PROBLEMS OF VALUE MEASUREMENT FOR A THEORY OF INDUCTION AND DECISIONS

C. WEST CHURCHMAN
CASE INSTITUTE OF TECHNOLOGY

1. Introduction

This paper is a philosophical evaluation of current "decision theory" and the pragmatic theory of induction. Its main argument is that there can be no theory without measurement, and that we have no method as yet of performing measurements relative to decisions. Statements made about "rationality" and "optimality" of decisions are premature. In order to perform measurements of values (or preferences) we will have to precommit ourselves to a general decision theory, because measurement is the most intricate and complex of all human decision processes. Indeed, we will be fortunate if we find one decision theory adequate to the task of generating controlled value measurements. Attempts to develop a minimum decision theory on the basis of "reasonably clear" assumptions are criticized on philosophical grounds; such attempts should be regarded as prolegomena to measurement, not as valid statements about rationality. Likewise, the paper criticizes the notion of factual indeterminacy arising out of the necessity to assume some statements about *a priori* probability distributions. In sum, we cannot expect that data about values will ultimately be "inserted" in a decision theory, simply because we require the strongest possible decision theory to generate the data.

2. The problem of pragmatic induction

The concern of this paper is with the pragmatic theory of truth. Roughly—very roughly—speaking, the pragmatic theory of truth states that truth is a property of actions that work out satisfactorily for the person or persons concerned. More specifically, the pragmatic problem of induction is to ground the justification of induction in terms of effectiveness of actions for objectives. Pragmatic "reconstructionism" is the reconstruction of science within a conceptual framework of decisions and their consequences [2].

The term "pragmatism" has a wide variety of meanings. The philosophical attitude of this paper is that philosophy must use present and future experimental sciences as sources of information and guidance in reflection on its problems. Perhaps the term "experimental" more closely reflects the intent of the writer, as a means of differentiating the present approach from that of pragmatists who have found their sources elsewhere than in the sciences.

The problem of the experimental theory of induction seems no different from the current problem of decision theory. Indeed, decision theory and experimental pragmatism are only two examples of a convergence of scientific interest in actions and goals: add to these operations research, social psychology, consumer research, psychoanalysis, law, to name a few. And if the mathematician in decision theory feels some discomfort at being put to bed with so many messy bedfellows, he can perhaps gather some solace from the anxiety neurosis of a psychoanalyst engendered by the sight of a mathematician.

Though many fields are converging on the problem of well-adjusted or rational actions, each still bears some mark of distinction of its own which causes it to view with a critical eye its peculiar colleagues.

That which marks the philosopher—which makes him philosopher and not anything else—is an unending curiosity about evidence. The theories of evidence and philosophy are—to this writer—one and the same. And when I say "unending" I mean this as much as I can—the philosophical mind never stops asking for the evidence (justification) that underlies a judgment or statement.

So, the experimental pragmatist brings to decision theory this fundamental curiosity about the evidence for statements that are made within this theory. He asks in a very general way how any statement in decision theory is possible, that is, can be justified by evidence. He asks also how specific methods of attacking problems in decision theory are justified.

Now in a sense this is not very much to say, since of course all good men and true want to justify what they say, even though they often do not have the leisure time of philosophers to make the necessary inquiries. Is it leisure that differentiates the philosopher from his fellow man? Rather than answer this, let us hasten to add one further qualification: the philosopher brings to bear a history of philosophy, a history of man's attempts to provide a theory of evidence. He thinks that attempts to use a theory of evidence should be judged in terms of their past performance—that is, in terms of the test of past philosophical inquiry. This is the wisdom of the philosopher.

3. Decision theory—is it a theory?

Suppose we try to illustrate this point of view, and, since the matter is of such concern to experimental pragmatism, suppose we turn our attention to current decision theory. Our reflections will keep a more even keel if we center attention not on specific contributors but on methodology—that is, on possible theories of evidence with respect to decision theory.

It will be most convenient for present purposes to define decision theory ostensively by pointing to a sample of references [11], [12], [13], and [14]. Generally, but not precisely, decision theory is concerned with a study and evaluation of actions in terms of objectives. Our chief interest is with that aspect of decision theory which tries to find criteria for optimum selection of actions relative to objectives and environmental circumstances.

