

CeMMAP Masterclass November 6-7, 2023

Charles F. Manski

Department of Economics and Institute for Policy Research
Northwestern University

Theme: **Social Planning under Uncertainty**

Monday November 6

Lecture 1: Credible Social Planning under Uncertainty

Lecture 2: Diversified Treatment under Ambiguity, with Application to Tuberculosis
Diagnosis and Treatment

Lecture 3: Statistical Decision Properties of Imprecise Trials Assessing COVID-19 Drugs

Tuesday November 7

Lecture 4: Identification of Income-Leisure Preferences and Evaluation of Income Tax
Policy

Lecture 5: Minimax-Regret Climate Policy with Deep Uncertainty in Climate
Modeling and Intergenerational Discounting

Lecture 6: Looking Ahead

CREDIBLE SOCIAL PLANNING UNDER UNCERTAINTY

Charles F. Manski

Department of Economics and Institute for Policy Research

Northwestern University

The Preamble of the Constitution of the United States states:

“**We the People** of the United States, in Order to form a more perfect Union, establish Justice, insure domestic Tranquility, provide for the common defence, **promote the general Welfare**, and secure the Blessings of Liberty to ourselves and our Posterity, do ordain and establish this Constitution for the United States of America.”

Alfred Marshall (1890) began his *Principles of Economics* with the sentence:

“Political economy or economics is a study of mankind in the ordinary business of life; it examines that part of individual and social action which is most closely connected with the attainment and with the use of the material requisites of **wellbeing**.”

A report on clinical practice guidelines by the U.S. Institute of Medicine (2010) states:

“Clinical practice guidelines are statements that include recommendations intended to **optimize patient care** that are informed by a systematic review of evidence and an assessment of the benefits and harms of alternative care options.”

The Constitutional premise that the United States should promote the general welfare, Marshall's concern with social action to promote wellbeing, and the IOM premise that clinicians should optimize patient care exemplify broad assertions that entities making societal decisions should aim to maximize social welfare.

Such assertions have rhetorical appeal, but they become meaningful only when several questions are answered:

What constitutes social welfare? What are the feasible actions?

What is known about the welfare consequences of alternative choices?

Maximization of welfare is a well-defined objective if enough is known about the welfare consequences of alternative choices to determine an unambiguous best action. It is ill-defined if the consequences are sufficiently uncertain that no action is clearly best.

My concern is societal decision making in such settings.

The Prevalent Study of Planning with Certainty

Economists have long studied policy choice by an actual or hypothetical social planner who aims to maximize welfare in democracies or other political systems where welfare is intended to express the values of a society rather than the preferences of a dictator.

The public may not be familiar with formal welfare economics, but basic ideas are familiar through the widespread use of the term *benefit-cost analysis*.

Economists often study planning with utilitarian welfare functions.

We sometimes specify ones that express a form of paternalism or principles of fairness.

The motivation for studying planning is most transparent when actual planners face specific decision problems.

A national government must design an income tax and develop a national defense.

Local governments maintain roads, perform policing, and organize public education.

Clinicians make medical choices on behalf of patients.

Parents act as planners for their families.

In these settings and others, the planning objective is to maximize some idea of welfare.

Welfare economics has also sought to shed light on noncooperative societal decision processes, where no actual planner exists.

Adam Smith metaphorically suggested that an *invisible hand* makes decentralized decision making in market economies promote social welfare.

Economists gradually formalized this notion to develop what have become known as the *fundamental theorems of welfare economics*.

These give idealized conditions under which equilibrium outcomes in markets have the desirable welfare property of Pareto efficiency, which would be sought by a planner using a utilitarian or other welfare function that aggregates personal welfare (aka utility).

A central concern of research in public economics has been to study planning when the idealized conditions of the fundamental theorems of welfare economics do not hold.

Pareto efficiency has served as a benchmark for measuring the inefficiency of outcomes.

The social welfare achieved by a hypothetical planner has also served as a benchmark in social-choice theory, which studies the outcomes produced by voting and other decentralized mechanisms that attempt to aggregate personal preferences.

Whether performing abstract theoretical studies or applied benefit-cost analyses, researchers have generally assumed that the planner knows enough about the choice environment to be able to determine an optimal action.

However, the consequences of decisions are often highly uncertain.

Aiming to circumvent this difficulty, researchers commonly invoke strong unsubstantiated assumptions and use them to study solvable optimization problems.

I have referred to this as policy analysis with *incredible certitude*.

(Manski, *EJ* 2011, *Public Policy in an Uncertain World*, 2013).

Planning with incredible certitude can harm society in multiple ways.

It seeks to maximize the social welfare that would prevail if untenable assumptions were to hold rather than actual social welfare.

If planners incorrectly believe that existing analysis provides an errorless description of society and accurate predictions of policy outcomes, they may make poor decisions.

They will not recognize the potential value of research aiming to improve knowledge.

Nor will they appreciate the potential usefulness of decision strategies that may help society cope with uncertainty and learn.

The dearth of study of planning under uncertainty is apparent in the comprehensive textbook on public economics of Atkinson and Stiglitz (1980), which mentions uncertainty only a few times and then only in passing.

Addressing the reluctance of research in welfare economics to come to grips with uncertainty has motivated my research program on credible social planning under uncertainty, which has developed over the past twenty years.

The word 'credible' is subjective and often difficult to pin down, but I use it nonetheless.

I will describe the themes of my work and applications.

As far as I am aware, only a small body of other research engages any of the themes that I will discuss.

Johansen (1978) called for research on macroeconomic planning under uncertainty, stating:

“Uncertainty is not something which should be considered as a theoretically interesting refinement or extension of standard theory and methodology, but a central factor of eminently practical importance. Sometimes uncertainty is itself the heart of the matter when decisions are to be taken.”

In the early 2000s, Hansen and Sargent initiated a program of work on robust macroeconomic policy, considering possible deviations of reality from the assumptions maintained in conventional macroeconomic models; see Hansen and Sargent (2008).

Barlevy (2011) reviews work on macroeconomic policy under ambiguity.

Uncertainty in Decision Theory

A fundamental difficulty with welfare maximization under uncertainty is apparent even in a simple setting with two feasible actions, say A and B, and two possible choice environments, say s_1 and s_2 .

Suppose that A yields higher welfare in s_1 and B yields higher welfare in s_2 .

If it is not known whether s_1 or s_2 is the actual environment, then it is not known which action is better.

Thus, maximization of welfare is logically impossible.

At most one can seek a reasonable way to make a choice. A basic issue is how to interpret and justify the word 'reasonable.'

Research in decision theory has posed and characterized various principles for reasonable decision making under uncertainty.

Decision theory is not specifically concerned with societal decisions.

It assume an abstract decision maker who must choice among a specified set of actions.

The decision maker could be an individual, a firm, or a planner.

The description of uncertainty in decision theory is abstract.

One supposes that outcomes are determined by the chosen action and by some feature of the environment, called the *state of nature*.

The decision maker is assumed able to list all states of nature that could possibly occur.

This list, the *state space*, is a primitive concept expressing uncertainty.

The larger the state space, the less the decision maker knows about the consequences of each action.

Much of decision theory adds a secondary expression of uncertainty in the form of a probability distribution over the state space.

Some studies view the probability distribution as a cognitive concept, expressing how decision makers might actually perceive uncertainty.

Others view it as a mathematical construct, whose existence might be inferred from analysis of choice behavior.

Two conceptually distinct but mathematically related approaches have been used to develop criteria for reasonable decision making.

One poses choice axioms as primitives. The other focuses on the substantive consequences of choices. The approaches are related by *representation theorems*.

Axiomatic Decision Theory

Axiomatic decision theory poses principles, called axioms, for consistency of hypothetical behavior across a class of potential choice problems.

Researchers may introspect and assert it to be reasonable, or rational, that a decision maker should adhere to these choice axioms.

The central research activity of axiomatic decision theory has been to pose and prove representation theorems establishing that adherence to a specified set of axioms is equivalent to acting as if one wants to maximize some welfare function, coping with uncertainty in some manner.

Particularly famous are theorems in Von Neumann-Morgenstern (1944) and Savage (1954).

Both establish that adherence to certain axioms is representable as maximizing expected utility.

They differ mainly in that the distribution on the state space used to form expected utility is pre-specified in VN-M and determined within the theory in Savage..

VN-M viewed the distribution as a primitive concept.

Savage viewed the distribution as a construct that may in principle be inferred from analysis of extensive choice behavior.

In neither theorem does the distribution have any necessary connection to objective reality.

When studying consistency axioms of the types posed by VN-M and Savage, decision theorists do not differentiate between private entities and social planners.

The presumption is that all decision makers should behave consistently in the same manner.

Choice axioms aim to characterize procedural reasonableness, or rationality, in the sense of consistency of hypothetical behavior across potential choice problems.

Axiomatic decision theory provides no description of substantively good decisions.

Consequentialist Decision Theory

Consequentialist decision theory specifies a welfare function and an expression of uncertainty as primitives. It then seeks reasonable criteria to make decisions.

The most prevalent recommendation has been maximization of expected utility.

One places a probability distribution on the state space and chooses an action that maximizes the expected value of welfare with respect to this distribution.

To assist decision makers who do not find it credible to express uncertainty through a probability distribution, decision theorists have studied criteria that, in some sense, works uniformly well over all of the state space.

Two prominent interpretations of this idea are the maximin and minimax-regret criteria.

The decision theory used in my research is consequentialist rather than axiomatic.

I suppose that the objective is to make substantively good decisions in particular settings.

I suppose that a planner specifies a welfare function, expresses uncertainty in a credible manner, and uses these primitives to make a decision.

The welfare function and expression of uncertainty (aka *expectations*) are context specific.

A fundamental problem with axiomatic decision theory is that it is unconcerned with the credibility of a decision maker's expectations.

The realism of expectations should matter to any decision maker.

It should matter particularly to a planner who represents a population.

Characterizing Uncertainty Regarding Objective Probability Distributions

To characterize uncertainty with enough concreteness to be useful to the study of planning, I draw on my econometric research on partial identification.

A fundamental difficulty when studying planning is the identification problem arising from the unobservability of counterfactual outcomes.

At most one can observe the outcomes that have occurred under realized policies. The outcomes of unrealized policies are logically unobservable.

Yet determination of an optimal policy requires comparison of all feasible policies.

For this and many other reasons, planners usually have only partial knowledge of the welfare achieved by alternative policies.

I first connected identification with decisions under uncertainty in Manski (2000), writing:

“This paper connects decisions under ambiguity with identification problems in econometrics. Considered abstractly, it is natural to make this connection. Ambiguity occurs when lack of knowledge of an objective probability distribution prevents a decision maker from solving an optimization problem. Empirical research seeks to draw conclusions about objective probability distributions by combining assumptions with observations. An identification problem occurs when a specified set of assumptions combined with unlimited observations drawn by a specified sampling process does not reveal a distribution of interest. Thus, identification problems generate ambiguity in decision making.”

I followed Ellsberg (1961) in using the word *ambiguity* to signify uncertainty when one specifies a set of feasible states of nature but does not place a probability distribution on the state space. Synonyms include *deep uncertainty* and *Knightian uncertainty*.

Statistical imprecision in empirical analysis is also relevant to planning, but identification is generally the deeper and more profound source of uncertainty.

What are the objective uncertainties with which actual social planning must cope?

They are many and varied. For now, I will simply list those I have studied.

These include numerous identification problems in medical risk assessment and prediction of treatment response. See Manski (*Patient Care under Uncertainty* 2019).

There is uncertainty in the epidemiological models used to predict the spread of infectious diseases, which inform choice of vaccination policy (Manski, *PNAS* 2010; *JPET* 2017).

There is uncertainty in the climate models used to predict climate change, which inform choice of climate policy (Manski, Sanstad, and DeCanio, *PNAS* 2021) and in the discount rate used to form a welfare function (DeCanio, Manski, and Sanstad, *EE* 2022).

Challenging identification problems arise when studying the preferences and behavior of human populations.

Knowledge of preferences is essential to policy evaluation when welfare is utilitarian.

An ability to predict behavior is required to evaluate policy consequences whatever the welfare function may be.

I have examined how uncertainties about preferences and behavior complicate evaluation of income tax policies, where a central consideration is the income-leisure preferences of potential workers (Manski, *QE* 2014; *EJ* 2014).

I have shown how uncertainty about the effect of policing on criminal behavior complicates evaluation of proactive policing programs (Manski, *EJ* 2006).

Planning with Incredible Certitude

Analyses of public policy regularly express certitude about the consequences of alternative policy choices. Expressions of uncertainty are rare. Yet predictions often are fragile.

Conclusions may rest on critical unsupported assumptions or on leaps of logic. Then the certitude of policy analysis is not credible.

One can resolve the tension between the credibility and power of assumptions by posing assumptions of varying strength and determining the conclusions that follow.

In practice, policy analysis tends to sacrifice credibility in return for strong conclusions.

Why so?

Analysts and policy makers respond to incentives.

The scientific community rewards strong novel findings.

The public wants unequivocal policy recommendations.

These incentives make it tempting to maintain assumptions far stronger than can be persuasively defended, in order to draw strong conclusions.

Expressing certitude also has been advocated in philosophy of science.

When there are multiple explanations for available data, philosophers recommend using a criterion such as 'simplicity' to choose one of them.

Manski (*EJ* 2011) introduced a typology of practices that contribute to incredible certitude. I have since elaborated in Manski (*PPUW* 2013; *JEL* 2015, *PNAS* 2019, *EaP* 2020):

The typology is

- * *conventional certitude*: A prediction that is generally accepted as true but is not necessarily true.
- * *dueling certitudes*: Contradictory predictions made with alternative assumptions.
- * *conflating science and advocacy*: Specifying assumptions to generate a predetermined conclusion.
- * *wishful extrapolation*: Using untenable assumptions to extrapolate.
- * *illogical certitude*: Drawing an unfounded conclusion based on logical errors.
- * *media overreach*: Premature or exaggerated public reporting of policy analysis.

I have provided illustrative examples and have offered suggestions to improve practices.

Perspectives on Social Welfare

Given a choice set and welfare function, analysis of optimal planning is straightforward in abstraction, although solution of the optimization problem may be difficult in practice.

