PLANNING UNDER AMBIGUITY

Charles F. Manski

Department of Economics and Institute for Policy Research
Northwestern University

CEMMAP: 19-20 MARCH 2009

Studying Treatment Response
To Inform Treatment Choice

Why perform research on treatment response?

An important objective is to provide information useful in choosing treatments.

Consider a planner who must choose treatments for a heterogeneous population. For example,

* a physician choosing medical treatments for patients.

* a judge choosing sentences for convicted offenders.

* a community offering preschool activities to children.

Economists have long asked how a planner should act.

A standard exercise specifies a set of feasible treatment policies and a social welfare function.

The objective is to characterize the optimal policy.

Planners often have partial knowledge of the welfare achieved by alternative policies.

Hence, they often cannot determine optimal policies.

The informativeness of research on treatment response is limited by

statistical uncertainty: inference on a study population form sample data.

identification problems: extrapolation from observable features of a study population to other features or other populations.

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

I use elements of decision theory to address this question.

Topics

1: Planning when Treatment Response is Partially Identified

2: Diversified Treatment under Ambiguity

3: Search Profiling with Partial Knowledge of Deterrence

4: Competitive Lending with Partial Knowledge of Loan Repayment

5: Treatment Choice with Sample Data

6. More on Treatment Choice with Sample Data

PLANNING WHEN TREATMENT RESPONSE IS PARTIALLY IDENTIFIED

Studies of treatment response aim to predict the outcomes that would occur if alternative treatment rules were applied to a population.

One cannot observe the outcomes that a person would experience under all treatments.

At most, one can observe a person's *realized outcome*; the one he experiences under the treatment he actually receives.

The *counterfactual outcomes* that would have experienced under other treatments are logically unobservable.

Example

Suppose that patients ill with a disease can be treated by drugs or surgery. The relevant outcome might be life span.

One may want to predict the life spans that would occur if all patients of a certain type were to be treated by drugs.

The available data may be observations of the realized life spans of patients in a study population, some of whom were treated by drugs and the rest by surgery.

The Selection Problem

Set T list all feasible treatments.

Each member j of a study population has covariates $x_j \in X$.

Each j has a *response function* $y_j(\cdot)$: $T \rightarrow Y$ that maps the mutually exclusive and exhaustive treatments $t \in T$ into outcomes $y_j(t) \in Y$.

Let $z_j \in T$ denote the treatment received by person j. Then $y_j \equiv y_j(z_j)$ is the realized outcome.

The outcomes $[y_j(t), t \neq z_j]$ he would have experienced under other treatments are counterfactual.

Observation may reveal the distribution $P[y(z)|x] = P(y|x)$ of realized outcomes for persons with covariates x.

Observation cannot reveal the distribution $P[y(t), t \neq z|x]$ of counterfactual outcomes.

Consider prediction of the outcomes that would occur if all persons with covariates x were to receive treatment t.

This requires inference on $P[y(t)|x]$.

The *selection problem* is the problem of identification of $P[y(t)|x]$ from knowledge of $P(y, z|x)$.

Prediction Using the Empirical Evidence Alone

The Law of Total Probability gives

$$P[y(t) \mid x]$$

$$= P[y(t) \mid x, z = t]P(z = t \mid x) + P[y(t) \mid x, z \neq t]P(z \neq t \mid x)$$

$$= P(y \mid x, z = t)P(z = t \mid x) + P[y(t) \mid x, z \neq t]P(z \neq t \mid x).$$

The *identification region* for $P[y(t) \mid x]$ is

$$H\{P[y(t) \mid x]\}$$
$$= \{P(y \mid x, z = t)P(z = t \mid x) + \gamma P(z \neq t \mid x), \gamma \in \Gamma_Y\},$$

where Γ_Y is the space of all distributions on Y.

Mean Treatment Response

The Law of Iterated Expectations gives

$$E[y(t)|x]$$

$$= E[y(t)|x, z = t]P(z = t|x) + E[y(t)|x, z \neq t]P(z \neq t|x)$$

$$= E(y|x, z = t)P(z = t|x) + E[y(t)|x, z \neq t]P(z \neq t|x).$$

Let Y have smallest and largest elements y_0 and y_1. The identification region for $E[y(t)|x]$ is the interval

$$H\{E[y(t)|x]\} =$$

$$[E(y|x, z = t)P(z = t|x) + y_0P(z \neq t|x),$$
$$E(y|x, z = t)P(z = t|x) + y_1P(z \neq t|x)].$$

Illustration: Sentencing and Recidivism

Consider how the sentencing of offenders may affect recidivism.

Data are available on the outcomes experienced by offenders given the sentences that they receive.

Researchers have debated the counterfactual outcomes that offenders would experience under other sentences.

Predicting the response of criminality to sentencing might reasonably be studied using the empirical evidence alone.

Manski and Nagin (1998) analyzed data on the sentencing and recidivism of males in the state of Utah who were born from 1970 through 1974 and who were convicted of offenses before they reached age 16.

We compared recidivism under the two main sentencing options available to judges:

confinement in residential facilities ($t = b$)

sentences not involving confinement ($t = a$).

The outcome of interest be

$y = 1$ if the offender is not convicted again for two years following sentencing

$y = 0$ if the offender is convicted of a subsequent crime

The data reveal that

Probability of residential treatment: $P(z = b) = 0.11$

Recidivism probability in sub-population receiving residential treatment: $P(y = 0 | z = b) = 0.77$

Recidivism probability in sub-population receiving nonresidential treatment: $P(y = 0 | z = a) = 0.59$.

Consider two policies, one mandating confinement and the other mandating no confinement.

The recidivism probabilities are $P[y(b) = 0]$ and $P[y(a) = 0]$.

Assuming treatment at random,

$$P[y(b) = 0] = P(y = 0 | z = b) = 0.77$$
$$P[y(a) = 0] = P(y = 0 | z = a) = 0.59.$$

Using the data alone,

$$H\{P[y(b) = 0]\} = [0.08, 0.97]$$
$$H\{P[y(a) = 0]\} = [0.53, 0.64].$$

The identification region for the average treatment effect is

$$H\{P[y(b) = 0] - P[y(a) = 0]\} = [-0.56, 0.44].$$

Partial Identification and Decisions under Ambiguity

Partial identification of treatment response may generate ambiguity about the identity of optimal treatment rules.

Criteria for Choice Under Ambiguity

Consider a choice set C and a decision maker who must choose an action from this set.

He wants to maximize an objective function $f(\cdot): C \rightarrow R$.

He faces an optimization problem if he knows C and $f(\cdot)$.

He faces a problem of choice under ambiguity if he knows C but knows only that $f(\cdot) \in F$, where F is a set of possible objective functions.

How should the decision maker behave?

He should not choose a *dominated* action.

Action $d \in C$ is dominated if there exists another feasible action, say c, that is at least as good as d for all objective functions in F and strictly better for some function in F.

How should he choose among the undominated actions?

One idea is to average the elements of F and maximize the resulting function. This yields Bayes rules.

Another is to choose an action that, in some sense, works uniformly well over all elements of F. This yields the maximin and minimax-regret criteria.

One-Period Planning with Individualistic Treatment

There are two treatments, a and b. Let $T \equiv \{a, b\}$.

Each member j of population J has a response function $y_j(\cdot): T \rightarrow Y$ that maps treatments t into outcomes $y_j(t)$.

$P[y(\cdot)]$ is the population distribution of treatment response.

The population is large, with $P(j) = 0$ for all $j \in J$.

The task is to allocate the population to treatments.

An allocation is a $\delta \in [0, 1]$ that randomly assigns a fraction δ of the population to b and the remaining $1 - \delta$ to a.

The planner wants to maximize mean welfare.

Let $u_j(t) \equiv u_j[y(t), t]$ be the net contribution to welfare if person j receives t and realizes $y_j(t)$.

Let $\alpha \equiv E[u(a)]$ and $\beta \equiv E[u(b)]$.

Welfare with allocation δ is

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

Treatment Choice Under Ambiguity

$\delta = 1$ is optimal if $\beta \geq \alpha$ and $\delta = 0$ if $\beta \leq \alpha$. The problem is treatment choice when (α, β) is partially known.

Let S index the feasible states of nature. The planner knows that (α, β) lies in the set $[(\alpha_s, \beta_s), s \in S]$. Assume that this set is bounded. Let

$$\alpha_L \equiv \min_{s \in S} \alpha_s, \quad \beta_L \equiv \min_{s \in S} \beta_s,$$

$$\alpha_U \equiv \max_{s \in S} \alpha_s, \quad \beta_U \equiv \max_{s \in S} \beta_s.$$

The planner faces ambiguity if $\alpha_s > \beta_s$ for some values of s and $\alpha_s < \beta_s$ for other values.

Bayes Rules

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

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

The Bayes decision assigns everyone to b if $E_\pi(\beta) > E_\pi(\alpha)$ and everyone to treatment a if $E_\pi(\alpha) > E_\pi(\beta)$. All allocations are Bayes decisions if $E_\pi(\beta) = E_\pi(\alpha)$.

Bayesian planning is conceptually straightforward, but it may not be straightforward to form a credible subjective distribution on the states of nature.

The allocation chosen by a Bayesian planner depends on the subjective distribution used.

The Bayesian paradigm is appealing only when a decision maker is able to form a subjective distribution that really expresses his beliefs.

The Maximin Criterion

Compute the minimum welfare attained by each allocation across all states of nature.

Choose an allocation that maximizes minimum welfare.

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

Suppose that (α_L, β_L) is a feasible value of (α, β).

Then the maximin allocation is $\delta = 0$ if $\alpha_L > \beta_L$ and $\delta = 1$ if $\alpha_L < \beta_L$.

The Minimax-Regret Criterion

The regret of allocation δ in state of nature s is the difference between the maximum achievable welfare and the welfare achieved with allocation δ.

Maximum welfare in state of nature s is max (α_s, β_s).

Hence,

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

In contrast to Bayes decisions and the maximin rule, the minimax-regret rule is *fractional*.

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

Let $M(a) \equiv \max_{s \in S(a)} (\alpha_s - \beta_s)$

$M(b) \equiv \max_{s \in S(b)} (\beta_s - \alpha_s)$.

The minimax-regret allocation is

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

Let (α_L, β_U) and (α_U, β_L) be feasible values of (α, β). Then

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

Proof:

The maximum regret of δ is max [R(δ, a), R(δ, b)], where

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

$$= \max_{s \in S(a)} \delta(\alpha_s - \beta_s) = \delta M(a),$$

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

$$= \max_{s \in S(b)} (1-\delta)(\beta_s - \alpha_s) = (1-\delta)M(b),$$

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

Both treatments are undominated, so $R(1, a) = M(a) > 0$ and $R(0, b) = M(b) > 0$.

As δ increases from 0 to 1, $R(\cdot, a)$ increases linearly from 0 to $M(a)$ and $R(\cdot, b)$ decreases linearly from $M(b)$ to 0.

Hence, the MR rule is the unique $\delta \in (0, 1)$ such that

$$R(\delta, a) = R(\delta, b).$$

This yields the result.

Illustration: Sentencing Juvenile Offenders in Utah

Consider a Utah judge who observes the treatments and outcomes of the study population and who must choose sentences for a new cohort of convicted offenders.