We begin with postulational methods, for sake of a beginning point. I assume no defense is needed for the relevance of our considerations, since postulational methods are common enough in current decision theory literature. See [1] and [9], for example. Now our philosophical question is simply concerned with the justification for the method itself. How are we to regard a set of postulates?

Suppose we say that the postulates "carve off' an exceptionally agreed-upon set of statements which then form the basis for a whole sequence of formal inquiries. That is, the initial set of assertions, the postulates, represents statements that any "rational" individual would accept. Indeed, one could argue—as classical rationalism did—that if one does not accept these statements, this is no evidence of their falsity. Rather it is evidence for one of two possibilities: the person failed to understand the meanings of the terms, or else he failed to have understanding at all. Confusion or idiocy is the only consequence to be drawn from disagreement.

Many realize, as the history of philosophy illustrates over and over, that there are subtle traps involved in such a theory of evidence. No mind can foresee all that other minds may logically infer from innocent first-beginnings. The test of what follows, says philosophical history, is far more significant than the test of current agreement. Descartes set down a program of science based on clearly perceived ideas and indubitable assertions [4]. He gambled on a theory of evidence based on current, overt agreement—and lost, if the subsequent history of philosophy means anything.

So one takes the next step. The postulates are agreements, not based on agreement. They are, in fact, an agreement to use terms in a certain manner. They are tautologies. This step follows a Leibnizian program—to ground all science in a fundamental set of tautologies [8]. This seems the most promising program to a rationalist mind, for its success depends on a much simpler test than could be applied to Descartes. The Cartesian indubitables would fall if one derived a single consequence that was doubtful. The Leibnizian tautologies seem to fall only if an inconsistency is deduced. Logic has shown that in many "strong" cases one can apparently guarantee a Leibnizian program—that is, guarantee consistency for all consequences. No one has shown how a Cartesian program can be similarly safeguarded.

One tends therefore to favor the Leibnizian point of view. Subsequent history tells us what is entailed in such a viewpoint. In the first place, we need to abandon much of ordinary language, and express our agreements in some formalized language, using as little English as we can. The point is that the so-called tautologies of English usually turn out to be much more than pure agreements. We'd like to say that anything of the form "A is A," or more generally, "A and B and ... is A" is a tautology. But English is a hectic language and about the only obvious thing that can be said about it is that any obvious thing said within it is almost always wrong. "Women are women, God bless 'em!"—a tautology? "Blackbirds are black birds"—really? "Hydrogen and oxygen—are oxygen?"

We may therefore proceed to phrase the agreements of decision theory in a formal language. If we do, we may in the long run be able to guarantee consistency in one of its several current senses. From this point of view, the "foundations" of decision theory (and I suppose statistics as well) consist of a formalized language adequate to the set of agreements that decision theorists wish to make.

But here arises a serious difficulty. Decision theory is often regarded not as a branch of mathematics or formal science, but as a theory of how people actually ought to behave. If so, agreements reached within a formal language will have little value unless we can also prove that the language is adequate to discuss actual cases. Indeed, pure formalization of decision theory seems to be the very last thing we want to do, not the first. For experimental pragmatism—and I suppose equally for operations research—we need to come out of the formal language again, and reach agreements on how observable behavior relates to the terms of the formal language. We need to know when something is a decision—that is, we need operational specifications for identifying decisions and their properties.

From this point of view, formal decision theory does not represent a "foundation" for a theory of decisions. One can now take a modified viewpoint based on an empirical philosophy: decision theory is or will be a science constructed from certain observables. Postulational methods in decision theory are agreements—or "preconstructions"—on what is to be observed, what the observables are to be called, and how they are to be used to define optimal decision patterns.

It would be better, I think, to go the whole distance at once, and drop the philosophically inadequate term "observable." Suppose we simply say that postulational decision theory is a set of agreements concerning the measurements of behavioral properties like "value," "probability," and the like. Postulational decision theory represents a prolegomenon to measurement, not a theory of decisions as such.

4. Decision rules and indeterminacy

Now any agreements reached prior to measurement are not binding and serve as devices to guide progress, not rules to limit it. We don't know how to measure the values of decisions today, and, until we do, it would be foolish to agree to any commitment once and for all.

