CHAPTER

7

A PROBABILISTIC THEORY OF UTILITY AND ITS RELATIONSHIP TO FECHNERIAN SCALING

R. Duncan Luce

PROFESSOR, DEPARTMENT OF PSYCHOLOGY,
UNIVERSITY OF PENNSYLVANIA

INTRODUCTION [1, 2, 3]

A measurement problem that pervades much, if not all, of social science in this: Given an individual's choices in a series of paired comparisons, where the stimuli vary along some dimension which is of concern to him, how to devise a relatively unique measure of their

[1] The last five sections of this article differ in but trivial ways from Sections A1.2 through A1.6 of Appendix 1 of *Games and Decisions* [4]. I have received Professor Raiffa's permission to reprint them here. The draft of my talk for the Measurement Symposium of the 1956 AAAS Christmas meetings, plus extensive critical comments by Professor Raiffa, served as a basis for the final draft of Appendix 1. When I prepared this paper a few months later, I did not see a sufficiently fresh approach to the material to warrant restating it.

[2] Many conversations and letters underlie the present version of this paper. I am particularly indebted to Professors P. F. Lazarsfeld and H. Raiffa, who devoted many hours to discussing earlier drafts of it, and to Professor John S. Chipman who, although he has never seen this paper in any of its drafts, much influenced it. He, in his capacity as associate editor of *Econometrica,* pointed out difficulties in the original technical exposition of the theory, which has since been revised and published [2].

[3] The preparation of this paper was done at Columbia University, and it was

subjective value, or worth, to him; or to infer a latent scale of values from the set of manifest data of the individual's pairwise choices. Although it is easy to point to the problem and to see that, in all likelihood, it is of widespread scientific interest, it is far from easy to see how to resolve it. In each of the behavioral disciplines, there is a more or less long history of work on the problem in one form or another, and, as one would expect, there are just about as many approaches and vocabularies as there are sciences. I mention this well-known fact because I shall be drawing upon portions of two of these approaches, that of psychology and of economics, without being able to afford the space to give an adequate history of either; and because any such amalgamation of ideas must also entail a mingling of terms which, at times, provides us with an abundance of riches. I must ask you to bear with my arbitrary terminological choices.

The general approach that I shall take to this problem is to define the scale of value implicitly by a set of axioms which assert some of its properties. These are, of course, properties that seem more intuitively acceptable or more basic than an explicit description of the scale might seem. The mathematical argument, which is too lengthy to present here, allows one to transform the implicit definition into an explicit form. The selection of the axioms is naturally a very subtle matter; we shall see this only too clearly later. One tends to depart but little from tradition in these choices, and the traditions I shall cling to can be found in psychology and in a part of economics that borders closely on both psychology and statistics.

THE PSYCHOLOGICAL TRADITION

Suppose that a and b are two alternatives confronting a subject who must select between them according to some relevant dimension. In psychophysics, it might be loudness; in a sociological study, attitude toward race; in much recent decision making work, it has been preference between alternatives. I shall use the latter interpretation. It is a widely accepted assumption of psychology, which for the most part has been idealized out of the corresponding economic models, that, generally, the probability is different from 0 or 1 of a particular person preferring a to b at a given instant, i.e., it is assumed that people do not generally exhibit perfect preference discrimination. I say that this is an assumption of psychology because no way that I know of

supported in part by a grant from the National Science Foundation for the study of "the Mathematics of Imperfect Discrimination."

has been devised to verify or refute it. Furthermore, no way has yet been found to estimate such probabilities without either assuming that each subject in some population has the same probability of preferring a to b or assuming that independent samples can be obtained from a single subject by offering him the same choice at several different and carefully spaced times. As neither of these assumptions seems uniformly valid—indeed, in many cases both seem erroneous—the task of empirical estimation is very delicate; however, I do not want to enter into that issue here. I shall simply suppose that such probabilities exist.

