13 R. D. Luce

Psychological Studies of Risky Decision Making

R. D. Luce, 'Psychological studies of risky decision making', in Strother, G. B. (ed.), *Social science approaches to business behavior*, R. D. Irwin, Inc., 1962, pp. 141–61, Tavistock Publications Ltd, 1962.

Introduction

At the outset, let me acknowledge my ignorance of the world of business. Professionally, I have had little to do with business or with its study, and so I write with no more than an educated layman's understanding of which intellectual endeavors are relevant to business problems. Because of this, it would be foolish and presumptuous to claim that the studies we shall examine here are directly applicable to them. I suspect that my general problem – the choices that people make among alternatives whose consequences are risky – is of quite general interest; however, specific theories and experiments may very well be another matter. But even if what I describe is not of direct relevance to business, hopefully it will be at least intriguing. I plan simply to sketch some of the ideas and findings that theoretical and experimental psychologists have made in this stimulating, confusing, and conflictful area of research.

The problem is, let me repeat, to describe the formal structure of a person's choices when the consequences following from a choice are risky. That is, we are searching for a satisfactory abstract – mathematical – description of a person choosing among two or more alternatives, each of which is, in more familiar terms, a lottery or gamble. Although the specific outcome to him depends to some extent on his choice, it also depends on chance events that are quite beyond his control.

Some early mathematicians were interested in choices among money gambles. It was Daniel Bernoulli, for example, who so neatly established that people do not choose on the basis of expected money returns alone. Out of these early studies grew, on the one hand, the mathematical notions of probability and, on the other hand, the economic notions of cardinal utility for nonnumerical commodities. Although separately evolved, these two groups of ideas have always had a close affinity, and less than 15 years ago they were reunited in the work of von Neumann and

Morgenstern, Savage, and Wald, to name only the most prominent contributors.

Even though history seems to suggest that our problem lies in the province of economics and statistics, it is evident that in part it also belongs to psychology. Indeed, I think the claim can be defended that the study of the choices which individuals and institutions should make to attain certain ends forms a part of – and quite possibly is equivalent to – statistics; that the study of the choices which individuals do in fact make forms a part of – but definitely is not equivalent to – psychology; and that the study of the choices which aggregates of people and institutions make forms a part of economics and sociology. That much theoretical work about the choices that individuals make when the alternatives are risky appears in the economic literature is, it seems to me, no more than a historical accident. Some approaches to economics seem to demand psychological underpinnings, and economists have never been loath to create their own psychologies when none could be borrowed.

Accepting that the descriptive half of the problem of individuals making choices belongs to psychology, what then have we done to solve it? Speaking broadly, there have been two mathematical tacks. One group of psychologists has developed models which, although different in detail, are similar in spirit to those proposed by the economists. These theories have been designed to account for the results of preference studies. For the most part, they have been applied only to choices among what we may call 'sure' alternatives, i.e., to those for which the outcome depends only on the subject's choice and not on any other events. Usually, there is no reason why preference theories cannot be applied to risky alternatives, but, until recently, such theories have failed to include in their structure the added richness provided by dealing with risky alternatives.

A second group of psychologists concerned with the processes of human learning have, almost inadvertently, studied choices among risky alternatives. Their end was not to understand the interrelations among the choices that subjects make in different situations of risk but rather to use these situations as devices to generate particular temporal patterns of responses. The subjects in these experiments initially know little or nothing about the events controlling the outcomes that result from their choices. But the same choices are offered over and over, and, in the course of time, the subjects acquire considerable statistical information about the events. Ultimately, their responses settle down to some stable pattern of behavior, much like that discussed by the utility and preference theorists. During the past decade, a sizable mathematical literature has evolved in the attempt to describe this and related sorts of learning.

So, in total, we have three general classes of theories which attempt to describe an individual's choices among risky alternatives – the economists' and statisticians' utility theories and the psychologists' preference and learning theories. The distinction between utility and preference theory is often somewhat fuzzy, for to a considerable extent it rests upon who did the work. I shall sharpen it for present purposes by restricting attention to those utility theories centered around the expected utility hypothesis and to those preference theories explicitly assuming stochastic behavior.

Subjective Expected Utility Theories

A characteristic of all utility theories, expected or otherwise, is that one can assign numerical quantities – utilities – to alternatives in such a way that alternative a is chosen from a set T of alternatives if and only if the utility of a is larger than that of any other alternative in T. When such an assignment is possible, we say that the person behaves optimally relative to his utility scale, that he maximizes utility.

When we study risky alternatives, we have a second guiding idea which is called the 'subjective expected utility hypothesis'. It holds that, in addition to utility assignments to all alternatives – risky as well as sure – one can also assign numbers to events. The numbers are interpreted as the subject's evaluation of the likelihood of the event's occurring; they are called 'subjective probabilities'. Like ordinary objective ones, they lie between 0 and 1. These two numerical scales are interlocked in the following way: the utility of a risky alternative is the sum of the utilities of its component outcomes, each weighted according to the subjective probability of its occurring. For example, if $a\alpha b$ denotes the risky alternative in which a is the outcome when the event α occurs and b when it fails to occur, then the subjective expected utility hypothesis asserts in this simple case that

$$u(a\alpha b) = u(a)\psi(\alpha) + u(b)[1 - \psi(\alpha)],$$

where u denotes the utility scale and ψ the subjective probability scale.

