

I

Credible Planning under Uncertainty

A foundational objective of the Constitution of the United States is to “promote the general Welfare.” The Preamble 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.

The Constitution does not define “general Welfare.”

A century later, Marshall (1890) began his *Principles of Economics* with this sentence (p. 1):

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.

The word “wellbeing” may be synonymous with welfare.

In this century, a report on clinical practice guidelines by the US Institute of Medicine (IOM) stated (Institute of Medicine, 2011, p. 4):

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 report did not specify what it means to optimize patient care.

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 may have rhetorical appeal but they lack substance. 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. Maximization is ill defined if the consequences are sufficiently uncertain that no action is clearly best. My concern is reasonable societal decision making in such settings.

What are the uncertainties with which planning must cope? They are too many and varied to summarize easily. For now, I will simply list those that I have studied, each of which will be discussed in this book. These include numerous uncertainties in medical risk assessment and prediction of treatment response; see Manski (2019a) for a broad exposition. There is much uncertainty in the epidemiological models used to predict the spread of infectious diseases, which inform choice of vaccination policy (Manski, 2010, 2017). There is also much uncertainty in the physical science climate models used to predict future climate change, which inform choice of climate policy (Manski, Sanstad, and DeCanio, 2021), and in the discount rate used to form a social welfare function (DeCanio, Manski, and Sanstad, 2022).

Challenging uncertainties 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. Manski (2007c) provides an abstract analysis. Manski (2014a, 2014b) examined how uncertainties about preferences and behavior complicate evaluation of income tax policies, where a central consideration is the relative preferences of potential workers for consumption goods and for availability of time to enable nonpaid activities. I have shown how uncertainty about the effect of policing on criminal behavior complicates evaluation of proactive policing programs (Manski, 2006).

ORGANIZATION OF THE BOOK

I lay out basic themes in abstraction in this opening chapter and flesh them out in what follows. Part I, constituting Chapters 2 through 4, is concerned with characterization of uncertainty. Part II, being Chapters 5

through 9, describes my research analyzing particular classes of planning problems. Chapter 10 looks ahead to performance of future research on social planning under uncertainty.

In this initial chapter, Section 1.1 calls attention to the prevalent research practice that studies planning with *incredible certitude*. Section 1.2 contrasts the conceptions of uncertainty in consequentialist and axiomatic decision theory. Section 1.3 presents the formal structure of consequentialist theory, which will be used throughout the book. Section 1.4 explains the prevalent econometric characterization of uncertainty, which distinguishes identification problems and statistical imprecision. Section 1.5 discusses the distinct perspectives on social welfare expressed in various strands of research on planning.

In Part I, Chapter 2 demonstrates how incredible certitude harms analysis of planning and assesses explanations that have been suggested for the prevalence of incredible certitude. Chapter 3 considers the central econometric problem of identification of treatment response. Chapter 4 discusses the comparably central problem of identification of choice behavior and the distribution of personal welfare in a society.

In Part II, Chapter 5 presents a core part of my work on treatment of individuals under ambiguity, developing the theme that diversification may be socially beneficial. Chapter 6 shows that use of statistical decision theory can improve treatment choice with data from statistically imprecise randomized trials, replacing the common use of hypothesis testing. Chapter 7 discusses my research on personalized treatment under uncertainty, where the planner wants to condition treatment on observed covariates but does not know how treatment response varies across persons.

Chapter 8 considers an important setting where treatment response has social interactions, this being vaccination to prevent transmission of infectious disease. Moving from treatment of individuals to global planning, Chapter 9 expositis my collaborative research on choice of a greenhouse gas abatement policy to reduce planetary warming when the physics of climate determination and the discount rate used in the social welfare function are uncertain. Chapter 10 looks ahead, calling for work that strengthens the foundations for planning under uncertainty, and touching on certain planning problems that need immediate and long-term attention.

As far as I am aware, only a small body of other research engages any of the themes that I will discuss. In the late 1970s, Johansen (1978) called for research on macroeconomic planning under uncertainty,

stating (pp. 263–264): “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 certain possible deviations of reality from the assumptions maintained in conventional macroeconomic models; see Hansen and Sargent (2008). Their work uses concepts of robust decision analysis, which I will explain in Section 1.3. Barlevy (2011) reviews work on macroeconomic policy under ambiguity.

1.1 THE PREVALENT STUDY OF PLANNING WITH INCREDIBLE CERTITUDE

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, in some sense, welfare is intended to express the values of society rather than the preferences of a dictator. The public at large may not be familiar with the formal structure of 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. They 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 structure and develop a system for national defense. Local governments choose how to maintain roads, perform policing, and organize public education. Planners need not be governmental. Clinicians make medical choices on behalf of patients. Parents act as planners for their families. In these settings and many more, the objective of the planner may be to maximize some idea of social welfare.