One tradition of scaling in psychology, which, in large measure, stems from psychophysical studies of the last century and which is associated with the name of Fechner, can be stated roughly as follows: One searches for a numerical scale having the property that the probability of preferring one alternative to another, provided the probability is neither 0 nor 1, depends only upon the difference of the alternatives' scale values and not upon their separate values. Thus, if u denotes the scale and $P(a, b)$ the probability that a is preferred to b, then if $P(a, b) \neq 0, 1, P(a, b)$ shall be a function only of $u(a) - u(b)$. Because one can assume that stimuli from any one psychophysical dimension form a continuum, it turns out that this condition specifies a psychophysical scale uniquely, except for its zero and unit; but it does not do so when the alternatives are discrete, as in a preference experiment. Then there are many inherently different scales compatible with the Fechnerian condition. To have a reasonably unique scale, it appears that both the probabilities P must satisfy some restrictive conditions and a set of alternatives must possess some "mathematical structure." I do not want to explain what I mean by this, except to point out that one abstracts a psychophysical continuum by the real number system which has a very rich mathematical structure. A finite set of "unconnected" stimuli does not. There also seems to be an exchange relation between the conditions on P and the structure on the alternatives: The more structure possessed by the set of alternatives, the less stringent need be the conditions on P; and the less structure on the set, the stronger must be the conditions on P for a relatively unique scale to exist. In the psychophysical case, the conditions needed on P are relatively weak. On the other hand, in most social psychological scaling, only a finite number of alternatives are assumed, such as a set of political candidates, and little or no mathematical structure can be assumed. Thus, to have a single acceptable scale, quite strong assumptions must be made about the probabilities of preference. These

include the familiar, but nonetheless controversial and largely untested, assumptions that certain variables are normally distributed and that the error terms are statistically independent of practically everything in sight. Such assumptions are most familiar from factor analysis and the Thurstone school of scaling.

When the problem is phrased in this way, the question immediately arises as to whether we can arrange a choice situation which renders preference scaling similar to psychophysical scaling in the sense that the underlying set of alternatives has a lot of mathematical structure. Certainly, one cannot just copy the psychophysical assumption of a continuum but possibly the spirit of that model can be reproduced.

THE MODERN DECISION MAKING TRADITION

The idea now current for model building comes from a part of economics and statistics known as the theory of decision making under risk. A central observation of these theorists, in particular von Neumann and Morgenstern [7] and Ramsey [5] (who, although less influential than von Neumann and Morgenstern, predated many of their utility ideas by several decades), is that it is rather more rare than common to make choices between pure prospects, as has been the case in most preference experiments. Generally, one is confronted with prospects that are built up of several outcomes which are conditional upon the occurrence or nonoccurrence of certain future events. Let me be more specific. Your wife suggests that you purchase tickets to a certain play for the evening of January 15. Is it fair to say that you must simply choose between spending $8 and seeing the play on that date versus spending nothing and, say, sitting home watching whatever television fare is available? Hardly. If nothing else, there is a certain possibility that you cannot go to the play that night even if you have the tickets—a pressing professional activity may arise or you may be ill. Let me call all such eventualities the event α. Thus, the choice will be among:

Spending $8 and not seeing the play if α occurs; or
Spending $8 and seeing the play if α does not occur.

versus

Spending nothing and not seeing television if α occurs; or
Spending nothing and seeing television if α does not occur.

This seems a little more realistic, and certainly the choice is more difficult now because it depends both upon how likely you think it is that α will occur and on how much you want to see the play.

Abstractly, such a mixed prospect will be symbolized as follows: if a and b are any two prospects and α is an event, the mixed prospect "a if α occurs or b if it does not" is denoted by $a\alpha b$.

There is, of course, no reason why the component prospects of a mixed prospect cannot themselves also be mixed prospects. To return to our theater example, suppose that one has taken the theater option and that α does not occur. Thus, one goes to the theater; but this is a mixed pleasure, depending upon the weather. One's enjoyment of the play is liable to be somewhat dampened if there is a sticky New York snowstorm to battle. So, in our symbolism, if b is actually the mixed prospect $c\beta d$, then the over-all mixed prospect is $a\alpha(c\beta d)$.