The judge believes that the study population and the new cohort have the same distribution of treatment response.

The judge does not find any other assumptions credible.

The distribution of realized treatments and outcomes in the study population is

$$P(z = b) = 0.11 \qquad P(y = 1 | z = b) = 0.23$$
$$P(y = 1 | z = a) = 0.41.$$

Let $q(t) \equiv E_\pi(y = 1 | z \neq t)$. A Bayes rule sets

$\delta = 1$ if $0.03 + (0.89)q(b) > 0.36 + (0.11)q(a)$

$\delta = 0$ if $0.03 + (0.89)q(b) < 0.36 + (0.11)q(a)$

The maximin rule sets $\delta = 0$.

The minimax-regret rule sets $\delta = 0.55$.

Illustration: Choosing Treatments for X-Pox

A new viral disease, x-pox, is sweeping the world.

Medical researchers have proposed two mutually exclusive treatments, a and b, reflecting alternative hypotheses, H_a and H_b, about the nature of the virus.

If H_t is correct, all persons who receive treatment t survive and all others die.

It is known that one of the two hypotheses is correct, but it is not known which. Thus, there are two states of natures, $\gamma = H_a$ and $\gamma = H_b$.

The objective is to maximize the survival rate of the population.

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

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

A Bayes rule generically sets $\delta = 0$ or 1.

The maximin and minimax-regret rules set $\delta = ½$.

DIVERSIFIED TREATMENT UNDER AMBIGUITY

Suppose that a planner can treat persons differentially.

He may make a *singleton* allocation, assigning all observationally identical persons to the same treatment.

He could choose a *fractional* allocation, randomly assigning positive fractions of these persons to both treatments.

Portfolio choice has long been framed as a choice among fractional allocations.

Social planning has commonly been viewed as a choice between singleton allocations.

Fractional allocations cope with ambiguity through diversification.

Diversification enables a decision maker to balance two types of potential error.

A Type A error occur when treatment a is chosen but is actually inferior to b, and a Type B error occurs when b is chosen but is inferior to a.

The singleton allocation assigning everyone to treatment a entirely avoids type B errors but may yield Type A errors, and vice versa for singleton assignment to treatment b.

Fractional allocations make both types of errors but reduce their potential magnitudes.

Topics

* One-Period Planning with Individualistic Treatment and Linear Welfare

* Nonlinear Welfare

* Interacting Treatments

* Dynamic Planning Problems

* Two-Planner Games

One-Period Problems with Individualistic Treatment and Linear Welfare

There are two treatments, labeled a and b. Let $T \equiv \{a, b\}$.

Each member j of population J has a response function $y_j(\cdot)$: mapping treatments t into outcomes $y_j(t)$.

$P[y(\cdot)]$ is the population distribution of treatment response. The population is large, with $P(j) = 0$ for all $j \in J$.

The minimax-regret allocation is

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

Planning with Observable Covariates

A planner may systematically differentiate among persons with different observed covariates $\xi \in X$.

He may segment persons by ξ and treat each group as the population. This works when the objective function is separable in covariates but not otherwise.

The Bayes objective function is always separable.

Maximin is separable if $(\alpha_{\xi L}, \beta_{\xi L})$, $\xi \in X$ is feasible.

Minimax-regret is separable if $(\alpha_{\xi L}, \beta_{\xi U})$, $\xi \in X$ and $(\alpha_{\xi U}, \beta_{\xi L})$, $\xi \in X$ are feasible.

Nonseparability occurs if the planner has information that relates the values of $(\alpha_{\xi s}, \beta_{\xi s})$ across $\xi \in X$.

Planning with Multiple Treatments

The MR allocation is not always fractional when a planner allocates the population among more than two treatments.

Stoye (2007) has studied a class of such problems and has found that the MR allocations are subtle to characterize.

They often are fractional, but he gives an example in which there exists a unique singleton allocation.

Nonlinear Welfare

Monotone Transformations of the Welfare Function

Let $W(\delta) = f[\alpha + (\beta - \alpha)\delta]$, where $f(\cdot)$ is strictly increasing.

The Bayes decision is generically singleton if $f(\cdot)$ is convex, but it may be fractional if $f(\cdot)$ has concave segments.

The shape of $f(\cdot)$ does not affect the maximin decision.

The minimax-regret allocation is fractional whenever $f(\cdot)$ is continuous and the optimal choice is ambiguous.

If $f(\cdot) = \log(\cdot)$ and $\{(\alpha_L, \beta_U), (\alpha_U, \beta_L)\}$ are feasible, then

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

Allocation of an Endowment Between a Safe and a Risky Asset

Consider an investor who must allocate an endowment between a safe and a risky asset. The safe asset is treatment a, with known return α. The risky asset is b, with return known to lie in $[\beta_L, \beta_U]$. Assume that $\alpha \in [\beta_L, \beta_U]$.

A Bayesian investor sets $\delta = 0$ if $E_\pi(\beta) < \alpha$.
He diversifies if $E_\pi(\beta) > \alpha$ and $\int f(\beta)d\pi < f(\alpha)$.
He sets $\delta > 0$ if $\int f(\beta)d\pi \geq f(\alpha)$.

A maximin investor sets $\delta = 0$.

A minimax-regret investor always diversifies, the specific allocation depending on the shape of $f(\cdot)$.

Fixed Costs

Let treatments a and b have known fixed costs C(a) and C(b). Let welfare be linear. The MR allocation is fractional if the fixed costs are small, but singleton if they are large.

If $C \equiv C(a) = C(b)$, then

$$\delta_{FMR} = 0 \qquad \text{if} \quad M(b) \le \min\{M(a),\ \delta_{MR}M(a) + C\}$$

$$\delta_{FMR} = \delta_{MR} \qquad \text{if} \quad \delta_{MR}M(a) + C \le \min\{M(a), M(b)\}$$

$$\delta_{FMR} = 1 \qquad \text{if} \quad M(a) \le \min\{M(b),\ \delta_{MR}M(a) + C\}.$$

Deontological Welfare Functions

Deontological ethics supposes that choices may have intrinsic value, apart from their consequences.

Fixed costs can be interpreted as expressing the deontological idea that any use of treatment a or b is bad.

Equal Treatment of Equals is a deontological principle

Fractional allocations adhere to the principle in the *ex ante* sense that all persons have equal probabilities of receiving particular treatments.

Fractional allocations are inconsistent with equal treatment in the *ex post* sense that all persons do not actually receive the same treatment.

From the ex ante perspective, all treatment allocations are deontologically equivalent.

From the ex post perspective, singleton allocations are advantageous relative to fractional ones.

Placing value C on equal ex post treatment does not alter the MR allocation if $C < \min \{M(a), M(b)\} - \delta_{MR}M(a)$. The MR allocation if singleton otherwise.

Interacting Treatments

Response function $y_j(\cdot, \cdot)$: $T \times [0, 1] \to Y$ maps own treatments t and allocations δ into outcomes $y_j(t, \delta)$.

$u_j(t, \delta) \equiv u_j[y(t, \delta), t, \delta]$ is the net contribution to welfare.

Let $\alpha(\delta) \equiv E[u(a, \delta)]$ and $\beta(\delta) \equiv E[u(b, \delta)]$.

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

The optimal allocation may be singleton or fractional, depending on how $\alpha(\cdot)$ and $\beta(\cdot)$ vary with δ.

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

Choosing Medical Treatments for an Infectious Disease

A population is susceptible to an infectious disease.

Treatment a is therapy after infection and b is vaccination before infection.

Assume that

the infection rate of unvaccinated persons decreases with the fraction of the population who are vaccinated. Thus, $\alpha(\delta)$ weakly increases with δ

vaccination of a person always prevents his infection. Thus, $\beta(\delta) \equiv \beta$ does not vary with δ.

Linear Treatment Interactions

Assume that $\alpha(\cdot)$ has the known linear form

$$\alpha(\delta) = \alpha(0)(1 - \delta) + \alpha(1)\delta.$$

Then the optimization problem is

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

The optimal allocation is

$$\delta^* = ½ + ½[\beta - \alpha(0)]/[\alpha(1) - \alpha(0)]$$

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

Treatment under Ambiguity

Assume $\alpha(\cdot)$ is bounded and weakly increasing, the bounds being real numbers L and U such that $L \le \alpha(0) \le \alpha(1) \le U$. β is known.

The minimax-regret allocation is

$$\delta_{MR} = 1 \qquad \text{if} \quad \beta > U,$$

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

δ_{MR} is always positive, no matter how small β is.

Dynamic Planning Problems

In each period $n = 0, \ldots, N$, a planner must choose treatments for the current cohort of a population.

Learning is possible, with observation of the outcomes experienced by earlier cohorts informing treatment choice for later cohorts.

Fractional treatment allocations generate randomized experiments yielding outcome data on both treatments.

Sampling variation is not an issue when cohorts are large. All fractional allocations yield the same information.

The Adaptive Minimax-Reget Criterion

In each period, the *adaptive minimax-regret* (*AMR*) criterion applies the static minimax-regret criterion using the information available at the time.

The AMR criterion is an appealing myopic rule. 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.

Unless fixed costs or deontological considerations make the AMR allocation singleton, it maximizes learning about treatment response.

Treating a Life-Threatening Disease

Close approximations to the AMR rule could be implemented in centralized health care systems where government or private agencies directly assign treatments.

Let y(t) be the number of years that a patient lives during the five years following receipt of treatment t.

The outcome gradually becomes observable as time passes. Initially, $y_j(t) \in [0, 1, 2, 3, 4, 5]$. A year later, one knows whether $y_j(t) = 0$ or $y_j(t) \geq 1$. And so on until year five.

Assume that u(t) = y(t). Assume no initial knowledge of β.

Table 1: Treating a Life-Threatening Disease						
cohort or year (n or k)	death rate in kth 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

The AMR Criterion and Randomized Clinical Trials

Randomized clinical trials (RCTs) are used to learn about medical innovations. The allocations produced by the AMR criterion differ from the practice of RCTs in many ways.

Fraction of the Population Receiving the Innovation

The AMR allocation can take any value in [0, 1]. The sample receiving the innovation in RCTs is typically a very small fraction of the population, with sample size determined by conventional calculations of statistical power.

Group Subject to Randomization

Under the AMR criterion, the persons receiving the innovation are randomly drawn from the full patient population. Clinical trials randomly draw subjects from pools of persons who volunteer to participate.

Measurement of Outcomes

Under the AMR criterion, one observes the health outcomes of interest as they unfold over time. RCTs typically have short durations of two to three years. Hence, medical researchers often measure *surrogate outcomes* rather than outcomes of real interest.

Blinding of Treatment Assignment

Under the AMR criterion, assigned treatments are known to patients and their physicians. Blinded treatment assignment has been the norm in clinical trials of new drugs.

Use of Empirical Evidence in Decision Making

Choosing a treatment allocation to minimize maximum regret is remote from the way that RCTs are used in decision making. The conventional approach is to perform a hypothesis test. The null hypothesis is that the innovation is no better than the status quo.

Two-Planner Games

Consider a two-planner setting where the planners may have different welfare functions and beliefs. Any departure from the status quo require agreement of the two planners.

When policy choice is framed as a binary decision, let the status quo be chosen if either planner prefers $\delta = 0$ to $\delta = 1$.