The subtleties in research on planning are conceptual rather than mathematical.

If analysis is to be useful as more than a theoretical exercise, the welfare function should express normative properties acceptable to some meaningful part of the relevant society.

The choice set should be realistic, comprising options that may actually be available.

Analysis should recognize that policy outcomes may be uncertain.

Specification of the welfare function has vexed economists and philosophers in broad terms, as well as policy analysts in particular contexts.

Most research by economists has supposed that the welfare function should aggregate the personal welfares of the individuals who compose society.

Yet it has long been understood that, in general, a heterogeneous society cannot develop a consensus social welfare function.

The Arrow (1950) Possibility Theorem diminished the residual hope that a heterogeneous society might be able to devise at least a coherent non-dictatorial welfare function.

How then should research on planning proceed? The literature is vast and varied.

The New Welfare Economics

One route was taken in the 1930s and 1940s by the economists who initiated study of the *new welfare economics* (Hicks, 1939).

Wary of any criterion to choose among policies that benefit some people but harm others, they retreated to the study of Pareto efficiency with extension to fictional redistributions proposed by Kaldor (1939) and Hicks (1939).

This restriction on their domain of concern drastically limited their ability to study actual planning problems. This led Chipman and Moore (1978) to write:

“In this paper we shall argue that, judged in relation to its basic objective of enabling economists to make welfare prescriptions without having to make value judgments and, in particular, interpersonal comparisons of utility, the New Welfare Economics must be considered a failure.”

Utilitarian and Maximin Welfare

In research that studies planning when policies benefit some people but harm others, it has been common among economists to specify a utilitarian welfare function.

The standard theory of rational individual behavior under certainty requires only an ordinal concept of personal welfare.

A utilitarian welfare function specifies interpersonally comparable cardinal personal welfares and sums them. Bentham (1776) may have had this in mind when he wrote:

“a fundamental axiom, it is the greatest happiness of the greatest number that is the measure of right and wrong.”

One need not sum personal welfares to develop welfare functions that respect Pareto efficiency. The Rawls (1971) maximin function has received attention outside economics.

Non-Personalist Welfare Functions

I have so far discussed research that assumes the welfare function aggregates personal welfare. Sen (1977) called this *welfarism*. I prefer the word “personalism.”

Non-personalist welfare functions place direct societal value on certain ethical concepts, beyond their possible manifestations as determinants of personal welfare.

These concepts have been given many seemingly simple labels, including justice, fairness, equity, and liberty.

These labels are difficult to interpret formally. See Manski, Mullahy, and Venkataramani (*PNAS* 2023).

Pragmatic Welfare Functions

Research in welfare economics and moral philosophy has mainly been abstract.

Studies of concrete planning problems have commonly used pragmatic welfare functions.

I use the word ‘pragmatic’ to mean that researchers motivate their welfare functions by some combination of conjecture regarding societal values, empirical study of population preferences, and concern for analytical tractability.

For example, the literature on optimal taxation stemming from Mirrlees (1971) has assumed a utilitarian welfare function and has placed various restrictions on the population distribution of income-leisure preferences.

Research on government spending to optimize macroeconomic growth has assumed utilitarian welfare and a representative infinite-lived household (e.g., Barro, 1990).

Integrated assessment studies of optimal climate policy has assumed that the objective is to maximize present-discounted gross world product (e.g., Nordhaus, 2008).

Analyses of optimal medical care often assume that the objective is to maximize the population mean quality-adjusted life years (QALYS) net of treatment cost.

When academic researchers specify pragmatic welfare functions, they may believe that these functions have sufficient social acceptability to make them worthy of study.

They usually do not argue that actual planners should necessarily use these welfare functions to make decisions.

The less ambitious goal is to learn what decisions would be optimal if specified welfare functions were to be used.

This perspective is maintained throughout my own work.

Uncertainty in Consequentialist Decision Theory

I now deepen the discussion of uncertainty in decision theory begun in the Introduction.

Maximization of welfare expresses a consequential perspective on decision making. Given a choice setting, the goal is to do as well as possible in achieving a specified objective.

The starting point is to suppose that the planner faces a predetermined choice set C and believes that the true state of nature s^* lies in a state space S .

The welfare function $w(\cdot, \cdot): C \times S \rightarrow \mathbb{R}^1$ maps actions and states into welfare.

The planner wants to maximize $w(\cdot, s^*)$ over C but does not know s^* .

Hence, maximization is infeasible except in special cases.

The state space S provides the basic decision theoretic expression of uncertainty. States of nature that are not elements of S are presumed impossible to occur.

The decision maker does not contemplate the possible existence of ‘unknown unknowns.’

Discussions of the state space often consider it to express uncertainty purely about the physical and social environment within which choice takes place.

A state space can also express uncertainty about the welfare function.

This may occur when the planner is utilitarian.

The planner must know the preferences of the population to maximize welfare, but this knowledge may not be available.

The state space is a subjective primitive of the decision problem.

However, being subjective does not imply that it is an arbitrary construction.

Credibility is fundamental in consequential decision theory in general and in the study of social planning specifically.

If planning decisions are to enhance well-being in the real world, the planner should specify a state space that expresses some reasonable sense of credibility.

Scientific research seeks to provide at least a partially objective basis for specification of the state space.

This basis is obtained by combining plausible theory with careful empirical analysis.

Decision Criteria

It is generally accepted that decisions should respect dominance.

Action $c \in C$ is weakly dominated if there exists a $d \in C$ such that $w(d, s) \geq w(c, s)$ for all $s \in S$ and $w(d, s) > w(c, s)$ for some $s \in S$.

To choose among undominated actions, decision theorists have proposed various ways of using $w(\cdot, \cdot)$ to form functions of actions alone, which can be optimized.

One should only consider undominated actions, but it often is difficult to determine which actions are undominated. Hence, it is common to optimize over all feasible actions.

Here I discuss settings without sample data. Wald (1950) extended the theory to settings where the planner observes sample data.

A familiar idea is to place a subjective probability distribution π on the state space, average state-dependent welfare with respect to π , and maximize subjective average welfare over C . The criterion solves

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

Another idea seeks an action that, in some sense, works uniformly well over all of S . This yields the maximin and minimax-regret (MMR) criteria.

The maximin criterion maximizes the minimum welfare attainable across S , solving the problem

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

The MMR criterion solves

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

Here $\max_{d \in C} w(d, s) - w(c, s)$ is the *regret* of action c in state s .

The true state being unknown, one evaluates c by its maximum regret over all states and selects an action that minimizes maximum regret.

The maximum regret of an action measures its maximum distance from optimality across states. Hence, maximum regret is uniform nearness to optimality.

Note: The above considers polar cases in which a planner asserts either a complete subjective distribution on the state space, or none. A planner might also assert a partial subjective distribution, placing lower and upper probabilities on states.

Learning Objective Probability Distributions

In many planning settings, it has become standard to specify the state space as a set of objective probability distributions that may possibly describe the system under study.

Haavelmo (1944) did so when he introduced *The Probability Approach in Econometrics*.

Studies of treatment choice do so when they consider the population to be treated to have a distribution of treatment response.

Research seeks to enhance the credibility of planning by constructive combination of theory and empirical analysis to provide information about the possible distributions.

The Koopmans (1949) formalization of identification analysis contemplated unlimited data collection that enables one to shrink the state space, eliminating distributions that are inconsistent with theory and with the information revealed by observation.

Note: Sample data generally are not informative enough to shrink the state space. Wald's statistical decision theory shows how sample data can be informative, nonetheless.

For most of the 20th century, econometricians commonly thought of identification as binary. A feature of an objective probability distribution is either identified or it is not.

Empirical researchers combined available data with assumptions that yield point identification, and they reported point estimates of parameters.

Many economists recognized that point identification often requires strong assumptions that are difficult to motivate. However, they saw no other way to perform inference.

Yet there is enormous scope for fruitful inference using weaker and more credible assumptions that partially identify population parameters.

A parameter is partially identified if the sampling process and maintained assumptions reveal that the parameter lies in a set, its *identification region* or *identified set*, that is smaller than the logical range of the parameter but larger than a single point.

There were isolated contributions to analysis of partial identification as early as the 1930s, but the subject remained at the fringes of econometric consciousness and did not spawn systematic study.

A coherent body of research took shape in the 1990s and has since grown rapidly. A textbook exposition is Manski (*Identification for Prediction and Decision* 2007).

The modern literature on partial identification emerged out of concern with traditional approaches to inference with missing outcome data.

Empirical researchers have commonly assumed that missingness is random, in the sense that the observability of an outcome is statistically independent of its value.

Yet this and other point-identifying assumptions have been criticized as implausible.

It was natural to ask what random sampling with partial observability of outcomes reveals about outcome distributions if nothing is known about the missingness process, or if assumptions weak enough to be widely credible are imposed.

Studying inference with missing outcome data led to analysis of treatment response.

A common objective of empirical research is to predict treatment response conditional on specified covariates, using data from a random sample of the population.

Analysis must contend with the problem that counterfactual outcomes are not observable; hence, findings on partial identification with missing outcome data are applicable.

Analysis of treatment response poses more than a generic missing-data problem.

One reason is that observations of realized outcomes, when combined with suitable assumptions, can provide information about counterfactual ones.

Another is that practical problems of treatment choice motivate research on treatment response and thereby determine what population parameters are of interest.

Whatever the specific subject under study, a common theme runs through research on partial identification.

One first asks what the sampling process alone reveals about the population of interest and then studies the identifying power of assumptions that aim to be credible in practice.

This conservative approach to inference makes clear the conclusions one can draw in empirical research without imposing untenable assumptions.

It establishes a domain of consensus among analysts who may hold disparate beliefs about what assumptions are appropriate.

It also makes plain the limitations of the available data. When credible identification regions turn out to be large, we should face up to the fact that the available data do not support inferences as tight as we might like to achieve.

Findings on partial identification imply that empirical research may shrink the state space for decision making but not reduce it to a single state of nature.

Let S be the state space without observation of the unlimited data assumed in an identification study.

Let $S_0 \subset S$ be the shrunken state space with these data.

Then the decision criteria posed earlier have the same forms, but with S_0 replacing S .

Minimax-Regret Planning

Maximization of subjective average welfare places a subjective distribution on the state space, whereas maximin and MMR do not.

Concern with the basis for specification of a subjective distribution motivated Wald (1950) to study the minimax criterion (maximin in my description), writing:

“a minimax solution seems, in general, to be a reasonable solution of the decision problem when an a priori distribution does not exist or is unknown.”

I am similarly concerned with decision making with no subjective distribution on states.

However, I have mainly measured performance of decisions by maximum regret rather than by minimum welfare.

The maximin and MMR criteria both provide ex ante evaluations of the worst result that a decision maker may experience ex post.

However, the criteria are equivalent only in special cases, particularly when optimal welfare is invariant across states. They differ more generally.

Whereas maximin considers the worst absolute outcome that an action may yield across states, MMR considers the worst outcome relative to what is achievable in a given state.

A conceptual appeal of using maximum regret to measure performance is that it quantifies how lack of knowledge of the true state of nature diminishes the quality of decisions.

The term “maximum regret” is shorthand for the maximum sub-optimality of a decision criterion across the feasible states of nature. A decision with small maximum regret is uniformly near optimal across all states. I think this a desirable property.

Diversified Treatment under Ambiguity

To show how study of planning under uncertainty can matter in practice, consider the study of diversified treatment under ambiguity initiated in Manski (*IPD* 2007), expanded in Manski (*IER* 2009), and applied in Cassidy and Manski (*PNAS* 2019).

I considered settings in which a planner can treat persons differentially. Examples include medical treatment, sentencing of offenders, and active labor-market programs.

The planner may make a *singleton* allocation, assigning all observationally identical persons to the same treatment. Or the planner may choose a *fractional* allocation, randomly assigning positive fractions of these persons to different treatments.

Fractional allocations cope with ambiguity through diversification.

Let there be two feasible treatments, A and B.

A Type A error occur when treatment A is chosen but is actually inferior to B. 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.

Bayesian and maximin planning may yield singleton or fractional allocations, depending on the specifics of the problem.

The MMR allocation of two treatments under ambiguity is always fractional.

The Institutional Separation of Research on Planning and Actual Planning

I conclude by extending remarks in *Public Policy in an Uncertain World* (2013).

I observed there that modern democratic societies have created an institutional separation between policy analysis and decision making, with professional analysts reporting findings to representative governments.

Separation of the tasks of analysis and decision making, the former aiming to inform the latter, appears advantageous from the perspective of division of labor.

Having researchers study planning problems and provide their findings to law makers and civil servants enables these planners to focus on the challenging task of policy choice, without having to perform their own research.

I also observed that the current practice of policy analysis with incredible certitude does not serve society well.

The problem is that the consumers of policy analysis cannot trust the producers.

To improve analysis and to increase trust, research on planning should transparently face up to uncertainty rather than hide it.

Some think this idea naïve or impractical.

I have repeatedly heard policy analysts assert that policy makers are either psychologically unwilling or cognitively unable to cope with uncertainty.

A more optimistic possibility is that incredible certitude is a modifiable social norm.

Salutary change can occur if awareness grows that incredible certitude is harmful.

Then society will want researchers to provide reasonable policy recommendations recognizing the subtlety of planning under uncertainty, not unequivocal ones lacking foundation.

**DIVERSIFIED TREATMENT UNDER AMBIGUITY,
With Application to Tuberculosis Diagnosis and Treatment**

Charles F. Manski

Department of Economics and Institute for Policy Research
Northwestern University

Diversified Treatment under Ambiguity

Charles F. Manski

International Economic Review, Vol. 50, No. 4, 2009, pp. 1013-1041.

I will discuss the study of diversified treatment under ambiguity initiated in Manski (*IPD* 2007), expanded in Manski (*IER* 2009), and applied in Cassidy and Manski (*PNAS* 2019).

I considered settings in which a planner can treat persons differentially. Examples include medical treatment, sentencing of offenders, and active labor-market programs.

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

Or the planner may choose a *fractional* allocation, randomly assigning positive fractions of these persons to different treatments.

Fractional allocations cope with ambiguity through diversification.

Note: Diversification has long been studied in the theory of financial portfolio allocation, assuming maximization of expected utility.

