Institute for Policy Research
Northwestern University
Working Paper Series

WP-07-08

Adaptive Partial Drug Approval

Charles F. Manski
Faculty Fellow, Institute for Policy Research
Professor of Economics
Northwestern University

Version: November 2007

DRAFT

Abstract

In the United States, the drug approval process of the Food and Drug Administration (FDA) is currently the main mechanism through which the government influences the production and dissemination of information on drug treatments. To obtain approval for a new drug, a pharmaceutical firm provides evidence on treatment response in randomized clinical trials that compare the new drug with an accepted treatment or a placebo. The FDA makes a binary approval decision after reviewing these trails' empirical findings. This paper brings welfare-economic and decision-theoretic thinking to bear on drug approval. Considering the matter from the minimax-regret perspective suggests an adaptive social planning process in which treatment with a new drug would vary— instead of being either fully allowed or denied as in current practice— as empirical evidence accumulates. The stronger the evidence on identified health outcomes, then the more the drug could be used. The adaptive process would improve on the current one by stimulating production of stronger information on treatment response and by reducing the welfare losses that arise from errors in approval decisions. Manski suggests a pragmatic version of the adaptive process that the FDA could implement.

This research was supported in part by National Science Foundation grant SES-0549544. I am grateful to David Meltzer, John Pepper, Bruce Spencer, and Alex Tetenov for comments. I have benefited from the opportunity to present this work in seminars at the University of Chicago and the University of Southern California. An earlier version of this paper was circulated under the title "Adaptive Minimax-Regret Treatment Choice, with Application to Drug Approval," NBER Working Paper W13312, August 2007.

1. Introduction

Partial knowledge of treatment response is the norm when choosing medical treatments. The various actors in the decision process— physicians, patients, insurers, and others—essentially never have perfect foresight. The most that decision makers can strive to know is the distribution of treatment response in certain settings. For example, when treating a life-threatening disease, one might strive to assess the probabilities of survival if a patient with specified characteristics were to receive alternative treatments.

The information about treatment response that is actually available to decision makers may combine experimental data from clinical trials and observational data from clinical practice. This information is a quasi-public good, so there is a prima facie case for some government role in its production and dissemination. The difficult policy question is to determine what this role should be.

In the United States, the drug approval process of the Food and Drug Administration (FDA) is presently the main mechanism through which the government influences the production and dissemination of information on drug treatments. The drug approval process determines whether a drug can legally be sold within the country. To obtain approval for a new drug, a pharmaceutical firm must provide to the FDA information on treatment response through the performance of randomized clinical trials that compare the new drug with an accepted treatment or a placebo. The FDA makes a binary (up or down) approval decision after reviewing the empirical findings of these trials.

This paper brings welfare-economic and decision-theoretic thinking to bear on drug approval. The result is a proposal to replace the present approval process with an adaptive process in which treatment with a new drug would vary smoothly as empirical evidence accumulates. The stronger the evidence on health outcomes of interest in the patient population of interest, the more that use of a new drug would occur. The adaptive process would improve on the present one by stimulating production of stronger information on treatment response and by reducing the welfare losses that arise from errors in approval decisions.

Section 2 provides background on the present FDA approval process. I observe that statistical

imprecision and identification problems may generate errors in approval decisions, with identification usually being the dominant concern. Important sources of identification problems are measurement of surrogate outcomes, experimentation on volunteers, noncompliance and attrition, and blinded treatment assignment.

Section 3 steps away from the present FDA process and considers how a hypothetical social planner might behave when a new drug is invented. At each point in time, the planner must choose how to allocate the current cohort of a heterogeneous patient population between the new drug and a status quo treatment. Over time, the planner wants to learn response to the new drug, in order to improve treatment of future cohorts. The identification problems discussed in Section 2 imply that the planner faces a joint task of learning and treatment choice under ambiguity.

I consider this task from the minimax-regret perspective and introduce the *adaptive minimax-regret* (*AMR*) treatment rule. The AMR rule treats each cohort as well as possible, in the minimax-regret sense, using the information available at the time. A generic result is that the minimax-regret treatment allocation is *fractional* whenever the available information does not suffice to determine whether a new treatment is superior or inferior to the status quo. That is, it is socially best to diversify treatment by assigning positive fractions of the cohort to both treatments. A consequence is that performance of randomized trials arises naturally from application of the AMR rule. The findings of these trials yield new information about treatment response, thereby improving treatment of future cohorts.

Section 4 discusses implementation of policies that embody features of the AMR rule. I first consider settings with centralized health care systems, where some public or private entity acts as a social planner and makes treatment choices for its patient population. Examples include the Veterans Health Administration (VA) in the United States, the National Health Service (NHS) in England, and private health maintenance organizations (HMOs). Such entities could implement close approximations to the AMR rule.

I then consider the decentralized American context within which the FDA functions. Full implementation of the AMR rule would require radical change in the American health care system, but the

FDA could embrace important features of the rule with relatively modest revision to the present drug approval process. I suggest a *quasi-AMR* process that (1) lengthens the duration of clinical trials to enable measurement of health outcomes of real interest and (2) permits limited sales of new drugs while clinical trials are underway, the limit varying with time as evidence accumulates.

Although this paper focuses on drug approval, the ideas developed here apply more broadly. They apply as much to FDA approval of medical devices as drugs. The AMR rule may also be applied to less regulated aspects of medical treatment, such as surgical procedures.

The ideas apply outside of medicine as well. Consider, for example, evaluation of educational interventions in early childhood. The outcomes of interest may be years of schooling completed by adulthood and job performance in adulthood. Not wanting to wait for these outcomes to unfold over time, researchers have often used performance in the early grades of school to judge the success of innovations, with binary implementation decisions in mind. Here, as with drug approval, it may be better to institute an adaptive process in which the scale of implementation of an intervention varies as evidence accumulates.

The analysis in this paper makes various simplifying assumptions. Perhaps the most weighty is that it takes the invention of a new drug as a predetermined event, without asking how the drug came to be. This perspective is appropriate when considering a social planner whose patient population is small relative to the potential market for new drugs. It may not be appropriate when considering such large actors as the FDA or the NHS. The drug approval process of the FDA and the treatment rule of the NHS may affect the innovation process that generate new drugs. It is natural to ask how implementation of the AMR rule by the NHS or of a quasi-AMR approval process by the FDA may affect the innovation process. This is an important question, but it may be a difficult one to answer. I do not address it here.

2. Errors in FDA Drug Approval due to Identification Problems

The present FDA process for drug approval begins with preclinical laboratory and animal testing of new compounds. Those that seem promising then go through three phases of clinical trials, in which the new drug is compared with an accepted treatment or placebo.[1] Phase 1 trials, which typically take about a year and are performed with twenty to eighty healthy volunteers, aim to determine the basic pharmacological action of the drug and the safety of different doses. Phase 2 trials, which usually take about two years and are performed with several hundred volunteers who are ill with a specific disease, give preliminary evidence on the effectiveness and short-term side effects of the drug. Phase 3 trials, which usually take about three years and are performed with several hundred to several thousand volunteers ill with the disease, give further evidence on effectiveness and side effects. Following completion of Phase 3, the firm files a New Drug Application and the FDA either approves or disapproves the drug.[2]

FDA evaluation of new drugs occurs with partial knowledge of treatment response. As a consequence, drug approval decision are susceptible to two types of errors. Type I errors occur when new drugs that actually are inferior to accepted treatments are approved because they appear superior when evaluated using the available information. Type II errors occur when new drugs that actually are superior to accepted treatments are disapproved because they appear inferior when evaluated using the available

[1] The circumstances in which the comparison should be to an accepted treatment or placebo are a matter of some controversy. See Rothman and Michels (1994).