An innovation replaces the status quo if both prefer $\delta = 1$ to $\delta = 0$.

Thus, the innovation bears the *burden of proof*.

When policy choice is framed as selection of a treatment allocation, there may exist fractional allocations that both planners prefer.

Noncooperative Application of the MR Criterion

Let the two planners be m = 1 and m = 2. Let each have a welfare function that is monotone in some mean outcome. Suppose that both planners use the MR criterion.

Let δ_{mMR} be the allocation that m would choose if he were able to dictate policy. Let $\delta_{1MR} \leq \delta_{2MR}$.

The set of pareto efficient allocations is $[\delta_{1MR}, \delta_{2MR}]$.

Consider a decision process where each planner announces his preferred allocation, after which the smaller of the reported values is chosen. Truthful revelation is the dominant strategy. Thus, δ_{1MR} is the chosen allocation.

Conclusion

Analysis of choice with incomplete information has usually presumed that the decision maker places a subjective distribution on the feasible states of nature and maximizes expected utility. A subjective distribution is a form of knowledge. There may be no credible basis to assert one.

Diversification is commonplace in private decisions with incomplete information. Yet discussions of social planning commonly presume that observationally identical persons should receive the same treatment. A planner need not limit consideration to singleton allocations.

I have focused attention on the minimax-regret criterion, which has received remarkably little attention from economists in the long period since it was proposed by Savage (1951).

The MR criterion diversifies treatment in a large class of planning problems, including ones with nonlinear welfare functions, interacting treatments, dynamics with learning, and non-cooperative aspects.

There are important problems in which the MR allocation is not fractional. It is easy to see why the MR allocation is singleton when fixed costs or deontological considerations loom large. It is harder to intuit the subtlety of the allocation in settings with more than two treatments.

SEARCH PROFILING WITH PARTIAL KNOWLEDGE OF DETERRENCE

Economists commonly assume that social planners know the deterrent effects of alternative law enforcement policies.

I explore how a planner with partial knowledge of deterrence might choose a policy.

I specifically consider profiling policies that make decisions to search for evidence of crime vary with observable covariates of the persons at risk of being searched.

I pose a planning problem whose objective is to minimize the social cost of crime and search.

Search is costly per se, and search that reveals a crime entails costs for punishment of offenders. Search is beneficial if it deters or prevents crime.

Deterrence is expressed through the *offense function*, describing how the offense rate of persons with given covariates varies with their search rate.

The planner has partial knowledge of the offense function and, hence, is unable to determine what policy is optimal.

The planner observes the offense rates of a study population whose search rule has previously been chosen.

He knows that the study population and the population of interest have the same offense function.

He knows that search weakly deters crime, but does not know the magnitude of the deterrent effect of search.

The planner should first eliminate dominated search rules.

He then may use a Bayes, minimax, or minimax-regret criterion to choose an undominated search rule.

The Planning Problem

There exists a large population J of potential offenders. Each member of J decides whether or not to commit an offense, taking into account the chance that he will be searched.

Let $y_j(t) = 1$ if person j chooses to commit an offense when the search probability is t, with $y_j(t) = 0$ otherwise.

Offenders are always apprehended through search but are not apprehended otherwise.

Ex ante search prevents an offense from causing social harm.

Ex post search apprehends an offender after completion of his offense.

The planner's problem is to choose the probabilities with which persons are searched.

Let person j have observable fixed covariates $x_j \in X$.

A feasible search rule is a function $z: X \rightarrow [0, 1]$ that assigns a single search rate to all persons with the same value of x, but possibly different search rates to persons with different covariates.

Let $p(t, x) \equiv P[y(t) = 1 | x]$ be the offense function.

Under search rule z, the offense rate among persons with covariates x is $p[z(x), x] = P\{y[z(x)] = 1 | x\}$.

The objective is to minimize a social cost function with three components:

(a) the harm caused by completed offenses

(b) the cost of punishing apprehended offenders

(c) the cost of performing searches.

Ex Ante Search

The social cost function is

$$S(z) =$$

$$\int a(x)\cdot p[z(x), x]\cdot[1 - z(x)]dP(x) + \int b(x)\cdot p[z(x), x]\cdot z(x)dP(x) + \int c(x)\cdot z(x)dP(x).$$

I assume that $a(x) > b(x) \geq 0$ and that $c(x) > 0$.

Th planning problem is separable in x.

An optimal search rate for persons with covariates x is

$$z^*(x) \equiv \underset{t \in [0, 1]}{\operatorname{argmin}} \; a(x)\cdot p(t, x)\cdot(1 - t) + b(x)\cdot p(t, x)\cdot t + c(x)\cdot t.$$

Linear Deterrence

Let $\rho(x) \equiv p(0, x)$.

Let $p(t, x) = \rho(x)$, $t \in [0, 1]$. Then the optimal search rate is

$$z^*(x) = \underset{t \in [0,1]}{\operatorname{argmin}} \; a(x)\cdot\rho(x)\cdot(1 - t) + b(x)\cdot\rho(x)\cdot t + c(x)\cdot t$$

$$= 0 \text{ if } c(x) > [a(x) - b(x)]\cdot\rho(x),$$

$$= 1 \text{ if } c(x) < [a(x) - b(x)]\cdot\rho(x).$$

Let $p(t, x) = \rho(x)\cdot(1 - t)$. Then the optimal search rate is

$$z^*(x) = \underset{t \in [0,1]}{\operatorname{argmin}} \; a(x)\cdot\rho(x)\cdot(1 - t)^2 + b(x)\cdot\rho(x)\cdot t(1 - t) + c(x)\cdot t.$$

The global minimum is at

$$t^{*}(x) = \frac{[2a(x) - b(x)]\cdot\rho(x) - c(x)}{2[a(x) - b(x)]\cdot\rho(x)} \quad .$$

Hence, the optimal search rate is

$$\begin{aligned} z^{*}(x) &= 0 && \text{if} \quad t^{*} < 0, \\ &= t^{*}(x) && \text{if} \quad 0 \le t^{*} \le 1, \\ &= 1 && \text{if} \quad t^{*} > 1. \end{aligned}$$

Partial Knowledge of Deterrence

Solution of the planning problem requires essentially complete knowledge of $p(\cdot, \cdot)$.

Suppose that the planner observes the offense rates of a study population.

He knows that the study population and the population of interest have the same offense function.

He knows that the offense rate weakly decreases as the search rate increases, but he does not know the magnitude of the deterrent effect of search.

Let $r(x)$ denote the search rate applied to persons with covariates x in the study population.

Let $q(x)$ denote the realized offense rate of these persons.

Assumption 1 (Study Population): The planner observes $[r(x), q(x)], x \in X$. He knows that $q(x) = p[r(x), x], x \in X$.

Assumption 2(Search Weakly Deters Crime): The planner knows that, for $x \in X$, $p(t, x)$ is weakly decreasing in t.

These assumptions imply that

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

Dominated Search Rules

Let Γ denote the set of offense functions that are feasible under Assumptions 1 and 2. For $\gamma \in \Gamma$, let

$$S(z, \gamma) \equiv$$

$$\int a(x)\cdot\gamma[z(x), x]\cdot[1 - z(x)]dP(x) + \int b(x)\cdot\gamma[z(x), x]\cdot z(x)dP(x) + \int c(x)\cdot z(x)dP(x)$$

be the social cost of rule z when the offense function is γ.

Rule z is *strictly dominated* if there exists another search rule $z' \in Z$ such that $S(z, \gamma) > S(z', \gamma)$ for all $\gamma \in \Gamma$.

The planning problem is separable in x, so it suffices to consider each covariate value separately.

Suppressing x, let

$$d(t, s; \gamma) \equiv [a(1 - t) + bt]\gamma(t) + ct \; - \; [a(1 - s) + bs]\gamma(s) - cs$$

be the difference in social cost between application of search rates t and s when the offense function is γ.

Let $D(t, s) \equiv \sup_{\gamma \in \Gamma} d(t, s; \gamma)$. Search rate s is strictly dominated if there exists a t such that $D(t, s) < 0$.

Lemma 2: Let Assumptions 1 and 2 hold. Search rate s is strictly dominated if any of these conditions hold:

(a) Let $c < (a - b)q$. Then s is strictly dominated if $s \leq r$.

(b) Let $c > (a - b)q$. Then s is strictly dominated if

$s > r + [aq(1 - r) + bqr]/c$.

(c) Let $c > a - b$. Then s is strictly dominated if

$a(1 - q)/[c - q(a - b)] < s \leq r$ or if $s > \max(r, a/c)$.

Decision theory offers various criteria for choice among undominated alternatives.

Bayes Rules

The planner places a subjective distribution Ψ on $p(\cdot, \cdot)$ and minimizes subjective expected social cost

$$E_\Psi[S(z)] =$$

$$\int a(x)\cdot\pi[z(x), x]\cdot[1 - z(x)]dP(x) + \int b(x)\cdot\pi[z(x), x]\cdot z(x)dP(x) + \int c(x)\cdot z(x)dP(x),$$

where $\pi(\cdot, x) \equiv E_\Psi[p(\cdot, x)]$ is the subjective mean of $p(\cdot, x)$.

A Bayesian planner acts as a pseudo-optimizer, using the subjective expected offense function π as if it were the actual offense function p.

Minimax Search

The minimax criterion is

$$\min_{z \in Z} \max_{\gamma \in \Gamma} S(z, \gamma).$$

The outer minimization problem is separable in x. Hence, the minimax search rate for persons with covariates x is

$$z^m(x) \equiv \underset{t \in [0, 1]}{\text{argmin}} \max_{\gamma \in \Gamma} [a(x)\cdot(1 - t) + b(x)\cdot t]\cdot\gamma(t, x) + c(x)\cdot t.$$

Suppressing x,

Lemma 3: Under Assumptions 1 and 2, the minimax search rate is

$$
\begin{aligned}
z^m &= 0 \quad \text{if} \quad c \geq a - b \text{ and } a \leq aq(1 - r) + bq{\cdot}r + cr, \\
&= r \quad \text{if} \quad c \geq a - b \text{ and } a \geq aq(1 - r) + bqr + cr \\
&\qquad\qquad \text{or if} \quad (a - b)q \leq c < a - b, \\
&= 1 \quad \text{if} \quad c \leq (a - b)q.
\end{aligned}
$$

Minimax-Regret Search

For $\gamma \in \Gamma$, let $S^*(\gamma) \equiv \min_{z \in Z} S(z, \gamma)$ be the lowest social cost achievable by any feasible search rule when the offense function is γ.

The regret of rule z in state of nature γ is $S(z, \gamma) - S^*(\gamma)$.

The minimax-regret criterion is

$$\min_{z \in Z} \sup_{\gamma \in \Gamma} S(z, \gamma) - S^*(\gamma).$$

Suppress x.

Lemma 4: Let $b = 0$ and $c < aq$. Under Assumptions 1 and 2, the minimax-regret search rate is

$$z^{mr} = (aq + cr)/(aq + c).$$

Ex Post Search

When search is ex post, the social cost function is

$$S'(z) = \int a(x)\cdot p[z(x), x]dP(x) + \int b(x)\cdot p[z(x), x]\cdot z(x)dP(x) + \int c(x)\cdot z(x)dP(x).$$

Linear Deterrence

If ex post search has no deterrent effect, the optimal search rate is zero.