Suppose that there are two feasible treatments, labelled A and B.

A Type A error occurs when treatment A is chosen but is actually inferior to B. A Type B error occurs when B is chosen but is inferior to A.

The singleton allocation assigning everyone to treatment A 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.

One-Period Problems with Individualistic Treatment and Linear Welfare

I begin with a simple leading case.

Each member j of a population J of observationally identical persons 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 task is to allocate the population to treatments A and B.

An allocation $\delta \in [0, 1]$ randomly assigns a fraction δ of the population to treatment B and $1 - \delta$ to treatment A.

A utilitarian planner wants to maximize mean personal welfare.

Let $u_j(t) \equiv u_j[y_j(t), t]$ be the personal welfare of person j when this person receives treatment t and realizes outcome $y_j(t)$. This welfare is cardinal and interpersonally comparable.

Let $\alpha \equiv E[u(A)]$ and $\beta \equiv E[u(B)]$ be mean welfare if all members of the population receive treatment A or B. Then mean welfare with allocation δ is

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

$\delta = 1$ is optimal if $\beta \geq \alpha$ and $\delta = 0$ if $\beta \leq \alpha$.

The problem is treatment choice when the planner has partial knowledge of (α, β) .

Let S index the feasible states of nature.

Let the planner know that (α, β) lies in a bounded set $[(\alpha_s, \beta_s), s \in S]$.

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

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

A Bayesian planner places a subjective distribution π on S and solves

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

where $E_{\pi}(\alpha) = \int \alpha_s d\pi$ and $E_{\pi}(\beta) = \int \beta_s d\pi$.

Thus, a Bayesian planner makes a singleton choice if $E_{\pi}(\beta) \neq E_{\pi}(\alpha)$.

A maximin planner solves

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

If (α_L, β_L) is feasible, the decision is $\delta = 0$ if $\beta_L < \alpha_L$, $\delta = 1$ if $\beta_L > \alpha_L$, and all δ if $\beta_L = \alpha_L$.

Thus, a maximin planner makes a singleton choice if (α_L, β_L) is feasible and $\beta_L \neq \alpha_L$.

The regret of allocation δ in state s is the difference between the maximum achievable welfare and the welfare achieved with δ . Maximum welfare in state s is $\max(\alpha_s, \beta_s)$. The minimax-regret criterion is

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

Regret with allocation δ in state s is $\max(\alpha_s, \beta_s) - [\alpha_s + (\beta_s - \alpha_s)\delta]$.

Manski (*IPD* 2007, Chapter 11) derived the MMR treatment allocation.

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

Let $M(A) \equiv \max_{s \in S(A)} (\alpha_s - \beta_s)$; $M(B) \equiv \max_{s \in S(B)} (\beta_s - \alpha_s)$.

If (α_L, β_U) and (α_U, β_L) are feasible, $M(A) = \alpha_U - \beta_L$; $M(B) = \beta_U - \alpha_L$.

The MMR allocation to treatment B is $\delta_{MR} = M(B)/[M(A) + M(B)]$.

Thus, the MMR allocation of two treatments is always fractional under ambiguity.

Note: The allocation is not always fractional with more than two treatments (Stoye, 2007).

Welfare Monotone in Mean Personal Welfare

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. In finance, this is the well-known finding that a risk-seeking investor, whose utility is convex in income, does not diversify but a risk-averse investor, whose utility is concave in income, may diversify.

The shape of $f(\cdot)$ does not affect the maximin decision. The reason is that the maximin criterion only uses ordinal, not cardinal properties of the welfare function.

Manski (*IER* 2009) shows that the MMR allocation is fractional whenever $f(\cdot)$ is continuous and the planner faces ambiguity. If $f(\cdot) = \log(\cdot)$ and $\{(\alpha_L, \beta_U), (\alpha_U, \beta_L)\}$ are feasible, then the MMR allocation is $\delta_{MR} = [\alpha_U(\beta_U - \alpha_L)] / [\alpha_U(\beta_U - \alpha_L) + \beta_U(\alpha_U - \beta_L)]$.

Deontological Welfare Functions

Deontological ethics supposes that choices may have intrinsic value. *Equal treatment of equals* is sometimes thought to be an important 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.

Manski (2009) studied welfare functions that express concern with ex post equal treatment by subtracting a fixed cost from welfare when the treatment allocation is not singleton. The MMR allocation may be singleton or fractional, depending on the specifics of the case.

Adaptive Diversification in Sequential Planning Problems

Manski (2009) considered sequential planning when, in each period, a planner chooses treatments for the current cohort of a population.

Now learning is possible. Observing outcomes for earlier cohorts informs treatment choice for later cohorts.

Fractional allocations generate randomized experiments, yielding outcome data on both treatments. Sampling variation is not an issue when cohorts are large.

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

This criterion treats each cohort as well as possible, given existing knowledge.

Tuberculosis Diagnosis and Treatment under Ambiguity

Rachel Cassidy and Charles F. Manski

Proceedings of the National Academy of Sciences, Vol. 116, No. 46, 2019, pp. 22990-22997.

Summary

1.6 million persons worldwide died from tuberculosis (TB) in 2017.

A challenge in fighting TB is to improve capacity for rapid and accurate diagnosis.

A new TB diagnostic test called *Xpert* was endorsed by the WHO in 2010.

Trials demonstrated that Xpert is faster and has greater sensitivity and specificity than smear microscopy – the most common diagnostic test.

However, subsequent trials found no impact of Xpert on morbidity and mortality.

We present a decision-theoretic model of how a clinician might decide whether to order XPert or other tests for TB and whether to treat a patient, with or without test results.

We study the conditions in which it is optimal to perform *empirical treatment*; that is, treatment without diagnostic testing.

We examine the implications for decision making of partial knowledge of TB prevalence or test accuracy, which generates ambiguity about the best testing and treatment policy.

In the presence of ambiguity, we show the usefulness of adaptive diversification of testing and treatment.

Optimal Diagnosis and Treatment Decisions

We first study decision-making when a clinician knows enough to make optimal diagnosis and treatment decisions.

We consider a clinician who cares for a population of patients.

We assume that this patient population is predetermined and that patients always comply with the clinician's decisions.

We assume that treatment response is individualistic. In particular, treatment decisions do not affect disease transmission.

When a patient appears at a clinic, the clinician initially observes covariates x that may include demographic attributes, medical history, and indicators of health status.

The clinician can prescribe a treatment immediately (empirical treatment) or order a test that may yield further evidence.

In the latter case, the clinician prescribes a treatment after observation of the test result.

In the context of TB, treatment $t = B$ means prescription of antibiotics and $t = A$ means a decision not to prescribe antibiotics.

Let s indicate whether the clinician orders the diagnostic test, with $s = 1$ or $s = 0$.

Let r denote the test result.

$r = p$ (positive) suggests that the patient has TB.

$r = n$ (negative) suggests the absence of TB.

The feasible actions may be expressed as a decision tree.

The clinician chooses $s = 0$ or $s = 1$ with knowledge of x .

If she chooses $s = 0$, she chooses $t = A$ or $t = B$ with knowledge of x .

If she chooses $s = 1$, she chooses $t = A$ or $t = B$ with knowledge of (x, r) .

Let $\delta_S(x)$ be the fraction of patients with covariates x who are tested. The clinician can choose $\delta_S(x)$ to be any fraction.

When considering treatment, we need to distinguish three types of patients.

Among patients with covariates x who are not tested, let $\delta_{T0}(x)$ be the fraction receiving B.

Among patients who are tested and have test outcome r , let $\delta_{T1}(x, r)$ be the fraction who receive B.

Welfare function

We assume that the clinician aggregates the benefits and harms of making a specific testing and treatment decision for a given patient into a scalar welfare measure.

Let $y(s, t)$ summarize the clinician's overall assessment of the benefits and harms that would occur if she were to make testing decision s and treatment decision t .

Patients may respond heterogeneously, so $y(s, t)$ may vary across patients.

We suppose that the objective of the clinician is to optimize care on average across the patients in her practice.

Mean welfare across the population of patients is determined by the fraction of those in each covariate group that the clinician assigns to each testing-treatment option.

Let $P(x)$ denote the fraction of patients with covariate value x .

For $r = p$ or n , let $f(r|x)$ denote the fraction of patients with covariates x who would have test result r if they were to be tested.

For each possible value of (s, t) , let $E[y(s, t)|x]$ be the mean welfare that would result if all patients with covariates x were to receive (s, t) .

Let $E[y(s, t)|x, r]$ be the mean welfare that would result if all patients with covariates x and test result r were to receive (s, t) .

Let $\delta = [\delta_s(x), \delta_{T0}(x), \delta_{T1}(x, r), x \in X, r \in \{p, n\}]$ denote any specified testing-treatment allocation.

The mean welfare $W(\delta)$ that results with allocation δ is obtained by averaging the mean welfare values $E[y(s, t)|x]$ and $E[y(s, t)|x, r]$ across the groups who receive them. Thus,

$$(1) \quad W(\delta) = \sum_{x \in X} P(x) \{ [1 - \delta_s(x)][1 - \delta_{T0}(x)]E[y(0, A)|x] + [1 - \delta_s(x)]\delta_{T0}(x)E[y(0, B)|x] \\ + \sum_{r \in \{p, n\}} f(r|x) \{ \delta_s(x)[1 - \delta_{T1}(x, r)]E[y(1, A)|x, r] + \delta_s(x)\delta_{T1}(x, r)E[y(1, B)|x, r] \} \}.$$

Optimal Testing and Treatment

An optimal testing and treatment decision sets

$$(2a) \quad \delta_S(\mathbf{x}) = 1 \quad \text{if} \quad \sum_{r \in \{p, n\}} f(r|\mathbf{x}) [\max \{E[y(1, A)|\mathbf{x}, r], E[y(1, B)|\mathbf{x}, r]\}] \\ \geq \max \{E[y(0, A)|\mathbf{x}], E[y(0, B)|\mathbf{x}]\}, \\ = 0 \text{ otherwise.}$$

$$(2b) \quad \delta_{T0}(\mathbf{x}) = 1 \quad \text{if} \quad E[y(0, B)|\mathbf{x}] \geq E[y(0, A)|\mathbf{x}], = 0 \text{ otherwise.}$$

$$(2c) \quad \delta_{T1}(\mathbf{x}, p) = 1 \quad \text{if} \quad E[y(1, B)|\mathbf{x}, p] \geq E[y(1, A)|\mathbf{x}, p], = 0 \text{ otherwise.}$$

$$(2d) \quad \delta_{T1}(\mathbf{x}, n) = 1 \quad \text{if} \quad E[y(1, B)|\mathbf{x}, n] \geq E[y(1, A)|\mathbf{x}, n], = 0 \text{ otherwise.}$$

Empirical treatment is optimal when the inequality in (2a) does not hold and the inequality in (2b) does hold. Empirical treatment is not optimal otherwise.

Risk of Illness and Treatment Decisions

The notation $y(s, t)$ leaves implicit how illness affects patient welfare.

Let $z = 1$ if the patient is ill and $z = 0$ otherwise.

Let $U(z, s, t)$ be the assessment of benefits and harms with knowledge of z .

Assume that the clinician does not know z when choosing (s, t) .

Then testing and treatment decisions depend on a patient's risk of illness rather than on realized illness outcomes.

To formalize this, replace $E[y(s, t)|x]$ and $E[y(s, t)|x, r]$ with

$$E[U(z, s, t)|x] = P(z = 0|x)E[U(0, s, t)|x] + P(z = 1|x)E[U(1, s, t)|x],$$

$$E[U(z, s, t)|x, r] = P(z = 0|x, r)E[U(0, s, t)|x, r] + P(z = 1|x, r)E[U(1, s, t)|x, r].$$

The earlier characterization of optimal testing and treatment holds, with $E[U(z, s, t)|x]$ and $E[U(z, s, t)|x, r]$ replacing $E[y(s, t)|x]$ and $E[y(s, t)|x, r]$.

To simplify further computations, we write

$$E[U(z, s, t)|x] = (1 - P_x)U_x(0, s, t) + P_x U_x(1, s, t),$$

$$E[U(z, s, t)|x, r] = (1 - P_{xr})U_{xr}(0, s, t) + P_{xr}U_{xr}(1, s, t).$$

With this notation, the treatment decision criteria are as follows:

(2b') treatment with no test result:

choose B if $(1 - P_x)U_x(0, 0, B) + P_xU_x(1, 0, B) \geq (1 - P_x)U_x(0, 0, A) + P_xU_x(1, 0, A)$,

choose A otherwise.

(2c') treatment with positive test result:

choose B if $(1 - P_{xp})U_x(0, 1, B) + P_{xp}U_x(1, 1, B) \geq (1 - P_x)U_{xp}(0, 1, A) + P_{xp}U_x(1, 1, A)$,

choose A otherwise.

(2d') treatment with negative test result:

choose B if $(1 - P_{xn})U_x(0, 1, B) + P_{xn}U_x(1, 1, B) \geq (1 - P_{xn})U_x(0, 1, A) + P_{xn}U_x(1, 1, A)$,

choose A otherwise.

Note on Measuring the Accuracy of Diagnostic Tests

P_x is the *base rate* or the *prevalence* of the illness.

P_{xp} is the *positive predictive value* of a test and $1 - P_{xn}$ is the *negative predictive value*.

An ideal test would have $P_{xp} = 1$ and $P_{xn} = 0$. In practice, $1 > P_{xp} > P_{xn} > 0$.

The medical literature commonly measures test accuracy by *sensitivity* and *specificity*.

Sensitivity is $P(r = p|x, z = 1)$. Specificity is $P(r = n|x, z = 0)$.

Sensitivity and specificity do not provide the information that a clinician would want to have to inform patient care.

Threshold Risk Assessments for Choice between Surveillance and Aggressive Treatment

Manski (*QE* 2018) shows that criteria (2b'-2d') yield simple solutions when treatment A is surveillance of a patient and B is aggressive treatment. This analysis fits TB.

