AMRL-TDR-64-77

PREDECISIONAL PROCESSES IN DECISION MAKING: PROCEEDINGS OF A SYMPOSIUM

DECEMBER 1964

BEHAVIORAL SCIENCES LABORATORY
AEROSPACE MEDICAL RESEARCH LABORATORIES
AEROSPACE MEDICAL DIVISION
AIR FORCE SYSTEMS COMMAND
WRIGHT-PATTERSON AIR FORCE BASE, OHIO

PREDECISIONAL PROCESSES IN DECISION MAKING: PROCEEDINGS OF A SYMPOSIUM

FOREWORD

This report was assembled and edited by Dr. Darwin P. Hunt and Captain Donald L. Zink. It is essentially a compilation of the proceedings of a symposium sponsored by the Presentation of Information Branch, Human Engineering Division, Behavioral Sciences Laboratory, Aerospace Medical Research Laboratories, under Project No. 7183, "Psychological Research on Human Performance," Task No. 718303, "Research on Human Intellectual Functions." The symposium was held at Wright-Patterson Air Force Base in April 1962.

Contributions to the symposium were made by: John W. Atkinson, Ph.D., University of Michigan; Daniel E. Berlyne, Ph.D., University of Toronto; Ward Edwards, Ph.D., University of Michigan; Janellen Huttenlocher, Ph.D., Harvard University; John T. Lanzetta, Ph.D., University of Delaware; Thornton B. Roby, Ph.D., Tufts University; Emir H. Shuford, Jr., Ph.D., University of North Carolina; Masanao Toda, Ph.D., Harvard University. Dr. Janellen Huttenlocher's contribution could not be included in these proceedings due to previous publication commitments.

The initial impetus for the symposium arose out of discussions with Dr. John T. Lanzetta, University of Delaware, who was at that time conducting research for the Behavioral Sciences Laboratory in the area of decision making. The editors are indebted to a number of individuals whose efforts contributed to the success of the symposium: Mr. A. J. Cannon, Special Activities Division, Plans and Operations Office, handled many necessary administrative details; Mrs. Mary Williams and Mr. Peter F. Lund, Office of Staff Judge Advocate, provided stenotype records of the discussions following each paper; Mrs. Barbara Calhoun, Presentation of Information Branch, typed the transcripts of the discussions and provided other secretarial assistance; Mrs. Ellen J. Jennings, Presentation of Information Branch, rendered valuable support in collating the report. The symposium was supported by Contract No. AF 18-600-1792, with the University of Virginia.

This technical documentary report has been review and is approved.

WALTER F. GRETHER, PhD
Technical Director
Behavioral Sciences Laboratory

ABSTRACT

The proceedings of a symposium on "Predecisional Processes in Decision Making" are presented. Predecisional processes may be characterized as the search for, the acquisition of, and the evaluation of information prior to the choice of a course of action. The objectives of the symposium were: an assessment of the adequacy of present decision theories in dealing with human decision making behavior; an assessment of other approaches to decision making situations; and an analysis of predecisional processes. Seven papers are presented that analyze this area from several different theoretical viewpoints. Transcripts of the discussions following each paper are also included.

TABLE OF CONTENTS

PREDECISIONAL PROCESSES IN DECISION MAKING: PROCEEDINGS OF A SYMPOSIUM

INTRODUCTION

The process of making a decision pervades human behavior; in fact, one of the contributors to this volume has declared that the phenomenon of choice provides the key to the study of "the higher mental processes". Whether or not one agrees with this point of view, it is true that the study of decision processes is an area that is receiving an increasing amount of attention, both as an area of research in its own right, and as an adjunct to other areas.

The formal description of a decision situation includes a statement of the possible courses of action available to the decision maker, a listing of the possible states of the world, and a matrix representing the outcome to the decision maker if he follows a particular course of action, and a particular state of the world obtains. Given a specific problem, probabilities may or may not be associated with the possible states of the world. The problem confronting the decision maker is to choose a course of action that will maximize the outcome that he realizes. Typical empirical research on decision making has usually provided the decision maker with information concerning the above components of a decision situation, and has greatly increased our knowledge of how values, probabilities, and other factors act as determiners of choices. (Edwards has provided reviews in refs. 1, 2, 3.)