Let $p(t, x) = \rho(x) \cdot (1 - t)$. Then the optimal search rate is

$$z'(x) = \underset{t \in [0, 1]}{\operatorname{argmin}}\ a(x) \cdot \rho(x) \cdot (1 - t) + b(x) \cdot \rho(x) \cdot t(1 - t) + c(x) \cdot t.$$

Solving this yields

$$\begin{aligned} z'(x) &= 0 \qquad \text{if } c(x) \geq a(x) \cdot \rho(x), \\ &= 1 \qquad \text{if } c(x) \leq a(x) \cdot \rho(x). \end{aligned}$$

Dominated Search Rules

Suppressing x, let

$$d'(t, s; \gamma) \equiv (a + bt)\gamma(t) + ct - (a + bs)\gamma(s) - cs$$

be the difference in social cost between application of search rates t and s when the offense function is γ.

Let $D'(t, s) \equiv \sup_{\gamma \in \Gamma} d'(t, s; \gamma)$. Search rate s is strictly dominated if there exists a t such that $D'(t, s) < 0$.

Lemma 6: Let S′ be the social cost function. Let Assumptions 1 and 2 hold. Search rate s is strictly dominated if any of these conditions hold:

(a) $a(1 - q)/(c + bq) < s \leq r$

(b) $s > \max(r, a/c)$

(c) $s > r + (aq + bqr)/c$.

Minimax Search

Lemma 7: Let S′ be the social cost function. Under Assumptions 1 and 2, the minimax search rate is

$$z^{m'} = 0 \qquad \text{if} \quad a \leq aq + bqr + cr,$$
$$= r \qquad \text{if} \quad a \geq aq + bqr + cr.$$

Minimax-Regret Search

Lemma 6: Let S′ be the social cost function. Let b = 0 and c < aq. Under Assumptions 1 and 2, the minimax-regret search rate is

$$z^{mr'} = 0 \qquad \text{if} \quad a \leq aq + cr\,,$$
$$= \text{all } t \geq r \qquad \text{if} \quad a > aq + cr.$$

Variations on the Planning Problem

* The available information may differ from that assumed here.

* Search may only sometimes apprehends offenders.

* A planner may be able to implement both ex ante and ex post search rules.

* a planner may be able to choose the severity with which apprehended offenders are punished.

* Offense decisions may be interdependent.

* Personal covariates may be malleable.

COMPETITIVE LENDING WITH PARTIAL KNOWLEDGE OF LOAN REPAYMENT

Consider a loan contract specifying that a borrower will receive one dollar today and repay r dollars tomorrow. A common problem is partial knowledge today of the amount that will actually be repaid tomorrow.

Economists usually assume that agents place subjective distributions on unknowns and maximize expected utility. It is usual to assume that expectations are rational and that agents use Bayes Rule to update expectations.

There are realistic circumstances in which an agent may have no credible basis for asserting any subjective distribution, never mind one that is objectively correct. Then agents face problems of decision making under ambiguity.

Ambiguity may be prevalent when a stable market experiences a significant unanticipated shock.

Market participants may be unsure how to interpret the shock. It may have been temporary, or it may indicate a regime change. Agents may have probabilistic expectations before the shock, but be unsure how to update them.

Consider the ongoing credit crisis. Lenders experienced an unanticipated shock when mortgage loan repayment rates fell well below recent norms. Lenders have subsequently been unsure how to predict future repayment rates.

Equilibrium with Partial Knowledge of Loan Repayment

We study a competitive credit market when lenders have partial knowledge of loan repayment.

The contracted repayment rate r equilibrates the supply and demand of loans.

Loan supply is determined by lenders who allocate monetary endowments between loans and a safe investment.

Borrowers demand loans to undertake investments.

Incomplete repayment occurs when borrowers have partial knowledge of the productivity of investments. When productivity is low ex post, they may lack the resources to fully repay their loans.

Bankruptcy law limits liability for repayment.

Market Dynamics Following a Temporary Productivity Shock

We explore market dynamics after an unanticipated productivity shock.

The market is initially in steady state, with lenders knowing the loan return. A shock temporarily lowers the productivity of borrower investments, reducing loan repayment.

Assuming that the shock does not affect loan demand, we consider the response of lenders.

Not knowing how to interpret the shock, lenders may use Bayesian, maximin, or minimax-regret criteria to determine loan supply.

The equilibrium contracted repayment rate will rise immediately following the shock, to a degree that depends on the decision criteria that lenders use.

The longer run market dynamic depends on how lenders interpret the shock, what decision criteria they use, and how they revise their beliefs as new empirical evidence accumulates.

Government Intervention in the Credit Market

We ask how a planner who knows that a shock is temporary might use policy instruments to restore the steady state.

Let the objective be maximization of the aggregate return to the investments financed by lender's endowments.

Considering a market whose steady state has a known loan return, we show that two policies can restore the steady state immediately:

* Manipulate the return on the safe investment.
* Guarantee a minimum loan return to lenders.

We conclude that the loan guarantee is preferable.

* It does not require detailed knowledge of lender behavior.
* The required safe return may be infeasibly low.

A Credit Market with Observationally Heterogeneous Borrowers

Finally, we study markets with borrowers who are observationally heterogeneous to lenders. Now lenders can price loans differentially, setting contracted repayment rates that vary with observed borrower covariates.

Relationship to the Literature

The literature studying credit markets is diverse and vast. This paper makes two main contributions.

1. We study competitive equilibrium when lenders face ambiguity about loan repayment.

2. We propose and study policy instruments that the government may use to restore a steady state following a temporary shock.

Easley and O'Hara (2008) and Caballero and Krishnamurthy (2008) study financial markets where agents face ambiguity. They also ask how government intervention might mitigate unpalatable market outcomes. These studies differ greatly from our work in their specifics.

Equilibrium in a Credit Market with Partial Knowledge of Loan Repayment

The credit market is small and competitive.

Loans are one-period contracts.

Inflation is anticipated.

Borrowers are observationally identical, as viewed by lenders.

Loans are securitized.

Lender knowledge of loan repayment is predetermined.

A population J of borrowers and a set K of lenders interact in period t.

A loan contract specifies that a borrower receives one dollar at time t and repays r_t dollars at t + 1.

r_t equates the aggregate demand for and supply of loans.

The Demand for and Repayment of Loans

Aggregate demand $D_t(\cdot)$ is continuous and strictly decreasing for r such that $D_t(r) > 0$, with $D_t(1) > 0$ and $\lim_{r \to \infty} D_t(r) = 0$.

Borrower j wants to solve

$$(1) \qquad \max_{x \geq 0} \quad u_{jt}(v_{jt}) + u_{j(t+1)} \{\max [0, v_{j(t+1)} + g_{jt}(x) - rx]\}.$$

v_{jt} and $v_{j(t+1)}$ are endowments. $g_{jt}(x)$ is the investment return to a loan of size x, with $g_{jt}(0) = 0$. $g_{jt}(\cdot)$ is increasing, differentiable, and concave.

If non-negative, $v_{j(t+1)} + g_{jt}(x) - rx$ is t + 1 consumption.

If negative, the borrower cannot fully repay. He declares bankruptcy, consumes zero, and repays $v_{j(t+1)} + g_{jt}(x)$.

A borrower who knows $g_{jt}(\cdot)$ can solve (1), which reduces to

$$\max_{x \geq 0} \; g_{jt}(x) - rx. \tag{2}$$

The maximum value of $g_{jt}(x) - rx$ over $x \geq 0$ is non-negative. Hence, the person repays his loan fully.

A borrower with partial knowledge of $g_{jt}(\cdot)$ may not be able to solve (1).

Borrowing $x > 0$ may be a reasonable decision ex ante but a poor one ex post. If the investment return is low, a borrower may not be able to fully repay his loan.

The Return on Loans

Securitization implies that lenders are concerned only with the aggregate repayment of loans. Let $x_{jt}(r)$ be loan demand by j at contracted repayment rate r. The loan return at this rate is

$$(3) \quad \lambda_t(r) \equiv \frac{\sum_{j \in J} \min\{rx_{jt}(r), v_{j(t+1)} + g_{jt}[x_{jt}(r)]\}}{\sum_{j \in J} x_{jt}(r)} .$$

A Simple Special Case

j knows that $v_{j(t+1)} = 0$, $g_{jt}(x) = \nu_{jt}h_{jt}(x)$, and $h_{jt}(0) = 0$.

He believes there are two feasible values of ν_{jt}:

$\nu_{0jt} = 0$ and $\nu_{1jt} > 0$.

He places subjective probabilities (p_{0jt}, p_{1jt}) on ν_{jt} and chooses x to maximize expected utility

$$p_{0jt} \cdot u_{j(t+1)}(0) + p_{1jt} \cdot u_{j(t+1)}\{\max[0, \nu_{1jt}h_{jt}(x) - rx]\}.$$

The optimization problem reduces to the perfect foresight problem $\max_{x \geq 0} \nu_{1jt}h_{jt}(x) - rx$.

Thus, loan demand with partial knowledge is the same as what demand would be if the borrower knew that $\nu_{jt} = \nu_{1jt}$.

The reason is bankruptcy protection. The borrower knows that if he receives a bad draw on ν, he will realize the utility of zero consumption regardless of what magnitude loan he demands. Hence, expected utility maximization ignores the possibility of a bad draw and optimizes for the case of a good draw.

In the absence of bankruptcy protection, expected utility would be

$$p_{0jt} \cdot u_{j(t+1)}(-rx) + p_{1jt} \cdot u_{j(t+1)}[\nu_{1jt} h_{jt}(x) - rx].$$

If the borrower is risk-neutral, expected utility is

$$p_{1jt} \cdot \nu_{1jt} h_{jt}(x) - rx.$$

Loan demand at rate r without bankruptcy protection equals demand at rate r/p_{1jt} with bankruptcy protection.

The loan return function $\lambda_t(\cdot)$ has a simple form if

* borrowers have rational expectations with common probabilities (p_{0t}, p_{1t})

* realizations of the parameter values are statistically independent across borrowers

* the population J of borrowers is "large."

Then $\lambda_t(r) = p_{1t}r$.

The Supply of Loans

Let lender k be endowed with m_{kt} dollars.

The lender allocates this asset between loans and a safe investment. A dollar invested in the safe investment at t returns a known value ρ_t at t + 1.

The lender chooses a fraction $\delta \in [0, 1]$, implying that he allocates δm_{kt} dollars to loans and $(1 - \delta)m_{kt}$ dollars to the safe investment.

If k chooses allocation δ, his asset endowment next period will be $[\delta\lambda_t(r) + (1 - \delta)\rho_t]m_{kt}$.

He wants to solve

$$(4) \qquad \max_{\delta \in [0, 1]} \quad f_k\{[\delta\lambda_t(r) + (1 - \delta)\rho_t]m_{kt}\},$$

where $f_k(\cdot)$ is strictly increasing.

The optimal allocation is

$\delta = 1$ if $\lambda_t(r) > \rho_t$ and $\delta = 0$ if $\lambda_t(r) < \rho_t$.

Our concern is asset allocation with partial knowledge of loan repayment.