It is clear that, even when both the set of pure prospects and the set of events are finite, the set of mixed prospects that can be generated recursively in this manner is infinite. It is only more so when the set of events is infinite, as we shall assume. The possibility of this set having a lot of reasonable mathematical structure is clear, and we shall impose some structure later in this paper.

The price that we pay—and from many points of view it is really a gain—for making the choice model a little more realistic is that we have two, not one, scaling tasks. In addition to scaling subjective values, we are pretty well forced to obtain a scale of subjective probability for the events. Depending upon one's interest, attention is generally focused on just one of the two scales and the other appears only as a necessary technical device.

One thing is immediately clear: the two scales are thoroughly interlocked since subjective values must be attached to mixed prospects, and these depend upon the events. If we let u denote the subjective value scale and ϕ the subjective probability scale, then one of the simplest ways they can be interlaced would be for the subjective value of a mixed prospect to be given by the expectation of the subjective value of its components, i.e., the subjective value of each component prospect is weighted according to its subjective probability of occurring. In the language of utility theory, this is described by saying: the individual behaves as if he were maximizing expected utility. A good deal of effort has gone into determining conditions under which this expected value property is met, for, without it, any mathematical model based upon a subjective value scale becomes dreadfully complicated. In this paper, however, we are concerned with a slightly different issue; we suppose that our scales do, in fact, possess this very desirable property and then inquire into its consequences when it is coupled with certain other intuitively plausible axioms.

PREFERENCE DISCRIMINATION AND INDUCED PREFERENCE

We shall need some notation. First, the set of pure alternatives, finite or infinite, will be denoted by A and the set (actually, Boolean algebra) of chance events by E. If α is an element of E, then $\bar{\alpha}$ will denote the complement of α. The set of mixed prospects, generated in the way just described, will be denoted by G.

Axiom 1. For every a in G, $a\alpha a = a$.

In words, the mixed prospect in which a is the outcome whether or not α occurs is not distinguished as different from a itself. It is hard to quarrel with this, although, when combined with Axiom 11, it implies that the subjective probabilities of an event and of its complement sum to 1, which Edwards [1] has questioned.

If a and b are two mixed prospects from G, we suppose that there exists an objective probability $P(a, b)$ that the given individual will prefer a to b. As I indicated earlier, it is not easy to see how to estimate such probabilities in practice, but we need not concern ourselves about that when describing the model.

Although it is true that imperfect preference discrimination has been introduced in part to avoid the strong transitivity requirements of the von Neumann and Morgenstern theory, it would be folly to ignore the empirical evidence suggesting that preferences are approximately transitive. It is easy to go astray at this point by assuming certain inequalities among the three quantities $P(a, b)$, $P(b, c)$, and $P(a, c)$; apparently this is not strong enough. Our tack is a bit different. Observe that, in an induced sense, a is "preferred or indifferent to" b if for every c in G both

$$P(a, c) \geq P(b, c) \qquad \text{and} \qquad P(c, b) \geq P(c, a).$$

Whenever these two sets of inequalities hold, we shall write $a \gtrsim b$. It is easy to see that $\gtrsim$ must always be transitive, but that, in general, there will be alternatives which are not comparable according to $\gtrsim$. A basic restriction we shall make about preference discrimination is that such comparisons are always possible, i.e.:

Axiom 2. For every a and b in G, either $a \gtrsim b$ or $b \gtrsim a$.

This is a strong assumption, but I do not believe it to be nearly so strong as the corresponding ones in the traditional nonprobabilistic utility models. There, comparability is operationally forced by the demand that the individual make a choice, but transitivity is in doubt. Here, transitivity is certain and comparability is in doubt. Although it is plausible that Axiom 2 is met in some empirical contexts, the following example, due to Howard Raiffa, strongly suggests that this is not always the case. Suppose that a and b are two alternatives of roughly comparable value to some person, e.g., trips from New York City to Paris and to Rome. Let c be alternative a plus \$20 and d be alternative b plus \$20. Clearly, in general,

$$P(a, c) = 0 \qquad \text{and} \qquad P(b, d) = 0.$$

It also seems perfectly plausible that, for some people,

$$P(b, c) > 0 \qquad \text{and} \qquad P(a, d) > 0,$$

in which event a and b are not comparable, and so Axiom 2 is violated. In one respect this example is special: c differs from a, and d from b, by the addition of an extra commodity which is always desirable; therefore, we may expect perfect discrimination within each of these two pairs. As we shall see, there are theoretical reasons for believing that the occurrence of perfect preference discrimination may require a somewhat different model than when it never occurs.