The formal description of a decision situation, however, greatly oversimplifies the problem: making a choice is merely the final step in a chain of psychological processes. The decision maker first acquires and organizes information about a decision situation; he evaluates the information available and determines if other information is necessary. Only after acquiring as much information as seems reasonable, does he make a choice, which he hopes is "optimal" in some sense. It is these processes of searching for and acquiring information, of organizing and evaluating information, that are here called "predecisional processes," and which were the subject of the symposium reported in this volume.

The symposium had the following stated objectives.

An assessment of the adequacy of present decision theories in dealing with human decision making behavior. The question may be put as to how well existing concepts of decision theory account for and predict decision behavior. Are existing concepts adequate for the inclusion of predecisional processes in an extended model of human behavior?

An assessment of other approaches to decision making situations. It is not unlikely that other existing theories may provide useful conceptual vehicles for understanding human choice behavior. At least it seems to the editors that failing to obtain a large favorable outcome would result in a degree of "frustration", or that a particular decision situation might prove to be especially "stressful". Certainly predecisional processes are closely related to the areas of perception and learning.

An analysis of predecisional processes. As indicated earlier, these are clearly an important part of the complete decision situation; yet, typical research on decision behavior seems to ignore these processes, or to assume that they have been completed. For a more complete theory of decision making, it would seem that the processes whereby an individual searches for, acquires, processes, and evaluates information prior to the point of making a choice should be clarified and studied.

Because these objectives are broad in nature, and especially because of the second objective stated above, the papers reported here are not what a reader might expect to find in a volume on decision making. The reader will be on familiar ground with the papers by Edwards and Shuford. The papers by Atkinson and Berlyne might seem especially foreign to the student of decision behavior; yet they are clearly relevant to a completer understanding of decision processes. Atkinson's emphasis on the activity already in progress at the time a stimulus is presented is relevant to any area of psychological research; Berlyne calls attention to the relationship between many concepts in behavior theory and decision making, and suggests, for example, how the concept of curiosity can be related to search for information. Lanzetta utilizes some new measures in his studies of decision situations, and reinforces many of the ideas suggested by Berlyne. Toda's analysis of decision making emphasizes cognitive aspects of the decision situation, especially perceptual ones; Roby analyzes the environment constructed by the decision maker, and suggests several potentially useful concepts for understanding information gathering.

The goals of the symposium, of course, were not fully realized; the editors would attach a high probability to the assertion that no symposium completely achieves its stated objectives. A more reasonable goal for any symposium might be to provide a starting point for further effort. The editors hope that at least this objective was achieved.

REFERENCES

1. Edwards, W., "A Perspective on Automation and Decision Making," in D. Willner, Editor, Decisions, Values, and Groups, New York, Pergamon Press, 1960.

2. Edwards, W., "Behavioral Decision Theory," Annual Review of Psychology, Vol 12, pp 473-490, 1961.

3. Edwards, W., "The Theory of Decision Making," Psychological Bulletin, Vol 51, pp 380-417, 1954.

BLANK PAGE

INFORMATION-SEEKING TO REDUCE THE RISK OF DECISION*

Ward Edwards
Engineering Psychology Laboratory
University of Michigan

The purposes for which I was invited to this meeting were to comment on the present state and future directions of research on human decision processes, and to say something about my own research. Most of my time will be devoted to reporting an experiment. Still, I can never quite resist the temptation to air my own opinions, especially when I'm invited to do so, so let me do a little of that first.

It is natural to begin a symposium on pre- and postdecisional factors in decision making by asking what precedes and succeeds decisions. One valid answer is that what precedes a decision is a previous decision; and what succeeds a decision is the next decision. Decisions come in ordered and, hopefully, in orderly sequences. An appropriate goal for research in this area would be to find out what the facts that decisions come in sequences and that these sequences are orderly might mean for the development of a theory.

I find it helpful to think of human intellectual processes as resembling those of a guide who must lead a wagon through a hitherto unexplored range of mountains. (Incidentally, my guide has many of the same characteristics as Dr. Masanao Toda's fungus-eater, and serves the same sort of thought-orienting purpose.) A sensible guide approaching the foothills of the mountain range would most certainly <u>not</u> make a blind foray into the range, in the spirit of the trial-and-error conception of even the most modern and sophisticated learning theories. He would be more likely to climb a tree so he can see better. Perhaps he learns by climbing the tree that half a mile off to the right there is a hill, and that if he were to climb the hill, he could see still more. From the hill, he may see two or three promising-looking canyons into the mountain range. Even then he does not lead his wagon train into one or another of the canyons, but instead sends a scout into each one. Only after a substantial number of such information-seeking moves, each of which may serve only to put him in a position where he can seek information more effectively than before, does he begin to make moves actually directed at getting the wagon train through the mountain range, and even then he intersperses more information-seeking actions among the actions which are directed more explicitly to the ultimate goal.