Let Γ_{kt} denote the states of nature that k thinks feasible for t+1, given the information he has at t.

Let $\lambda_{0kt} \equiv \min_{\gamma \in \Gamma_{kt}} \lambda_{\gamma}$ and $\lambda_{1kt} \equiv \max_{\gamma \in \Gamma_{kt}} \lambda_{\gamma}$. The optimal allocation is indeterminate if $\lambda_{0kt} < \rho_t < \lambda_{1kt}$.

We do not argue that lenders "should" use a particular decision criterion. If the optimal allocation is indeterminate, there is no "right" way for lenders to make decisions.

Bayesian Lending

Let π denote the subjective distribution. The lender solves

$$(5) \qquad \max_{\delta \in [0, 1]} \int f\{[\delta\lambda_\gamma + (1 - \delta)\rho]m\} d\pi.$$

Let $\lambda_M \equiv \int \lambda_\gamma d\pi$ be the subjective mean of λ. If $f(\cdot)$ is convex, the solution is $\delta = 0$ when $\lambda_M < \rho$ and $\delta = 1$ when $\lambda_M > \rho$. All $\delta \in [0, 1]$ are solutions if $\lambda_M = \rho$.

The solution may be fractional if $f(\cdot)$ has strictly concave segments.

The Maximin Criterion

The criterion is

$$(6) \qquad \max_{\delta \in [0,\, 1]} \; \min_{\gamma \in \Gamma} \; f\{[\delta\lambda_\gamma + (1 - \delta)\rho]m\}.$$

The solution is $\delta = 0$ if $\lambda_0 < \rho$ and $\delta = 1$ if $\lambda_0 > \rho$.

All $\delta \in [0, 1]$ are solutions if $\lambda_0 = \rho$.

The Minimax-Regret Criterion

The regret of allocation δ in state of nature γ is

$$\max\,[f(\lambda_\gamma m), f(\rho m)] - f\{[\delta\lambda_\gamma + (1-\delta)\rho]m\}.$$

The minimax-regret criterion is

$$(7) \quad \min_{\delta \in [0,1]} \; \max_{\gamma \in \Gamma} \; \max\,[f(\lambda_\gamma m), f(\rho m)] - f\{[\delta\lambda_\gamma + (1-\delta)\rho]m\}.$$

The solution to (7) is always fractional under ambiguity; that is, when $\lambda_0 < \rho < \lambda_1$. The minimax-regret allocation takes a simple form if $f(\cdot)$ is linear. Then

$$(8) \quad \delta_{MR} = \min\,[\max\,(\frac{\lambda_1 - \rho}{\lambda_1 - \lambda_0}, 0), 1].$$

Equilibrium Contracted Repayment Rates and Loan Transactions

The supply of loans at rate r is the possibly set-valued mapping

$$(10) \quad S_t(r) = \sum_{k \in K} \delta_{kt}(r) m_{kt}.$$

Rate r equilibrates supply with demand if

$$(11) \quad D_t(r) \in \sum_{k \in K} \delta_{kt}(r) m_{kt}.$$

We often assume that all lenders believe the highest feasible loan return is r and the lowest feasible return is $\alpha_t r$, for some $\alpha_t \in [0, 1]$.

Common Knowledge of the Loan Return

Let $\lambda_t(\cdot)$ be common knowledge. Let $m_t \equiv \sum_{k \in K} m_{kt}$.

Lender k sets $\delta_{kt}(r) = 0$ if $\lambda_t(r) < \rho_t$ and $\delta_{kt}(r) = 1$ if $\lambda_t(r) > \rho_t$. He is indifferent among all $\delta \in [0, 1]$ if $\lambda_t(r) = \rho_t$.

r equalizes supply and demand if

(12) $$1[\lambda_t(r) > \rho_t]m_t \leq D_t(r) \leq 1[\lambda_t(r) \geq \rho_t]m_t,$$

A *full-supply* equilibrium has $\lambda_t(r) > \rho_t$ and $D_t(r) = m_t$.

An *indifferent-supply* equilibrium has $\lambda_t(r) = \rho_t$ and $D_t(r) \in [0, m_t]$.

A *zero-supply* equilibrium has $\lambda_t(r) < \rho_t$ and $D_t(r) = 0$.

There exists a unique equilibrium with positive loan transactions if $\lambda_t(\cdot)$ satisfies the *single-crossing property* with respect to ρ_t:

That is, there exists a unique r_t^* such that

$$\lambda_t(r) < \rho_t \text{ for } r < r_t^*$$

$$\lambda_t(r_t^*) = \rho_t,$$

$$\lambda_t(r) > \rho_t \text{ for } r > r_t^*.$$

Let $D_t(r_t^*) > 0$.

If $D_t(r_t^*) \leq m_t$, then r_t^* is the unique equilibrium.

If $D_t(r_t^*) > m_t$, equilibrium is at $r > r_t^*$ such that $D_t(r) = m_t$.

In the special case discussed earlier, $\lambda_t(r) = p_{1t}r$ and $r_t^* = \rho_t/p_{1t}$.

Lender k sets $\delta_{kt}(r) = 0$ if $\lambda_{Mkt}(r) < \rho_t$ and $\delta_{kt}(r) = 1$ if $\lambda_{Mkt}(r) > \rho_t$. He is indifferent among all $\delta \in [0, 1]$ if $\lambda_{Mkt}(r) = \rho_t$.

An equilibrium value of r satisfies

$$(13) \quad \sum_{k \in K} 1[\lambda_{Mkt}(r) > \rho_t]m_{kt} \leq D_t(r) \leq \sum_{k \in K} 1[\lambda_{Mkt}(r) \geq \rho_t]m_{kt}.$$

Suppose every lender gives λ a uniform distribution on an interval $[\alpha_t r, r]$. Then $\lambda_{Mkt}(r) = (\alpha_t + 1)r/2$ and (13) becomes

$$(14) \quad 1[r > 2\rho_t/(\alpha_t + 1)]m_t \leq D_t(r) \leq 1[r \geq 2\rho_t/(\alpha_t + 1)]m_t.$$

The equilibrium r is

$r = 2\rho_t/(\alpha_t + 1)$ if $0 < D_t[2\rho_t/(\alpha_t + 1)] \leq m_t$

r such that $D_t(r) = m_t$ if $D_t[2\rho_t/(\alpha_t + 1)] > m_t$.

No loans are transacted if $D_t[2\rho_t/(\alpha_t + 1)] = 0$.

Maximin Decision Making

The maximin allocation for lender k at repayment rate r is $\delta_{kt}(r) = 0$ if $\lambda_{0kt}(r) < \rho_t$ and $\delta_{kt}(r) = 1$ if $\lambda_{0kt}(r) > \rho_t$.
All $\delta \in [0, 1]$ are maximin solutions if $\lambda_{0kt}(r) = \rho_t$.

An equilibrium value of r satisfies

$$(15) \qquad \sum_{k \in K} 1[\lambda_{0kt}(r) > \rho_t]m_{kt} \;\leq\; D_t(r) \;\leq\; \sum_{k \in K} 1[\lambda_{0kt}(r) \geq \rho_t]m_{kt}.$$

Suppose each lender sets $\lambda_{0kt}(r) = \alpha_t r$. Then (15) becomes

$$(16) \qquad 1[r > \rho_t/\alpha_t]m_t \;\leq\; D_t(r) \;\leq\; 1[r \geq \rho_t/\alpha_t]m_t.$$

The equilibrium is

$r = \rho_t/\alpha_t$ if $0 < D_t(\rho_t/\alpha_t) \leq m_t$

r such that $D_t(r) = m_t$ if $D_t(\rho_t/\alpha_t) > m_t$.

No loans are transacted if $D_t(\rho_t/\alpha_t) = 0$.

The MR allocation for lender k at r is

$$(17) \quad \delta_{MRkt}(r) = \min\left[\max\left(\frac{\lambda_{1kt}(r) - \rho_t}{\lambda_{1kt}(r) - \lambda_{0kt}(r)}, 0\right), 1\right].$$

An equilibrium value of r solves the equation

$$(18) \qquad D_t(r) = \sum_{k \in K} \min\left[\max\left(\frac{\lambda_{1kt}(r) - \rho_t}{\lambda_{1kt}(r) - \lambda_{0kt}(r)}, 0\right), 1\right] m_{kt}.$$

Suppose that every lender k sets $\lambda_{0kt}(r) = \alpha_t r$ and $\lambda_{1kt}(r) = r$. Then (18) becomes

$$(19) \qquad D_t(r) = \min\left[\max\left(\frac{r - \rho_t}{(1 - \alpha_t) r}, 0\right), 1\right] m_t.$$

Market Dynamics Following a Temporary Productivity Shock

The market initially is in steady state, with common knowledge of the loan return.

The steady-state is the indifferent-supply equilibrium, with r^* such that $\lambda(r^*) = \rho$.

A temporary negative productivity shock occurs at $t = 1$, yielding loan return $\lambda_1(r^*) < \rho$.

The shock does not affect borrower behavior.

Lenders learn $\lambda_1(r^*)$ at $t = 2$. They then form beliefs about $\lambda_2(\cdot)$ and choose allocations $\delta_{k2}(\cdot)$, $k \in K$.

The market dynamics depend on how lenders interpret the shock, what decision criteria they use, and how they revise their beliefs as new empirical evidence accumulates.

The Immediate Response to the Shock

We conjecture that a negative shock to the loan return will induce lenders to revise downward their beliefs about future loan returns.

Whatever decision criteria lenders use (Bayes, maximin, or minimax-regret), the supply of loans will fall at $t = 2$.

Hence, $r_2 > r^*$ and $D(r_2) < D(r^*)$.

Let the steady state have common knowledge of full loan repayment; thus, $\lambda(r) = r$ and $r^* = \rho$.

After the shock, let all lenders believe that the highest feasible loan return is r and the lowest feasible return is $\alpha_2 r$.

If lenders have linear utility and are Bayesian with uniform subjective distributions on $\lambda_2(\cdot)$, then $r_2 = 2\rho/(\alpha_2 + 1)$.

If lenders use the maximin criterion, then $r_2 = \rho/\alpha_2$.

If lenders have linear utility and use the minimax-regret criterion, then r_2 solves the equation

$$D(r_2) = m_2(r_2 - \rho)/[(1 - \alpha_2)r_2].$$

It is the case that $r_2 \in (\rho, \rho/\alpha_2)$.

Subsequent Market Outcomes

We cannot make a sharp prediction on long run market dynamics.

There are many plausible ways in which lenders could revise beliefs and make decisions as empirical evidence accumulates. Hence, there are many plausible sequences of market outcomes.

A Favorable Scenario

Let $D(r_2) > 0$.

At $t = 3$, lenders who supplied positive loans at $t = 2$ observe that their loans are fully repaid; that is, $\lambda_2(r_2) = r_2$.

Lenders who supplied no loans at $t = 2$ do not directly observe the loan return. However, suppose they learn it indirectly through communication across lenders.

We conjecture that full repayment of the loans transacted at $t = 2$ is interpreted as evidence that the shock was temporary. Hence, lenders revise their beliefs upward to some degree and increase the supply of loans relative to $t = 2$.

The result is $r^* < r_3 < r_2$ and $D(r_2) < D(r_3) < D(r^*)$.