Let us say that a and b are indifferent in the induced sense, and write $a \sim b$, whenever both $a \gtrsim b$ and $b \gtrsim a$. We next argue that certain two-stage gambles should be indifferent.

Consider the mixed prospect $(a\alpha b)\beta c$, where a, b, and c are pure alternatives. If one analyzes what this means, one sees that outcome a results if both α and β occur, i.e., if the event $\alpha \cap \beta$ occurs; b results if both $\bar{\alpha}$ and β occur, i.e., if $\bar{\alpha} \cap \beta$ occurs; and c results if $\bar{\beta}$ occurs. A similar analysis of the prospect $a(\alpha \cap \beta)(b\beta c)$ shows that a, b, and c occur under exactly the same conditions. Thus, there is no difference between the two mixed prospects and it is reasonable to argue that a person should be indifferent between them. We shall demand that this holds not strictly but only in the weaker sense of induced preference.

Axiom 3. If a, b, and c are in A and α and β are in E, then

$$(a\alpha b)\beta c \sim a(\alpha \cap \beta)(b\beta c).$$

Actually, the results that we shall state depend only upon the weaker assumption

$$(a\alpha b)\beta b \sim a(\alpha \cap \beta)b,$$

which follows from Axiom 3 by setting $c = b$ and then using Axiom 1.

LIKELIHOOD DISCRIMINATION AND QUALITATIVE PROBABILITY

Suppose that our subject must decide between the two prospects $a\alpha b$ and $a\beta b$. He can simplify his choice by asking himself which alternative, a or b, he prefers, and which event, α or β, he considers more likely to occur. Of the four combinations, two should lead to preference for $a\alpha b$ over $a\beta b$:

1. a is preferred to b, and α is deemed more likely to occur than β.
2. b is preferred to a, and β is deemed more likely to occur than α.

By assumption, the probability that he will prefer a to b is $P(a, b)$. If we suppose that his discrimination as to the likelihood of events is statistically independent of his preference discriminations, and that it is governed by a probability $Q(\alpha, \beta)$, then the probability that he will both prefer a to b and deem α more likely to occur than β is $P(a, b)\, Q(\alpha, \beta)$. Similarly, the probability that he will both prefer b to a and deem β more likely to occur than α is $P(b, a)\, Q(\beta, \alpha)$. Since these two cases are exclusive of each other, the sum of the two numbers should give the probability that he will prefer $a\alpha b$ to $a\beta b$.

The important assumption made in this argument is that the two discrimination processes are statistically independent. This seems reasonable when and only when the subject believes the two prospects a and b to be "independent" of the events α and β, for, if alternative a depends on α and he believes α is likely to occur, then he is really forced to compare the outcome of a which arises when α occurs with $a\beta b$, in which case his preference between $a\alpha b$ and $a\beta b$ may be different from what it would be if a were independent of α. There is at least one case when it is plausible that the subject should deem a and b to be independent of α and β, namely, when a and b are pure alternatives having nothing to do with chance events. We shall assume that our conclusion holds in that case.

Axiom 4. There is a probability $Q(\alpha, \beta)$ for every α and β in E such that, if a and b are in A,

$$P(a\alpha b, a\beta b) = P(a, b)\, Q(\alpha, \beta) + P(b, a)\, Q(\beta, \alpha).$$

There is, as yet, no direct evidence as to whether these two discriminations actually are statistically independent. Conceptually, we clearly separate preferences among alternatives from likelihood among events, and it seems reasonable that people attempt to deal with these as distinct, independent dimensions. On the other hand, casual observation indicates that people do play long shots, and such behavior appears to violate the axiom. At the least, the axiom seems sufficiently compelling as a dictum of sensible behavior to warrant its investigation, and it can be looked on as a generalization of related, but nonprobabilistic, assumptions found in other work, e.g., in Ramsey [5] and in Savage [6].