[2] See www.fda.gov/cder/handbook/ for an FDA description of the drug approval process. Although The FDA was created over a century ago in the Pure Food and Drug Act of 1906, the present drug approval process is a much more recent invention. From 1906 to 1938, the agency was unable to disapprove the sale of purported medicines. It only was able to outlaw labeling and other advertising that made unsupported claims of treatment safety and effectiveness. The Food, Drug, and Cosmetics Act (FDCA) of 1938 gave the FDA power to prohibit the sale of unsafe drugs, but without a requirement to assess effectiveness. The 1962 Amendments to the FDCA established the modern drug approval process, which requires pharmaceutical firms to demonstrate that new drugs are safe and effective through a series of randomized clinical trials. See Peltzman (1973) and Temin (1980) for further discussion.

information; see Viscusi, Harrington, and Vernon (2005, Chapter 22) for a textbook discussion. Some Type I errors eventually are corrected after approval through the FDA's post-market surveillance program, which analyzes data on the outcomes experienced when the drug is used in clinical practice.[3] Type II errors commonly are permanent because, after a drug is disapproved, use of the drug ceases and no further data on treatment response are produced.

A well-recognized potential source of errors in drug approval is the statistical imprecision of empirical findings from clinical trials with finite samples of subjects. The FDA limits the frequency of statistical errors by requiring that sample sizes suffice to perform conventional hypothesis tests with specified power, and by using the results of these tests to make approval decisions. There is much reason to question the logic of using hypothesis tests as a criterion for drug approval.[4] Nevertheless, conventional power calculations do in practice ensure that sample sizes are large enough to make statistical error a relatively minor concern.

The dominant determinants of errors in drug approval are a host of identification problems.

---

[3] See www.fda.gov/cder/regulatory/applications/postmarketing/surveillancepost.htm for an FDA description of this program. At present, the post-market surveillance program only aims to detect adverse side effects of approved drugs, not to better assess their effectiveness in treating the conditions for which they are intended. The data available for post-market surveillance is limited by the fact that FDA cannot compel a firm to perform new clinical trials after a drug has been approved. This being the case, the main instrument of post-market surveillance is the Adverse Event Reporting System, which encourages patients and physicians to submit reports of adverse side effects related to drug administration. A recent Institute of Medicine study makes recommendations for strengthening the post-market surveillance program. See Committee on the Assessment of the US Drug Safety System (2006).

[4] Approval of a new drug normally requires one-sided rejection of the null hypothesis of zero average treatment effect when a new drug is compared with an accepted treatment or placebo (Fisher and Moyé, 1999). This sets a high bar for approval, requiring that pharmaceutical firms demonstrate "substantial evidence of effect" for their products (Gould, 2002).

The use of hypothesis tests in drug approval is difficult to motivate from the perspective of treatment choice. First, there is no decision-theoretic rationale for the standard practice of handling the null and alternative hypotheses asymmetrically, fixing the probability of a type I statistical error and seeking to minimize the probability of a type II error. One should instead handle the two errors symmetrically. Second, error probabilities only measure the chance of making an approval error. They do not measure the loss in welfare resulting from an error.

Identification problems are the inferential difficulties that would persist even if statistical imprecision were eliminated by letting the sample sizes in clinical trials go to infinity; see Manski (1995, 2007a) for exposition and analysis of many such problems. The clinical trials used in the FDA drug approval process suffer from identification problems because, even if their sample sizes were to grow without bound, these trials would not reveal the distribution of treatment response in relevant patient populations. Some important reasons are

Measurement of Surrogate Outcomes: The clinical trials used to support New Drug Applications have relatively short durations. When trials are not long enough to observe the health outcomes of real interest, the practice is to measure so-called surrogate outcomes and base drug approval decisions on their values. For example, treatments for heart disease may be evaluated using data on patient cholesterol levels and blood pressure rather than data on heart attacks and life span. In such cases, which occur regularly, the clinical trials used in drug approval at most reveal the distribution of surrogate outcomes in the patient population, not the distribution of outcomes of real health interest.

Experimentation on Volunteers: Participation in experiments ordinarily cannot be mandated in democracies. Hence, the clinical trials performed for drug approval draw subjects at random 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.

Noncompliance and Attrition: Within the sample who agree to participate in a trial, some subjects may not comply with their assigned treatments or may leave the trial early, before their outcomes can be measured. Hence, a trial at most reveals the distribution of response to an assigned treatment within the sub-population

of subjects who comply with treatment assignment and who do not leave the trial early.[5]

Blinded Treatment Assignment: Blinded treatment assignment has been the norm in clinical trials of new drugs. Hence, a trial at most reveals the distribution of response in a setting where patients and physicians are uncertain what treatment has been assigned. It does not reveal the distribution of response in a real clinical setting where patients and physicians would know the assigned treatment.

Although the FDA has adopted formal procedures to limit the frequency of statistical errors in the drug approval process, it has not adopted procedures to cope coherently with identification problems. Instead, the approval process informally extrapolates from available trial data to the distribution of treatment response in real clinical settings. This practice has been criticized, but it persists.

Extrapolation from Surrogate Outcomes

Health researchers have called particular attention to the difficulty of extrapolating from surrogate outcomes to health outcomes of interest. Fleming and Demets (1996), who review the prevalent use of surrogate outcomes in Phase 3 trials evaluating drug treatments for heart disease, cancer, HIV/AIDS, osteoporosis, and other diseases, write (p. 605):

> "Surrogate end points are rarely, if ever, adequate substitutes for the definitive clinical outcome in phase 3 trials."

Sculpher and Claxton (2005), who consider decisions about whether new drugs are sufficiently cost-effective

[5] Much medical research seeks to circumvent the issue of noncompliance through performance of *intention-to-treat* analysis, which studies the outcomes associated with assigned rather than received treatments. Noncompliance is logically impossible in intention-to-treat analysis. This form of analysis is well-motivated if it is credible to assume that noncompliance patterns in real clinical settings will be similar to those observed in the available trial data. If this assumption is not credible, intention-to-treat analysis can be misleading for prediction of treatment response in real clinical settings.

for reimbursement in collectively funded health-care systems, write (p. 441):

> "Arguably the biggest challenge that reimbursement agencies have to face in terms of the uncertainty surrounding existing evidence relates to costs and outcomes which have not been observed directly in trials. There are two frequent manifestations of this: linking intermediate outcomes to ultimate measures of health gain, and extrapolating costs and benefits over a longer-term time horizon."