This point can be illustrated by an example, before making more specific remarks on the measurement aspects of decision theory. Suppose we turn to the aspect of postulational decision theory that is of chief interest to statisticians, that is, the study of decision rules relative to a set of data. One might be tempted to urge the following kind of agreement. First we note that the problem of decision rules has a determinate solution if there exists one decision rule which is "best" in terms of the postulates of the decision theory and the "given" set of observations. The problem is epistemologically indeterminate if (1) there exists a subset of all decision rules, having more than one member, (2) no member of the subset is "best," by the criteria of the postulates and data, but every member is better than any decision rule lying outside the subset, (3) there exists a set of statements, no two of which can be true together, that are consistent with and independent of the postulates and of such a nature that if any one of the statements is added to the postulates, one decision rule of the subset becomes "best," and (4) there is no better evidence for one member of the statement set than for another in terms of the postulates or data. Such a set of statements might be called "metaphysical." They seem to be outside of "science," and yet in some sense provide answers to questions which science cannot answer. They might also be called "personalistic," since presumably each individual decision-maker would make a personal selection from among the statements in order to select one decision rule. One might dramatize the indeterminacy by calling the set of statements possible "decisions" of nature, and paint into an otherwise rather drab discussion of decisions in the abstract the awful tragedy of man's inability to understand his Mother. Nature "plays" moves which we observe. What is her over-all strategy? If it is random, we get one decision rule. If it is "agin' us," we get another.

Epistemological indeterminacy is not new and it is certainly worth our effort to see what has been its fate in the past. Two centuries ago, Hume set himself the problem of finding the origin of the idea of causal necessity (see section V [6]). Specifically, whence comes our idea that given an event A, event B *necessarily* follows? Hume found that the idea of causal necessity could not be grounded in the data (which at best reveal repetitions of A-then-B, not necessity), or in "reason" (there is nothing in the meaning of A as such that necessitates B). He came to the conclusion that the idea must come from a psychological habit—that is, a personal predilection, not a "scientific" result.

It's worth seeing what happened to Hume's argument. Kant (see I, part II, in second analogy [7]) argued that causal necessity does have a foundation in science. Why? Well, in effect Hume asked the questions in the wrong way. He asked how one could derive causal necessity from data. He should have asked how the data were possible in the first place. If he had asked his question in this manner, thought Kant, he would have seen that the data are possible only because the mind views them within a conceptual framework of causal necessity.

Kant was too rigid in his thinking—but he was not wrong in principle. Those who want to go beyond Kant—not around him—must admit that the really difficult problem of science is the data—or, in general, measurement. How we get the "givens"—this is the real difficulty. How we know we have reliable data—that is, how we arrive at "stability"—is the heart of scientific method. Hume's argument for epistemological indeterminacy was premature. So is the argument for indeterminacy in decision theory. Indeed, in decision theory the Kantian viewpoint is all the more obvious. For here we require certain value measurements, or at least certain preference rankings. What are the decision rules by which we as scientists accept a value ranking as valid? If it should turn out—as it will—that to measure and control data on values we need an elaborate framework of decision rules, it is certainly premature to discuss epistemological indeterminacy within a framework in which the data on values are assumed to be given.

The post-Kantians argued in effect that there is no mind without another mind. We can paraphrase their thesis by saying that there is no decision theory without another decision theory. A decision theory based on given value information and other relevant data requires another decision theory to tell us how the information was obtained. The indeterminacy that arises in the first may be explained and removed in the second. Can it be that the recurrence of epistemological indeterminacy in philosophy and science is always the sign of a need for a more reflective viewpoint—another mind to view what these minds have done?

Reflect then on what it means to measure or rank a man's preferences. Reflect first of all on what it means to tell a man what he ought to do. Kant thought the hypothetical imperative, "If you want A, you ought to do B," was a scientifically provable statement. Indeed, the statement, he thought, was equivalent to "A follows, given B." But his argument missed the real difficulty. When told I ought to do B, I not only need to know that A follows, but I also need to know that I *want* A. How do I know this? Rather, how does the observing scientist know it?

What do we mean when we ask a person to rank preferences? What are we asking him to do? If his answers are inconsistent, does this mean that he changes his mind or that he didn't know what our question meant?

One can adopt the methodological principle that there are certain primitive operations that can be performed which form the foundation of value measurements. These primitives usually consist of "simple" choices which are posed to the person. Thus when a "reasonable" person is asked to rank certain choices, he will always rank in one way: the question is "clear" to him and his responses are not subject to "serious" error. If one could find this kind of elementary data, and if one could construct value measurements on the basis of it, then the methodological circularity of decision theories would disappear into one fundamental and unquestioned decision: the acceptance of the elementary data. But this pious hope is almost sure to be doomed; even the relatively simple experiments on utility show that no such consistency will be obtained [10]. One could also hope that "macroscopic" data on economic behavior can be obtained of sufficient reliability to form a "test" of certain decision theoretic assertions; see [5]. But the difficulty here is to judge whether the data or the theory is wrong if an inconsistency arises.