Our next axiom is comparatively innocent. Let me state it first and then discuss its import.

Axiom 5. For every a and b in G,

$$P(a, b) \geq 0 \quad \text{and} \quad P(a, b) + P(b, a) = 1.$$

For every α and β in E,

$$Q(\alpha, \beta) \geq 0 \quad \text{and} \quad Q(\alpha, \beta) + Q(\beta, \alpha) = 1.$$

There exist at least two alternatives a^* and b^* in A such that $P(a^*, b^*) > \frac{1}{2}$.

First, we have supposed that the P's and Q's are actually probabilities in the sense that they lie between 0 and 1 inclusive, and we have supposed that the subject is forced to make choices between alternatives and between events. That is, he cannot report that he is indifferent between a and b. Experimentally, this is known as the "forced-choice" technique, and it is in standard use. It may be worth mentioning that, if one allows indifference reports in the sense of only demanding $P(a, b) + P(b, a) \leq 1$, then the mathematics leads to two quite distinct cases —one we shall describe here, and another one somewhat like it but apparently less realistic. The final condition simply demands that the situation be nontrivial in the sense that not all pure alternatives are equally confused with respect to preference.

From Axioms 4 and 5, it is trivial to show that

$$Q(\alpha, \beta) = \frac{P(a\alpha b, a\beta b) + P(a, b) - 1}{2P(a, b) - 1},$$

for every a and b in A such that $P(a, b) \neq \frac{1}{2}$ [by Axiom 5, at least one such pair (a^*, b^*) exists]. This expression is useful because it permits one to determine whether a given set of preference data do satisfy the independence assumption and, if they do, to estimate $Q(\alpha, \beta)$.

In complete analogy to "induced preference," we may define a relation on the set of events E. We write $\alpha \gtrsim \beta$ if

$$Q(\alpha, \delta) \geq Q(\beta, \delta) \quad \text{and} \quad Q(\delta, \alpha) \geq Q(\delta, \beta)$$

for every δ in E. We shall refer to this as the "qualitative probability" (induced by Q) on E. One might expect us now to impose a comparability axiom like Axiom 2 on qualitative probability, but this is unnecessary as it is a consequence of our other axioms. Rather, an entirely different assumption, peculiar to the notion of probability, is required. We shall suppose that the subject is certain that the universal event e of the Boolean algebra E will occur. For the moment, we will demand that no event have a qualitative probability in excess of e or less than its complement.

Axiom 6. If e is the universal event in E, then

$$e \gtrsim \alpha \gtrsim \bar{e} \quad \text{for every } \alpha \text{ in } E.$$

THE UTILITY AND SUBJECTIVE PROBABILITY FUNCTIONS

So far, our technique of study has been similar to that normally employed in utility theory, but now we depart from that tradition by assuming that utility and subjective probability functions exist having, among others, properties like those that are traditionally established. Of course, neither of these two functions, however we choose them, can be a complete representation of the assumed data in the same sense that traditional utility functions are. We no longer have a simple transitive relation to be represented numerically but rather a set of probabilities. The role of what we shall continue to call the utility and subjective probability functions will be a partial and—as we shall see—comparatively simple representation of the probabilities. It is analogous to using a statistic such as the mean or standard deviation to give a partial description of a probability distribution.

We shall suppose that there exists at least one real-valued function u on G, called the utility function, and at least one real-valued function ϕ on E, called the subjective probability function, and that the following axioms are met.

Axiom 7. u preserves the induced preference relation on G, and ϕ preserves the qualitative probability on E, i.e.,

$$u(a) \geq u(b) \text{ if and only if } a \gtrsim b, \quad \text{for } a \text{ and } b \text{ in } G$$

and

$$\phi(\alpha) \geq \phi(\beta) \text{ if and only if } \alpha \gtrsim \beta \quad \text{for } \alpha \text{ and } \beta \text{ in } E.$$

As this sort of condition is very familiar in all of utility theory, I need not comment on it.