The primary theoretical problem has been to justify this representation. Typically, a series of empirically testable assumptions

is stated which relate choices one to another, and from these it is shown by mathematical argument that the expected utility representation follows. The assumptions made are always plausible as canons of rational behavior. For example, they usually include one something like this: when a is preferred to b and when event α is more likely than event β, then $a\alpha b$ is preferred to $a\beta b$. Usually it is easy to persuade oneself and one's more rational friends that one should abide by these axioms, but rather fewer of us are willing to claim that they actually describe our behavior.

There are three major theories of this type. In 1947 von Neumann and Morgenstern (35) stated the first such system of axioms, which served as a normative underpinning for their theory of games. As a normative theory it may well be satisfactory, but as description it was badly marred by the supposition that objective probability can substitute for subjective probability. In 1954, Savage (31) rectified this weakness by incorporating into a single, elegant axiom system both their utility ideas and de Finetti's (19, 20) subjective probability notions. Although this is probably the definitive axiomatization, other more special ones are often useful. For example, Davidson and Suppes (9), working out an idea suggested in 1931 by the philosopher Ramsey (30), developed a utility theory that required only data from choices among gambles generated by but one event – one having a subjective probability of $\frac{1}{2}$. This considerably simplifies certain experimental problems.

Only two major empirical studies have been carried out, one to test the von Neumann–Morgenstern theory and the other the Davidson–Suppes theory. Without directly examining the correctness of the von Neumann–Morgenstern axiom system, Mosteller and Nogee (29) applied the expected utility equation, using objective probabilities, to certain observed choices and attempted to construct individual utility functions for money. They found that such functions could, in fact, be calculated and that they were reasonably smooth, increasing functions of money. With these functions in hand and again assuming the expected utility hypothesis, they then predicted the subject's choices between risky alternatives different from those used to construct the functions. These predictions, although somewhat more accurate than those calculated from expected money returns, were far from perfectly accurate. By strict standards, the model failed, but the failure was ambiguous. Did it stem from the expected utility hypothesis itself or from the equating of subjective to objective probability? No one could be sure.

The second study, performed by Davidson, Suppes, and Siegel (10), required only an event having a subjective probability of $\frac{1}{2}$. With that in hand, the utility function was generated by carrying out an extremely careful exploration of the choices made when the money outcomes were changed by 1-cent amounts. Although a 1-cent change is small in absolute value, it was actually a sizable percentage change of the sums involved: from 3 to 25 per cent. The most important consequence of this procedure was that it prevented an exact determination of the utility function; only upper and lower bounds could be found. Once the bounds were determined, the experimenters tested the theory in much the same way as Mosteller and Nogee by predicting choices in situations different from those used to generate the utility functions. Because of the indeterminacy of the utility functions, however, a number of the more sensitive predictions could not be made. Of those that were unambiguous, an extremely high proportion were correct.

It is generally agreed that this experiment provides the strongest support for the subjective expected utility hypothesis as a description of behavior, but, even so, its success is equivocal. One doubts that the hypothesis has been adequately taxed both because of the indeterminacy of the utility function and because the theory is restricted to only one chance event.

Several other experimental studies have been interpreted as unfavorable to the expected utility hypothesis (Coombs [6]; Coombs and Pruitt [7]; and Edwards [11–14, 16]). Some of these have been concerned with preferences among gambles having different money variances but the same expected values, and others with preferences among the probabilities of the events themselves. For example, Edwards' subjects chose between bets having the same expected money values. When that expectation was positive, they consistently preferred those bets with a 50:50 chance of winning and avoided those with a 75:25 chance; when the expected value was negative, these preferences were reversed.

It is not easy to know what to make of such studies. For one thing, the gambles all involve only two outcomes, and so when one variable, such as variance, is changed, others, such as range, must also automatically change. Thus it is impossible to know which of several variables is actually relevant to the behavior. For another, most, if not all, of the results are explicable in terms of the subjective expected utility hypothesis, provided that we are willing to accept utility and subjective probability functions with enough twists and bends.

Where do we stand? It is impossible at present to cite a study either clearly supporting or clearly rejecting the subjective expected utility hypothesis. Because of the freedom to select both the utility and the subjective probability functions, I do not find it surprising that we have been unable clearly to reject it. What would be truly surprising would be for the experimental evidence to fail to give strong support, were the hypothesis in fact correct. This, coupled with the peculiar results that Edwards and others have obtained, has made some psychologists suspicious that, at least in detail, the subjective expected utility hypothesis is wrong. A number of us suspect that the structure of choices is somewhat more subtle than can be satisfactorily encompassed by this very simple and appealing model. Of course, it may well turn out that under some conditions this model is a good approximation to a more nearly correct one, but we probably will not know what the conditions are until we have discovered the more correct one. Consequently, many of us have turned to other approaches.