An Unfavorable Scenario

The least favorable scenario occurs if lender beliefs and actions make r_2 so high that $D(r_2) = 0$, shutting down the credit market.

Then lenders do not observe a loan return at $t = 2$ and, hence, obtain no empirical evidence that the shock was temporary. The situation repeats itself at $t = 3$.

Thus, the credit market could disappear permanently unless some lenders decide to experiment, making loans that currently appear unprofitable in order to update their beliefs about the loan return.

Government Intervention in the Credit Market

We consider intervention to stabilize the credit market after a temporary productivity shock.

We measure welfare by the aggregate return to the investments financed by lenders' endowments.

The government knows the shock is temporary.

The safe investment is a government security with guaranteed return ρ_t.

The social return ρ^* measures the productivity of investments made with the funds the government obtains from lenders.

The Social Welfare Function

Each period, an Authority chooses from a set of feasible policies.

Let (x_{jtc}, $j \in J$) be the equilibrium loan allocation if the Authority implements policy c.

Welfare under policy c is

$$W_t(c) \equiv \sum_{j \in J} g_{jt}(x_{jtc}) + \rho^*(m_t - \sum_{j \in J} x_{jtc}) = \sum_{j \in J} [g_{jt}(x_{jtc}) - \rho^* x_{jtc}] + \rho^* m_t.$$

$\sum_{j \in J} g_{jt}(x_{jtc})$ is the return on the investments made with loan financing.

$\rho^*(m_t - \sum_{j \in J} x_{jtc})$ is the social return on the assets that lenders allocate to the safe investment.

Restoration of a Steady State with Common Knowledge of Full Repayment

Let the market initially be in steady state, with common knowledge of full loan repayment. The safe investment is offered at rate ρ^*. Let $0 < D(\rho^*) \leq m_t$. Then the steady-state has $r^* = \rho^*$.

A shock occurs at $t = 1$. The Authority knows the shock is temporary, but it cannot credibly communicate this to lenders. It can only provide incentives to lenders.

The Authority should choose a policy to restore the steady state equilibrium. The reason is that the steady state maximizes the aggregate return to lender assets.

With perfect foresight and $r^* = \rho^*$, each borrower j chooses x to maximize $g_j(x) - \rho^* x$. Thus, decentralized borrowing maximizes the aggregate return.

Manipulation of the Return on the Safe Investment

Suppose the Authority can manipulate ρ_2, making it deviate from ρ^*. If the Authority knows how lenders behave, it can set ρ_2 to restore the credit market to the steady state.

Uniform-Prior Bayesian Decision Making with Linear Utility: With no intervention, the equilibrium is

$$r_2 = 2\rho_2/(\alpha_2 + 1).$$

To make $r_2 = \rho^*$, the Authority should set $\rho_2 = \rho^*(\alpha_2 + 1)/2$.

Maximin Decision Making: With no intervention, $r_2 = \rho_2/\alpha_t$. To make $r_2 = \rho^*$, the Authority should set $\rho_2 = \alpha_2\rho^*$.

Minimax-Regret Decision Making with Linear Utility: With no intervention, r_2 solves $D(r_2) = m_2(r_2 - \rho_2)/[(1 - \alpha_2)r_2]$. To make $r_2 = \rho^*$, the Authority should set

$$\rho_2 = \rho^* - [(1 - \alpha_2)\rho^*]D(\rho^*)/m_2.$$

Guaranteeing a Minimum Loan Return

Suppose that the Authority can guarantee to lenders that the loan return will equal the contracted repayment rate.

Then lenders behave as before the shock. Borrowers are observationally identical, so the guarantee does not give an incentive to make bad loans.

The guarantee restores the steady state. Borrowers fully repay their loans, so the Authority does not have to pay off on the guarantee.

Implementation of this policy does not require the Authority to know how lenders behave. This contrasts with manipulation of the return on the safe investment.

Intervention in a Credit Market with Partial Loan Repayment

Let the market initially be in steady state with common knowledge of the loan return, which yields less than full repayment.

In steady state, the safe investment is offered at return ρ, which need not equal ρ^*.

Assume that $0 < D(\rho) \leq m$ and that the loan-return function $\lambda(\cdot)$ satisfies the single-crossing property with respect to ρ.

Then the steady-state is r^* such that $\lambda(r^*) = \rho$.

Let the market experience a temporary productivity shock.

Guaranteeing a Minimum Loan Return

Let the Authority guarantee that the loan return at rate r will at least equal the steady-state return $\lambda(r)$.

Then lenders behave as they did before the shock. The guarantee restores the steady state.

The realized repayment rate is $\lambda(r^*)$, so the Authority does not have to pay off on the guarantee.

This guarantee requires the Authority to know the steady-state loan return function $\lambda(\cdot)$.

Observation of the pre-shock steady state reveals that $\lambda(r^*) = \rho$, but it does not reveal $\lambda(r)$ at other r.

Fortunately, the Authority does not need to know the full structure of $\lambda(\cdot)$ to make an effective guarantee.

Consider any guarantee function $\mu(\cdot)$ that satisfies the single-crossing property with respect to ρ; that is,

$$\mu(r) < \rho \text{ for } r < r^*, \ \mu(r^*) = \rho, \text{ and } \mu(r) > \rho \text{ for } r > r^*.$$

Any such guarantee induces steady state behavior by lenders.

Setting the Steady-State Return on the Safe Investment

How should the Authority set ρ?

Welfare is maximized if the investment made by each borrower j solves $\max_{x \geq 0} g_j(x) - \rho^* x$.

A borrower with partial knowledge of $g_j(\cdot)$ cannot solve the optimization problem.

This opens the possibility that the Authority can improve welfare by setting ρ to a value other than ρ^*.

Let C be the set of feasible values for ρ.

Let $x_{j\rho}$ be borrower j's loan demand in the market equilibrium that would occur if the Authority were to set the return on the safe investment equal to ρ.

Steady-state welfare under policy ρ is

$$W(\rho) \equiv \sum_{j \in J} g_j(x_{j\rho}) + \rho^*(m - \sum_{j \in J} x_{j\rho}) = \sum_{j \in J} [g_j(x_{j\rho}) - \rho^* x_{j\rho}] + \rho^* m.$$

The Authority should choose ρ to maximize $W(\cdot)$.

However, this presumes the Authority knows $\{[x_{j\rho}, g_j(x_{j\rho})], j \in J\}, \rho \in C$.

A Credit Market with Observationally Heterogeneous Borrowers

Let borrowers have observable covariates. Let lenders be able to price loans differentially to borrowers with different covariates.

Steady State Equilibrium with Common Knowledge of the Loan Return

Analysis is a straightforward extension of our earlier work.

A Simple Special Case

Let each class z of borrowers satisfy the special assumptions described earlier, with z-specific parameters.

Then $\lambda_z(r) = p_{1z}r$ and $r_z^* = \rho/p_{1z}$. The equilibrium is of the indifferent-supply type if $\sum_{z \in Z} D_z(\rho/p_{1z}) \leq m$.

Market Dynamics and Government Intervention after Temporary Shocks

A Temporary Productivity Shock

The Authority can use a loan guarantee to restore the steady state. The argument is as earlier, except that now the Authority makes z-specific guarantees.

A Temporary Securitization Shock

Discussions of the American credit crisis have cited deceptive securitization of mortgage loans as an antecedent event.

The idea is that the mortgage market was initially in steady state, with mortgages priced differentially to borrowers with different characteristics—prime and subprime borrowers.

A shock may have occurred when loan originators began to offer prime-rate mortgages to subprime borrowers, bundling these loans with those made to prime borrowers.

Loan originators sold the bundles to lenders who believed that the bundles would have the high repayment rate of prime borrowers.

Ex post, the loan bundles had lower repayment rates, reflecting the mix of prime and sub-prime loans.

The deceptive bundling process was ended when regulators stopped it; hence, the shock was temporary. However, lenders did not know how to interpret the shock and reduced loan supply.

An Authority who understands the shock is temporary and who is able to prevent further deception can use a loan guarantee to restore the steady state.

Concluding Comments

1. We think it highly desirable to study intervention when the Authority, like lenders, does not know how to interpret an unanticipated shock.

2. Our model does not fit the mortgage market.

* We assume that the return on a borrower's investment depends only on the magnitude of his investment. However, the return on a home purchase depends on future home prices, which are determined by aggregate home purchase decisions.

* Our social welfare function is inappropriate for study of the mortgage market. We have assumed that the private and social returns to loan-financed investments coincide.

TREATMENT CHOICE WITH SAMPLE DATA

The ambiguity in treatment choice studied thus far arose purely out of identification problems.

In practice, a planner may observe only a random sample of the study population. This generates further ambiguity.

The standard practice is to estimate point-identified population features by sample analogs. Econometricians appeal to asymptotic theory to justify the practice.

Asymptotic theory gives at most approximate guidance to a planner who must make treatment choices using sample data.

The limit theorems of asymptotic theory describe the behavior of estimates as sample size increases to infinity.

They do not reveal how estimates perform in specific finite-sample settings.

A planner’s objective is not to obtain estimates with good asymptotic properties but rather to choose a good treatment rule with the data available.

The Wald (1950) development of statistical decision theory addresses treatment choice with sample data directly, without recourse to asymptotic approximations.

It seamlessly integrates the study of identification and statistical inference.

Wald developed the principles of statistical decision theory in the 1930s and 1940s. Important extensions and applications followed in the 1950s, but this period of rapid development came to a close by the 1960s.

Why did statistical decision theory lose momentum?

* One reason may have been the technical difficulty of the subject. Wald's ideas are fairly easy to describe in the abstract, but applying them tends to be analytically and numerically demanding.

* Another reason may have been diminishing interest in decision making as the motivation for statistical analysis. Modern statisticians and econometricians tend to view their objectives as estimation and hypothesis testing rather than decision making.

* A third contributing factor may have been the criticism of Wald's thinking put forward by those decision theorists who espouse the *conditionality principle* as a sine qua non of statistical decision making.

Wald's Development of Statistical Decision Theory

Wald considered the broad problem of using sample data to make a decision.

His world view eliminates the common separation of activities between empirical research and decision making. The researcher and decision maker are the same person.

Wald posed the task as choice of a *statistical decision function*, which maps the available data into a choice among the feasible actions.

In a treatment-choice setting, a statistical decision function is a rule for using the data to choose a treatment allocation.

I call such a rule a *statistical treatment rule.*

No statistical decision function that makes non-trivial use of sample data can perform best in every realization of a sampling process.

Hence, Wald recommended evaluation of statistical decision functions as *procedures* applied as the sampling process is engaged repeatedly to draw independent data samples.

The idea of a procedure transforms the original statistical problem of induction from a single sample into the deductive problem of assessing the probabilistic performance of a statistical decision function across realizations of the sampling process.

Admissibility

Wald suggested comparison of statistical decision functions by mean performance across realizations of the sampling process.

Wald termed this *risk*. Here, where the goal is maximization of a social welfare function, I call it *expected welfare*.

The most basic prescription of Wald's statistical decision theory is that a decision maker should not choose an action that is dominated in risk.

Such an action is *inadmissible*. An action that is not dominated in risk is *admissible*.

Admissibility is a desirable property, but its operational implications are limited for two main reasons.