Axiom 8. $\phi(e) = 1$ and $\phi(\bar{e}) = 0$.

This prescribes more clearly the role of the universal event e. It is an event which is subjectively certain to occur, and its complement is subjectively certain not to occur.

Given a subjective probability function ϕ, we may follow the usual terminology for objective probabilities and say that two events α and β are *(subjectively) independent* if and only if $\phi(\alpha \cap \beta) = \phi(\alpha)\,\phi(\beta)$. It is clear that we cannot ascertain which events are independent until we know the subjective probability function ϕ, and thus it would appear as though we were rapidly getting ourselves into a circle. However, it turns out that all of our final conclusions can be stated without reference to independent events provided only that Axiom 4 can be extended in a certain way and that there are enough independent events—so many that no exhaustive check would be possible anyhow. These conditions will be formulated as Axioms 9 and 10.

Earlier, when we introduced Axiom 4, describing the statistical independence of the two discrimination processes, we held that it should be met whenever the two prospects a and b are "independent" of the events α and β, without, however, specifying what we might mean by this except that it should hold for all pure alternatives. We now extend Axiom 4 as follows:

Axiom 9. If a and b are in A, and α and β are events which are subjectively independent of event γ, then

$$P[(a\gamma b)\alpha b, (a\gamma b)\beta b] = P(a\gamma b, b)\,Q(\alpha, \beta) + P(b, a\gamma b)\,Q(\beta, \alpha).$$

Axiom 10. The subjective probability function ϕ shall have the property that, for all numbers x, y, and z, where $0 \le x, y, z \le 1$, there are events α, β, and γ in E such that:

(a) $\phi(\alpha) = x$, $\phi(\beta) = y$, and $\phi(\gamma) = z$.

(b) α and β are both subjectively independent of γ.

This axiom postulates a very dense set of independent events, so dense that every conceivable subjective probability is exhibited at least twice. Put another way, we are making a continuum assumption about the individual being described via the axioms. Although this type of assumption is not often made so explicit, it is nevertheless implicit whenever the assumption is made that any objective probability can be realized.

Axiom 11. These two subjective scales satisfy the expected utility hypothesis in the sense that, for a and b in A, and α in E,

$$u(a\alpha b) = \phi(\alpha)\,u(a) + \phi(\bar{\alpha})\,u(b).$$

This, except for the restriction to pure alternatives, is a familiar feature of utility theory. Although no restrictions are usually stated when the expected utility hypothesis is made, it is always tacitly assumed that it only holds for mixed prospects whose component events are independent of the event α of the hypothesis. In utility theory, of course, independence is meant in the usual objective sense. For our purposes, it is sufficient to assume the hypothesis only for pure alternatives which are trivially independent of events.

CONCLUSIONS ABOUT THE SUBJECTIVE SCALES

On the basis of these eleven axioms, the following conclusions can be established as to the form of the discrimination functions and the subjective scales. First of all, Q must depend only upon the difference of the subjective probabilities of its two events. Put more formally, there exists a real-valued function Q^* of one real variable such that

$$Q(\alpha, \beta) = Q^*[\phi(\alpha) - \phi(\beta)].$$

This result is interesting because of its connection with the old psychological problem mentioned on p. 146.

Actually, we can give a much more explicit result than that ϕ is a Fechnerian sensation scale: we can describe the mathematical form of Q. There are three cases. In the first, there is a positive constant ϵ and Q is of the form

$$Q(\alpha, \beta) = \begin{cases} \frac{1}{2} + \frac{1}{2}[\phi(\alpha) - \phi(\beta)]^{\epsilon} & \text{if } \alpha \succ \beta \\ \frac{1}{2} & \text{if } \alpha \sim \beta \\ \frac{1}{2} - \frac{1}{2}[\phi(\beta) - \phi(\alpha)]^{\epsilon} & \text{if } \beta \succ \alpha. \end{cases}$$

The second is the discontinuous function

$$Q(\alpha, \beta) = \begin{cases} 1 & \text{if } \alpha \succ \beta \\ \frac{1}{2} & \text{if } \alpha \sim \beta \\ 0 & \text{if } \beta \succ \alpha, \end{cases}$$

which results from the first case by taking the limit as ϵ approaches 0. This represents perfect likelihood discrimination. The third is the function obtained by taking the limit as ϵ approaches infinity, and it represents almost total lack of discrimination.