Preference Theories

An incidental result of the Mosteller–Nogee experiment which has since been seen in a number of other studies is that subjects are not always consistent in their choices. If a pair of alternatives is presented many times, successive presentations being well separated by other choices, a given subject does not necessarily choose the same alternative each time. At first, one is tempted to attribute such inconsistency to changes of state in the subject or to other errors of measurement, much as one would in testing a physical theory. There are, however, two features of the data which lead one to suspect that the phenomenon may in fact be basic to the choice process. The pattern of inconsistency is very regular when it occurs. For example, suppose the two alternatives are the gamble $x\alpha - 5$ ¢ and the pure outcome nothing. When x is small, say 1 ¢, the gamble is never chosen: when it is larger, say \$1, it is always chosen. And as x varies from 1 ¢ to \$1, the probability of choosing the gamble increases in a smooth S-shaped curve, such as that shown in Figure 1. In contrast, there are other pairs, such as x versus nothing, where the behavior is perfectly discontinuous: when $x < 0$, x is never chosen, and when $x > 0$, it is always chosen. It appears that a wonderfully complex error theory will be needed to account for such different results.

As an alternative to treating these results as errors, we can try to construct inherently probabilistic choice models. The postulates so far suggested are numerous and their interrelations complex; much of this net of implications was worked out by Block and Marschak (2) and by Marschak (28). Of the various postulates, the two that have received most empirical attention are weak and strong stochastic transitivity. In both it is assumed that we have three alternatives – *a*, *b*, and *c* – such that when *a* and *b* are presented, the subject selects *a* at least half the time and that when *b* and *c* are presented, he chooses *b* at least half

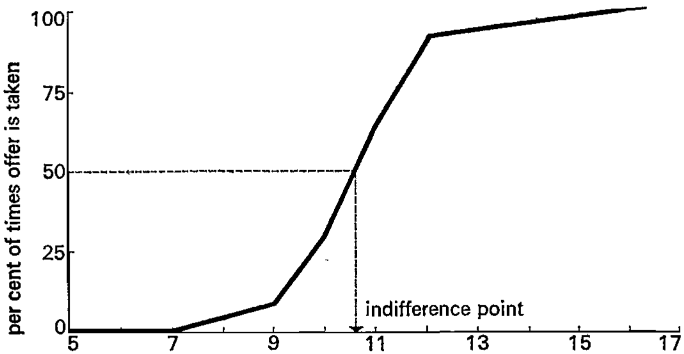

Figure 1 Percentage of times a gamble of the form xα – 5¢, where event α has a probability 0·332 of occurring, was chosen over nothing as reported by Mosteller and Nogee (29, Figure 2). [See page 122.]

the time. The question is what happens when *a* and *c* are presented. Weak stochastic transitivity says that *a* will be chosen at least half the time. Strong stochastic transitivity says that the proportion of times *a* is selected over *c* will be at least as large as the proportion of times *a* is selected over *b* and at least as large as the proportion of times *b* is selected over *c*. In symbols, if $P(a, b) \geqslant \frac{1}{2}$ and $P(b, c) \geqslant \frac{1}{2}$, then weak transitivity says simply that $P(a, c) \geqslant \frac{1}{2}$, and strong says that $P(a, c) \geqslant P(a, b)$ and $P(b, c)$. Strong clearly implies weak transitivity, and it reduces to ordinary algebraic transitivity when the probabilities are 0 or 1.

Neither assumption is strong enough to stand alone as a theory of behavior – certainly not as a theory about choices among risky alternatives, because nothing is assumed about the nature of the alternatives. Aside from a certain *a priori* reasonableness, they are

of interest mainly because weak transitivity is a consequence of practically every more elaborate stochastic theory that has been proposed and strong transitivity, of many. Thus, should we collect data that rejected either, with it would be swept away many other proposed theories.

Almost any preference study that one might perform affords an opportunity to test these hypotheses, but few, if any, studies that have been performed provide, in my view, adequate tests. Davidson and Marschak (8) obtained single observations from each subject for each pair of alternatives in several triplets. Because single observations do not permit very subtle estimates of probabilities, they were forced to a subtle statistical analysis to test the two hypotheses. Although they interpret their results as supporting both, it would have required such extreme data to reject either that I do not find their evidence very convincing. Chipman (5) and Coombs and Pruitt (7) made several observations for each pair of alternatives, estimated the choice probabilities from these, and tested the two hypotheses. The authors in both cases concluded that strong transitivity may well be wrong. For example, Coombs and Pruitt found that strong transitivity was violated in 25 per cent of the triples where it could be tested. However, as they point out, one is bound to have some apparent rejections as a result of sampling variability, and when each probability is estimated from only 6 or 8 observations, as was the case in these two studies, one suspects that this may be quite a serious problem. No one has yet worked out the necessary statistical analysis to know whether their results should be accepted at face value. Griswold and Luce (22) employed 30 to 50 observations per pair to estimate the choice probabilities, and among 103 triples of money gambles they found a total of 13 violations of strong transitivity. Moreover, most of these violations were within one standard deviation of the quantities estimated and so may very well not really indicate violations. In their opinion, there was no substantial reason to doubt strong transitivity.