I speculate that my guide uses precisely three intellectual processes, in addition, of course, to the tools of memory, perception, etc. The first of these intellectual processes is the acquisition and organization of information relevant to the ultimate goal and to the route from here to there. The second is the formulation, on the basis of the information at hand, of a set of alternative courses of action. Usually, this set will not be exhaustive in any mathematical sense.

---

*This work was sponsored by Project MICHIGAN under Department of the Army Contract DA-36-089 SC-78801, administered by the United States Army Signal Corps. I am grateful to Robert Ollman for assistance in running <u>Ss</u> and processing data, to Harold Lindman for writing a computer program, to Mrs. Carol Orwant, who invented the punchboard, helped design the study, ran <u>Ss</u> and did some of the data processing, and to S. Paul Slovic, who finished the data analysis.

Furthermore, it will usually be dominated, in the early stages at least, by information seeking actions rather than by actions looking directly toward attainment of the ultimate goal. The third intellectual process is the selection of a course of action from among those invented in the course of the second step. Ordinarily, selection of a course of action is followed by its execution. I speculate further that these three stages always occur prior to any considered action, and that they occur in that order. Naturally, however, I am not suggesting that they occur only once. Instead, as the mountain-traversing metaphor suggests, I assume that each step occurs not once but many times in the course of solution of any significant problem, with the information obtained as the result of any one move contributing to the re-analysis of the relevant data that precedes the next move.

Of the three intellectual processes I have just suggested, - acquisition and organization of information, invention of courses of action, and selection of one among them - we get to observe only the output of the third. Clearly, the other two, which might be called concept formation and thinking, must precede the third, which I call decision making; so in a sense, they -- concept formation and thinking -- are the predecisional processes in decision making. Still, the observable events that precede a given decision are usually decisions whose informational outcomes are relevant to it, and hopefully also to the ultimate decision, if any, that terminates the problem.

From this point of view, only a very heroic simplification of the overall problem can reduce it to the dimensions of the utility - subjective probability model of static decision theory. In particular, in order to use that model it is necessary to assume that the informational outcomes of earlier decisions are irrelevant to later decisions. I know of no experiment in the decision theory literature that gives any adequate basis for that assumption, and I have never been able to figure out how to run an experiment for which I would condisder such an assumption to be valid or sensible. Of course, the assumption is less idiotic for some experiments than for others, and one can hope that the utility and subjective probability effects may be so strong, or the load on memory so great, that the effects of informational outcomes of earlier decisions on later ones may be of secondary importance relative to other effects. If there is any sense in the traditional decision theoretical literature at all, and I think there is, it is based on that hope.

This phrasing of the problem suggests that decision theory should study sequences of decisions, especially from the point of view of the relevance of the informational outcomes of earlier decisions to later ones. Such a change of emphasis is a choice point. I can see two directions in which decision theory could develop. One, already somewhat explored, is the direction of simplicity-seeking. Simplicity-seekers attempt to invent very simple axioms about human behavior in sequential situations which nevertheless turn out to have sufficiently powerful consequences to be worth considering. The outstanding examples of this strategy for theorizing, of course, are the stochastic learning models. Simple they certainly are, at least in their assumptions. For example, the current trend toward the one-element sampling models is almost completely canfed by the fact that such models are analytically more tractable than their multi-element predecessors.

This direction is inherently unattractive compared with other alternatives. For one thing, simplicity is often more apparent than real. For another thing, these models are unresponsive to many of the considerations which seem important to decision theorists, and especially unresponsive to cost-payoff considerations. An attempt is in progress to introduce cost-payoff considerations into some of the models. So far, I have felt this is a somewhat artificial grafting operation. Most important of all, at least to me, is the fact that these stochastic models are patently absurd. I cannot understand how a man can spend an afternoon in a committee meeting, striving with all of his knowledge, experience, and ingenuity to maximize utility in a task such as deciding what graduate students to admit to the department, and then go home that evening and write yet another article attributing the afternoon's activities to the operation of a random stimulus sampling mechanism! I don't base my decisions on such random processes, and I think it very unlikely that you do. So I would advocate considering alternatives to simplicity-seeking.