Welfare economics has also sought to shed light on noncooperative societal decision processes, where no actual planner exists. In the late 1700s, Adam Smith metaphorically suggested that an *invisible hand* makes decentralized decision making in market economies promote social welfare. Between then and the mid 1900s, 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 societal outcomes when the idealized conditions of the fundamental theorems of welfare economics do not hold. The social welfare achieved by a hypothetical planner has 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. Even when actual societal decisions are made by processes distant from planning, study of hypothetical planning problems has been valuable to clarify the respects in which the members of society agree and to make explicit the nature of disagreements.

I wrote previously that welfare economics has studied maximization of welfare. 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 practice as policy analysis with *incredible certitude* (Manski, 2011b, 2013c).

Planning with incredible certitude can harm society in multiple ways. Most obviously, 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 the current state of society and accurate predictions of policy outcomes, they may make substantively poor decisions. Moreover, they will not recognize the value of new research aiming to improve knowledge. Nor will they appreciate the potential usefulness of decision strategies that may help society cope with uncertainty and learn. In Chapter 2, I will present a typology of research practices that generate incredible certitude and discuss many specific cases.

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. Mongin and Pivato (2016) began their review article with this sentence (p. 711): “PERHAPS surprisingly, uncertainty plays no role whatsoever in the classical works on social welfare.”

Addressing the failure 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-five years. The word “credible” is inevitably subjective and difficult to pin down, but I use it nonetheless.

1.2 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 action A yields higher welfare in environment s_1 and action B yields higher welfare in s_2 . If it is not known whether s_1 or s_2 is the actual choice environment, 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 presumes the existence of an abstract decision maker who must choose among a specified set of actions. The decision maker could be an individual, a firm, or another institution. When the decision maker is an entity making societal decisions, it is a social planner. Thus, decision theory provides the formal basis for the study of social planning under uncertainty.

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, called the *state space*, is a primitive concept which provides the most basic expression of uncertainty. The larger the state space, the less the decision maker knows about the consequences of each action. Decision theorists usually describe the state space mathematically, without reference to an actual choice problem. For example, they might describe it as a finite or a convex set.

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. Arguing for the psychological realism of subjective probabilities, Tversky and Kahneman (1974) made plain the difference between the two perspectives, writing (p. 1130):

It should perhaps be noted that, while subjective probabilities can sometimes be inferred from preferences among bets, they are normally not formed in this fashion. A person bets on team A rather than on team B because he believes that team A is more likely to win; he does not infer this belief from his betting preferences. Thus, in reality, subjective probabilities determine preferences among bets and are not derived from them.

Two conceptually distinct but mathematically related approaches have been used to develop criteria for reasonable decision making. Consequentialist theory focuses on the substantive consequences of choices. Axiomatic theory poses choice axioms that characterize consistency of behavior across choice settings and proves *representation theorems* relating choice axioms to consequentialist decision criteria. My research has applied consequentialist rather than axiomatic theory. I explain why in Sections 1.2.1 and 1.2.2.

1.2.1 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 mainly studied criteria that, in some sense, works uniformly well over all of the state space. Two prominent interpretations of this broad idea are the maximin and minimax regret criteria. I will formalize these criteria in Section 1.3 and apply them throughout the book, particularly minimax regret.

The decision theory used in my research on planning is consequentialist. I suppose that the objective is to make substantively good societal decisions in particular settings. To accomplish this, I suppose that a planner specifies a suitable welfare function, expresses uncertainty in a credible manner, and uses these primitives to make a decision. The suitability of a welfare function and the credibility of an expression of uncertainty are context specific. These matters will be discussed in general terms in Sections 1.4 and 1.5 and in specific contexts in Part II.

1.2.2 Axiomatic Decision Theory

Axiomatic decision theory poses principles, called axioms, for consistency of hypothetical behavior across a class of potential choice problems. Researchers 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 theorists 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 use some consequentialist decision criterion, coping with uncertainty in some manner.

Perhaps the most famous representation theorems are those of Von Neumann and Morgenstern (1944) and Savage (1954). Both theorems establish that adherence to certain axioms is equivalent to maximization of expected utility. They differ mainly in that the probability distribution on the state space used to form expected utility is pre-specified in the former work and determined within the theory in the latter. Von Neumann and Morgenstern (VN-M) viewed the probability distribution as a primitive concept. Savage viewed the distribution as a construct that may in principle be inferred from analysis of choice behavior. I explain this distinction further on. I emphasize that in neither theorem does the probability distribution have any necessary connection to an objective reality.

Axiomatic theorists have long debated which axioms have normative appeal. Appraisal of normative appeal rests on introspection, so there should be no expectation that consensus will emerge. Indeed, decision theorists exhibit considerable difference in opinion. Binmore (2009) catalogues and assesses a wide spectrum of consistency axioms.

Why should one consider the VN-M, Savage, or other axioms to be compelling? No theorem answers this question. Instead, decision theorists call for introspection. In lecture notes for a Ph.D. course in decision theory, Kreps (1988) counseled a decision maker contemplating application of the VN-M theorem that he must first (p. 5): “Decide that you want to obey the axioms because they seem reasonable guides to behavior.”