First, when it is possible to determine which statistical decision functions are admissible, there often turn out to be many such functions.

Second, there are many settings of practical interest where analysis of admissibility is technically challenging.

Implementable Criteria for Treatment Choice

To develop implementable criteria for decision making with sample data, statistical decision theorists have studied the same broad ideas that were discussed earlier, but now applied to risk.

The Bayesian prescription uses a statistical decision function that works well on average across the feasible states of nature.

The maximin and minimax-regret prescriptions use a decision function that, in one of two senses, works well uniformly over Γ.

Using a Randomized Experiment to Evaluate an Innovation

Outcomes are binary, there are no observed covariates, and there are two treatments, one being the status quo and the other being an innovation.

The planner knows the response distribution of the status quo treatment, but not that of the innovation.

To learn about the innovation, a classical randomized experiment is performed.

The problem is to use the experimental data to inform treatment choice.

The Setting

Let $t = a$ be the status quo treatment and $t = b$ be the innovation.

The planner knows $\alpha \equiv P[y(a) = 1]$, but not $\beta \equiv P[y(b) = 1]$.

The planner wants to choose a treatment allocation that maximizes the success probability.

An experiment is performed to learn about outcomes under the innovation, with N subjects randomly drawn from the population and assigned to treatment b. There is full compliance with the assigned treatment.

n subjects realize $y = 1$. The remaining $N - n$ realize $y = 0$. These outcomes are observed.

This classical experiment point-identifies β. The planner only faces a problem of statistical inference.

The sample size N indexes the sampling process and the number n of experimental successes is a sufficient statistic for the sample data.

The feasible statistical treatment rules are functions

$$\delta(\cdot): [0, . . . , N] \rightarrow [0, 1]$$

that map the number of experimental successes into a treatment allocation.

For each value of n, rule $\delta(\cdot)$ randomly allocates a fraction $\delta(n)$ of the population to treatment b and $1 - \delta(n)$ to a.

The expected welfare of rule δ is

$$W(\delta, P, N) = \alpha\cdot E[1 - \delta(n)] + \beta\cdot E[\delta(n)] = \alpha + (\beta - \alpha)\cdot E[\delta(n)].$$

The number of experimental successes is distributed binomial **B**[β, N], so

$$E[\delta(n)] = \sum_{i=0}^{N} \delta(i) \cdot f(n = i; \beta, N),$$

where $f(n = i; \beta, N) \equiv N![i! \cdot (N - i)!]^{-1} \beta^i (1 - \beta)^{N-i}$ is the Binomial probability of i successes.

The only unknown determinant of expected welfare is β. Hence, Γ indexes the feasible values of β.

Let $\beta_\gamma \equiv P_\gamma[y(b) = 1]$. Suppose that $(\beta_\gamma, \gamma \in \Gamma)$ contains values that are smaller and larger than α.

The Admissible Treatment Rules

It is reasonable to conjecture that admissible treatment rules should be ones in which the fraction of the population allocated to treatment b increases with n. The admissible treatment rules are a simple subclass of these rules.

Karlin and Rubin (1956) define a *monotone treatment rule* to be one of the form

$$\delta(n) = 0 \quad \text{for } n < n_0,$$

$$\delta(n) = \lambda \quad \text{for } n = n_0,$$

$$\delta(n) = 1 \quad \text{for } n > n_0,$$

where $0 \leq n_0 \leq N$ and $0 \leq \lambda \leq 1$ are constants specified by the planner.

Let $0 < \alpha < 1$ and let the feasible set $(\beta_\gamma, \gamma \in \Gamma)$ exclude the values 0 and 1. Karlin and Rubin shows that the collection of monotone treatment rules is the set of admissible rules.

Some Monotone Rules

The collection of monotone treatment rules is a mathematically "small" subset of the space of all feasible treatment rules, but it still contains a broad range of rules.

Data-Invariant Rules: The rules $\delta(\cdot) = 0$ and $\delta(\cdot) = 1$, which assign all persons to treatment a or b respectively, whatever the realization of n may be.

Empirical Success Rules: An optimal treatment rule allocates all persons to treatment a if $\beta < \alpha$ and all to treatment b if $\beta > \alpha$. An empirical success rule replaces β with its sample analog, the empirical success rate n/N. Thus,

$$\delta(n) = 0 \quad \text{for } n < \alpha N,$$
$$\delta(n) = \lambda \quad \text{for } n = \alpha N, \quad \text{where } 0 \leq \lambda \leq 1,$$
$$\delta(n) = 1 \quad \text{for } n > \alpha N.$$

Bayes Rules: Consider the class of Beta priors, the conjugate family for a Binomial likelihood. Let $(\beta_\gamma, \gamma \in \Gamma) = (0, 1)$ and let the prior be Beta with parameters (c, d). The posterior mean for β is $(c + n)/(c + d + N)$. The resulting Bayes rule is

$$\delta(n) = 0 \quad \text{for } (c + n)/(c + d + N) < \alpha,$$
$$\delta(n) = \lambda \quad \text{for } (c + n)/(c + d + N) = \alpha, \quad \text{where } 0 \le \lambda \le 1,$$
$$\delta(n) = 1 \quad \text{for } (c + n)/(c + d + N) > \alpha.$$

As (c, d) tend to zero, the Bayes rule approaches an empirical success rule. The class of Bayes rules includes the data-invariant rules $\delta(\cdot) = 0$ and $\delta(b, \cdot) = 1$. The former occurs if the parameters (c, d) of the Beta prior distribution satisfy $(c + N)/(c + d + N) < \alpha$. The latter occurs if $c/(c + d + N) > \alpha$.

Statistical Significance Rules: These rules use a one-sided hypothesis test to choose between the status quo treatment and the innovation. The null hypothesis is that both treatments yield the same social welfare; that is, $\beta = \alpha$. The alternative is that treatment b is superior to treatment a; that is, $\beta > \alpha$. Treatment b is chosen if the null is rejected, and treatment a otherwise. Thus, the rule is

$$\delta(n) = 0 \quad \text{for} \quad n \le d(\alpha, s, N),$$

$$\delta(n) = 1 \quad \text{for} \quad n > d(\alpha, s, N),$$

where s is the specified size of the test and $d(\alpha, s, N)$ is the associated critical value. With n is binomial, $d(\alpha, s, N) = \min i: f(n > i; \alpha, N) \le s$.

Although statistical significance rules are monotone treatment rules, the conventional practice of hypothesis testing is remote from the problem of treatment choice with sample data.

If the null hypothesis [$\beta = \alpha$] is correct, all feasible treatment rules yield the same expected welfare. If not, alternative rules may yield different expected welfare.

A statistical test indicates only whether the sample data are inconsistent (in the usual sense of having low probability of being realized under the null) with the hypothesis that all feasible rules yield the same expected welfare.

The Maximin Rule: Minimum expected welfare for rule δ is

$$\min_{\gamma \in \Gamma} W(\delta, P_\gamma, N) = \alpha + \min_{\gamma \in \Gamma} (\beta_\gamma - \alpha)E_\gamma[\delta(n)],$$

$E_\gamma[\delta(n)] > 0$ for all $\beta_\gamma > 0$ and all monotone treatment rules except for $\delta(\cdot) = 0$, the rule that always chooses treatment a.

Hence, the maximin rule is the data-invariant rule $\delta(\cdot) = 0$.

The Minimax-Regret Rule: The regret of rule δ in state of nature γ is

$$\max(\alpha, \beta_\gamma) - \{\alpha + (\beta_\gamma - \alpha)\cdot E_\gamma[\delta(n)]\} =$$

$$(\beta_\gamma - \alpha)\{1 - E_\gamma[\delta(n)]\}\cdot 1[\beta_\gamma \geq \alpha] + (\alpha - \beta_\gamma)E_\gamma[\delta(n)]\cdot 1[\alpha \geq \beta_\gamma].$$

Thus, regret is the mean welfare loss when a member of the population is assigned the inferior treatment, multiplied by the expected fraction of the population assigned this treatment.

The minimax-regret rule does not have an analytical solution but it can be determined numerically.

It is very close to the empirical success rule.

Savage on the Maximin and Minimax-Regret Criteria

Treatment choice using the minimax-regret rule differs fundamentally from treatment choice using the maximin rule.

Savage (1951), whose review of Wald (1950) first explicitly distinguished between these criteria for decision making, argued strongly against application of the minimax (here maximin) criterion, writing (p. 63):

"Application of the minimax rule . . . . is indeed ultra-pessimistic; no serious justification for it has ever been suggested, and it can lead to the absurd conclusion in some cases that no amount of relevant experimentation should deter the actor from behaving as though he were in complete ignorance."

Our finding that the maximin treatment rule is data-invariant illustrates this “absurd conclusion.”

Savage emphasized that although the minimax criterion is “ultra-pessimistic,” the minimax-regret criterion is not.

Our finding that the minimax-regret rule approximates the empirical success rule illustrates that the minimax-regret criterion is not particularly pessimistic.

ADMISSIBLE TREATMENT RULES FOR A RISK-AVERSE PLANNER WITH EXPERIMENTAL DATA ON AN INNOVATION

Consider the treatment choice problem of Lecture 5 except that the objective is maximization of a concave-monotone function f(·) of the success rate.

Thus, the planner is "risk-averse."

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

The Planning Problem

There are two treatments, t = a being the status quo and t = b the innovation.

The population success rates if everyone were to receive the same treatment are $\alpha \equiv P[y(a) = 1]$ and $\beta = P[y(b) = 1]$.

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

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

Welfare is $f[\alpha + (\beta - \alpha)\delta]$, where $f(\cdot)$ is an increasing, concave transformation of the success rate.

The Empirical Evidence and Admissible Treatment Rules

α is known. Evidence on β comes from a randomized experiment, where N subjects are drawn at random and assigned to treatment b. Of these subjects, n experience y(b) = 1 and N − n experience y(b) = 0. The outcomes of all subjects are observed.

N indexes the sampling process and n is a sufficient statistic for the sample data. The feasible rules are the functions $z(\cdot): [0, \ldots, N] \rightarrow [0, 1]$ that map the number of experimental successes into a treatment allocation.

Thus, if there are n successes, rule z allocates a fraction z(n) of the population to treatment b and the remaining 1 − z(n) to treatment a.

Following Wald, we evaluate a rule by its expected performance across repeated samples.

Let $p(n; \beta) \equiv N![n!\cdot(N - n)!]^{-1}\beta^n(1 - \beta)^{N-n}$ denote the Binomial probability of n successes in N trials.

The expected welfare of rule $z(\cdot)$ across repeated samples is

$$(2) \qquad W(z; \beta) \equiv \sum_{n=0}^{N} p(n; \beta)\cdot f[\alpha + (\beta - \alpha)\cdot z(n)].$$

Some Definitions

Let B index the set of *states of nature* (or parameter space).

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

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

A class of rules is *complete* if, given any rule not in the class, there exists a member of the class that performs at least as well in all states of nature and better in some state of nature.