What can we do instead? One way of getting at that question would be to have a look at what is good in the static approach to decision theory, and to see if it might be possible to import some of the good things in that approach into this more complex and realistic conception of the problem. I think there are three such good things. One of them is the notion, inherent in the concept of utility, that trade-off relationships among inconsistent value dimensions are essential to any conception of decision making. Such trade-off relationships are fully as important in the dynamic as in the static conception of the problem. The second good thing that I think we can adopt from static decision theory might be phrased as the notion that all events have probabilities, whether they do or not. More sophisticatedly, the notion asserts that men impute probabilities to events, and base their decisions on these probabilities, regardless of whether relative frequencies for the events in question can or cannot be meaningfully defined. The third good thing to borrow from traditional static decision theory is that men do combine these utilities and probabilities according to a principle which either is or closely resembles the calculation of an expected value for each course of action considered, and then base choices on the ordering of these expected values.

This may sound as though I have imported static decision theory completely into the dynamic problem, and in a way I have. The difference, however, is that in the sequential conception of decision processes, opinions about probabilities and even about payoffs are a function of the information available at the time, and so change as the information changes. Furthermore, expected value is to be maximized over a sequence of decisions, embedded in this changing information milieu, rather than for any single decision considered alone. To this kind of utility-subjective probability maximization theory for sequences of decisions, I have given a name chosen for high content of OK words: Dynamic Decision Theory. The situations to which it is applicable, I call dynamic decision situations, or dynamic situations.

In dynamic situations, a new complication not found in the static situations arises. The environment in which the decision is set may be changing, either as a function of the sequence of decisions, or independently of them, or both. This possibility of an environment that changes while you collect information about it makes the task of dynamic decision theory difficult and fun. Most of the manageable cases are cases in which the environment is not changing, or is changing only in systematic response to the decision maker's decisions. It is often necessary to assume that the environment is what the mathematicians call stationary. This means, roughly speaking, that the environment's behavior is controlled by some variety of statistical process, and that the characteristics of that statistical process do not change with time. A roulette wheel or a pair of dice are stationary in this sense. Reasonable mathematical treatment of most decision problems involving stationary environments is possible; reasonable mathematical treatment of decision problems involving non-stationary environments is often unavailable.

How should we go about building a dynamic decision theory? Borrowing an idea from Tanner, we might define an ideal dynamic decision maker, comparable in nature, abstractness, and realism or lack thereof to Tanner's ideal psychophysical observer. I am under the impression that much of the mathematics required to think about an ideal dynamic decision maker is already well understood. In particular, information theory, Bayesian statistical conceptions, and dynamic programming conceptions are all relevant to the definition of such an ideal dynamic decision maker. These areas of mathematical thinking, initially separate, are currently being put together in very interesting and relevant ways by Kullback, Watanabe, and Bellman, and some of the interrelations, I think, suggest extremely interesting research problems. Two essential ideas about information are going to have to be distinguished in any dynamic theory of this sort. One of these has to do with how much information there is in any observation. We are accustomed to using information theory for questions like that. The other has to do with the amount of relevance that information may have to a decision (and therefore its utility or economic value). For simple cases, both of these ideas are already well defined, and their interrelations are known. Bayesian statisticians have been especially concerned with developing such ideas.

Strictly speaking, most of the mathematical techniques I refer to are applicable only to stationary environments. But there is much less to this limitation than meets the eye. In exactly the same sense, the Bush-Mosteller model is concerned only with stationary processes. A number of devices for using stationary mathematics to fit non-stationary real life are found in engineering and mathematics, and the Bayesian approach lends itself well to some of these devices. Any lawful phenomenon is stationary in the sense that the laws controlling it do not change with time, or that the super-laws controlling these laws do not change.

Of course, the ideal dynamic decision maker does not always closely resemble a real decision maker, but real and ideal dynamic decision makers resemble each other relatively closely in situations considerably more complex than one might initially suppose. The experiment reported later in this paper is one such finding. Another experiment was done in the Michigan laboratory by Gordon Robinson. He required subjects to keep track, by means of a tracking handle, of changing probabilities displayed by means of two flashing lights. He found that subjects can keep track of these probabilities as they change without notice exceedingly well, better than the best linear model we could invent, and as well as the best non-linear optimal model we could invent. This suggests the thought that people can function as very good transducers for probabilities, and that thought, if correct, has very considerable implications for systems design, particularly of Bayesian systems.