It is easy to see that, in the first case, but not in the other two, one can express ϕ in terms of Q, namely, as

$$\phi(\alpha) = [Q(\alpha, \bar{e}) - Q(\bar{e}, \alpha)]^{1/\epsilon}$$

or, more usefully, as

$$\phi(\alpha) = \begin{cases} \frac{1}{2} + \frac{1}{2}[2Q(\alpha, \bar{\alpha}) - 1]^{1/\epsilon} & \text{if } Q(\alpha, \bar{\alpha}) > \frac{1}{2} \\ \frac{1}{2} & \text{if } Q(\alpha, \bar{\alpha}) = \frac{1}{2} \\ \frac{1}{2} - \frac{1}{2}[1 - 2Q(\alpha, \bar{\alpha})]^{1/\epsilon} & \text{if } Q(\alpha, \bar{\alpha}) < \frac{1}{2}. \end{cases}$$

Similar results hold for u and P over the set A of pure alternatives. First, P can be shown to be a function only of $u(a) - u(b)$ for a and b in A. Second, assuming a Q of the first type above and letting ϵ be the constant determined there, then

$$P(a, b) = \begin{cases} \frac{1}{2} + \frac{1}{2}[P(a^*, b^*) - P(b^*, a^*)][u(a) - u(b)]^{\epsilon} & \text{if } a > b \\ \frac{1}{2} & \text{if } a \sim b \\ \frac{1}{2} - \frac{1}{2}[P(a^*, b^*) - P(b^*, a^*)][u(b) - u(a)]^{\epsilon} & \text{if } b > a, \end{cases}$$

and

$$u(a) = \begin{cases} \left[\dfrac{P(a, b^*) - P(b^*, a)}{P(a^*, b^*) - P(b^*, a^*)}\right]^{1/\epsilon} & \text{if } a > b^* \\ 1 - \left[\dfrac{P(a^*, a) - P(a, a^*)}{P(a^*, b^*) - P(b^*, a^*)}\right]^{1/\epsilon} & \text{if } b^* > a, \end{cases}$$

where a^* and b^* are mentioned in Axiom 5. Any positive linear transformation of u is equally acceptable.

Thus, we have the following situation. If the axioms are accepted and if it is assumed that discrimination of events is neither perfect nor totally absent, then the mathematical form of the model is completely specified except for a single parameter ϵ which appears to reflect the individual's sensitivity of discrimination; and the two subjective scales can be inferred from the empirical estimates of the probabilities P. The subjective probability scale is unique, and the utility scale is unique except for its zero and unit. There is only one trouble with all of this: it is extremely doubtful that people satisfy all the axioms.

An example, again due to Howard Raiffa, and a theorem will formulate our doubts. Although the mathematical argument used to establish our results rests heavily on steps involving independent events, the final results can be shown to hold for events whether or not they are independent, so we need not worry about independence in a counterintuitive example. Consider the two chance events: rain on Wall Street at time t, and rain on both Wall Street and 34th Street at time t. Since the locations are not widely separated, both being in New York City, it is highly likely that if it rains on Wall Street if will also rain on 34th Street, so the subjective probability of rain on Wall Street alone will only be slightly larger than rain at both places. Yet, if one is asked which is more likely, it seems silly ever to say the latter. If so, we have $\phi(\alpha)$ and $\phi(\beta)$ very close and $Q(\alpha, \beta) = 1$. If people actually behave in this way when making choices, then at least one of our axioms must be false.