As I have said, neither transitivity assumption is a complete theory of behavior, and neither has anything in particular to do with risky alternatives; so let me consider next a theory which is restricted to that case. It is described in detail in *Individual Choice Behavior* (Luce [24]). The basic assumptions are three forms of statistical independence. The first, called the 'choice axiom', is a probabilistic version of the ubiquitous independence of irrelevant alternatives notion in decision theory (Luce and

Raiffa [26]). It says that the probability of choosing alternative a from a set S of alternatives is identical with the conditional probability of choosing a from a larger set T, provided that we consider only those choices which, in fact, lie in S, i.e.,

$$P_S(a) = P_T(a \mid S),$$

assuming that the conditional probability exists. For example, suppose that two thirds of the time a person selects steak over lamb chops when they are the only two alternatives. Suppose we now confront him with chicken and liver as well as steak and lamb chops, and let us look only at those occasions when he selects either steak or chops. The axiom asserts that the relative frequency of steak choices will be exactly the same as when the chicken and liver were not present, namely, two thirds.

Our second major assumption limits the theory to risky alternatives. Suppose one must choose between $a\alpha b$ and $a\beta b$. There are two conditions when one should prefer the former, namely, when one prefers a to b and judges α as more likely than β and also when one prefers b to a and judges β more likely than α. If we suppose that a person's preference decisions are statistically independent of his judgments of likelihood, then

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

where P denotes the preference probability and Q the judgment probability. This has been called the 'decomposition axiom'.

The third assumption is the choice axiom applied to likelihood selections among events; it is the same equation as before with Q replacing P and events replacing alternatives.

Of the two results I shall cite, the second leads to an experimental test of the theory. Recall that there are data which suggest that the probability paint is spread somewhat unevenly over the pairs of risky alternatives. The decomposition theory not only allows this to happen but in a sense requires it. Specifically, it can be shown either that the choice probabilities for pairs of sure alternatives are 0, ½, or 1, as happens with money, or that the probabilities of judging one event as more likely than another can assume only three possible values. The probabilities of choice among gambles are not so severely restricted. This somewhat surprising and strong result has, on the whole, made psychologists skeptical of the theory. There is a common view that weaker theories are better than strong ones, for the strong ones are bound to be wrong. My view is just the opposite: the stronger, the better, for, on the one hand, a correct theory will be strong

and, on the other hand, an incorrect one is much more quickly discovered to be incorrect if it is strong than if it is weak.

Our testable conclusion concerns choices in the following situation:

$$\begin{array}{ccc} & \text{I} & \text{II} \\ \alpha & a & b \\ \alpha & d & c \end{array}$$

The subject selects a column, the event α selects the row, and the payoff is the corresponding sum of money. On the assumption that $a > b > c > d$, it is evident that if α is an event which never occurs, the subject should never choose column I; whereas if α is certain to occur, then he should always choose column I. Thus,

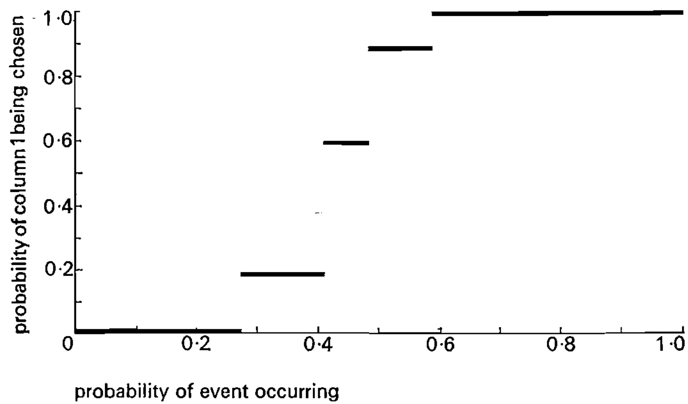

Figure 2 A typical example of the prediction of the decomposition theory that the choice probability is a step function of the event probability. This figure is reproduced from Luce (24) by permission of John Wiley & Sons.

as the probability of α occurring changes from 0 to 1, the probability of choosing column I should also go from 0 to 1. Judging by the data shown in Figure 1 and from related results in other parts of psychology, one might expect this to be a continuous change. Not so, if the decomposition theory is correct. It must be a step function of the sort shown in Figure 2. Just how many steps to expect or where to find them is not specified by the theory, but it says that the function is not continuous.

Obviously, it will be possible to confirm this prediction only if the steps are large and conspicuous. Taking the risk that they might not be, Elizabeth Shipley and I (27) ran the experiment. We used six different payoff matrices for each subject and 15 closely spaced events whose probabilities were known to the subjects. The events were located in the region where the choice probability changed from 0 to 1. We made 50 observations for each event-payoff condition. Of our five subjects, two exhibited