Even when the real dynamic decision maker does not closely resemble his ideal counterpart, the comparison between them can teach us quite a lot. For one thing, the discrepancies may be simpler to discover than the laws that relate real performance to other less abstract characterizations of the task. This, too is illustrated in the experiment to be reported. Finally, at the very lowest level of usefulness, the notion of the dynamic decision maker is very fruitful in suggesting new experiments.

My faith in the theoretical and practical usefulness of the comparison-with-ideal-performance approach to the task of dynamic decision theroy can be summed up by saying that one single, simple description of human behavior far excels all others for widespread truth and relevance to all decision tasks. It is contained in the following question and answer. Question: What is he doing? Answer: He's doing the best he can!

Now, consider an experiment. This experiment was sponsored by Project MICHIGAN, which is an Army contract concerned with battlefield surveillance. I hope I may be forgiven for bringing Army sponsored work into an Air Force symposium, but I hear that collaboration of this sort is sometimes encouraged. Collaborating with me in conducting it and analysing the data were S. Paul Slovic, Carol Orwant, and Robert Dillman.

To what degree can optimization notions, like those defining an ideal dynamic decision maker, be used to give an account of how people seek information in order to reduce the risk of an eventual decision, when the information is costly? Consider a commander who possesses a limited supply, say one, of tactical nuclear missiles, and receives somewhat unreliable information locating a large tank concentration at a certain set of coordinates. Either he fires, hoping that the information was right, that the coordinates were correct, and so on, or he sends out a special sensor mission, accepting the resulting delay in time in order to make sure before he spends that precious missile. If he does send out a special sensor mission in order to make sure, he runs the risk that while it's going and coming back the concentration may have moved elsewhere or may have dispersed and he may have missed his opportunity. We simulated some aspects of such a problem by means of a very simple device, a punch board, illustrated in Fig. 1.