It is often credible to make various assumptions about patient welfare when comparing surveillance and aggressive treatment. In particular,

- (i) Health is better than illness: $U_x(0, s, t) > U_x(1, s, t)$ for all (s, t) .
- (ii) Testing is costly/harmful: $U_x(z, 0, t) > U_x(z, 1, t)$ for all (z, t) .
- (iii) Surveillance is better when healthy: $U_x(0, s, A) > U_x(0, s, B)$, all s .
- (iv) Aggressive treatment is better when ill: $U_x(1, s, B) > U_x(1, s, A)$, all s .

These assumptions are realistic in the TB context.

Under these assumptions, aggressive treatment is optimal if the risk of illness equals or exceeds a threshold that equalizes mean welfare under treatments A and B.

Surveillance is better if risk is less than or equal to the threshold.

In the absence of testing, risk of illness is measured by P_x and the threshold is

$$P_{x0}^* \equiv \frac{U_x(0, 0, A) - U_x(0, 0, B)}{[U_x(0, 0, A) - U_x(0, 0, B)] + [U_x(1, 0, B) - U_x(1, 0, A)]} .$$

With testing, risk of illness is measured by P_{xp} or P_{xn} respectively. The threshold is

$$P_{x1}^* \equiv \frac{U_x(0, 1, A) - U_x(0, 1, B)}{[U_x(0, 1, A) - U_x(0, 1, B)] + [U_x(1, 1, B) - U_x(1, 1, A)]} .$$

How Testing Affects Treatment

In general, response to this question may be complex because P_{x0}^* and P_{x1}^* may differ.

Substantial simplification occurs if the thresholds are equal. Let the common value be P_x^* .

Note: A sufficient condition for equality is the assumption that testing imposes an additive treatment-invariant cost on welfare; that is, if $U_x(z, 0, t) - U_x(z, 1, t) = K > 0$ for some K .

Then testing affects optimal treatment if and only if $P_{xn} < P_x^* < P_{xp}$.

Given this inequality, a patient with a positive test result receives treatment B and one with a negative test result receives A.

In the absence of testing, the patient might receive either A or B.

Testing and Treatment under Ambiguity

Optimal testing and treatment is feasible if one knows $U_x(\cdot, \cdot, \cdot)$, (P_x, P_{xn}, P_{xp}) , and f_x .

A clinician with incomplete knowledge may not be able to optimize.

One may in principle study decision making using standard criteria, including maximization of subjective expected welfare, maximin, and minimax-regret.

Maximization of subjective expected welfare is a standard dynamic programming problem, but it requires specification of a subjective distribution on the state space, which we find difficult to motivate.

Study of maximin and minimax regret appears to require complex new analysis.

Piecemeal Minimax-Regret Decision Making

We propose a piecemeal minimax-regret criterion.

We consider each of the four component decisions in isolation from one another.

These are (1) choice to test or not to test, (2) choice between A and B without testing, (3) choice between A and B with testing and a positive result, (4) choice between A and B with testing and a negative result.

Each is a decision between two options, making piecemeal decision making relatively simple to study.

Piecemeal decisions are realistic in settings where each component decision may be performed by a different clinician.

We extend the study of minimax-regret decision making in Manski (2009).

The extension is especially simple if we suppose that $U_x(\cdot, \cdot, \cdot)$ is known. Then the threshold risk assessment P_x^* is known.

Considerable ambiguity may remain due to incomplete knowledge of (P_x, P_{xn}, P_{xp}) and f_x .

Piecemeal analysis applies minimax-regret separately to each component decision.

In each case, the result is a singleton allocation of patients in the absence of ambiguity and a fractional allocation with ambiguity.

A fractional allocation means diversification.

Adaptive Diversification

Finally, consider adaptive use of the piecemeal criterion across a sequence of cohorts.

Suppose that the distributions of test results and treatment response among patients remain stable over time.

Suppose that observation of patients eventually reveals whether they are ill.

Then complete learning eventually occurs if $\delta_S(\mathbf{x}) > 0$ for some cohort.

Randomized testing reveals f_x . Randomized treatment after testing reveals P_{xp} and P_{xn} .

Ambiguity in the TB context

To optimize, the clinician must know a patient's risk of illness conditional on each test result and the probability of a positive test result.

There are many reasons why these parameters are subject to ambiguity in the TB context.

We emphasize incomplete conditioning of research on patient covariates, which prevents use of available evidence to make personalized testing and treatment decisions.

When epidemiological studies estimate prevalence, they report findings conditional on a small subset of the patient attributes that a clinician observes. The same issue arises when trials of diagnostic tests report predictive values.

We use data for patients in Cape Town, South Africa to illustrate.

Statistical Decision Properties of Imprecise Trials Assessing COVID-19 Drugs

Charles F. Manski and Aleksey Tetenov
Value in Health, Vol. 24, No. 5, pp. 641-647.

Background

- “Statistical Treatment Rules for Heterogeneous Populations,” *Econometrica* 72, 2004, 221-246.
- “Sufficient Trial Size to Inform Clinical Practice,” *Proceedings of the National Academy of Sciences* 113, 2016, 10518-10523.
- “Trial Size for Near-Optimal Choice Between Surveillance and Aggressive Treatment: Reconsidering MSLT-II,” *The American Statistician* 73:sup1, 2019, 305-311.

Example of current practice

- Cao et al., “A Trial of Lopinavir–Ritonavir in Adults Hospitalized with Severe Covid-19,” *NEJM*, 18 March 2020.
- RCT
 - 99 patients assigned to receive lopinavir–ritonavir + “standard care”
 - 100 patients assigned to “standard care” alone
 - Measured outcomes up to 28 days after randomization

Reported trial outcomes

- Primary Finding: “In a modified intention-to-treat analysis, lopinavir–ritonavir led to a median time to clinical improvement that was shorter by 1 day than that observed with standard care (hazard ratio, 1.39; 95% CI, 1.00 to 1.91).”
- Secondary Finding: “Mortality at 28 days was similar in the lopinavir–ritonavir group and the standard-care group (19.2% vs. 25.0%; difference, –5.8 percentage points; 95% CI, –17.3 to 5.7).”

Conclusions from the trial

- Cao et al. “no benefit was observed with lopinavir–ritonavir treatment beyond standard care.”
- U.S. NIH panel guidelines then recommended against the use of lopinavir/ritonavir writing: “lopinavir/ritonavir was studied in a small randomized controlled trial in patients with COVID-19 with negative results.”
 - This trial was the main piece of evidence, summarized as: “No difference in primary outcome (time to clinical improvement) was observed, and 28-day mortality was similar between groups.”

Questions

- How should we measure precision of an RCT?
 - Maximum expected loss in patient welfare for treatment chosen based on an RCT relative to the unknown best treatment. (*maximum regret*)
 - This depends on how the trial results are translated into clinical decisions. (*statistical treatment rule*)
- How should we use the results of clinical trials to decide which treatment to use?
 - Prevailing practice is to use a two-sided 5% hypothesis test to reach a binary conclusion: Standard care if the null isn't rejected; innovation if the null is rejected with a significant positive estimate of average treatment effect.
 - We argue for the **Empirical Success** rule: choose the treatment with better average outcome and measure the outcome that patients want to maximize!

What happens in a trial with 100:99 patients using 28-day mortality as the outcome?

- Let mortality rate with standard care = 0.25 and use the standard t-test rule:

	Mortality rate with new treatment				
	0.35	0.30	0.25	0.20	0.15
% of trials after which standard care will be prescribed:	99.98%	99.7%	97.5%	86.76%	57.36%
Loss from choosing standard care:	0	0	0	0.05	0.10
% of trials after which new treatment will be prescribed:	0.02%	0.3%	2.5%	13.24%	42.64%
Loss from choosing treatment:	0.10	0.05	0	0	0
Expected loss:	0.0000	0.0002	0	0.0434	0.0574

- Maximum expected loss occurs when the new treatment has mortality rate 0.548 and standard care has rate 0.661. Then expected loss is $(0.661 - 0.548) \times \text{error probability } 0.624 = 0.071$.

- Same scenarios, using the empirical success rule

Mortality rate with new treatment					
	0.35	0.30	0.25	0.20	0.15
% of trials after which standard care will be prescribed:	94.28%	79.61%	51.64%	21.18%	4.22%
Loss from choosing standard care:	0	0	0	0.05	0.10
% of trials after which treatment will be prescribed:	5.72%	20.39%	48.36%	78.82%	95.78%
Loss from choosing treatment:	0.10	0.05	0	0	0
Expected loss:	0.0057	0.0102	0	0.0106	0.0042

- Maximum expected loss occurs when the new treatment has mortality rate 0.527 and standard care has rate 0.473. Then expected loss is $(0.527 - 0.473) \times \text{error probability } 0.226 = 0.012$. The same expected loss occurs when standard care has mortality 0.527 and the new treatment 0.473.

Why use the empirical success rule?

- Theoretical study proves that it exactly or approximately minimizes maximum expected loss
 - Exactly optimal in balanced trials with binary outcomes (Stoye, *JoE*, 2009)
 - Asymptotically optimal in other two-arm trials (Hirano & Porter, *ECMA*, 2009)
- Treats Type I and Type II errors symmetrically
- Hypothesis testing treats the two errors asymmetrically. Maximum loss when the innovation is better is 250 times greater than maximum loss when it is worse.

Why we shouldn't treat the two options asymmetrically

- “Standard care” for COVID-19 has been postulated without evidence that it is better than other options.
- If we were to start with a different definition of standard care, we would be stuck with it for a long time.
- Clinical equipoise
 - EMA: “There should be equipoise (“genuine uncertainty within the expert medical community [...] about the preferred treatment”) at the beginning of a randomised trial.”
 - One might motivate asymmetric decision-making after trials by having asymmetric Bayesian priors,
 - but interpreting ethical guidelines for starting trials through a Bayesian lens suggests that experts must
 1. Have disagreeing priors
 2. Some priors must favor one treatment, some the other

Multiple outcomes (side effects)

- Hypothesis testing does not protect against side effects outweighing benefits in primary outcome:
 - In sufficiently large trials, even small differences in “primary outcomes” will be detected, leading to headline conclusions that a new therapy is “effective”
 - Researcher definitions of primary outcomes often differ from patient-relevant outcomes (e.g., mortality)
- Empirical success rule can be applied to weighted averages of **all** patient-relevant outcomes observed in the trial
 - provided that patient-relevant outcomes are reported.

What sample sizes are sufficient?

- For two-armed trials with binary outcomes, using the empirical success rule yields these maximum expected losses:

Sample sizes	Near-optimality
20:20	0.0269
50:50	0.017
100:100	0.012
200:200	0.0085
500:500	0.0054
1000:1000	0.0038
4000:4000	0.0019
15000:15000	0.001

Downside of large sample sizes required by conventional testing rules

- Delay: It takes longer to recruit patients; hence, longer to reach conclusions.
- Crowds out trials of other treatments.
- Statistical significance requirement impedes subgroup analyses
 - There may be substantial heterogeneity in treatment effectiveness and the prevalence of side effects (e.g., by age)
 - The welfare weights attached to different outcome measures may vary with patient attributes.

Clinical trial landscape

- There are many alternative treatments in trials now
 - Each trial has a different set of inclusion criteria, a different PI, and only tests 1 innovation against its own definition of standard care.
 - Study populations differ across trials.
- It will be difficult to compare alternative treatments across trials.

Mutli-arm trials for Covid-19

- UK nationwide "Recovery" trial started with 5 arms
 - Standard care
 - Lopinavir-Ritonavir
 - Low dose corticosteroids (dexamethasone)
 - Hydroxychloroquine
 - Azithromycin
- Patients were assigned to treatments in a 2:1:1:1:1 ratio
- WHO organized an international "Solidarity" trial with 5 arms
 - Standard care
 - Remdesivir
 - Lopinavir-Ritonavir
 - Lopinavir-Ritonavir plus Interferon beta-1a
 - Chloroquine or hydroxychloroquine
- These trials allow comparisons of multiple treatments on same population.

Evaluating multi-arm trials such as Recovery

- The Recovery protocol calls for results to be analyzed using Dunnett's test. This is a multiple t-test procedure, with 0.05 Type I error probability that at least one test yields a positive statistically significant ATE. Presumably, the innovation with highest average outcome will be selected among those that pass the significance test. Otherwise, standard care will be selected.
- We contrast this with the empirical success rule, which selects the treatment with the highest average outcome, regardless of statistical significance.
- In practice, 3 treatment arms were stopped at different times. The results for each treatment were analyzed separately as if coming from a two-arm trial.

	Standard care	A	B	C	D
Sample size in each arm	500	250	250	250	250
Mortality rate of each treatment	0.25	0.15	0.20	0.30	0.35
Panel A: What happens if treatment decisions are made using two-sided Dunnett's test at 5% significance					
% of trials after which new treatment will be prescribed	25.65%	70.60%	3.75%	0	0
Loss from prescribing each treatment	0.1	0	0.05	0.15	0.2
Probability of error times the magnitude of loss	0.0257	0	0.0019	0	0
Expected loss given these mortality rates					0.0275
Panel B: What happens if treatment decisions are made using the empirical success rule					
% of trials after which new treatment will be prescribed	0.02%	92.95%	7.03%	0	0
Loss from prescribing each treatment	0.1	0	0.05	0.15	0.2
Probability of error times the magnitude of loss	0	0	0.0035	0	0
Expected loss given these mortality rates					0.0035

What sample sizes are sufficient?

- For five-armed trials with binary outcomes and 2:1:1:1:1 sample ratio, choosing the treatment using the empirical success rule and Dunnett's test rule imply the following maximum expected losses:

Sample sizes per arm	Near-optimality using Empirical Success rule	Near-optimality using Dunnett's test rule
100:50:50:50:50	0.0362	0.1224
200:100:100:100:100	0.0256	0.0855
500:250:250:250:250	0.0160	0.0532
1000:500:500:500:500	0.0112	0.0380
2000:1000:1000:1000:1000	0.0080	0.0274

- It is slightly better to divide the sample into equal-sized arms for the ES rule.

Near-optimality of empirical success rule with patient-specific treatment and multiple outcomes

- The above calculations concern settings where patients are observationally identical and trial outcomes are binary.
- In clinical practice, trial outcomes may take multiple values. Trials of COVID-19 drugs may report mortality outcomes and time to recovery for patients who survive. Patients may vary in treatment response by age, gender, and comorbidities.
- Methodological research has shown how to compute or bound the near-optimality of the empirical success rule when applied in a broad range of settings.