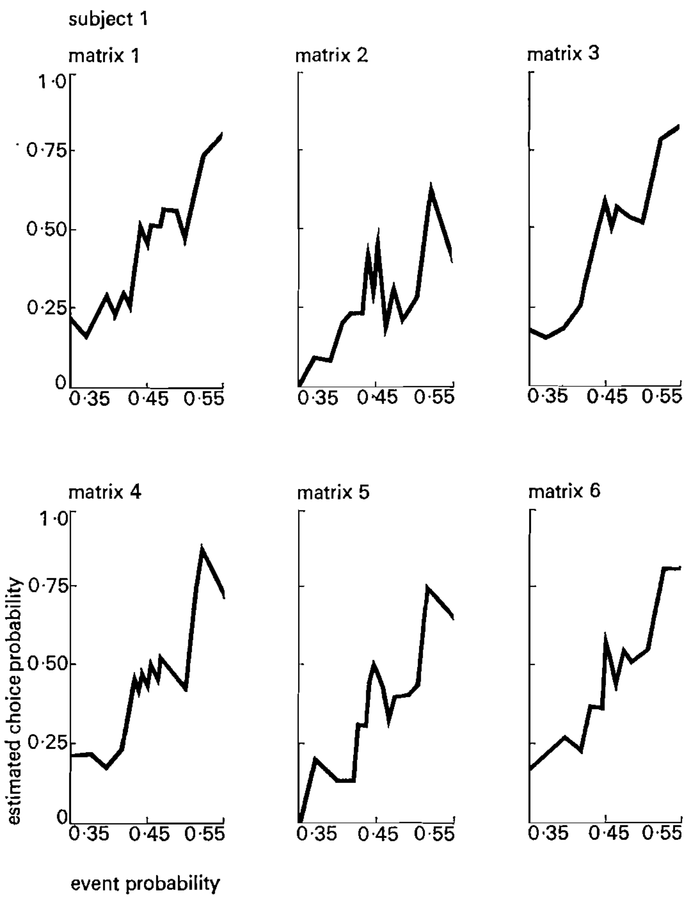

Figure 3 Estimated choice probabilities as a function of event probability for six payoff matrices and three subjects (Luce and Shipley [27]).