The obvious solution is to perform clinical trials of sufficient length to measure the health outcomes of real interest. However, this has been thought politically infeasible. Pasty *et al.* (1999) write (p. 789):

> "One systematic approach is a requirement that, prior to their approval, new drug therapies for cardiovascular risk factors should be evaluated in large, long-term clinical trials to assess their effects on major disease end points. The use of surrogate outcomes is avoided, and the major health outcomes are known prior to marketing. Such an approach would slow the time to drug approval and may meet with resistance from pharmaceutical manufacturers."

Indeed, pharmaceutical firms eager for returns on investments and patient groups wanting access to new drugs have often advocated shortening rather than lengthening the present time to approval.[6]

3. Treatment Choice by A Social Planner

To simplify analysis of complex collective decision problems, economists often find it useful to consider how a hypothetical social planner would behave. A standard exercise begins by specifying a set

[6] In 1992, pressure for quicker approval decisions led to passage of the Prescription Drug User Fee Act. This legislation shortened the time that the FDA takes to review New Drug Applications, the expedited review being funded by user fees assessed on the firms seeking approval. The Act did not shorten the duration of the clinical trials performed in support of New Drug Applications.

of policy alternatives and a social welfare function. The planner is presumed to know the welfare achieved by each policy alternative. With this knowledge, one can derive the optimal policy. The Mirrlees (1971) study of optimal income taxation is a leading example.

In this vein, I consider how a planner might behave when a new drug is invented. However, the informational character of this planning problem differs fundamentally from that assumed in the standard exercise. The standard approach would be to presume complete knowledge of the distribution of patient response to the new drug. But the essential feature of the FDA decision problem is that the agency begins with partial knowledge of treatment response and uses the drug approval process to induce production of information. If consideration of planner behavior is to carry lessons for drug approval in practice, we should presume that the planner begins with partial knowledge of treatment response.

How might a planner behave with partial knowledge of treatment response? He might, in the Bayesian manner, assert a subjective probability distribution over the unknown distribution of treatment response and maximize subjective expected welfare. See Meltzer (2001) for applications to medical decision making. However, a subjective probability distribution is itself a form of knowledge, and the planner may have no good basis for asserting one. How then might he behave? I shall study adaptive application of the *minimax-regret (MR)* criterion, a general principle for decision making with partial knowledge that was first suggested by Savage (1951).

Section 3.1 describes the minimax-regret criterion in general terms, gives an illustration, and explains its adaptive extension to multi-period planning problems. Section 3.2 formalizes the criterion in the new-drug context and derives the adaptive minimax-regret treatment allocation in a relatively simple setting. Section 3.3 gives numerical illustrations. Section 3.4 extends the analysis to some more complex settings.

3.1. The Adaptive Minimax-Regret Rule

First consider a one-period setting, where a planner must allocate a single cohort of patients between two treatments. In this setting, the minimax-regret rule chooses a treatment allocation that minimizes the maximum loss to welfare resulting from not having complete knowledge of the distribution of treatment response. Specifically, the *regret* of a treatment allocation under a possible distribution of treatment response is the difference between the maximum welfare that would be achievable given complete knowledge and the welfare that is achieved by this allocation. Given complete knowledge, the best decision obviously would be to choose an allocation that minimizes regret, setting it equal to zero. In the absence of complete knowledge, the MR rule chooses a treatment allocation that minimizes maximum regret across all possible distributions of treatment response.

This one-period planning problem has been studied previously in Manski (2005, 2007a, 2007b) and Stoye (2007a). A general finding is that when there are two treatments and the available knowledge of treatment response does not suffice to determine which treatment is better, the MR rule does not assign all observationally identical persons to the same treatment. Instead, it fractionally allocates these persons across the two treatments, the fraction receiving each treatment being determined by the available knowledge.

With its fractional treatment allocation, the minimax-regret criterion enables a planner to socially diversify risks that are privately indivisible. A dramatic illustration occurs in this hypothetical problem of treatment choice considered in Manski (2007a, Section 11.7).

Choosing Treatments for X-Pox: Suppose that a new viral disease called x-pox is sweeping the world. Researchers have proposed two mutually exclusive treatments, say a and b, which reflect alternative hypotheses, say H_a and H_b, about the nature of the virus. If H_a (H_b) is correct, all persons who receive treatment a (b) survive and all others die. A planner knows that one of the two hypotheses is correct, but

does not know which one. The objective is to maximize the survival rate of the population.

In this setting, the risk of death is privately indivisible. An individual receives either treatment a or b, and this person either lives or dies. Yet society can diversify by having positive fractions of the population receive each treatment. Consider the rule in which a fraction $\delta \in [0, 1]$ of the population receives treatment b and the remaining $1 - \delta$ receives treatment a. Then the fraction who survive is either δ or $1 - \delta$. A planner who uses the MR criterion would set $\delta = 0.5$, implying that half of the population survives and half dies. □

Social diversification is a central qualitative feature of minimax-regret choice between two treatments.[7] It is not a general feature of Bayesian decision making. For example, in the x-pox illustration, a Bayesian planner allocates the entire population to the treatment with the higher subjective probability of success.[8]

Now consider a multi-period setting, where the planner must allocate a sequence of cohorts of patients between two treatments. In contrast to the one-period planning problem, there now is an opportunity for learning, with the observed outcomes of treatments assigned in early periods being used to inform treatment choice in later periods. The adaptive minimax-regret criterion applies to each cohort the minimax-regret criterion using the knowledge of treatment response available at the time of treatment. The result is a fractional treatment allocation whenever the available knowledge does not suffice to determine which treatment is better. The rule is adaptive because knowledge of treatment response accumulates over time, so successive cohorts may receive different fractional allocations. Eventually, the planner may learn which treatment is better. From this point on, he assigns new cohorts entirely to the better treatment.

---

[7] The situation is more subtle when there are more than two treatments. Then the MR criterion typically yields a fractional treatment allocation, but in some circumstances it assigns everyone to the same treatment. See Stoye (2007a).

[8] A Bayesian planner does select a fractional treatment allocation if social welfare is a strictly concave function of the population survival rate. Let λ denote the planner's subjective probability that hypothesis H_b is correct. Let $s(\delta)$ denote the survival rate with allocation δ. If social welfare is $\log[s(\delta)]$, it can be shown that a Bayesian planner sets $\delta = \lambda$. See Keeney (1980) for further Bayesian analysis of "catastrophe" problems similar to the x-pox example.

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

The adaptive minimax-regret criterion shares the broad objective of using observed treatment responses to inform subsequent treatment choices. However, the AMR criterion is intended for application within the entire patient population rather than only within the sample drawn into a traditional clinical trial. Morever, the AMR criterion is intended to cope with identification problems. In contrast, proposals for adaptive trials have been concerned only with statistical imprecision.

3.2. Allocating a Sequence of Cohorts to a Status Quo Treatment and an Innovation

Basic Concepts, Notation, and Assumptions

I now formalize the above discussion. To begin, there are two treatments for a condition, labeled a and b. Treatment a is the *status quo*, a pre-existing accepted treatment whose properties are known from historical experience. Treatment b is an *innovation*, whose properties are not known initially. In particular, the innovation may be a new drug treatment.