Near-optimality with binary primary and secondary outcomes

- Manski and Tetenov (2019) study near-optimality of the empirical success rule when there are two treatments and patient welfare is a weighted sum of binary primary and secondary outcomes. The primary outcome is survival. The secondary one denotes whether the patient suffers a specified side effect.
- When a patient does not suffer the side effect, we let welfare equal 1 if a patient survives and equal 0 if he does not survive. When a patient experiences the side effect, welfare is lowered by a specified fraction h . Thus, a patient with the side effect has welfare $1 - h$ if he survives and $-h$ if he does not survive.
- We develop an algorithm to compute the near-optimality of the empirical success rule.

Near-optimality with bounded outcomes

- Exact computation of near-optimality becomes onerous when outcomes can take many discrete values or are continuous.
- When outcomes are bounded, large-deviations inequalities yield upper bounds on the near-optimality of the empirical success rule. These bounds are simple to compute and are sufficiently informative to provide useful guidance to clinicians.
- Manski (2004) used the Hoeffding inequality for sample averages to derive an upper bound on near-optimality when there are two treatments.
- Manski and Tetenov (2016) extended the analysis to multi-arm trials. Let L be the number of treatments and V be the range of the outcome. When the trial has a balanced design, with n subjects per arm, the upper bounds on near-optimality are $(2e)^{-1/2}V(L-1)n^{-1/2}$ and $V(\ln L)^{1/2}n^{-1/2}$. The former is tighter than the latter for two or three treatments. The latter is tighter for four or more treatments.

Near-optimality with heterogeneous patients

- Patient response to treatments may vary with observed covariates. A clinician can assess the near-optimality of a decision criterion when applied to patients with similar covariates.
- In principle, a clinician may view each group of patients with similar covariates as a separate population and may apply the empirical success rule separately to each group.
- In practice, the ability to differentially treat patients with different covariates is limited by the failure of medical researchers to report how trial findings vary with covariates. Research articles often report only subgroup findings that are statistically significant.
- Information is lost when reporting research findings is tied to statistical significance. The analysis of this paper makes clear that estimates of treatment effects need not be statistically significant to be clinically useful.

Topics for future research

- We have considered treatment choice using data from one trial with full validity.
- Internal validity may be compromised by non-compliance and loss to follow up. External validity may be compromised by measurement of surrogate outcomes and study of patients who differ from those that clinicians treat in practice. The concept of near-optimality is applicable when analyzing data from trials with limited validity, but the calculations made in this paper require modification.
- A clinician may learn the findings of multiple trials and may be informed by observational data. Near-optimality is well-defined in these settings, but methods for application are yet to be developed.
- A further issue concerns dynamic treatment choice when new trials and observational data may emerge in the future. The concept of near-optimality should be extendable, but methodology is yet to be developed.
 - Dynamic analysis of treatment choice made with hypothesis tests may be especially difficult, because testing views standard care and new treatments asymmetrically. As new data accumulate, the designation of standard care may change, leading to a change in the null hypothesis when new trials are evaluated.

Technical Appendix

We use concepts and notation in Manski (2004) and Manski and Tetenov (2016, 2019).

The clinician must assign one of L treatments studied in the trial to each member of treatment population J .

Denote treatments by $T = \{1, 2, \dots, L\}$, with $t = 1$ being standard care.

Each $j \in J$ has a response function $y_j(\cdot): T \rightarrow Y$ mapping treatments $t \in T$ into patient-relevant outcomes $y_j(t) \in Y$. Outcomes can be multi-valued and multi-dimensional. Treatment response is individualistic.

The distribution $P[y(\cdot)]$ of the random function $y(\cdot): T \rightarrow Y$ describes treatment response across the population. The set of feasible distributions is $\{P_s, s \in S\}$, S indexing feasible *states of nature*.

In Tables 2 and 4, we include in S all logically possible outcome distributions.

Patient welfare is a known function $u: Y \rightarrow \mathbf{R}$ of individual outcomes.

For binary outcomes $Y = \{0, 1\}$, with 1 denoting success, and $u(y) = y$. For two-dimensional outcomes $y = (y_p, y_{se})$, where y_p denotes the primary outcome and y_{se} a side effect, Manski and Tetenov (2019) considered welfare function $u(y) = y_p - h y_{se}$.

Consider data generation. Ψ denotes the sample space. Q_s denotes the sampling distribution on Ψ in state of nature s . Q_s is the distribution of trial outcomes.

We consider trials that randomize a predetermined number of subjects n_t to each treatment t . The set $n_T \equiv [n_t, t \in T]$ of sample sizes defines the design. The total number of subjects is $N \equiv \sum_{t \in T} n_t$. The data ψ are the N pairs of individual treatment assignments t_i and outcomes y_i : $\psi = [(t_i, y_i), i = 1, 2, \dots, N]$.

Q_s is determined by the distribution of treatment response P_s and the trial design, with $Q_s(y_i | t_i) = P_s(y(t_i))$.

A statistical treatment rule maps sample data into a treatment allocation. A feasible rule is a function that randomly allocates persons across the different treatments. Let Δ denote the space of functions that map T into the unit interval and that satisfy the adding-up condition: $\delta \in \Delta \Rightarrow \sum_{t \in T} \delta(t, \psi) = 1, \forall \psi \in \Psi$. Each function $\delta \in \Delta$ defines a statistical treatment rule.

The mean welfare of treatment t in state of nature s is denoted by $\mu_{st} \equiv E_s[u(y(t))]$. The maximum mean welfare achievable in state s is $\max_{t \in T} \mu_{st}$.

After data ψ are observed, the fraction $\delta(t, \psi)$ of patients will be treated with treatment t , resulting in mean patient welfare $\sum_{t \in T} (\mu_{st} \delta(t, \psi))$. The mean welfare of patients across repeated realizations of the trial is

$$\int_{\Psi} \sum_{t \in T} (\mu_{st} \delta(t, \psi)) dQ_s(\psi) = \sum_{t \in T} \mu_{st} E_s[\delta(t, \psi)],$$

where $E_s[\delta(t, \psi)] = \int_{\Psi} \delta(t, \psi) dQ_s(\psi)$ is the expected (across samples) fraction of persons assigned to t .

Application of rule δ in state of nature s yields expected loss (regret)

$$(A1) \quad \max_{t \in T} \mu_{st} - \sum_{t \in T} \mu_{st} E_s[\delta(t, \psi)].$$

The near-optimality (maximum regret) of rule δ is the maximum of (A1) over all feasible states of nature:

$$(A2) \quad \max_s \left(\max_{t \in T} \mu_{st} - \sum_{t \in T} \mu_{st} E_s[\delta(t, \psi)] \right).$$

Hypothesis Testing Rules

First consider rules based on hypothesis tests for univariate outcomes y . Denote the sample mean of y observed in arm t of the trial by $\bar{y}_t = \frac{1}{n_t} \sum_{i:t_i=t} y_i$. To test the null hypothesis that all treatments have the same outcome distribution, use $\hat{\sigma}^2 = \frac{1}{N-L} \sum_{t \in T} \sum_{i:t_i=t} (y_i - \bar{y}_t)^2$ as the estimator of common variance. The t-statistic for comparing the mean outcome of treatment $t = 2, \dots, L$ with that of standard care equals $\tau_t = \frac{\bar{y}_t - \bar{y}_1}{\hat{\sigma} \sqrt{1/n_t + 1/n_1}}$. Let c be the critical value adjusted for multiplicity. We use the t-distribution for two-arm trials and the Dunnett's test critical value for multiple comparisons for multi-arm trials.

The test rule prescribes treatment 1 (standard care) to everyone if all t-statistics are below the critical value.:

$$\delta_H(1, \psi) \equiv 1_{\left\{ \max_{t \in \{2, \dots, L\}} \tau_t \leq c \right\}}.$$

If some t-statistics comparing treatments $2, \dots, L$ to standard care exceed the critical value, these treatments are considered statistically significantly better than standard care. We assume that among these treatments the one with the largest mean outcome in the trial will be prescribed (with equal probability if there is a tie).

$$\delta_H(t, \psi) \equiv \frac{1_{\left\{ \tau_t > c, \bar{y}_t = \max_{t' \in \{2, \dots, L\}} \bar{y}_{t'} \right\}}}{\sum_{t' \in \{2, \dots, L\}} 1_{\left\{ \tau_{t'} > c, \bar{y}_{t'} = \max_{t'' \in \{2, \dots, L\}} \bar{y}_{t''} \right\}}}.$$

The Empirical Success Rule

Let $\bar{u}_t = \frac{1}{n_t} \sum_{i:t_i=t} u(y_i)$ denote the average welfare observed in treatment arm $t = 1, 2, \dots, L$.

The empirical success rule prescribes the treatment with the largest observed average patient welfare. If there is a tie, all treatments with the largest observed average patient welfare are prescribed with equal probability.

$$\delta_{ES}(t, \psi) \equiv \frac{1_{\{\bar{u}_t = \max_{t' \in \{1, \dots, L\}} \bar{u}_{t'}\}}}{\sum_{t' \in \{1, \dots, L\}} 1_{\{\bar{u}_{t'} = \max_{t'' \in \{1, \dots, L\}} \bar{u}_{t''}\}}}.$$

Computing near-optimality for two-arm trials with binary outcomes

When computing the results in Table 2, S is the set of all distributions of binary outcomes with means $p_1 \equiv E[y(1)]$, $p_2 \equiv E[y(2)]$, $(p_1, p_2) \in [0, 1]^2$.

Let m_1 and m_2 denote the number of positive outcomes in each arm of the trial. $\psi = (m_1, m_2)$ is a sufficient statistic for the sample. Hence, it is sufficient to consider the sample space $\Psi = \{0, 1, \dots, n_1\} \times \{0, 1, \dots, n_2\}$. The probability density function of ψ is a product of two binomial density functions.

The function (A1) is continuous in (p_1, p_2) but may have multiple global and local maxima. We approximate the maximum in (A2) by grid search using 1000 possible values for each parameter equally spaced on $[0,1]$: $\{0.0005, 0.0015, \dots, 0.9995\}$.

Computing near-optimality for multi-arm trials with binary outcomes

In Table 4, S is the set of all distributions of binary outcomes with means $p_t \equiv E[y(t)]$, $t = 1, \dots, L$, $(p_1, \dots, p_L) \in [0, 1]^L$. Let m_t denote the number of positive outcomes in arm t of the trial. $\psi = (m_1, \dots, m_L)$ is a sufficient statistic for the sample. Hence, we consider the sample space $\Psi = \{0, 1, \dots, n_1\} \times \dots \times \{0, 1, \dots, n_L\}$.

The large size of the sample space makes it impractical to evaluate (A1) exactly. Given each value of (p_1, \dots, p_L) we simulate a large number of trial outcomes to approximate Q_s . Our computations of the maximum of (A2) proceed in three steps.

(1) We conduct a grid search using 51 possible values for each parameter $p_t \in [0, 0.02, \dots, 1]$. For each combination of parameters, we approximate the sampling distribution Q_s by simulating 100,000 trial outcomes. The results of this grid search suggest that the largest expected loss for the empirical success rule occurs when the parameters have the form $p_1 = a$, $p_2 = p_3 = p_4 = p_5 = b$, $a > b$. The largest expected loss for the Dunnett's test rule occurs when $p_1 = a$, $p_2 = b$, $p_3 = p_4 = p_5 = c$, $b > a$, $b > c$.

(2) We conduct a grid search over these two lower-dimensional parameter spaces using 101 possible parameter values from $[0, 0.01, \dots, 1]$ for a , b , and c . In this step we approximate Q_s by simulating 1,000,000 trial outcomes.

(3) We take 10 parameter combinations yielding the largest estimated expected loss for each decision rule in step 2 and recompute expected loss by simulating 100,000,000 trial outcomes. We do this to verify that our results are not affected by bias resulting from approximating Q_s by simulation.

Identification of Income-Leisure Preferences and Evaluation of Income Tax Policy

Charles F. Manski

Quantitative Economics, Vol. 5, No. 1, 2014, pp. 145-174.

Economists have long recognized that the relative merits of alternative income tax policies depend on the preferences of persons for income and leisure.

Income-leisure preferences play both positive and normative roles in analysis of tax policy.

The positive role is that preferences yield labor supply and other decisions that determine tax revenue.

The normative role is that social welfare aggregates personal preferences in utilitarian policy evaluation.

Among the simplifying assumptions that Mirrlees (1971) made in his seminal study of optimal utilitarian income taxation, he stated:

“The State is supposed to have perfect information about the individuals in the economy, their utilities and, consequently, their actions.”

He recognized the difficulty of inference on population preferences, writing:

“The examples discussed confirm, as one would expect, that the shape of the optimum earned-income tax schedule is rather sensitive to the distribution of skills within the population, and to the income-leisure preferences postulated. Neither is easy to estimate for real economies.”

I studied identification of income-leisure preferences using revealed-preference analysis of labor supply in Manski (2014). I reached this pessimistic conclusion:

“As I see it, we lack the knowledge of preferences necessary to credibly evaluate income tax policies.”

Taxation and Labor Supply

Standard economic theory does not predict the direction or magnitude of the response of labor supply to income taxation. To the contrary, it shows that a worker may rationally respond in disparate ways.

As tax rates increase, a person may rationally decide to work less, work more, or not change his labor supply at all. The silence of theory on labor supply has long been appreciated; see Robbins (1930).

Modern labor economics envisions labor supply as a complex sequence of schooling, occupation, and work effort decisions made under uncertainty over the life course, perhaps with bounded rationality. However, we need only consider a simple static scenario to see that a person may respond to income taxes in disparate ways.

Suppose that a person with a predetermined wage and no unearned income allocates each day between paid work and the various non-paid activities that economists have traditionally called leisure.

Let a proportional income tax reduce his wage by the prevailing tax rate, yielding his net wage.

Assume that the person allocates time to maximize utility, which is an increasing function of net income and leisure.

Different utility functions imply different relationships between the tax rate and labor supply.

The labor supply implied by utility functions in the CES family increases or decreases with the tax rate depending on the elasticity of substitution.

Other utility functions imply that labor supply is *backward-bending*.

Still other utility functions yield more complex non-monotone relationships between net wage and labor supply. See Stern (1986).