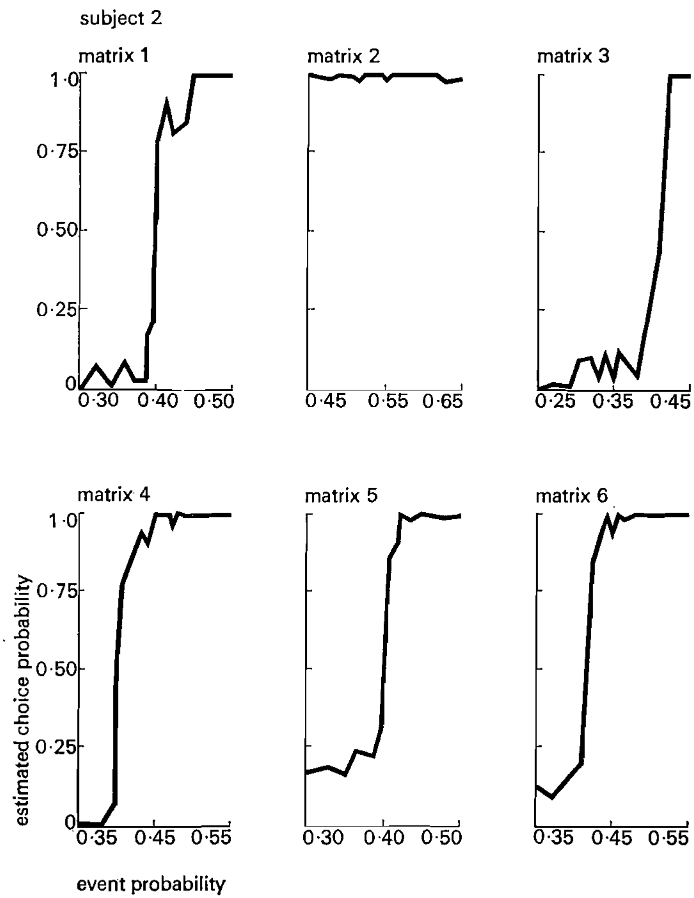

Figure 3 continued

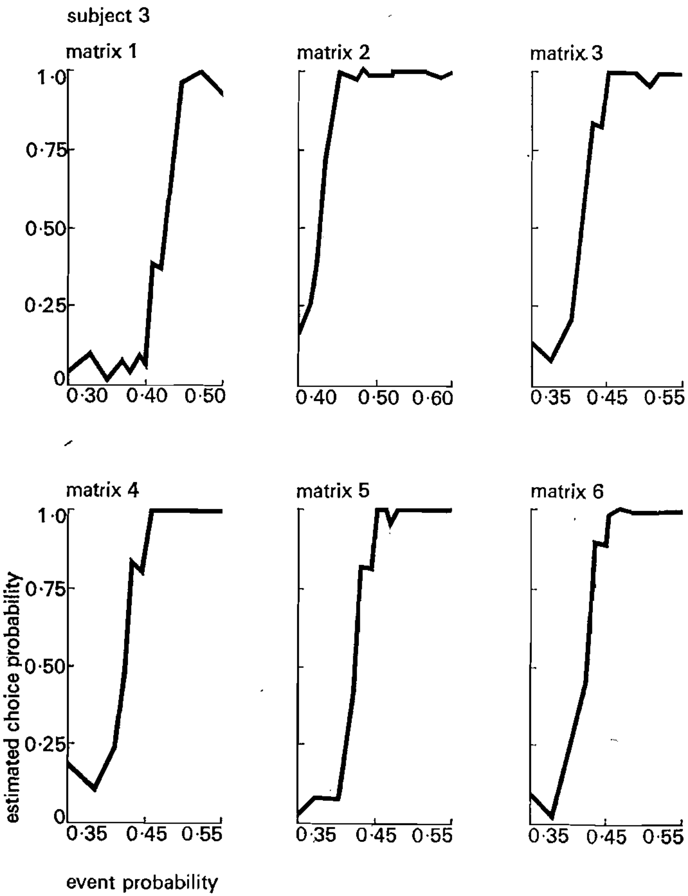

Figure 3 continued

almost perfectly discontinuous functions, and so were consistent with the prediction, but not in a very interesting way. The other three subjects yielded the data shown in Figure 3. You will note that there are curious plateaus of a sort not found in, say, psychophysical data. To see whether these results might have arisen from sampling fluctuations, we fit the data as best we could with smooth *S*-shaped curves (logistic functions). Assuming that

these continuous functions described the probabilities involved, we carried out Monte Carlo computer runs which mimicked the experimental runs. The question was how often in a set of 200 Monte Carlo runs did a pattern of reversals occur that was at least as marked as that exhibited by the data. It turned out that such patterns were very rare indeed. As a result, we concluded that the data could not have arisen from such smooth probability functions. Whether this means that the decomposition theory is approximately correct is another matter; I am not prepared to argue that it is until we have more studies which confirm it.

Learning Theories

Our third group of theorists who study human choice behavior when the alternatives are risky have, like the preference theorists, assumed that responses are stochastically controlled. But rather than seek out relations among the response probabilities in related choice situations, learning theorists have focused upon the mechanisms whereby these choice probabilities change with repeated experience in one situation. Without entering into the mathematical details of the learning models,[1] two of their most important features can be mentioned. First, they all describe an organism with an extremely limited memory. This means that they are totally incapable of accounting for human learning of, for example, periodic binary sequences. Second, the learning models assume that when a subject makes a response and it is rewarded, his probability for making that response again is increased a little bit and that when he makes one and it is unrewarded or punished, the probability is decreased a little.

Naïve as these postulates may seem, some data suggest that they are not always grievously incorrect. Suppose, for example, that on each trial a subject must predict which of two lights will appear, the one being correct a random 75 per cent of the time and the other only 25 per cent. Depending on just what specific model one assumes, the subject's asymptotic response probability is predicted to be the same or nearly the same as the event probability of 0·75. A rational analysis says that, to maximize the number of correct predictions, he should always select the 0·75 light. The data, even after hundreds of trials, indicate that human subjects overshoot the event probability but that, in general, they fail to adopt the rational solution.

1. The details can be found in Bush and Mosteller (4) and Bush and Estes (3).

Or consider another kind of experiment. On each trial, a faint light increment may or may not be presented in a background patch of light, and the subject reports whether or not he thinks it is there. Following his response, he is told whether or not he was correct, and he wins or loses money according to a given payoff schedule. Suppose that there are sensory discrimination thresholds such that internal fluctuations cause the patch to exceed the threshold with one probability, whereas the increment in the patch exceeds it with another, somewhat larger probability. Moreover, suppose not only that the subject bases his responses

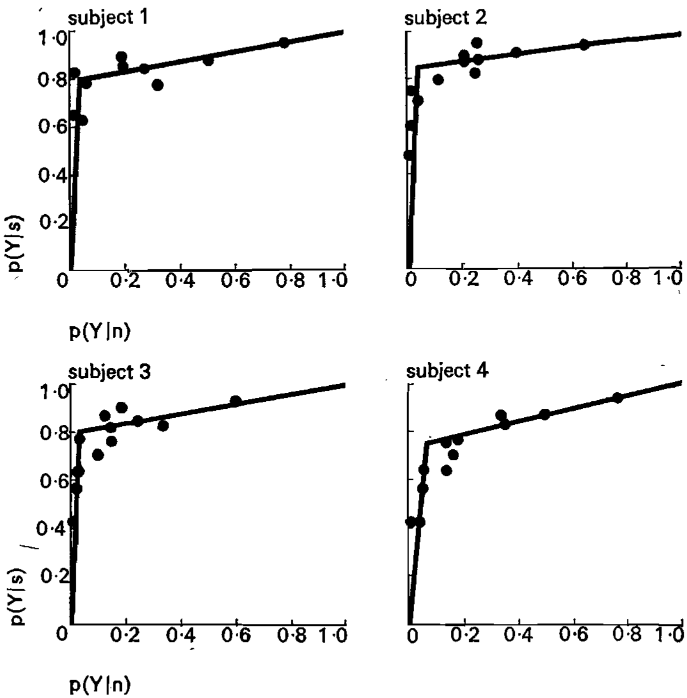

Figure 4 Plots of the estimated conditional probability of detection response when a signal is presented versus *conditional probability of detection when no signal is presented. The data points were generated by using the same visual signal and background but different payoff matrices (Swets, Tanner, and Birdsall [33]), and the theoretical curves are those of a threshold theory (Luce [25]).*

on what he thinks he sees but that he biases what he says in terms of the payoff matrix used. Then it can be shown that a maximization of expected utility leads to just one of three possible modes of behavior (Luce [25]). Either on all trials he says no increment is present, or he always says it is present, or he simply reports what he detects. Which he does depends on the proportion of trials having increments present and on the exact payoffs used. In contrast, if we suppose that the subject learns to bias his responses as a result of the information feedback, then all the intermediate points on the straight lines connecting the above three points are possible. This prediction is shown as the theoretical lines in Figure 4. The actual value of the response probabilities depends on the payoffs and frequency of stimulus presentations. Of these two models, the learning one is favored, judging by the data shown in Figure 4.