A cohort of patients appears each period and requires treatment. The periods are labeled n = 0, 1, ... , N. Here n = 0 is the first period in which the innovation is available. In each period $n \geq 0$, the set of feasible treatments is T = {a, b}.

Each member j of cohort n, denoted J_n, has a response function $y_j(\cdot)$: $T \to Y$ mapping treatments $t \in T$ into outcomes $y_j(t) \in Y$. Subscripting $y_j(\cdot)$ by j indicates that treatment response may vary across the

members of the cohort. Whereas the members of cohort n are assigned treatments in period n, their outcomes unfold over the subsequent K periods. In particular, $y_j(t)$ might have the time-additive form

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

where $y_{jk}(t)$ is the component of the outcome that is realized k periods after person j receives treatment.

For example, the treatments may be alternative cancer therapies and the outcome may be life span. Then a period may be a year, with K being a specified horizon of interest. Here $y_{jk}(t) = 1$ if person j would, in the event of receiving treatment t, be alive k years after treatment, and $y_{jk}(t) = 0$ otherwise.

Let $P[y(\cdot)]$ denote the distribution of treatment response across a cohort. Observe that I do not index P by n. Thus, all cohorts share the same distribution of treatment response. This assumption enables social learning. I assume that the planner observes outcomes as they unfold. Hence, he can use data on outcomes in early cohorts to inform treatment choice for later cohorts.

I maintain two further simplifying assumptions. First, the members of the population are observationally identical. In practice, persons may have observable covariates, and a planner may be able to differentially treat persons with different covariates. In such cases, the present analysis can be applied separately to each sub-population of observationally identical persons.

Second, each cohort J_n is large in the formal sense of being atomless; that is, $P(j) = 0$ for all $j \in J_n$, This idealization implies that whenever the planner implements a fractional treatment allocation, the sub-population of persons who are assigned each treatment is infinite. This formally eliminates statistical imprecision as an issue in inference on treatment response.

The Treatment Choice Problem

The planner's problem is to allocate each cohort between the two treatments. A treatment rule 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 feasible treatment rules are the elements of the hyper-rectangle $[0, 1]^{(N+1)}$.

Let $u(t) \equiv u[y(t), t]$ denote the net contribution to social welfare that occurs when a person receives treatment t and realizes outcome y(t). I assume that the planner wants to choose a treatment rule that maximizes mean welfare summed across cohorts. Let $\alpha \equiv E[u(a)]$ and $\beta \equiv E[u(b)]$ be the mean welfare that would result if all members of a cohort were to receive treatment a or b respectively. The quantities α and β are not indexed by n because, by assumption, the distribution of treatment response is the same for all cohorts. The social welfare achieved by rule δ is

$$(2) \quad W(\delta) \equiv \sum_{n=0}^{N} \beta\delta_n + \alpha(1 - \delta_n) = (N + 1)\alpha + (\beta - \alpha) \sum_{n=0}^{N} \delta_n.$$

W(·) is an ordinary consequentialist social welfare function that aggregates individual contributions to welfare in an additive manner. A notable special case occurs when the function u(·) expresses private preferences. Then W(·) is the utilitarian social welfare function that weights all cohorts equally. A slightly broader definition of W(·) would permit the social welfare function to differentially weight cohorts that vary in size or to express time discounting. The present analysis extends easily to such cases.

The optimal treatment rule is obvious if (α, β) are known. The planner should choose $\delta_n = 1$ for all n if $\beta > \alpha$ and $\delta_n = 0$ if $\beta < \alpha$. All values of δ yield the same welfare if $\beta = \alpha$. The problem of interest is treatment choice when (α, β) is only partially known. In particular, I shall consider situations in which α is known but β is not.

It is often reasonable to suppose that α is known from historical experience. All members of cohorts $n < 0$ received the status quo treatment. Hence, a planner can learn α empirically if he is able to observe the outcomes experienced by cohort −K or an earlier cohort.

The innovation having been introduced at period 0, empirical evidence cannot reveal the value of β before period K. When the planner treats a cohort $n < K$, he can only observe the components of the outcomes experienced to date by members of earlier cohorts who were assigned the innovation. Thus, at $n = 0$, the planner has no empirical evidence. At $n = 1$ he can observe first-period outcomes for those members of cohort 0 who were assigned the innovation. At $n = 2$, he can observe second-period outcomes for members of cohort 0 who were assigned the innovation, as well as first-period outcomes for members of cohort 1. And so on.

Minimax-Regret Treatment Choice in a One-Period Planning Problem

First consider a one-period planning problem; thus, $N = 0$. If the planner assigns a fraction δ_0 of cohort 0 to the innovation and the remainder to the status quo, social welfare is

$$(3) \quad W(\delta_0) = \alpha + (\beta - \alpha)\delta_0.$$

The problem is to choose δ_0 in the absence of empirical evidence on β.

Application of the MR criterion requires only that the planner be able to place β within some bounded interval $[\beta_L, \beta_U]$. Such an interval always exists when the outcome is itself bounded. Then, if the planner knows nothing about the innovation, he can set β_L and β_U equal to the smallest and largest logically possible outcome values. Or $[\beta_L, \beta_U]$ may be a subset of the logically possible outcomes, excluding values the planner deems infeasible.

The MR rule is a function of α, β_L, and β_U. It assumes nothing about the position of β within the interval $[\beta_L, \beta_U]$. This contrasts with Bayesian planning, which requires assertion of a subjective probability distribution on the interval of feasible values.

By definition, regret is the difference between the maximum achievable welfare and the welfare

achieved with a specified treatment rule. The maximum achievable welfare is max (α, β). Hence, the regret of allocation δ_0 is max (α, β) − [α + (β − α)δ_0]. Regret depends on the unknown value of β. The MR rule computes maximum regret over all feasible values of β and chooses a treatment allocation to minimize maximum regret. Thus, the MR criterion is

$$(4) \quad \min_{\delta_0 \in [0,\,1]} \max_{\beta \in [\beta_L,\,\beta_U]} \max\,(\alpha, \beta) - [\alpha + (\beta - \alpha)\delta_0].$$

It is easy to see that the MR decision, denoted δ_{MR}, is $\delta_{MR} = 0$ if $\beta_U < \alpha$ and $\delta_{MR} = 1$ if $\beta_L > \alpha$. In the former (latter) case, the planner knows that the innovation is worse (better) than the status quo. Our concern is with situations where the planner does not know which treatment is better; that is, where $\beta_L \le \alpha \le \beta_U$. Manski (2007a, Section 11.3) shows that the MR decision then is

$$(5) \quad \delta_{MR} = (\beta_U - \alpha)/(\beta_U - \beta_L).$$

The proof is simple, so I reproduce it here.