Given that theory does not predict how income taxation affects labor supply, prediction requires empirical analysis.

Robbins (1930) emphasized this, writing:

“we are left with the conclusion that any attempt to predict the effect of a change in the terms on which income is earned must proceed by inductive investigation of elasticities.”

Economists have performed a huge number of empirical studies of labor supply.

Reading the literature concerned with uncompensated (Marshallian) elasticities of labor supply, I was struck to find that while authors may differ on the magnitude of elasticities, they largely agree on the sign.

The consensus has been that increasing tax rates usually reduces work effort. Keane (2011) stated the directionality of the effect without reservation, writing:

“the use of labor income taxation to raise revenue causes people to work less.”

Considering the effect of a rise in a proportional tax, Meghir and Phillips (2010) wrote:

“in most cases this will lead to less work, but when the income effect dominates the substitution effect at high hours of work it may increase effort.”

Here and elsewhere, researchers may recognize the theoretical possibility that effort may increase with tax rates but view this as an empirical rarity rather than a regularity.

This view has been accepted in official government forecasts of the response of labor supply to income taxation. See Congressional Budget Office (2007).

Curiously, the opposite consensus prevailed early in the twentieth century. Gilbert and Pfouts (1958) cite assertions by Pigou and Knight in the 1920s that increasing tax rates increases work effort.

Examining the models of labor supply used in empirical research, I became concerned that the prevailing consensus on the sign of uncompensated elasticities may be an artifact of model specification rather than an expression of reality.

Models of labor supply differ across studies, but they have generally shared two key restrictive assumptions.

First, they suppose that labor supply varies monotonically with net wages.

Thus, model specifications do not generally permit backward-bending labor supply functions or other non-monotone relationships.

Second, they suppose that the response of labor supply to net wage is homogeneous within broad demographic groups.

With occasional exceptions, researchers specify hours-of-work equations that permit hours to vary additively across group members but that assume constant treatment response.

The literature contains some precedent for my concern that empirical findings on labor supply may be artifacts of model specification. Concluding his detailed comparison of alternative labor supply functions, Stern (1986) wrote:

“Our general conclusion must be in favour of diversity of functions and great caution in drawing policy conclusions on results based on a particular form.”

Stern and other writers such as Blundell and MaCurdy (1999) have called attention to the potential detrimental consequences of restrictive functional-form assumptions.

The reality may be that persons have heterogeneous income-leisure preferences and, consequently, heterogeneous labor-supply functions.

If so, estimates of models that assume monotonicity and homogeneity of labor supply can at most characterize the behavior of an artificial “representative” person. The estimates may not have even this limited interpretation.

In light of the above, I examined identification of income-leisure preferences.

I studied inference when data on time allocation under status-quo tax policies are interpreted through the lens of standard theory.

I found it productive to study the classical static model in which persons with separable preferences for private and public goods must allocate one unit of time to work and leisure.

I considered the use of revealed preference analysis to predict labor supply and tax revenue under a proposed policy that would alter persons' status-quo tax schedules.

The policies that I had in mind use tax revenue to produce public goods or to redistribute income from persons who pay positive income tax to ones who pay negative tax.

Basic Revealed-Preference Analysis

I first assumed only that persons prefer to have more income and leisure.

Basic revealed-preference analysis of the type pioneered by Samuelson (1938) shows that observation of a person's time allocation under a status-quo tax policy may bound his allocation under a proposed policy or may have no implications, depending on the tax schedules and the person's status-quo time allocation.

Basic analysis assuming only that more-is-better generically does not predict the sign of labor-supply response to change in the tax schedule.

I call this a “basic” analysis of revealed preference because it maintains no assumptions about preferences except that individual utility is an increasing function of net income and leisure. In short, more is better.

Tax Policy and Labor Supply

In the absence of assumptions restricting the population distribution of preferences, predicting population labor supply under a proposed tax policy simply requires aggregation of individual predictions. Hence, the analysis focuses on one person.

Let person j be endowed with wage w_j , unearned income z_j , and one unit of time. The person must allocate the time endowment between leisure and work. If he allocates a fraction $L \in [0, 1]$ to leisure and $1 - L$ to work, he receives gross income $w_j(1 - L) + z_j$.

The status-quo tax policy, denoted S , subtracts the work-dependent tax revenue $R_{jS}(L)$ from gross income, leaving j with net income $Y_{jS}(L) \equiv w_j(1 - L) + z_j - R_{jS}(L)$.

Taxes may be positive or negative. The $R_{jS}(\cdot)$ notation allows the status-quo tax schedule to be specific to person j .

Person j chooses a value of L from a set $\Lambda_j \subset [0, 1]$ of feasible leisure alternatives. For example, $\Lambda_j = \{0, \frac{1}{2}, 1\}$ means that the feasible options are full-time work ($L = 0$), half-time work ($L = \frac{1}{2}$), and no work ($L = 1$).

Preferences are expressed in the utility function $U_j(\cdot, \cdot)$, whose arguments are (net income, leisure). Utility is strictly increasing in both arguments. Let $L_{js} \in \Lambda_j$ denote the amount of leisure that j chooses under tax schedule $R_{js}(\cdot)$.

Utility maximization implies

$$U_j[Y_{js}(L_{js}), L_{js}] \geq U_j[Y_{js}(L), L], \quad \text{all } L \in \Lambda_j.$$

Using $U_j(\cdot, \cdot)$ to express preferences suppresses the possible dependence of preferences on public goods produced with tax revenue under policy S . This is innocuous if preferences are separable in private and public goods.

Predicting Labor Supply under a Proposed Tax Schedule

Suppose that one observes the wage, unearned income, and other tax-relevant attributes of person j . One also observes the leisure L_{jS} chosen by j under tax schedule $R_{jS}(\cdot)$.

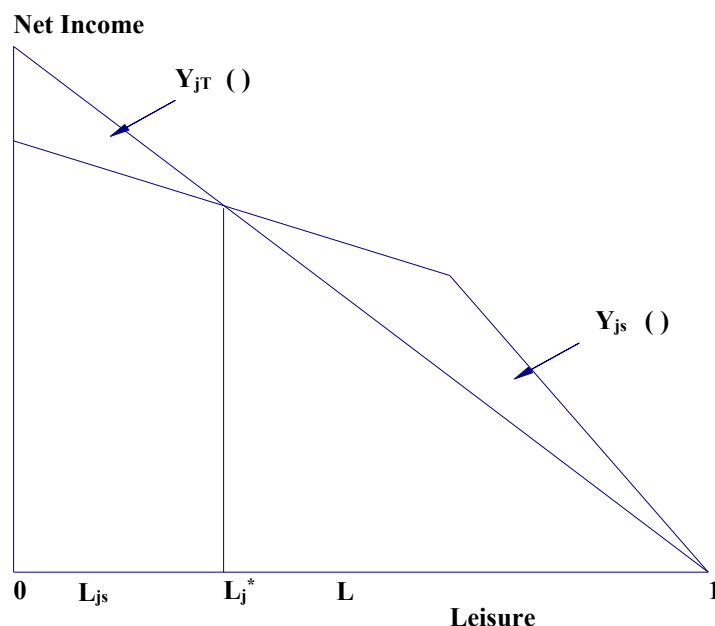
Let $R_{jT}(\cdot)$ denote the tax schedule if j were to face a proposed tax policy T . What can one predict about time allocation under $R_{jT}(\cdot)$?

The answer depends on the value of L_{jS} and on the budget sets $\{[Y_{jS}(L), L], L \in \Lambda_j\}$ and $\{[Y_{jT}(L), L], L \in \Lambda_j\}$ that j faces under the status-quo and proposed tax schedules.

I give an illustration, modestly extending the revealed-preference argument of Samuelson.

Policy S has a two-rate progressive schedule and T has a proportional one, the latter crossing the former from above when leisure equals L_j^* . Person j has no unearned income.

Figure 1: Net Income with Progressive and Proportional Tax Schedules



Let $L_{jS} \in [0, L_j^*]$, and consider any $L > L_j^*$.

$[Y_{jT}(L_{jS}), L_{jS}]$ is feasible under policy T. This pair is preferred to $[Y_{jS}(L_{jS}), L_{jS}]$.

$[Y_{jS}(L), L]$ is feasible under policy S. This pair is preferred to $[Y_{jT}(L), L]$.

Observation that j chose $[Y_{jS}(L_{jS}), L_{jS}]$ reveals that j prefers this pair to $[Y_{jS}(L), L]$.

Hence, he prefers $[Y_{jT}(L_{jS}), L_{jS}]$ to $[Y_{jT}(L), L]$. Thus, under T, j would not choose $L > L_j^*$.

Prediction When Downward-Sloping Net-Income Functions Cross Once

Basic revealed-preference analysis is complex in general, but it is simple when, as in Figure 1, the net income functions $Y_{jS}(\cdot)$ and $Y_{jT}(\cdot)$ implied by the status-quo and proposed tax schedules are both downward sloping.

Let $\Lambda_{j<}$ denote the feasible leisure values such that the (income, leisure) pairs in the $R_{jT}(\cdot)$ budget set are strictly smaller than some pair in the $R_{jS}(\cdot)$ budget set. That is,

$$\Lambda_{j<} \equiv \{L_{<} \in \Lambda_j: [Y_{jT}(L_{<}), L_{<}] < [Y_{jS}(L), L] \text{ for some } L \in \Lambda_j\}.$$

Let $\Lambda_{j>}$ denote the leisure values such that the (income, leisure) pairs in the $R_{jT}(\cdot)$ budget set are strictly larger than $[Y_{jS}(L_{jS}), L_{jS}]$, the pair that j chooses under policy S . That is,

$$\Lambda_{j>} \equiv \{L_{>} \in \Lambda_j: [Y_{jT}(L_{>}), L_{>}] > [Y_{jS}(L_{jS}), L_{jS}]\}.$$

Define $\Lambda_{j\leq}$ and $\Lambda_{j\geq}$ analogously, with weak inequalities replacing the strict ones.

When tax schedules are downward sloping, basic analysis can have predictive power for labor supply only if the two net-income functions cross at least once.

To see that a crossing is necessary for predictive power, consider a policy T such that $Y_{jT}(L) \geq Y_{jS}(L)$ for all $L \in [0, 1]$. Thus, at all L a person pays weakly less tax under T than S.

Then set $\Lambda_{j<}$ is empty because $Y_{jT}(L_{<}) < Y_{jS}(L) \Rightarrow L_{<} > L$.

Similarly, consider T such that $Y_{jT}(L) \leq Y_{jS}(L)$ for $L \in [0, 1]$. Thus, at all L a person pays weakly more tax under T than S. Then $\Lambda_{j\geq}$ is empty because $Y_{jT}(L_{<}) > Y_{jS}(L_{jS}) \Rightarrow L_{<} < L_{jS}$.

Policy comparison is simple when schedules S and T cross once. Observation that $L_{jS} \leq L_j^*$ implies that j would not choose $L > L_j^*$. Observation that $L_{jS} > L_j^*$ has no predictive power. The sign of labor supply response is predictable only if L_{jS} is at a corner or equals L_j^* .

Restrictions on the Preference Distribution

A huge distance separates basic revealed preference analysis from the practice of empirical analysis of labor supply.

The models used in empirical studies usually suppose that labor supply responds monotonically to variation in net wage. Moreover, it is common to assume that time allocation differs across persons only via a person-specific additive constant.

I explored the identifying power of assumptions restricting the distribution of preferences across persons.

I supposed that one observes the time allocation of each person in a population whose members may have heterogeneous preferences, wages, and face various status-quo tax schedules.

I found it analytically helpful to suppose that persons choose among a finite set of feasible (income, leisure) values rather than the continuum often assumed in the literature.

I used the discrete-choice framework of Manski (*IER* 2007) to characterize preferences and to predict aggregate labor supply and tax revenue when various assumptions restrict the distribution of preferences.

I studied the identifying power of two classes of assumptions.

The first assumes exogenous variation in choice sets. That is, groups of persons who face different choice sets have the same distribution of preferences. Exogenous variation makes the choice set an instrumental variable.

The second restricted the shape of the preference distribution. For example, one may assume all persons have CES preferences, with heterogeneous parameters.

These assumptions have identifying power, yielding partial rather than point identification of the preference distribution.

Computation of the identification region is challenging but tractable in some cases.

The classical static model is simplistic to warrant application in substantive study of labor supply. Nevertheless, I found it revealing to perform computational experiments that show the identifying power of alternative assumptions on the preference distribution.

Considering a setting with choice data on labor supply observed under a progressive tax schedule, the task was to predict tax revenue per capita under a proposed proportional schedule. I showed the identifying power of these assumptions:

(1) more is better + (2) persons in specified wage groups have the same preference distribution + (3) preferences are CES + (4) CES utility functions in a wage group have the same elasticity of substitution.

Enriching the Data for Identification of Preferences

I do not expect new theory will sharpen knowledge of preferences. Richer data may.

1. Resume the performance of randomized experiments with tax policy that began with the negative income tax experiments. Randomized experiments can make it credible to assume that groups who face different tax schedules have the same distribution of preferences.

2. Obtain repeated observations from individuals. With a static model of labor supply, it is useful to observe an individual with varying wages, unearned income, or tax schedules. Transitivity implies that basic revealed preference analysis has increasing predictive power as more choices are observed. However, this interpretation of data rests on the static model.

3. Pose choice scenarios with hypothetical wages or tax schedules, and ask persons to predict their choice behavior in these scenarios.

Implications for Utilitarian Policy Evaluation

A familiar exercise in normative public economics poses a utilitarian social welfare function and ranks tax policies by the welfare they achieve.

Performing this ranking requires knowledge of income-leisure preferences both to predict tax revenues and to compute the welfare achieved by alternative policies.

My analysis reached highly cautionary findings about present knowledge of income-leisure preferences, carrying implications for evaluation of tax policy.

Partial knowledge of preferences implies that one can only partially predict tax revenue and one can only partially evaluate the utilitarian welfare of policies.

Thus, choice of tax policy becomes a problem of planning under ambiguity.

Choosing Size of Government under Ambiguity: Infrastructure Spending and Income Taxation

Charles F. Manski

The Economic Journal, Vol. 124, May 2014, pp. 359-376.