I should add, however, that if we assume a different psychophysical model, such as the signal detection model in which there are no thresholds, then the maximization of expected value may yield sensible results. For expositions of this approach, see Green (21), Licklider (23), and Tanner and Swets (34).

Both the guessing and the threshold analyses suggest that a simple adaptive learning process is not, in general, consistent with an optimizing model. This seems, indeed, to be a general proposition. So far as I know, no one has yet stated a learning process which, when the behavior stabilizes at its asymptotic values, results in optimal choices. If, instead of trying to find mechanisms that lead to optimizing behavior, we look for ones leading to the stochastic preference models, the relations are more complex, Suppes (32) has shown that an existing linear learning model, due originally to Estes (18), predicts that the choice axiom holds asymptotically, although it may not hold during the learning phase itself. On the other hand, there do not seem to be any indications that the learning models lead asymptotically to the decomposition axiom. Probably this is related to their failure to lead to a maximization of expected utility, for the decomposition axiom and the expected utility decomposition are certainly very similar in spirit.

Concluding Remarks

These, then, are some of the main outlines of experimental and theoretical psychological research on the question of choices among risky alternatives. Much the most dominant theoretical ideas have been the subjective expected utility hypothesis and the notion that a person chooses optimally relative to his utility function. Although I cannot cite an unambiguous experimental refutation of these utility notions, their *a priori* flexibility, coupled with the failure to achieve really smashing experimental successes by using them, leads me to feel that other formulations must be seriously considered.

Of these, we examined two. Both postulate stochastic choice mechanisms. Preference theory is concerned with static constraints existing among the choice probabilities in different, but related, choice situations. Learning theory is concerned with the dynamic constraints relating the choice probabilities on successive exposures to the same choice situation. Ultimately, we must find a way to fuse these two approaches into a single model, specializations of which lead to the restricted theories we are now trying to develop. At present, however, we are far from that ideal state. Our learning models, which have received some experimental support, are not able to account for all our static assumptions, which also have received some experimental support. Moreover, there are some simple learning situations, e.g., prediction of simple periodic sequences, for which the learning models are completely inadequate.

It is not difficult to indicate where some of the trouble lies, but it is quite another matter to recast the models in such a way as to overcome it. Human beings appear to be both 'adaptive' and 'cognitive'; they sometimes adjust their behavior gradually to experience, and they sometimes 'understand' and analyze choice situations. Furthermore, both processes often seem to go on at the same time. The current learning theories are exclusively adaptive, whereas, almost by definition, the static assumptions of the preference theories are cognitive. By appropriately designing our experiments so as to draw upon just one of these two aspects of behavior, we are able to find support for each class of models. But other experiments in which both processes occur can also be designed, and these are bound to reject both classes of models. Such studies are interesting beyond being mere demonstrations, for it is from them that we shall begin to understand which features of a choice situation control the degree to which the behavior is adaptive or cognitive. Only when we develop such insights, will we be able to construct models that effectively take both into account.

If I am not mistaken, students of business have been much more deeply influenced by those theories that are primarily cognitive, especially those in which the decision maker is supposed to behave optimally. At best, this is only a part of what goes on, and certainly in some situations adaptive models are much more to the point. Regrettably, we cannot yet begin even to indicate what differentiates the two classes of situations or to suggest how to synthesize the two classes of models. It is certain, however, that these problems will receive a good deal of attention in the near future, and, judging by the rapid developments during the past decade, we may have a much clearer idea about the relations between these two aspects of behavior in another ten years.