Proof: Maximum regret across the feasible values of β is

$$\max_{\beta \in [\beta_L,\,\beta_U]} (\alpha - \beta)\delta_0 \cdot 1[\beta < \alpha] + (\beta - \alpha)(1 - \delta_0) \cdot 1[\beta > \alpha] = \max\,[(\alpha - \beta_L)\delta_0,\, (\beta_U - \alpha)(1 - \delta_0)].$$

Thus, the MR rule solves the optimization problem

$$\min_{\delta_0 \in [0,\,1]} \max\,[(\alpha - \beta_L)\delta_0,\, (\beta_U - \alpha)(1 - \delta_0)].$$

The quantity $(\alpha - \beta_L)\delta_0$ is increasing in δ_0, whereas $(\beta_U - \alpha)(1 - \delta_0)$ is decreasing in δ_0. The MR allocation is obtained by choosing δ_0 to equalize these two quantities. This gives (5). □

Observe that the MR rule yields a fractional allocation when $\beta_L < \alpha < \beta_U$. The fraction of the cohort assigned to the innovation depends on the location of α within the interval $[\beta_L, \beta_U]$, with δ_{MR} increasing linearly from 0 to 1 as α decreases from β_U to β_L. This behavior is sensible. As α decreases from β_U to β_L, the potential gain from choosing the innovation rises and the potential loss falls.

The value of maximum regret achieved by the MR decision depends on the values of α, β_L, and β_U. The proof of (5) shows that the maximum regret of δ_{MR} is $(\beta_U - \alpha)(\alpha - \beta_L)/(\beta_U - \beta_L)$. Ceteris paribus, this expression increases with β_U and decreases with β_L. Thus, maximum regret increases with the extent of ambiguity about the value of β.

Observe that leading alternatives to the MR rule, including the maximin rule and Bayes rules, generically do not deliver fractional treatment allocations when applied to this treatment-choice problem. The maximin criterion chooses an allocation that maximizes welfare when β take its lowest feasible value. Thus, the maximin criterion is

$$(6) \qquad \max_{\delta_0 \in [0,\,1]} \; \min_{\beta \in [\beta_L,\,\beta_U]} \; [\alpha + (\beta - \alpha)\delta_0].$$

Solution of this problem yields $\delta_0 = 0$ if $\beta_L < \alpha$ and $\delta_0 = 1$ if $\beta_L > \alpha$.

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

$$(7) \qquad \max_{\delta_0 \in [0,\,1]} \; \alpha + [E_\pi(\beta) - \alpha]\delta_0,$$

where π is the subjective distribution on β and $E_\pi(\beta) = \int \beta d\pi$ is its subjective mean. Solution of this problem yields $\delta_0 = 0$ if $E_\pi(\beta) < \alpha$ and $\delta_0 = 1$ if $E_\pi(\beta) > \alpha$. [9]

The Multi-Period Extension

Now let $N > 0$. Extending the notation introduced above, suppose that at period n, the planner knows that β lies in a bounded interval $[\beta_{Ln}, \beta_{Un}]$. This interval typically will shrink with n, as empirical evidence accumulates on the outcomes experienced by members of earlier cohorts who were treated with the innovation. Section 3.3 gives illustrations.

The adaptive minimax-regret rule applies the MR rule to each successive cohort, using the knowledge of β available at the time. Thus, the AMR decision at each n is

$$\begin{aligned} (8)\quad \delta_{AMR(n)} &= (\beta_{Un} - \alpha)/(\beta_{Un} - \beta_{Ln}) && \text{if } \beta_{Ln} \le \alpha \le \beta_{Un}, \\ &= 0 && \text{if } \beta_{Un} < \alpha, \\ &= 1 && \text{if } \beta_{Ln} > \alpha. \end{aligned}$$

Given allocation $\delta_{AMR(n)}$, maximum regret for cohort n is $\max[0, (\beta_{Un} - \alpha)(\alpha - \beta_{Ln})/(\beta_{Un} - \beta_{Ln})]$ and mean welfare is $\alpha + (\beta - \alpha)\delta_{AMR(n)}$. Maximum regret is computed using the information available at the time of treatment. Mean welfare depends on the value of β, which is not known before period K.

The AMR rule achieves the dual objectives of social learning and diversification. Inspection of (8)

---

[9] Manski and Tetenov (2007, Proposition 5) show that a Bayesian planner may make a fractional treatment allocation if the social welfare function is changed to $f[\alpha + (\beta - \alpha)\delta_0]$, where $f(\cdot)$ is monotone and continuously differentiable. Then the Bayes problem is

$$\max_{\delta_0 \in [0,\,1]} \int f[\alpha + (\beta - \alpha)\delta_0] d\pi.$$

Solutions to this problem are in the interior of the unit interval if $E_\pi(\beta) > \alpha$ and $\int f(\beta) d\pi < f(\alpha)$. In particular, this occurs if the function $f(\cdot)$ is sufficiently concave.

shows that the necessary and sufficient condition for learning to occur is $\beta_{U0} > \alpha$. This condition is sufficient for learning because it implies that the planner assigns a positive fraction of cohort 0 to the innovation. Thus, performance of a randomized trial at n = 0 is an inherent consequence of the AMR rule when $\beta_{U0} > \alpha$.

The condition $\beta_{U0} > \alpha$ is necessary for learning because, if $\beta_{U0} \leq \alpha$, the planner assigns no one to the innovation and, hence, never learns its outcomes. The absence of learning has no consequence for welfare in this case. The planner knows from the beginning that the innovation cannot be better than the status quo treatment. Hence, there is no need for a randomized trial.

The AMR rule diversifies each cohort's treatment allocation in a manner that reflects the available knowledge of treatment response. As empirical evidence accumulates and the interval $[\beta_{Ln}, \beta_{Un}]$ shrinks, the value of $\delta_{AMR(n)}$ varies accordingly. Eventually, observation of outcomes under the innovation reveals whether β is larger or smaller than α. From that point on, diversification is no longer warranted and the AMR rule assigns all persons to the better treatment.

Normative Properties of the AMR Rule

The AMR rule has normative appeal for the information that it produces and, with one caveat, for the manner in which it allocates treatments.

The informational appeal is clear. The evidence on treatment response produced under the AMR rule is much stronger than the evidence now produced in clinical trials performed for FDA drug approval. Consequently, treatment decisions under the AMR rule are less prone to error. Important reasons include

Measurement of Outcomes of Real Health Interest: As outcomes of real interest unfold over time, they are observed and used to inform subsequent treatment decisions. In contrast, the FDA currently bases approval decision on surrogate outcomes measured in clinical trials of short duration.

Randomization of the Full Patient Population: The patients assigned to the innovation are randomly drawn from the full patient population. In contrast, present clinical trials randomly draw subjects from a sub-population of volunteers.

Unblinded Treatment Assignment: Treatment assignment is unblinded. In contrast, present clinical trials make blinded assignments, which diminish the information that physicians and patients have relative to real clinical settings.

The AMR treatment allocation has clear appeal to each successive cohort. Each period, the minimax-regret rule is applied using the knowledge available at the time. In this sense, adaptive implementation of the MR rule treats each cohort as well as possible given the available knowledge. It does not ask the members of one cohort to sacrifice its own welfare for the benefit of other cohorts.