Considering the matter at length, Savage (1954) put it this way (p. 7):

I am about to build up a highly idealized theory of the behavior of a “rational” person with respect to decisions. In doing so I will, of course, have to ask you to agree with me that such and such maxims of behavior are “rational.” In so far as “rational” means logical, there is no live question; and, if I ask your leave there at all, it is only as a matter of form. But our person is going to have to make up his mind in situations in which criteria beyond the ordinary ones of logic will be necessary. So, when certain maxims are presented for your consideration, you must ask yourself whether you try to behave in accordance with them, or, to put it differently, how you would react if you noticed yourself violating them.

After discussing the positive role of logic in guiding actual human behavior, Savage wrote (p. 20):

The principal value of logic, however, is in connection with its normative interpretation, that is, as a set of criteria by which to detect, with sufficient trouble, any inconsistencies there may be among our beliefs, and to derive from the beliefs we already hold such new ones as consistency demands. It does not seem appropriate here to attempt an analysis of why and in what contexts we wish to be consistent; it is sufficient to allude to the fact that we often do wish to be so.

Then, addressing his basic axiom P1, which assumes that the decision maker places a complete binary preference ordering on all potential actions, he wrote:

Pursuing the analogy with logic, the main use I would make of P1 and its successors is normative, to police my own decisions for consistency and, where possible, to make complicated decisions depend on simpler ones. Here it is more pertinent than it was in connection with logic that something be said or why and when consistency is a desideratum, though I cannot say much.

Thus, Savage opined that humans may want their behavior to be consistent beyond the degree required by logic, but he was unable to explain why.

In a famous critique of the Savage axioms, Ellsberg (1961) sharply questioned the Savage conclusion that a rational decision maker must behave as if he places a subjective probability distribution on the state space. He observed that thoughtful persons sometimes exhibit behavioral patterns that violate the Savage axioms in ways implying that they do not hold subjective distributions. Considering this behavior, he wrote (p. 669):

Are they foolish? It is not the object of this paper to judge that. I have been concerned rather to advance the testable propositions: (1) certain information states can be meaningfully identified as highly ambiguous; (2) in these states, many reasonable people tend to violate the Savage axioms with respect to certain choices; (3) their behavior is deliberate and not readily reversed upon reflection; (4) certain patterns of “violating” behavior can be distinguished and described in terms of a specified decision rule.

If these propositions should prove valid, the question of the optimality of this behavior would gain more interest. The mere fact that it conflicts with certain axioms of choice that at first glance appear reasonable does not seem to me to foreclose this question; empirical research, and even preliminary speculation, about the nature of actual or “successful” decision making under uncertainty is still too young to give us confidence that these axioms are not abstracting away from vital considerations. It would seem incautious to rule peremptorily that the people in question should not allow their perception of ambiguity, their unease with their best estimates of probability, to influence their decision: or to assert that the manner in which they respond to it is against their long-run interest

and that they would be in some sense better off if they should go against their deep-felt preferences. If their rationale for their decision behavior is not uniquely compelling ..., neither, it seems to me, are the counterarguments. Indeed, it seems out of the question summarily to judge their behavior as irrational: I am included among them.

In any case, it follows from the propositions above that for their behavior in the situations in question, the Bayesian or Savage approach gives wrong predictions and, by their lights, bad advice. They act in conflict with the axioms deliberately, without apology, because it seems to them the sensible way to behave. Are they clearly mistaken?

When studying consistency axioms of the types posed by VN-M and Savage, decision theorists ordinarily do not differentiate between private entities and social planners. The presumption is that all decision makers should behave consistently in the same manner. However, some theorists have proposed that social planners should adhere to additional ethical axioms that require them, in some sense, to respect the preferences of their populations and/or behave fairly. Review articles include Fleurbaey (2018) and Mongin and Pivato (2016).

Representation Theorems

I now remark further on representation theorems. The staple formalism of axiomatic decision theory considers a collection of hypothetical choice settings and proposes axioms that mandate specific forms of consistency of behavior across settings. A representation theorem proves that adherence to the axioms is necessary and sufficient for behavior across settings to be representable as solution of some consequentialist optimization problem.

Consider the VN-M and Savage representation theorems. Both begin with a basic axiom stipulating that a decision maker has a complete binary preference ordering over a universe A of actions. They then propose further axioms mandating certain consistency properties for the preference ordering. The theorems prove that adherence to the axioms is necessary and sufficient for representation of behavior when facing any hypothetical choice set $D \subset A$ as maximization of expected utility.

Consequentialist decision theory takes the utility function to be a primitive specified by the decision maker to express what he wants to achieve. In contrast, the representation theorems of axiomatic theory view the utility function as a mathematical construct implied by hypothetical choice behavior. In neither the VN-M nor the Savage theorem does the distribution on the state space have any necessary connection to an objective reality. Considering this distribution, Berger (1985) cautioned that (p. 121)