The optimal size of government has been a subject of continuing debate.

Disagreements may stem in part from the fact that “size of government” is an imprecise term—persons using it may not interpret it the same way.

Persons with a common understanding of the term may disagree on what size is optimal.

They may have different normative perspectives on social welfare or different beliefs about the outcomes yielded by alternative policy choices.

Attempting to shed light on the optimal size of government, economists have studied planning problems.

A standard exercise specifies a set of feasible policies and a social welfare function, typically utilitarian. The planner is assumed to know the welfare achieved by each policy. The analysis characterizes the optimal policy.

A body of research stimulated by Mirrlees (1971) has studied use of income taxation to redistribute income. Economists have derived optimal tax schedules assuming the planner knows the income-leisure preferences of the population.

Another, following Barro (1990), has considered the use of public spending to promote growth. Economists have derived optimal tax-financed spending levels under the assumption that the planner knows the consumption preferences of the population and how public spending affects aggregate output.

Lack of knowledge of the welfare achieved by alternative policies limits the relevance of optimization studies to actual policy choice.

I examined choice of size of government as a problem of planning under ambiguity.

I focused on tax-financed public spending for infrastructure that aims to enhance private productivity. My focus on infrastructure spending was similar to the growth literature.

My formalization of the planning problem stemmed from the one used by Mirrlees to study optimal income taxation. I posed a static setting where each person allocates time to paid work and other activities (leisure). Persons have predetermined heterogeneous wages. An income tax schedule is applied to gross income, yielding net income.

Persons allocate time to maximize utility. Social welfare is utilitarian.

I departed from the Mirrlees setup in three main ways.

First, government chooses how much to spend on infrastructure and on activities that directly affect personal utility.

Second, persons may have heterogeneous preferences.

Third, the planner may have partial knowledge of population preferences and of the productivity of infrastructure spending.

Partial knowledge generates two distinct difficulties.

First, the planner may be unable to predict tax revenue with certitude and, thus, may not know if a policy will yield a balanced budget. Second, he may be unable to determine the welfare achieved by a policy. I bypassed the first issue and focused on the second.

The first issue is difficult to address in generality. Satisfactory evaluation of policies that may not yield balanced budgets requires study of a dynamic planning problem that permits surpluses and deficits to occur and recognizes their intertemporal welfare implications.

To bypass the complexity of dynamic policy evaluation under ambiguity, I considered settings in which the planner can ensure budget balance by choosing components of policies sequentially rather than simultaneously.

When budget balance can be ensured by sequential choice of policy components, planning may be studied using established criteria for static decision making under ambiguity.

I did so in a setting that is simple enough to yield easily interpretable closed-form findings.

In this setting, the planner only considers tax schedules that make the tax proportional to income.

Persons have Cobb-Douglas income-leisure preferences and no non-labor income.

These assumptions imply that time-allocation choices are invariant to policy and they enable the planner to achieve budget balance by first choosing the level of public spending, then observing the resulting population income, and finally choosing the tax rate to balance the budget.

For further simplicity, I assumed that all public spending is on infrastructure and that wages are person-specific positive constants multiplied by an aggregate production function expressing the wage-enhancing effect of spending on infrastructure.

Finally, I assumed that the planner has partial knowledge of the aggregate production function, obtained by observing the outcome of a status quo policy and by assuming that public spending enhances wages but with diminishing marginal returns.

Then the space of possible states of nature indexes all concave-monotone aggregate production functions that yield the outcome of the status quo policy.

In this setting, I showed that the planner can reasonably choose a wide range of spending levels. Thus, a society can rationalize having a small or large government.

The choice made depends on the decision criterion that the planner uses to cope with ambiguity. I considered planning that maximizes subjective expected welfare or that uses one of several criteria—maximin, minimax-regret, or a Hurwicz criterion—that do not place a subjective probability distribution on unknown quantities.

I drew conclusions that are methodologically constructive and substantively cautionary.

The methodologically constructive conclusion was that, when performing normative research on size of government, economists need not impose assumptions strong enough to yield optimal policies.

Decision theory provides a formal framework for study of planning under ambiguity.

The substantively cautionary conclusion was that study of planning with credible assumptions shows that a wide range of policy choices can be rationalized.

The only way to achieve credible conclusions about the desirable size of government is to vastly improve current knowledge of population preferences and the productivity of public spending.

There is no immediate way to achieve this, but a research program with a suitably long-run perspective may make progress possible.

**Minimax-Regret Climate Policy with Deep Uncertainty in Climate Modeling
and Intergenerational Discounting**

Stephen J. DeCanio, Charles F. Manski, and Alan H. Sanstad

Ecological Economics, Vol. 201, 2022, <https://doi.org/10.1016/j.ecolecon.2022.107552>

Introduction

Integrated assessment (IA) models enable quantitative evaluation of the benefits and costs of alternative climate policies.

Policy comparisons are performed by considering a planner who seeks to make optimal trade-offs between the costs of carbon abatement and the damages from climate change.

The planning problem has been formalized as a control problem with these components:

- (1) equations coupling GHG emissions and abatement to the accumulation of GHGs in the atmosphere and resulting temperature increases.
- (2) a damage function that quantifies economic effects of climate change in terms of the loss of global economic output as a function of temperature increases.
- (3) an abatement cost function that expresses the cost of actions to reduce GHG emissions relative to a stipulated baseline emissions trajectory.

Costs and damages are expressed as percentage reductions in gross world product.

The problem is to minimize the costs of abatement and damages over a time horizon.

Studying climate policy as a control problem presumes that a planner knows enough to make optimization feasible, but physical and economic uncertainties abound.

Physical scientists have performed multi-model ensemble (MME) analysis. Lacking a consensus climate model, they have developed multiple models. To cope with inter-model *structural uncertainty*, they compute simple or weighted averages of the outputs of MMEs. Choosing appropriate weights has been problematic.

Economists have estimated multiple damage functions and abatement cost functions. In general, economists have not performed MME analyses that combine multiple functions by averaging. They have reported disparate findings from separate studies.

Manski, Sanstad, and DeCanio (*PNAS* 2021) framed structural uncertainty in climate modeling as a problem of *partial identification*, generating *deep uncertainty*.

This problem refers to situations in which the underlying mechanisms, dynamics, or laws governing a system are not completely known and cannot be credibly modeled definitively even in the absence of data limitations in a statistical sense.

We proposed use of the *minimax-regret* (MMR) decision criterion to account for deep climate uncertainty in integrated assessment without weighting climate model forecasts.

We developed a theoretical framework for cost-benefit analysis of climate policy based on MMR and we applied it computationally with a simple illustrative IA model.

It is important to recognize deep uncertainty in both the physical and economic components of IA models.

Perhaps the most contentious economic issue has been how a planner should assess the costs and benefits of policies across generations.

In our new paper, we study choice of climate policy that minimizes maximum regret with deep uncertainty regarding both the correct climate model and the appropriate intergenerational assessment of policy consequences.

Economists have long framed intergenerational policy assessment using a discount rate.

They have evaluated climate policies by the present discounted value of the sum of abatement costs and the corresponding damages.

There has been considerable debate about what discount rate to use. The choice is consequential.

Low rates favor policies that reduce GHG emissions aggressively and rapidly.

High rates favor policies that act modestly and slowly.

To express deep uncertainty, we suppose that the appropriate discount rate lies within an interval that covers the spectrum of rates used in the literature.

Our mathematical analysis is a straightforward generalization of M-S-D.

There we supposed that the correct climate model is one of six prominent models in the literature on climate science, whereas the correct economic model is known.

We supposed that a planner compares six policies, each of which chooses an emissions abatement path that is optimal under one and only one of the six climate models.

Regret is the loss in welfare if the model used in policy making is not correct and, consequently, the chosen abatement path is actually sub-optimal.

The MMR rule chooses a policy that minimizes the maximum regret, or largest degree of sub-optimality, across all six climate models.

Here we suppose that the correct climate model is one of the six examined in M-S-D.

We characterize uncertainty about the discount rate by supposing that it takes one of the seven values $\{0.01, 0.02, \dots, 0.07\}$, a range that covers the rates commonly used.

This range reflects both empirical uncertainty about the future of the economy and normative uncertainty (or perhaps disagreement) about how the current population values the welfare of future generations.

Given joint uncertainty about the climate model and the discount rate, we suppose that a planner compares forty-three policies.

Forty-two policies entail choosing an emissions abatement path that is optimal under one of the $\{\text{discount rate, climate model}\}$ pairs. The remaining one is a passive policy in which the planner chooses no abatement.

The MMR criterion chooses a policy that minimizes maximum regret across all forty-three potential policies.

The MMR analysis points to use of a discount rate of 0.02 for climate policy.

The MMR decision rule keeps the maximum future temperature increase below 2°C for most of the parameter values used to weight costs and damages.

Prevalent Approaches to Climate and Discount-Rate Uncertainty

Averaging Outputs of MMEs of Climate Models

All climate models are based on a specific set of deterministic nonlinear partial differential equations describing large-scale atmospheric dynamics.

Implementation of the equations is subject to numerous practical choices involving discretization, solution methods, and other details.

Some components of the system – such as cloud formation and heat transfer between land surfaces and the atmosphere – are not yet fully understood and must be approximated.

For these reasons, multiple climate models have been developed and are in use, each reflecting different but credible choices in model design and implementation.

Existing models yield different projections of the global climate.

The range of projections produced by different models is a gauge of deep uncertainty about the climate system given the current state-of-the-science.

Virtually all methods of MME analysis combine model outputs into single projections of future climate variables.

However, climate researchers have recognized persistent methodological problems in combining model projections.

A common technique is to take the simple average across model projections of policy-relevant variables.

Researchers may compute weighted average projections when they believe that models can be ranked with respect to relative accuracy.

However, model performance with respect to historical data does not imply skill in predicting the future climate.

Combining MME outputs into single projected trajectories of the future global climate remains a challenging and unresolved problem.

The recent IPCC physical sciences report states:

“...despite some progress, no universal, robust method for weighting a multi-model projection ensemble is available...”

Uncertainties and Disagreements Regarding the Discount Rate

The economic losses from climate change are represented by damage functions that give the decreases in world-wide output resulting from increases in mean global temperature.

Economists study dynamic optimization by a planner, which entails discounting to quantify the present value of future economic costs and benefits.

The appropriate magnitude of the discount rate has been contentious.

Controversy persists in part because choice of a discount rate is not only an empirical question regarding the future of the economy.

It is also a normative question, concerning social preferences for equity across future generations.

A simple version of the famous *Ramsey formula* provides a transparent expression of the interplay of normative and empirical considerations in choosing a discount rate.

Let the social welfare function be additively separable in the utility of future generations.

Let ρ be the rate at which the planner discounts the utility of future generations.

Let the utility of a representative consumer be an increasing and concave function of consumption, with constant elasticity ($-\eta$) of marginal utility.

Let g be the annualized growth rate of consumption between time 0 and a future time t .

Ramsey showed that it is optimal to discount future consumption between the present (time 0) and time t at the rate $\delta = \rho + \eta g$.

From the perspective of the present, the empirical value of g may be uncertain. This uncertainty is similar conceptually to the uncertainty that climate modelers face as they attempt to project the future trajectory of climate variables.

ρ formalizes how the planner views intergenerational equity, with $\rho = 0$ if the planner gives equal weight to the welfare of all future generations and $\rho > 0$ if the planner weights welfare more heavily in the near future than in the distant future.

η formalizes the desirability of intergenerational consumption equity.

A planner may feel normative uncertainty about what values of ρ and η to use.

Supposing that the planner aims to represent society, a source of this uncertainty may be normative disagreements within the present population.

Such disagreements were evident in a dispute between Nordhaus (2007), who used the value $\rho = 0.03$, and Stern (2006), who used $\rho = 0.001$.

Stern concluded that policy should seek to reduce GHG emissions aggressively and rapidly. Nordhaus favored policies that act more modestly and slowly.

We argue against any attempt to cope with empirical and normative uncertainty by choosing a single discount rate.

Instead, we study formation of climate policy recognizing a set of possibly appropriate discount rates.

Minimax-Regret Policy Evaluation

To begin, we specify the control problem that a planner would solve with no uncertainty.

The Optimal-Control Problem

Let B_t represent baseline GHG emissions at time t , A_t be GHG abatement actions at time t under some climate policy, measured in the same units as emissions, $C(A_t)$ be the cost of these actions, and $E_t^{A_t} = B_t - A_t$ be the resulting net emissions.

We refer to A_t and $E_t^{A_t}$ as “paths” or “trajectories,” and we assume that abatement paths are chosen from some space of feasible paths.

Emissions paths are used as inputs to a climate model M .

We focus on the global mean temperatures projected by M as a function of these paths.

Let $T(E_t^{A_t}, M)$ be the global mean temperature at time t determined by the GHG trajectory $E_t^{A_t}$ when it is predicted by the climate model M .

Then a damage function can be written as $D\left(T(E_t^{A_t}, M)\right)$.

For abatement path A_t and climate model M , denote the associated total cost (abatement plus damages) at time t as

$$\mathbb{C}(A_t, M) \equiv C(A_t) + D\left(T(E_t^{A_t}, M)\right)$$

A policymaker seeks to minimize the present value of cost over a planning horizon. As usual in the climate economics literature, we assume an infinite horizon.

The control problem given climate model M is to solve

$$\min_{A_t} \int_0^{\infty} \mathbb{C}(A_t, M) e^{-\delta t} dt$$

where δ is the discount rate.

We suppose that the optimal A_t is chosen with commitment at time zero. That is, it is not updated over time as new climate or cost information is obtained.

Under certain assumptions, this optimization problem has a unique solution.

The Minimax-Regret Decision Rule

Let $\Delta = \{\delta_1, \dots, \delta_K\}$ be a set of discount rates and $\mathbf{M} = \{M_1, \dots, M_N\}$ be a model ensemble.

The planner now faces the problem of minimizing cost over the horizon while recognizing joint {discount rate, model} uncertainty.