The caveat is that the AMR allocation may not be best from the perspective of a planner who wants to maximize welfare aggregated across cohorts. The AMR rule may not minimize maximum regret in terms of the multi-period objective function (2). Minimization of maximum regret in multi-period decision problems is a subtle matter that requires joint consideration of all of cohorts rather than sequential consideration of them one at a time. The nature of the multi-period minimax-regret treatment allocation is an open question.

3.3. Numerical Illustrations

This section illustrates the AMR rule. I present two hypothetical treatment-choice problems. In each case the presumed outcome of interest unfolds over multiple periods. As empirical evidence accumulates, the AMR treatment allocation changes accordingly.

Treating a Life-Threatening Disease

When treating a life-threatening disease, the outcome of interest may be the number of years that a patient survives within some time horizon. For this illustration, I let the horizon be five years and I define y(t) to be the number of years that a patient lives during the five years following receipt of treatment t, where t is the status quo or the innovation. Thus, y(t) has the time-additive form (1), with $y_{jk}(t) = 1$ if patient j is alive k years after treatment, $y_{jk}(t) = 0$ otherwise, and K = 5.

The outcome gradually becomes observable as time passes. At the time of treatment, $y_j(t)$ can take any of the values [0, 1, 2, 3, 4, 5]. A year later, one can observe whether patient j is still alive and hence can determine whether $y_j(t) = 0$ or $y_j(t) \geq 1$. 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 subsequent years, the death rates are 5 and 2 percent respectively. Overall, the entries imply that 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.

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 $[\beta_{L0}, \beta_{U0}] = [0, 5]$. Applying equation (8), 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 $\delta_5 = 1$.

The final two columns of Table 1 gives the maximum regret and mean life span of each cohort, both computed using the AMR treatment allocation. The regret values are 1.5, 0. 72, 0.60, 0.43, 0.02, and 0. The mean life spans are 3.74, 3.72, 3.78, 3.90, 4.28, and 4.30.

Treating a Chronic Disease of Aging

When treating a chronic disease of aging, the outcome of interest may be quality adjusted life years (QALYs) within a specified time horizon. For this illustration, let the horizon be twenty years and let y(t) be the number of QALYs experienced during the twenty years following receipt of treatment t, where t is the status quo or the innovation. Thus, y(t) is time-additive, with $y_{jk}(t) \in [0, 1]$, and K = 20.

The welfare of a treatment is its benefit minus its cost. Let the status quo be a no-program setting with zero cost and let the innovation cost \$5000 per person. The benefit of a treatment is the benefit of one QALY multiplied by the number of QALYs that a person experiences. I consider two values for the social benefit of one QALY, \$10,000 and \$20,000. Then $u(a) = v \cdot y(a)$ and $u(b) = v \cdot y(b) - 5000$, where v = 10,000 or 20,000. Moreover, $\alpha = v \cdot E[y(a)]$ and $\beta = v \cdot E[y(b)] - 5000$.

Table 2 presents hypothetical data on mean QALYs following each treatment. Two columns show $E[y_k(t)]$ for each k = 1, . . . , K and each value of t. The entries describe a situation in which the innovation never does harm relative to the status quo and raises the mean in in some years. Overall, the entries imply that the mean numbers of QALYs experienced during the twenty-year horizon are E[y(a)] = 13.28 and E[y(b)] = 13.57. It follows that if v = 10,000, then α = 132,800 and β = 130,700; hence, the status quo is the better treatment. If v = 20,000, then α = 265,600 and β = 266,400; hence, the innovation is the better treatment in this case.

The value of E[y(a)] is known at the outset from historical experience, but E[y(b)] becomes

observable only gradually. To compute the AMR treatment allocation, I suppose the planner initially knows that treatment b never yields a result lower than $E[y_k(a)]$ and that it may raise it by a maximum of 0.10. Consider, for example, the first and twelfth years after treatment. The table shows that $E[y_1(a)] = 0.98$ and $E[y_{12}(a)] = 0.80$. Hence, the corresponding values under the innovation are initially known to lie in the intervals [0.98, 1] and [0.80, 0.90]. These bounds generate the bounds on E[y(b)] shown in Table 2.

The final columns of Table 2 show the AMR treatment allocations. The patterns for the two values of v are quite different. If v = 10,000, then $\delta_n = 0.65$ in years 0 through 4, δ_n slowly decreases to 0.59 in year 10, then quickly rises to 0.74 at year 12, and finally falls to zero in year 18, when it becomes known that the status quo is better than the innovation. If v = 20,000, then δ_n stays close to 0.83 through year 10 and then rises to one in year 12, when it becomes known that the innovation is better than the status quo.

Due to space constraints, Table 2 does not give the maximum regret and mean welfare of each cohort. However, these quantities are easy to compute, as shown in Section 3.2.

3.4. Extensions of the Analysis

The planning problem analyzed in Section 3.2 is simple in various respects. It assumes that the planner knows α at the outset and knows β eventually. The only inferential problem is that outcomes unfold over time, so K periods must pass to learn β fully.

This section extends the analysis to situations in which the planner faces joint and possibly persistent ambiguity about (α, β). I also discuss treatment choice when the patient population is small enough to make statistical precision an issue when studying treatment response.

Joint Ambiguity About (α, β)

Let us drop the assumption that α is known. Instead suppose that, in period n, (α, β) is known to lie

in some set of feasible pairs $[(\alpha_s, \beta_s), s \in S_n]$. Here S_n indexes the set of feasible *states of nature* in period n.

Let $S_n(a)$ and $S_n(b)$ be the subsets of S_n on which treatments a and b are superior. That is, let $S_n(a) \equiv \{s \in S_n: \alpha_s \geq \beta_s\}$ and $S_n(b) \equiv \{s \in S_n: \beta_s \geq \alpha_s\}$. Manski (2007a, Complement 11A) proves that the AMR rule assigns all patients to the status quo (innovation) if it is superior to the innovation (status quo) in all feasible states of nature and the rule makes a fractional treatment allocation otherwise. The result is

$$
\begin{aligned}
(9)\quad \delta_{AMR(n)} &= 0 && \text{if } \beta_s \leq \alpha_s, \quad \text{all } s \in S_n \\
&= 1 && \text{if } \beta_s \geq \alpha_s, \quad \text{all } s \in S_n \\
&= \frac{\max\limits_{s \in S_n(b)} (\beta_s - \alpha_s)}{\max\limits_{s \in S_n(a)} (\alpha_s - \beta_s) + \max\limits_{s \in S_n(b)} (\beta_s - \alpha_s)} && \text{otherwise.}
\end{aligned}
$$

This general result simplifies considerably when the set of feasible values of (α, β) has the rectangular form $[(\alpha_s, \beta_s), s \in S_n] = [\alpha_{Ln}, \alpha_{Un}] \times [\beta_{Ln}, \beta_{Un}]$, where $[\alpha_{Ln}, \alpha_{Un}]$ and $[\beta_{Ln}, \beta_{Un}]$ are intervals of feasible values for α and β. Then (9) reduces to

$$
\begin{aligned}
(10)\quad \delta_{AMR(n)} &= 0 && \text{if } \beta_{Un} \leq \alpha_{Ln}, \\
&= 1 && \text{if } \beta_{Ln} \geq \alpha_{Un}, \\
&= \frac{\beta_{Un} - \alpha_{Ln}}{(\alpha_{Un} - \beta_{Ln}) + (\beta_{Un} - \alpha_{Ln})} && \text{otherwise.}
\end{aligned}
$$

Result (10) further reduces to (8) when α is fully known; that is, when $\alpha_{Ln} = \alpha_{Un} = \alpha$.

