievable. Researchers and policy makers 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. Adaptive diversification of treatments enables a planner to cope with ambiguity in the short run and reduce it in the long run.

There are many potential applications for the ideas discussed in this paper, to educational policy and other domains. Nevertheless, the reader should be aware that diversified treatment is not always feasible. In some circumstances, society may wish to insist on ex post equal treatment of equals, giving up the advantages of diversification. Moreover, society faces collective threats ranging from strategic nuclear war to global warming to asteroid strikes. As long as humanity inhabits just one world, these are undiversifiable risks.

References

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

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

Manski, C. (1990), "Nonparametric Bounds on Treatment Effects," *American Economic Review Papers and Proceedings*, 80, 319-323.

Manski, C. (1995), "Identification Problems in the Social Sciences," Cambridge, Mass.: Harvard University Press.

Manski, C. (1997), "The Mixing Problem in Programme Evaluation," *Review of Economic Studies*, 64, 537-553.

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. (2002), "Treatment Choice Under Ambiguity Induced by Inferential Problems," *Journal of Statistical Planning and Inference*, 105, 67-82.

Manski, C. (2003), "Partial Identification of Probability Distributions," New York: Springer-Verlag.

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

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. (2009a), "Social Planning with Partial Knowledge of Social Interactions," *Journal of Institutional and Theoretical Economics*, 165, 146-158.

Manski, C. (2009b), "Diversified Treatment under Ambiguity," Department of Economics, Northwestern University.

McCarthy, C. (1912), *The Wisconsin Idea*, New York: Macmillan.

Rosenbaum, P. (1999), "Choice as an Alternative to Control in Observational Studies," *Statistical Science*, 14, 259-304.

Savage, L. (1951), "The Theory of Statistical Decision," *Journal of the American Statistical Association*, 46, 55-67.