References

1. ADAMS, E. W., 'Survey of Bernoullian Utility Theory', in H. Solomon (ed.), *Mathematical Thinking in the Measurement of Behavior*, pp. 151–268. Glencoe, Ill.: Free Press, 1960.
2. BLOCK, H. D., and MARSCHAK, J., 'Random Orderings and Stochastic Theories of Responses', in I. Olkin *et al.* (eds.), *Contributions to Probability and Statistics*, pp. 97–132. Stanford: Stanford University Press, 1960.
3. BUSH, R. R., and ESTES, W. K. (eds.), *Studies in Mathematical Learning Theory*. Stanford: Stanford University Press, 1959.
4. BUSH, R. R., and MOSTELLER, F., *Stochastic Models for Learning*. New York: John Wiley & Sons, 1955.
5. CHIPMAN, J. S., 'Stochastic Choice and Subjective Probability', in D. Willner (ed.), *Decisions, Values, and Groups*, Vol. I, pp. 70–95. New York: Pergamon Press, 1960.
6. COOMBS, C. H., 'On the Use of Inconsistency of Preferences in Psychological Measurement', *Journal of Experimental Psychology*, Vol. LV (1958), pp. 1–7.
7. COOMBS, C. H., and PRUITT, D. G., 'Components of Risk in Decision Making: Probability and Variance Preferences', *Journal of Experimental Psychology*, Vol. LX (1960), pp. 265–77.
8. DAVIDSON, D., and MARSCHAK, J., 'Experimental Tests of a Stochastic Decision Theory', in C. W. Churchman and P. Ratoosh (eds.), *Measurement: Definitions and Theory*, pp. 238–69. New York: John Wiley & Sons, 1959.
9. DAVIDSON, D., and SUPPES, P., 'A Finitistic Axiomatization of Subjective Probability and Utility', *Econometrica*, Vol. XXIV (1956), pp. 264–75.
10. DAVIDSON, D., SUPPES, P., and SIEGEL, S., *Decision Making*. Stanford: Stanford University Press, 1957.
11. EDWARDS, W., 'Probability-Preferences in Gambling', *American Journal of Psychology*, Vol. LXVI (1953), pp. 349–64.
12. ——. 'Probability Preferences among Bets with Differing Expected Values', *ibid.*, Vol. LXVII (1954), pp. 56–67.
13. ——. 'The Reliability of Probability Preferences', *ibid.*, pp. 68–95.
14. ——. 'Variance Preferences in Gambling', *ibid.*, pp. 441–52.
15. ——. 'The Theory of Decision Making', *Psychological Bulletin*, Vol. LI (1954), pp. 380–417.
16. ——. 'The Prediction of Decisions among Bets', *Journal o Experimental Psychology*, Vol. LI (1955), pp. 201–14.

17. ——. 'Behavioral Decision Theory', in *Annual Review of Psychology*, pp. 473–98. Palo Alto, Calif.: Annual Reviews, Inc., 1961.
18. ESTES, W. K., 'Toward a Statistical Theory of Learning', *Psychological Review*, Vol. LVII (1950), pp. 94–107.
19. FINETTI, B. DE., 'La prévision: ses lois logiques, ses sources subjectives', *Annales de l'Institut Henri Poincaré*, Vol. VII (1937), pp. 1–68.
20. ——. 'Recent Suggestions for the Reconciliation of Theories of Probability', in *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability.* Berkeley, Calif.: University of California Press, 1951.
21. GREEN, D., 'Psychoacoustics and Detection Theory', *Journal of the Acoustical Society of America*, Vol. XXXII (1960), pp. 1189–1203.
22. GRISWOLD, B., and LUCE, R. D., 'Choices among Uncertain Outcomes: A Test of the Decomposition, Transitivity, and Strong Stochastic Transitivity Assumptions', *American Journal of Psychology*, in press, 1961.
23. LICKLIDER, J. C. R., 'Three Auditory Theories', in S. Koch (ed.), *Psychology: A Study of a Science*, Vol. I, pp. 41–144. New York: McGraw-Hill Book Co., 1959.
24. LUCE, R. D., *Individual Choice Behavior.* New York: John Wiley & Sons, 1959.
25. ——. 'A Threshold Theory for Simple Detection and Recognition Experiments'. Department of Psychology Memorandum MP-12. University of Pennsylvania, 1961. Mimeographed.
26. LUCE, R. D., and RAIFFA, H., *Games and Decisions.* New York: John Wiley & Sons, 1957.
27. LUCE, R. D., and SHIPLEY, E., 'Preference Probability between Gambles as a Step Function of Event Probability', *Journal of Experimental Psychology*, in press, 1961.
28. MARSCHAK, J., 'Binary-Choice Constraints and Random Utility Indicators', in K. Arrow, S. Karlin, and P. Suppes (eds.), *Mathematical Methods in the Social Sciences*, 1959, pp. 312–29. Stanford: Stanford University Press, 1960.
29. MOSTELLER, F., and NOGEE, P., 'An Experimental Measurement of Utility', *Journal of Political Economy*, Vol. LIX (1951), pp. 371–404.
30. RAMSEY, F. P., *The Foundations of Mathematics and Other Logical Essays.* New York: Harcourt, Brace & Co., 1931.
31. SAVAGE, L. J., *The Foundations of Statistics.* New York: John Wiley & Sons, 1954.
32. SUPPES, P., *Behavioristic Foundations of Utility.* Technical Rept. No. 23, Institute for Mathematical Studies in Social Sciences, Stanford University, 1959.
33. SWETS, J. A., TANNER, W. P., JR., and BIRDSALL, T. G., *The Evidence for a Decision-Making Theory of Visual Detection.* Technical Rept. No. 40, Engineering Research Institute, University of Michigan, 1955.
34. TANNER, W. P., JR., and SWETS, J. A., 'A Decision-Making Theory of Visual Detection', *Psychological Review*, Vol. LXI (1954), pp. 401–9.
35. VON NEUMANN, J., and MORGENSTERN, O., *Theory of Games and Economic Behavior.* Princeton: Princeton University Press, 1947.