Incomplete Observation of Treatment Outcomes

Let us drop the assumption that the planner fully observes treatment outcomes and, hence, fully knows

(α, β) after K periods. In practice some outcome data may be missing or mismeasured, or perhaps only surrogate outcomes may be observable. Whatever the data problem may be, in period n the available empirical evidence combined with assumptions the planner finds credible will yield some set $[(\alpha_s, \beta_s), s \in S_n]$ of feasible values for (α, β); see Manski (2007a) for exposition of the form of this set when the data are incomplete in various ways.

Incomplete data create no problem for application of the AMR rule, whose form continues to be given by (9). The rule is applicable even if the planner never accumulates enough empirical evidence to know which treatment is superior. In this event, the AMR treatment allocation is fractional in all periods.

Small Populations

I have assumed an atomless patient population, in order to keep attention focused on the identification problems that create most errors in drug approval. This idealization approximates well the actual environment for treatment of widespread conditions such as diabetes, heart disease, and various cancers. However, statistical imprecision in empirical findings on treatment response may be a non-negligible cause of errors when the patient population is small.

In a one-period planning problem with no identification problems, a planner with clinical trial data can apply findings on finite-sample minimax-regret treatment choice developed in Manski (2004, 2005), Manski and Tetenov (2007), Stoye (2006), and Schlag (2007). Manski (2007b) and Stoye (2007b) consider one-period planning problems with certain identification problems.

Multi-period planning problems are more complex. In multi-period problems, there is a tension at each point in time between achievement of two desirable objectives. One is to choose a treatment allocation that minimizes the maximum regret of the current cohort, given the information currently available. The other is to produce as much new information on treatment response as possible, in order to improve the treatment of future cohorts. These objectives do not conflict in large patient populations because even a small fractional

allocation to the innovation produces a large treatment group. They may conflict in small populations.

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

4. Revising Drug Approval Policy in Practice

In this concluding section, I discuss the potential for implementing policies that embody features of the AMR rule. Section 4.1 addresses settings with centralized health care systems. Section 4.2 suggests a quasi-AMR drug approval process that could be implemented by the FDA.

4.1. Centralized Health Care Systems

Close approximations to the AMR rule could be implemented in centralized health care systems where government agencies directly assign treatments. Examples include the Veterans Health Administration in the United States and the National Health Service in England. Implementation could also occur in employer-based and other private systems where health maintenance organizations directly provide medical care.

The VA, NHS, and HMOs could only implement approximate versions of the AMR rule because these social planners do not have fixed patient populations as assumed in Section 3. Patients who are unhappy with the care provided by the VA, NHS, or HMOs can opt-out and seek medical care elsewhere. However, strong financial disincentives typically limit opting-out to the relatively wealthy.

The Ethics of Fractional Treatment Rules

I have written that approximate AMR rules "could" be implemented in centralized health care systems. An important open question is whether the relevant social planners and patient populations would accept the idea of fractional treatment allocation. A possible objection is that fractional rules violate the ethical principle calling for "equal treatment of equals." Fractional rules are consistent with this principle in the *ex ante* sense that observationally identical people have the same probability of receiving a particular treatment. They violate it in the *ex post* sense that observationally identical persons ultimately receive different treatments.

There is precedent for societies to implement major policies that use fractional treatment rules. American examples include random drug testing and airport screening, calls for jury service, and the Green Card and Vietnam draft lotteries. In addition, present-day randomized clinical trials implement fractional rules. The norm has been to randomize volunteers, but the FDA recently approved a set of studies in which critically ill patients are randomized without consent; see Stein (2007).

Fractional treatment allocations derived from the AMR rule are consistent with prevailing standards of medical ethics. Medical ethics permit randomized clinical trials only under conditions of *equipoise*; that is, when partial knowledge of treatment response prevents a determination that one treatment is superior to another. These are exactly the circumstances in which the AMR rule yields a fractional allocation.

4.2. Quasi-AMR Drug Approval by the FDA

As indicated in Section 4.1, AMR treatment rules could be implemented by some parts of the American health care system without change in the FDA drug approval process. The VA, the Military Health System, other public agencies, and many private HMOs could institute such rules right now, considering innovations to be drugs that are newly approved by the FDA.

Implementation of AMR rules after FDA approval of new drugs should reduce the impact of Type I errors in drug approval, both by diversifying treatment choice and by producing new information on treatment response. However, it would not reduce Type II errors. At present, the VA, HMOs and other American health care providers can only choose among the drugs that are approved by the FDA.

Two policy changes could reduce the incidence and impact of Type II errors. First, the federal government could amend the Food, Drug, and Cosmetics Act to enable specified health care providers to treat their patient populations with drugs that have not received FDA approval. In return for exemption from the FDA approval process, these providers could be required to apply appropriate adaptive treatment rules.

Second, the FDA could replace its present binary approval process with an adaptive partial approval process. The permitted use of a new drug now has a sharp discontinuity at the date of the FDA approval decision. Beforehand, a typically tiny fraction of the patient population receives the new drug in clinical trials. Afterwards, use of the drug is unconstrained if approval is granted and zero if approval is not granted. An adaptive approval process would eliminate this discontinuity and instead permit use of a new drug to vary smoothly as empirical evidence accumulates.

Although full FDA implementation of the AMR rule would require radical change in the American health care system, the agency could embrace some important features of the rule with relatively modest revision to the present drug approval process. I sketch below such a quasi-AMR process.

Adaptive Partial Approval through Limited-Term Sales Licenses

The revised drug approval process would begin, as at present, with a pharmaceutical firm performing preclinical testing followed by Phase 1 and 2 trials. It seems prudent to retain these preliminary stages of the approval process in close to their current form. The changes would appear in the subsequent Phase 3 trials and in the FDA decision process. First, the duration of Phase 3 trials would be lengthened sufficiently to measure health outcomes of real interest, not just surrogate outcomes. Second, the present binary approval

decision following a Phase 3 trial would be replaced by an adaptive process that monitors the trial while in progress and that periodically grants limited-term sales licenses.

A limited-term sales license would permit a firm to sell no more than a specified quantity of the new drug over a specified time period. To enforce the upper bound requires a means to monitor sales of a new drug. This is most straightforward if the FDA places only an overall upper bound on sales, rather than distinct bounds on sales to patient groups with different characteristics. If only an overall upper bound is imposed, then it suffices to monitor gross sales in pharmacies. If group-specific bounds are imposed, it is necessary to monitor the distribution of sales.