AN IMPOSSIBILITY THEOREM

Casual observation suggests that there are many situations, e.g., those involving gambles of money, in which these conditions can be satisfied. First, there are at least three prospects a, b, and c which are perfectly discriminated with respect to preference, i.e., $P(a, b) = P(b, c) = P(a, c) = 1$. This will hold, we are sure, when all other things are equal and $a = \$10$, $b = \$5$, and $c = \$1$. Second, there are at least two events, α and β, which are neither perfectly discriminated nor equally confused, i.e., such that $Q(\alpha, \beta) \neq 0$, ½ or 1. The impossibility theorem asserts that these two assumptions are inconsistent with the eleven axioms we have previously stated.

This result seems disturbing, for most of the assumptions on which it is based have, by now, acquired a considerable respectability. Yet, clearly, they cannot all be satisfied. The task of reappraising them is quite delicate, for there are numerous reasons for supposing that they are not terribly far from the truth. For example, the derived form of the discrimination function for events is sufficiently similar to much discrimination data to suggest that we are not completely afield.

It would appear that six of our assumptions are subject to the greatest doubt. Of these, three (Axiom 2, requiring that every pair of mixed prospects be comparable by the induced preference relation; Axiom 3, requiring that two prospects which decompose in the same way be indifferent in the induced sense; and Axiom 4, requiring that the two discrimination processes be statistically independent for pure alternatives) are subject to direct experimental study. The other three (Axiom 9, requiring that Axiom 4 hold for certain mixed prospects involving subjectively independent events; Axiom 10, requiring that certain triples of independent events be extremely dense; and Axiom 11, requiring that the expected utility hypothesis be true for pure alternatives) are impossible to study directly. Because of this, one can expect that most attempts to get out of the bind will be concentrated on the second three.

Since much of decision theory is so dependent upon the expected-utility hypothesis, special attention will undoubtedly be given to

Axioms 9 and 10. There is the intriguing possibility that these subjective scales are discrete rather than continuous, as has generally been assumed, which would make them more in accord with the way people seem to classify, say, events: impossible, not very likely, etc. In that case, Axiom 10 might be abandoned. On the other hand, Axiom 9 when coupled with our definition of independence may be the source of difficulty. As the axiom seems reasonable for one's intuitive idea of subjectively independent events, it may be the definition that should be altered.

As it stands, two conceptual features of this theory are of interest. First, by making the assumption that the two discrimination processes are statistically independent, it has been possible to deal simultaneously with both subjective value (utility) and subjective probability. Second, by using axioms which are closely related to those of traditional utility theory and the independence assumption (Axiom 4), it has been possible to demonstrate that both utility and subjective probability form sensation scales in the Fechnerian sense. In psychophysics it has been argued, though never fully accepted, that subjective experience must be represented by such scales; however, the defining condition is neither simple nor has it been derived from other assumptions. The traditional practice has been to postulate this condition as an *a priori* definition of subjective sensation, and, of course, many have objected that it is much too sophisticated to be accepted as a basic axiom. Whether a model that parallels this one and that arrives at sensation scales as a consequence, not as a postulate, can be developed for psychophysical problems is not known.[4]

For a fuller statement of this theory and for proofs of the assertions, see [2].

REFERENCES

1. Edwards, Ward, "The Theory of Decision Making," *Psychological Bulletin*, **51**, 1954, pp. 380–417.
2. Luce, R. D., "A Probabilistic Theory of Utility," *Econometrica*, **26**, 1958, pp. 193–224.
3. ——, *Individual Choice Behavior*, John Wiley & Sons, New York, 1959.

[4] Since this was written, I have developed a model, based upon a single plausible axiom (aside from those of ordinary probability theory), which is a probabilistic generalization of transitivity, that establishes the existence of sensation scales for arbitrary sets of alternatives whenever discrimination is not perfect; see [3].

4. Luce, R. D., and H. Raiffa, *Games and Decisions*, John Wiley & Sons, New York, 1957.
5. Ramsey, F. P., *The Foundations of Mathematics*, Harcourt, Brace and Co., New York, 1931.
6. Savage, L. J., *Foundations of Statistics*, John Wiley & Sons, New York, 1954.
7. von Neumann, John, and Oskar Morgenstern, *Theory of Games and Economic Behavior*, 2nd ed., Princeton University Press, Princeton, 1947.