For rate δ_i and model M_j , let $A_{t;\delta_i,M_j}^*$ be the optimal abatement path defined by

$$A_{t;\delta_i,M_j}^* = \arg \min_{A_t} \int_0^{\infty} \mathbb{C}(A_t, M_j) e^{-\delta_i t} dt$$

Let $\mathbb{C}^*(A_{t;\delta_i,M_j}^*, \delta_i, M_j)$ be the associated minimum cost:

$$\mathbb{C}^*(A_{t;\delta_i,M_j}^*, \delta_i, M_j) = \int_0^{\infty} \mathbb{C}(A_t^*, M_j) e^{-\delta_i t} dt$$

Now consider any feasible abatement trajectory A_t . The *regret* $\mathbb{R}(A_t, \delta_i, M_j)$ associated with A_t , when discount rate δ_i and climate model M_j describe the world, is the difference between the cost of A_t and the cost of the *optimal* policy associated with δ_i and M_j :

$$\mathbb{R}(A_t, \delta_i, M_j) = \int_0^{\infty} \mathbb{C}(A_t, M_j) e^{-\delta_i t} dt - \mathbb{C}^*(A_{t;\delta_i M_j}^*, \delta_i, M_j)$$

To apply the MMR rule, the planner considers each feasible abatement path A_t and finds the model and discount rate combination that maximizes regret, solving the problem

$$\begin{aligned} & \max_{\delta_i, M_j} \mathbb{R}(A_t, \delta_i, M_j) \\ &= \max_{\delta_i, M_j} \left[\int_0^{\infty} \mathbb{C}(A_t, M_j) e^{-\delta_i t} dt - \mathbb{C}^*(A_{t;\delta_i, M_j}^*, \delta_i, M_j) \right] \end{aligned}$$

The MMR solution is to find A_t to solve the problem

$$\min_{A_t} \left[\max_{\delta_i, M_j} \mathbb{R}(A_t, \delta_i, M_j) \right]$$

Rather than use the MMR rule, one might use the minimax rule, which embodies the principle of preparing for the worst case.

MMR analysis uses information in a more nuanced and less conservative way.

If a climate policy maker selects one model and discount rate from an ensemble and chooses an emissions abatement path that is optimal for that {model, rate} pair, regret is the excess cost of that abatement path if a different (model, rate) pair is the correct one.

Thus, regret measures the potential sub-optimality of policies.

Choosing a policy to minimize maximum regret means choosing one to minimize the maximum degree of sub-optimality across the set of policies under consideration.

Use of Δ to Express Empirical and Normative Uncertainty

The term “uncertainty” has usually referred to incomplete knowledge of the empirical environment of a decision maker, called the “state of nature” or the “state of the world.”

This notion of uncertainty applies to incomplete knowledge of the future global temperature, abatement costs, and damages under alternative climate policies.

We also consider uncertainty about the discount rate.

Our use of the set Δ to express both empirical and normative uncertainty regarding the discount rate departs from the usual decision-theoretic focus on empirical uncertainty.

Normative uncertainty may have an empirical source, namely incomplete knowledge of the population preferences that a utilitarian planner would seek to maximize.

The planner may face the difficult task of representing a population whose members may not be clear about their time preferences or concern with intergenerational inequalities.

Using Δ to express normative uncertainty is a more radical departure from the decision-theoretic norm if normative disagreements exist within the present population.

A segment of the population may strongly value intergenerational equity whereas another segment may be less concerned with the fate of future generations.

Then one may think it necessary to abandon the idealization of a utilitarian planner and replace it with conceptualization of policy making as a non-cooperative political game.

We nonetheless find it attractive to study MMR decision making in this setting.

The MMR rule has some appeal as a broadly acceptable mechanism for policy choice.

Recall that the *regret* of a policy in a specified state of nature measures its degree of sub-optimality in that state, and maximum regret measures the maximum degree of sub-optimality across all states.

Suppose that the members of a heterogeneous present population disagree on what {discount rate, model} should be considered the “true” state of nature.

Then use of the MMR rule to choose policy minimizes the maximum degree of sub-optimality that will be experienced across the population.

Computational Model

To show the consequences of adoption of the MMR decision rule, we present a simple IA model that summarizes the essential economic and physical mechanisms.

The standard in the literature has been to report results about a century into the future.

Analyzing the uncertainty associated with discount rates necessitates attention to longer time horizons.

Phenomena in the more distant future that are negligible in economic terms with high discount rates become salient with low rates.

Model Details

As a simple expression of complex climate dynamics, we use Matthews *et al.* (2009).

They showed that the “carbon-climate response” (CCR), the change in global mean temperature over periods of decades or longer, varies approximately linearly with the increase in cumulative carbon emissions over the same period.

Net cumulative emissions is

$$\mathbf{E}_t^{A_t} = \int_0^t E_t^{A_t} dt = \int_0^t (B_t - A_t) dt$$

There is no requirement that $(B_t - A_t)$ be non-negative. A_t exceeding B_t implies adoption of mitigation measures that yield negative net emissions.

The CCR vary across climate models. The CCR parameter m for model M_j is estimated by determining the model's projected temperature response when driven by a carbon emissions path according to

$$T_t = m_j \mathbf{E}_t^{A_t}, \quad j = 1, \dots, 6$$

where T henceforth indicates the temperature increase over its initial value at time $t = 0$.

We estimate m_j with historic and projected emissions and temperature data from model j .

Our model ensemble \mathbf{M} comprises six Earth System Models (ESMs). These ESMs were used in the Climate Model Intercomparison Project Phase 5 (CMIP5).

Table 1
Earth system models used to estimate Carbon-Climate Response (CCR)
parameters, with estimated CCR values (°C per teraton carbon)

<i>Model and model number</i>	<i>CCR</i>
1. GFDL-ESM-2G - Geophysical Fluid Dynamics Laboratory Earth System Model version 2G	0.00157
2. BCC-CSM-1 - Beijing Climate Center Climate System Model version 1.1	0.00186
3. FIO-ESM - FIO-ESM - First Institute of Oceanography Earth System Model	0.00194
4. Had-GEM2-ES - Hadley Global Environmental Model 2 - Earth System	0.00229
5. IPSL-CM5A-MR - Institut Pierre Simon Laplace Coupled Model 5A - Medium Resolution	0.00236
6. MIROC-ESM - Model for Interdisciplinary Research on Climate - Earth System Model	0.00244

We specify abatement cost and climate damage functions in quadratic form to implement the IA model as an optimal control problem, allowing for plausible non-linearity in these functions as the abatement effort A_t and the global temperature increase T_t at time t :

$$C(A_t) = \frac{1}{2} \alpha A_t^2$$

$$D(T_t) = \frac{1}{2} \beta T_t^2$$

The quadratic form and value of α are derived from Dietz and Venmans (2019).

The quadratic form and the value of β are taken from Nordhaus and Moffat (2017).

A baseline emissions trajectory B_t is derived from the “Representative Concentration Pathway (RCP) 8.5” scenario in its extended version to year 2500.

This envisions a relatively high growth rate of global carbon emissions from fossil fuel use through the 21st century, followed by a peak plateau period of constant emissions until 2150, and then a decline to a very low level by 2250.

A smoothed functional form having the same general shape as the RCP 8.5 was fitted by nonlinear least squares. The fitted equation for B_t is

$$B_t = \left(\theta t + \frac{B_0}{\exp(\theta \varphi)} \right) \exp(-\theta(t - \varphi)) .$$

The control problem is to minimize, for a given discount rate and model, the present value of abatement costs plus climate damages over an infinite horizon, subject to the dynamic relationship between cumulative emissions and temperature:

$$\min_{A_t} \int_0^{\infty} \frac{1}{2} (\alpha A_t^2 + \beta T_t^2) e^{-\delta t} dt$$

subject to

$$\frac{d}{dt} \mathbf{E}_t^{A_t} = E_t^{A_t} = B_t - A_t$$

$$T_t = m \mathbf{E}_t^{A_t}$$

$$\mathbf{E}_0^{A_t} = \mathbf{E}_0$$

The last equation specifies an initial condition for net cumulative emissions.

First-order necessary conditions include two coupled differential equations in abatement and the atmospheric greenhouse gas concentration associated with the optimal abatement:

$$\frac{dA_t}{dt} = \delta A_t - \frac{\beta m^2}{\alpha} \mathbf{E}_t^{A_t}$$
$$\frac{d\mathbf{E}_t^{A_t}}{dt} = B_t - A_t$$


These equations can be solved in closed form for A_t and $\mathbf{E}_t^{A_t}$.

The model satisfies convexity properties implying that the first-order conditions are sufficient for these to be unique optimal solutions to the control problem.

The left-hand panel of Figure 1 shows the baseline B_t and the optimal abatement A_t for a particular set of parameters.

The right-hand panel shows net cumulative emissions under A_t and under a policy of no abatement ($A_t = 0$ for all t).

Figure 1 – Trajectories of B_t , optimal A_t , $\mathbf{E}_t^{A_t}$, and $\mathbf{E}_t^{\text{noabatement}}$ for $m = 0.002286$, $\alpha = 0.000125$, $\beta = 0.018$, and $\delta = 0.05$



MMR Analysis

We discussed our climate model ensemble \mathbf{M} above. We specify the possible discount rates as $\Delta = \{0.01, 0.02, 0.03, 0.04, 0.05, 0.06, 0.07\}$. There are thus forty-two combinations of (δ, m) expressing the range of deep uncertainty.

Regrets can be calculated for any feasible abatement path A_t .

To keep calculation tractable, and because a planner may restrict attention to policies that are optimal in some state of nature, we consider policies that are optimal for some (δ, m) in $\Delta \times \mathbf{M}$. There are 42 such policies, and a 43rd when “No Abatement” is a possibility.

To explore the sensitivity of the MMR policy to the α and β parameters, we calculated the MMR for nine combinations of α and β .

Table 2

Values of MMR, uncertain Model and δ , for combinations of α and β
 Potential values of δ : {0.01, 0.02, 0.03, 0.04, 0.05, 0.06, 0.07}

Baseline RCP 8.5 (fitted): $B_t = (\theta t + \frac{B_0}{\exp(\theta\varphi)} \exp(-\theta(t - \varphi)))$; $\theta \rightarrow 0.0123125$, $\varphi \rightarrow 339.565$

$\alpha = 0.000075$ $\beta = 0.014$			$\alpha = 0.000075$ $\beta = 0.018$			$\alpha = 0.000075$ $\beta = 0.022$		
Model	δ	MMR	Model	δ	MMR	Model	δ	MMR
IPSL	0.02	0.172	HAD	0.02	0.172	HAD	0.02	0.178
$\alpha = 0.000125$ $\beta = 0.014$			$\alpha = 0.000125$ $\beta = 0.018$			$\alpha = 0.000125$ $\beta = 0.022$		
Model	δ	MMR	Model	δ	MMR	Model	δ	MMR
MIROC	0.02	0.266	IPSL	0.02	0.273	IPSL	0.02	0.284
$\alpha = 0.0002$ $\beta = 0.014$			$\alpha = 0.0002$ $\beta = 0.018$			$\alpha = 0.0002$ $\beta = 0.022$		
Model	δ	MMR	Model	δ	MMR	Model	δ	MMR
MIROC	0.02	0.478	MIROC	0.02	0.436	MIROC	0.02	0.423

For all (α, β) combinations, the discount rate corresponding to the MMR solution is 0.02.

The IA model allows for calculation of the maximum temperature increase that will be reached for any policy path, and how long it will take to reach that temperature.

Because the actual state of the world is unknown, the temperature increase under the MMR policy cannot be known at the time the policy decision is made.

What is known is that it will be less than or equal to the maximum over all six models, which will occur if MIROC is the true model because m_6 is the greatest of the CCRs.

We find that the MMR decision rule keeps the maximum future temperature increase below 2°C above the 1900-09 level for most parameter values.

Table 3**Values of Maximum Temperature Increase (Tmax) in °C and Years after 2000 when reached,**

<i>$\alpha = 0.000075$ $\beta = 0.014$</i>			<i>$\alpha = 0.000075$ $\beta = 0.018$</i>			<i>$\alpha = 0.000075$ $\beta = 0.022$</i>		
MMR Model	Years	Tmax	MMR Model	Years	Tmax	MMR Model	Years	Tmax
IPSL	124	1.248	HAD	121	1.056	HAD	118	0.879

<i>$\alpha = 0.000125$ $\beta = 0.014$</i>			<i>$\alpha = 0.000125$ $\beta = 0.018$</i>			<i>$\alpha = 0.000125$ $\beta = 0.022$</i>		
MMR Model	Years	Tmax	MMR Model	Years	Tmax	MMR Model	Years	Tmax
MIROC	134	1.831	IPSL	130	1.564	IPSL	125	1.315

<i>$\alpha = 0.0002$ $\beta = 0.014$</i>			<i>$\alpha = 0.0002$ $\beta = 0.018$</i>			<i>$\alpha = 0.0002$ $\beta = 0.022$</i>		
MMR Model	Years	Tmax	MMR Model	Years	Tmax	MMR Model	Years	Tmax
MIROC	149	2.660	MIROC	141	2.187	MIROC	135	1.859

Discussion

The MMR rule provides a reasonable way to form climate policy with empirical uncertainty about the climate and normative uncertainty regarding the discount rate.

Our computational analysis offers a new reason for using a low discount rate in climate policy analysis, on the order of 2% per annum.

This discount rate encompasses the pure rate of time preference, intergenerational inequality aversion, projection of the economy's future rate of growth, and other factors that potentially can affect the discount rate.

MMR decision making copes with deep uncertainty without adopting the extreme conservatism of minimax decisions.

MMR enables a planner to deal with heterogeneous populations, who may not themselves be clear about their time preferences or concern with intergenerational equity.

There is no scientific or economic reason that everyone should hold the same normative values. Some people may have only a vague understanding of discounting.

We also find it appealing to view MMR as a consensus-building mechanism.

Calculating regrets enables people with different values to see how implementation of alternative policies might play out from their perspectives.

Our IA model is simple and computationally tractable.

This is partially because we have not considered all possible sources of uncertainty.

The appropriate baseline emissions path is highly uncertain.

The abatement cost and climate damage functions are also uncertain.

We have addressed this partially by sensitivity analysis, calculating MMR solutions with various parameters (α, β) on abatement cost and climate damages.

It would be desirable to expand the analysis to encompass deep uncertainty about the correct values for these weights, a formidable computational task.