The duration of the license would depend on the schedule for reporting new findings in the trial. For example, if the firm reports updated outcome data to the FDA annually, then the licensing decision could be updated annually as well. On each iteration of the decision, the maximum quantity of drug that the firm is permitted to sell would be set by the FDA with the assistance of an expert advisory board, similar to those now used in drug approval. This is where the AMR rule comes in. To give the licensing decision transparency and coherence, the FDA could be mandated to compute the AMR treatment allocation with a specified social welfare function. The role of the expert advisory board would be to use the available empirical evidence to determine the set $[(\alpha_s, \beta_s), s \in S_n]$ of feasible values for (α, β).

When the lengthened Phase 3 trial is complete and the outcomes of health interest have been observed, the FDA would make a longer-term approval decision. If the drug is deemed safe and effective, the firm would be permitted to sell it with no quantity restriction. Further use would be prohibited otherwise. As in the current environment, the FDA would retain the right to rescind approval should new evidence warrant. Post-market surveillance would be necessary because lengthening Phase 3 trials to measure health outcomes of interest may not suffice to determine with certainty whether the innovation is superior to the status quo. As with present Phase 3 trials, the lengthened trials would only reveal treatment response for volunteer subjects who comply with treatment and do not attrit from the trial. Moreover, unless the FDA changes its

norms on blinding treatment assignment, the trials would not reveal treatment response in real clinical settings where patients and physicians know the assigned treatments.

A Pragmatic Compromise

The quasi-AMR approval process suggested here would not achieve all of the benefits of the AMR rule. The AMR rule calls for randomly allocating the entire patient population between the status quo treatment and the innovation. Randomization optimally diversifies the treatment of each cohort from the minimax-regret perspective, and it produces as much information as possible about treatment response.

The quasi-AMR process would only randomize the sample of persons who volunteer for Phase 3 trials, as at present. The FDA would influence treatment allocation in the full patient population only through the upper bound it sets on sales. Subject to this bound, treatment allocation would be determined by the decentralized pricing decisions of the pharmaceutical firm, coverage decisions of insurers, and treatment decisions of physicians and patients.

Thus, the new process would diversify treatment allocation, but not necessarily in the optimal manner. It would produce observational data on treatment response in the full patient population that is not available in the present approval process. However, these data would not be as informative as data produced by randomizing treatments.

The rationale for these compromises is pragmatism. The quasi-AMR process would preserve the decentralized health care system that many Americans prize. While Phase 3 trials are underway, it would give the general patient population some access to new drugs and it would give firms some revenue from the sale of new drugs. I conjecture that these features of the new process will make it acceptable to lengthen the duration of Phase 3 trials to enable observation of health outcomes of real interest.

Variations on the Theme

Granting limited-term sales licenses is not the only way that the FDA could implement a quasi-AMR rule. The agency could empower a subset of physicians to prescribe the drug. It could restrict treatment to patients with specified characteristics. Or it could tax sales, the tax rate being set to induce the desired usage. These other mechanisms are less direct than granting sales licenses, but they may merit consideration.

References

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

Committee on the Assessment of the US Drug Safety System, A. Baciu, K. Stratton, and S. Burke (editors) (2006), *The Future of Drug Safety: Promoting and Protecting the Health of the Public*, Washington, DC: National Academies Press.

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

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

Gould, A. (2002), "Substantial Evidence of Effect," *Journal of Biopharmaceutical Statistics*, 12, 53-77.

Keeney, R. (1980), "Equity and Public Risk," *Operations Research*, 28, 527-534.

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

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

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

Manski, C. (2007a), *Identification for Prediction and Decision*, Cambridge, Mass.: Harvard University Press.

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

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

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.

Peltzman, S. (1973), "An Evaluation of Consumer Protection Legislation: The 1962 Drug Amendments," *Journal of Political Economy*, 81, 1049-1091.

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.

Rothman, S. and K. Michels (1994), "The Continued Unethical Use of Placebo Controls," *New England Journal of Medicine*, 331, 394-398.

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

Sculpher, M. and K. Claxton (2005), "Establishing the Cost-Effectiveness of New Pharmaceuticals under Conditions of Uncertainty—When Is There Sufficient Evidence?" *Value in Health*, 8, 433-446.

Stein, R. (2007), "Critical Care without Consent," *Washington Post*, May 27, A01.

Schlag, K. (2006), "Eleven," Technical Report, European University Institute.

Stoye, J. (2006), "Minimax-Regret Treatment Choice with Finite Samples," Department of Economics, New York University.

Stoye, J. (2007a), "Minimax Regret Treatment Choice with Incomplete Data and Many Treatments," *Econometric Theory*, 23, 190-199.

Stoye, J. (2007b), "Minimax-Regret Treatment Choice with Finite Samples and Missing Outcome Data," Proceedings of the 5th International Symposium on Imprecise Probability, Prague, Czech Republic.

Tamura, R., D. Faries, J. Andersen, and J. Heiligenstein (1994), "A Case Study of an Adaptive Clinical Trial in the Treatment of Out-Patients with Depressive Disorder," *Journal of the American Statistical Association*, 89, 768-776.

Temin, P. (1980), *Taking Your Medicine: Drug Regulation in the United States*, Cambridge, Mass.: Harvard University Press.

Viscusi, K., J. Harrington, and J. Vernon (2005), *Economics of Regulation and Antitrust*, Cambridge, Mass.: MIT Press.

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	maximum regret of AMR allocation for cohort n	mean life span achieved by 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

Table 2: Treating a Chronic Disease of Aging					
year (n or k)	mean QALY in k^{th} year after treatment		bound on E[y(b)] for cohort n	AMR allocation for cohort n, by social benefit of one QALY	
	Status Quo	Innovation		$10,000	$20,000
0			[13.28, 14.70]	0.65	0.82
1	0.98	0.99	[13.29, 14.69]	0.65	0.83
2	0.98	0.99	[13.30, 14.68]	0.65	0.83
3	0.98	0.99	[13.31, 14.67]	0.65	0.84
4	0.98	0.98	[13.31, 14.65]	0.65	0.84
5	0.98	0.98	[13.31, 14.63]	0.64	0.83
6	0.98	0.98	[13.31, 14.61]	0.64	0.83
7	0.95	0.96	[13.32, 14.57]	0.63	0.83
8	0.95	0.96	[13.33, 14.53]	0.62	0.83
9	0.90	0.92	[13.35, 14.45]	0.61	0.84
10	0.90	0.92	[13.37, 14.37]	0.59	0.84
11	0.80	0.90	[13.47, 14.37]	0.66	0.93
12	0.80	0.90	[13.57, 14.37]	0.74	1
13	0.50	0.50	[13.57, 14.27]	0.70	1
14	0.50	0.50	[13.57, 14.17]	0.65	1
15	0.40	0.40	[13.57, 14.07]	0.58	1
16	0.40	0.40	[13.57, 13.97]	0.48	1
17	0.10	0.10	[13.57, 13.87]	0.30	1
18	0.10	0.10	[13.57, 13.77]	0	1
19	0.05	0.05	[13.57, 13.67]	0	1
20	0.05	0.05	[13.57, 13.57]	0	1