It seems clear enough, relative to actual preference rankings, that the choices the individual must make are not in any sense primitive—by themselves they mean very little. The big difficulty seems to be that the person must have complete knowledge and awareness. That is, a choice made by a person who is ignorant of the outcome or unaware of one or more of the possibilities does not seem to be a choice one can effectively use to measure his values. But here we are posed with the prevalent circularity of scientific measurement. There seems to be no clear way of ascertaining a person's knowledge without knowing his interests. To me you may seem to be an ignorant fool if you eat at Greasy Joe's when next door Pierre serves a fine meal at the same price. But perhaps good food does not interest you.

We can say in a very general way what a set of value measurements for a person is: it is a set of numbers that predict the probability that the person will make certain behavior choices in certain environments, given perfect awareness and knowledge on his part. Knowledge and awareness behave like the extraneous variables in the definition of physical standards: the standard of length requires specification of other physical variables.

What we cannot say at present is how to generate a controlled set of value measurements. By a controlled set of measurements I mean a set that will check "satisfactorily" against measurements derived by operationally different means. At the present time we are generating information on values in a wide number of different fields: economics, accounting, operations research, education, sociology, anthropology, psychoanalysis, psychology, ethics, industrial relations, consumer research [3]. We have no method of evaluating the results obtained in these different fields. That is, we have no way of adjusting one set of results back to a standard in terms of which we can make comparisons with results obtained by other means. Until we do, we have no decision theory—but only tentative suggestions as to where we might go from here. Until we do, the agreements we reach in postulational decision theory are not foundational. I happen to be one who thinks that epistemological indeterminacy does not stand up under an experimental approach to problems. There are always questions that are unanswered—but there is also always a theory of evidence supplied by another approach to the problem which will show how to generate information concerning the question.

In sum, the experimental approach to the problem of decision-making lacks a theory of data collection—a theory of stability of information. Until such a theory is at least tentatively formed, we lack any foundations for decision theory. Indeed, we might safely admit that the foundations of any field are never to be found—they are as much the ideals of experimental science as are true, exact, and riskless answers to questions.

In acknowledgment, I should like to state that while S. B. Littauer cannot escape the responsibility of having inspired the central theme of this paper—data stability—he can evade responsibility for the actual statements made herein.

REFERENCES

[1] H. Chernoff, "Rational selection of decision functions," *Econometrica*, Vol. 22 (1954), pp. 422–443.
[2] C. W. Churchman, "Notes on a pragmatic theory of induction," *Scientific Monthly*, Vol. 79 (1954), pp. 149–151.
[3] ———, "Measurement of values—a survey" (part of a monograph on *Decision Processes and Value Theory*, to be issued by the *Journal of the Operations Research Society of America*).
[4] R. Descartes, *Discourse on Method*.

[5] M. Friedman and L. J. Savage, "The utility analysis of choices involving risk," *Jour. of Political Economy*, Vol. 56 (1948), pp. 279–304.
[6] D. Hume, *An Enquiry Concerning Human Understanding*.
[7] I. Kant, *Critique of Pure Reason*.
[8] G. W. Leibnitz, *Animadversiones in partem generalem Principiorum Cartesianorum*, Gerhardt IV, Berlin, Weidmann, 1875–1890, pp. 350–354.
[9] J. Marschak, "Rational behavior, uncertain prospects and measurable utility," *Econometrica*, Vol. 18 (1950), pp. 111–141.
[10] F. Mosteller and P. Nogee, "An experimental measure of utility," *Jour. of Political Economy*, Vol. 59 (1951), pp. 371–404.
[11] J. von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior*, Princeton, Princeton University Press, 1944.
[12] L. J. Savage, *The Foundations of Statistics*, New York, John Wiley and Sons, 1954.
[13] R. M. Thrall, C. H. Coombs, and R. L. Davis, *Decision Processes*, New York, John Wiley and Sons, 1954.
[14] A. Wald, *Statistical Decision Functions*, New York, John Wiley and Sons, 1950.