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

Admissible Rules for a Risk-Neutral Planner

Let f(·) be the identity function. Then the expected welfare of rule z is

(3) $$W(z; \beta) = \alpha + (\beta - \alpha)E_\beta[z(n)],$$

where $E_\beta[z(n)] = \sum_n p(n; \beta)z(n)$.

z is admissible if there exists no z′ such that

$$(\beta - \alpha)E_\beta[z(n) - z'(n)] \leq 0 \text{ for all } \beta \in B,$$

$$(\beta - \alpha)E_\beta[z(n) - z'(n)] < 0 \text{ for some } \beta \in B.$$

A KR-monotone treatment rule has the form

(4)
$$z(n) = 0 \quad \text{for } n < k,$$
$$z(n) = \lambda \quad \text{for } n = k,$$
$$z(n) = 1 \quad \text{for } n > k,$$

where $0 \leq k \leq N$ and $0 \leq \lambda \leq 1$.

Karlin and Rubin (1956, Theorem 4) shows that the admissible and KR-monotone rules coincide if B excludes the extreme values 0 and 1.

Admissible Rules for Risk-Averse Planners

The Binomial density function possesses the strict form of the monotone-likelihood ratio property:

$$(n > n', \beta > \beta') \Rightarrow p(n; \beta)/p(n; \beta') > p(n'; \beta)/p(n'; \beta').$$

Thus, larger values of n are unambiguously evidence for larger values of β.

It is reasonable to conjecture that good treatment rules are ones that make the fraction of the population allocated to treatment b increase with n.

Define z to be *fractional monotone* if $n > n' \Rightarrow z(n) \geq z(n')$.

Proposition 1: If $f(\cdot)$ is weakly increasing and concave, the class of fractional monotone rules is essentially complete. If $f(\cdot)$ is also strictly increasing, the class of fractional monotone rules is complete.

Proposition 1 can be strengthened considerably if $f(\cdot)$ is differentiable with derivative function $g(\cdot)$ that does not decrease too rapidly.

Call z *M-step monotone* if $n < n' \Rightarrow z(n) \leq z(n')$ and, for a given positive integer M, $n + M \leq n' \Rightarrow z(n) = 0$ or $z(n') = 1$.

Proposition 2: Let $f(\cdot)$ be weakly increasing, concave and differentiable on $(\inf \{B\}, \sup \{B\})$, with derivative $g(\cdot)$.

If $[x(1-x)^{-1}]^M g(x)$ weakly increases with x, then the M-step monotone rules are an essentially complete class.

If $f(\cdot)$ is also strictly concave or if $[x(1-x)^{-1}]^M g(x)$ strictly increases with x, then the M-step monotone rules are a complete class.

Properties of KR-Monotone Rules

KR-monotone rules are M-step monotone rules with $M = 1$.

Proposition 2 shows that the classs of KR-monotone rules is essentially complete if $[x(1-x)^{-1}]g(x)$ weakly increases in x.

This class is complete if $f(\cdot)$ is strictly concave or if $[x(1-x)^{-1}]g(x)$ strictly increases with x.

The KR-monotone rules are a complete class if $f(x) = \log(x)$.

$f(x) = \log(x) - x$ has the threshold derivative $x^{-1}(1-x)$.

Proposition 3: Let f(·) be strictly increasing. If the class of KR-monotone rules is essentially complete, then every KR-monotone rule is admissible. If the class of KR-monotone rules is complete, then it is minimal complete.

Combining Propositions 2 and 3 yields

Proposition 3, Corollary: Let f(·) be strictly increasing and differentiable on (inf {B}, sup {B}), with $[x(1-x)^{-1}]g(x)$ weakly increasing in x. Then the class of KR-monotone rules is minimal complete if f(·) is strictly concave or if $[x(1-x)^{-1}]g(x)$ is strictly increasing.

However, it can be shown that some KR-monotone rules are inadmissible when f(·) has sufficiently strong curvature.

Bayes Rules

A Bayesian planner places a prior distribution Π on B.

Observing the number n of experimental successes in the randomized trial, he forms a posterior distribution $\Pi(\beta|n)$.

Treating β as a random variable with distribution $\Pi(\beta|n)$, the planner solves the problem

(16) $$\max_{\delta \in [0, 1]} \int f[\alpha + (\beta - \alpha)\delta]d\Pi(\beta|n).$$

Proposition 5: Consider (16). Let $\Pi(\beta|n)$ be non-degenerate. Let $E_{\Pi(\beta|n)}[\beta]$ denote the posterior mean of β.

(a) Let $f(\cdot)$ be strictly concave. Then the Bayes rule is unique, therefore admissible. It is $\delta = 0$ if $E_{\Pi(\beta|n)}[\beta] \leq \alpha$.

(b) Let $f(\cdot)$ be continuously differentiable. Let $f(\cdot)$ and $\Pi(\beta|n)$ be sufficiently regular that

$$\partial\{\int f[\alpha + (\beta - \alpha)\delta]d\Pi(\beta|n)\}/\partial\delta$$

$$= \int \{\partial f[\alpha + (\beta - \alpha)\delta]/\partial\delta\}d\Pi(\beta|n)$$

in a neighborhood of $\delta = 0$. Then a Bayes rule has

$\delta > 0$ if $E_{\Pi(\beta|n)}[\beta] > \alpha$.

$\delta \in (0, 1)$ if $E_{\Pi(\beta|n)}[\beta] > \alpha$ and $\int f(\beta)d\Pi(\beta|n) < f(\alpha)$.

Minimax-Regret Rule

For $\beta \in B$, max $[f(\alpha), f(\beta)]$ is the maximum welfare achievable given knowledge of β.

$W(z; \beta)$ is the expected welfare achieved by rule $z(\cdot)$. The difference between these quantities is regret

(17) $$R(z; \beta) \equiv \max [f(\alpha), f(\beta)] - W(z; \beta).$$

A minimax-regret rule z_{mmr} solves

(18) $$\inf_{z \in Z} \sup_{\beta \in B} R(z; \beta).$$

The minimax-regret rule is always fractional if $f(\cdot)$ has sufficient curvature.

Proposition 6: Let $\alpha > 0$. Let B contain a sequence of positive values that converges to zero. If

$$(19) \qquad \lim_{\beta \to 0_+} \beta^M \cdot f(\beta) = -\infty$$

for some $M \geq 0$, then $z_{mmr}(n) < 1$ for all $n \leq M$ regardless of sample size N.

Let $f(x) = -x^{-K}$, where $K > 1$. Then (19) holds for $M < K$ and $z_{mmr}(n) < 1$ for $n < K$.

Let $f(x) = -\exp(1/x)$. Then (19) holds for all M and $z_{mmr}(n) < 1$ for any n.

Treatment Choice in Practice

Use of Test-Based Rules in Medicine

Choice between a status quo treatment and an innovation occurs often in practice, but explicit use of statistical decision theory is rare.

In medicine, the branch of statistics that has strongly influenced practice has been hypothesis testing rather than decision theory.

In the U. S., testing the hypothesis of zero average treatment effect is institutionalized in the Food and Drug Administration drug approval process, which calls for comparison of a new treatment under study (t = b) with a placebo or an approved treatment (t = a).

FDA approval of the new treatment requires one-sided rejection of the null hypothesis of zero average treatment effect $\{H_0: E[y(b)] = E[y(a)]\}$ in clinical trials.

In the context of treatments with binary outcomes, this means performance of a test with null hypothesis $\{H_0: \beta = \alpha\}$ and alternative $\{H_1: \beta > \alpha\}$.

The classical practice of fixing the probability of a type I error and seeking to minimize the probability of a type II error is difficult to motivate 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 loss in welfare resulting from a sub-optimal choice.

See Tetenov (2007) for interpretation of test-based rules from an "asymmetric minimax-regret" perspective.

Even if statistical decision theory does not motivate test-based treatment rules, we can use decision theory to evaluate such rules.

A conventional test-based rule assigns treatment b to the entire population if the number of experimental successes is large enough to reject H_0 and assigns treatment a otherwise. Thus, a test-based rule has the form

$$(20) \quad z(n) = 0 \quad \text{for} \quad n \le d(s, \alpha),$$
$$z(n) = 1 \quad \text{for} \quad n > d(s, \alpha),$$

where s is the specified size of the test and $d(s, \alpha)$ is the associated critical value. Given that n is binomial, $d(s, \alpha) = \min i: p(n > i; \alpha) \le s$.

Test-based rules are KR-monotone, hence admissible if the welfare function has sufficiently weak curvature.

Numerical analysis for linear and log welfare functions shows that the maximum regret of the rule based on the exact binomial test with size $s = 0.05$ is much larger than that of the minimax-regret rule.

When the sample size is larger than ten, the ratio of the former maximum regret to the latter is typically about 5 to 1. This quantifies the inferiority of the test-based rule from the vantage of maximum regret.

However, the test-based rule is not inferior in all states of nature. Being admissible, this rule must yield smaller regret in some states of nature. The test-based rule delivers smaller regret than the minimax-regret rule in states of nature with $\beta < \alpha$ and larger regret in states with $\beta > \alpha$. The inferiority of the rule in terms of maximum regret arises because the latter losses are much larger than the former gains.

The maximum regret of the empirical success rule is close to that of the minimax-regret rule. The former rule is a close step-function approximation to the latter.

The minimax-regret rule is relatively easy to compute, but the empirical success rule is simpler yet.

Hence, a practitioner who is not equipped to compute the minimax-regret rule would suffer little by using the empirical success rule as an approximation.

NOTE: The above conclusion assumes binary outcomes. Considering linear welfare, Schlag (2007) and Stoye (2009) show that it does not extend to other cases where outcomes have bounded support. Tetenov (2007) shows that it does extend to the case where outcomes are normally distributed.

For Further Study:

Use of Covariate Information in Treatment Choice

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

Abstract

An important objective of empirical research on treatment response is to provide decision makers with information useful in choosing treatments. This paper studies minimax-regret treatment choice using the sample data generated by a classical randomized experiment.

Consider a utilitarian social planner who must choose among the feasible *statistical treatment rules*, these being functions that map the sample data and observed covariates of population members into a treatment allocation.

If the planner knew the population distribution of treatment response, the optimal treatment rule would maximize mean welfare conditional on all observed covariates.

The appropriate use of covariate information is a more subtle matter when only sample data on treatment response are available.

I consider the class of *conditional empirical success* rules; that is, rules assigning persons to treatments that yield the best experimental outcomes conditional on alternative subsets of the observed covariates.

I derive a closed-form bound on the maximum regret of any such rule. Comparison of the bounds for rules that conditional on smaller and larger subsets of the covariates yields *sufficient sample sizes* for productive use of covariate information.

When the available sample size exceeds the sufficiency boundary, a planner can be certain that conditioning treatment choice on more covariates is preferable (in terms of minimax regret) to conditioning on fewer covariates.

Stoye, J., "Minimax Regret Treatment Choice with Finite Samples," *Journal of Econometrics*, forthcoming.

Abstract

This paper applies the minimax regret criterion to choice between two treatments conditional on observation of a finite sample. The analysis is based on exact small sample regret and does not use asymptotic approximations nor finite-sample bounds.

Core results are:

(i) Minimax regret treatment rules are well approximated by empirical success rules in many cases, but differ from them significantly — both in terms of how the rules look and in terms of maximal regret incurred —for small sample sizes and certain sample designs.

(ii) Absent prior cross-covariate restrictions on treatment outcomes, they prescribe inference that is completely separate across covariates, leading to no-data rules as the support of a covariate grows.

I conclude by offering an assessment of these results.