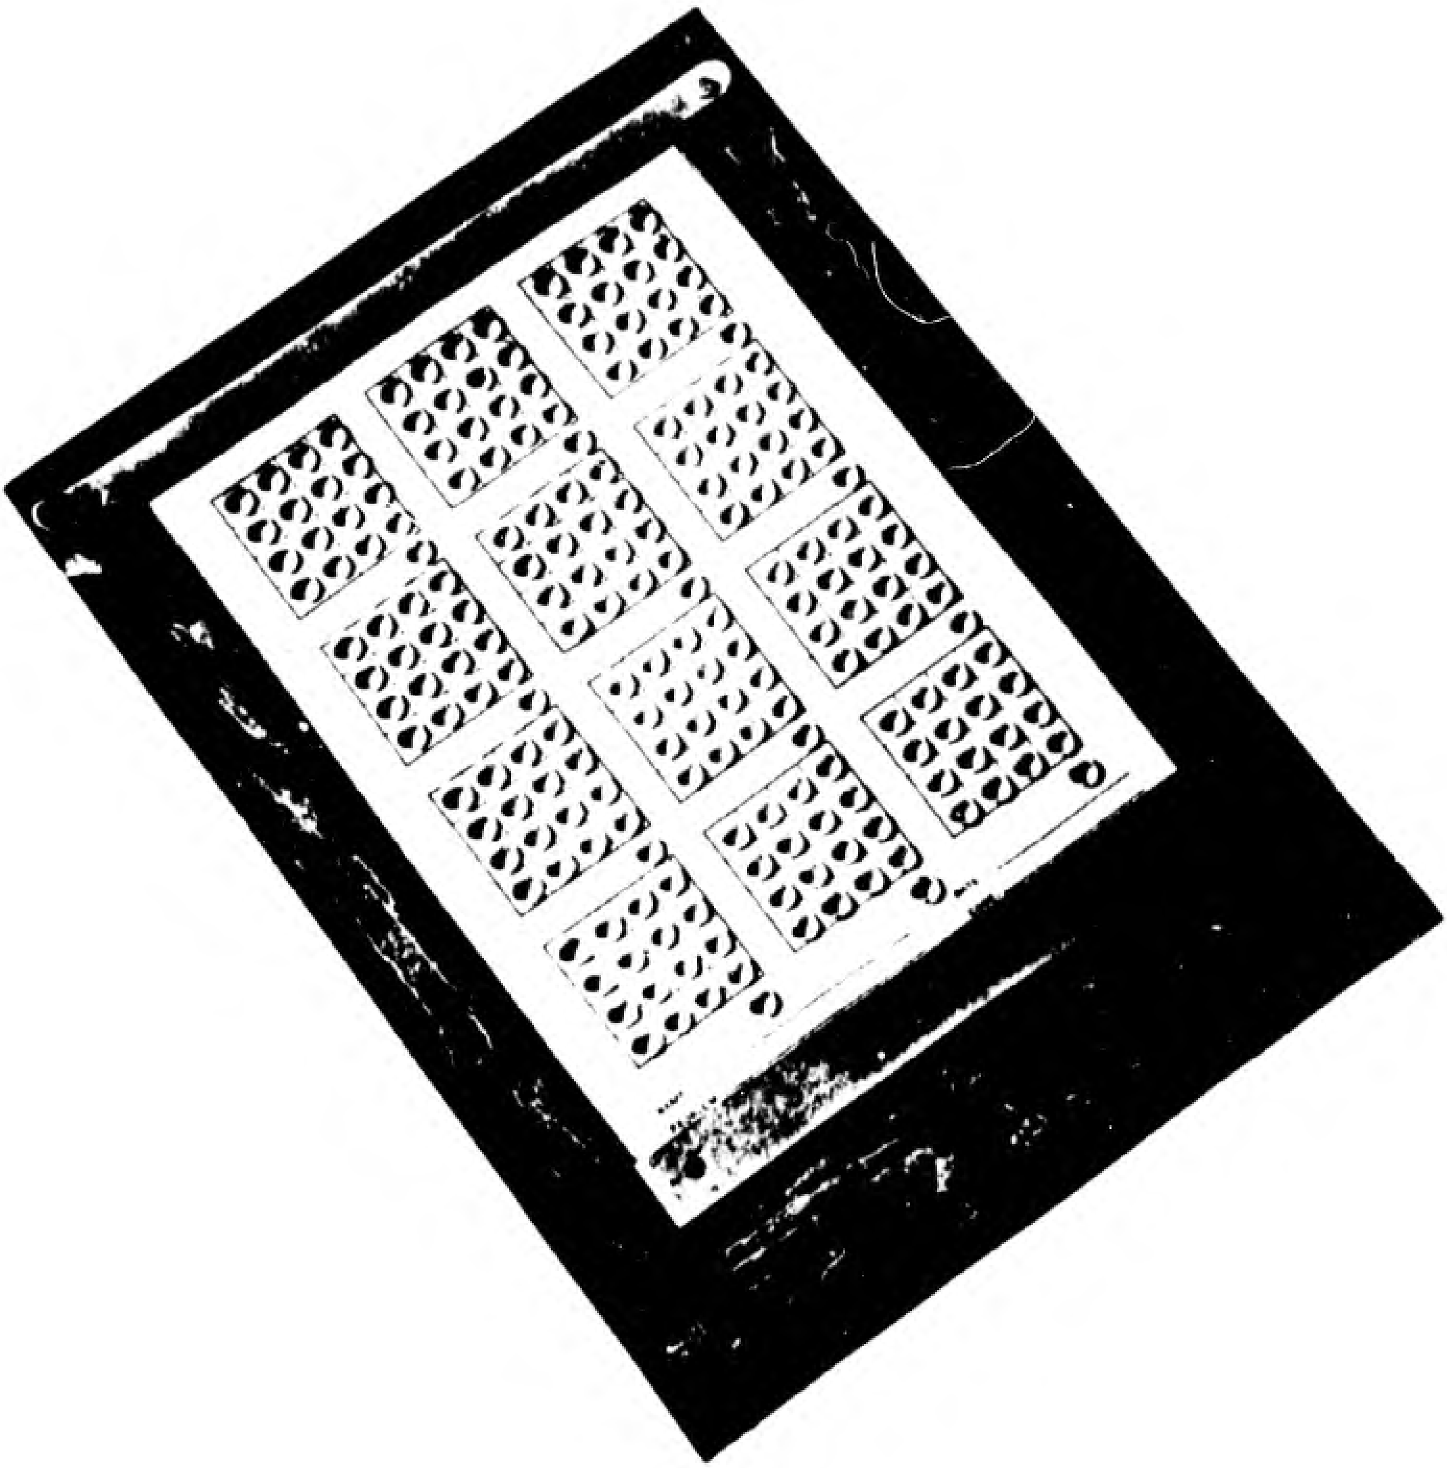

Figure 1. The Punchboard

This is an assembly consisting of a holder, a backing plate which is underneath a piece of paper, the piece of paper, then a surface plate of transparent plastic with holes in it. The assembly has 12 16-hole matrices. In each 16 cell matrix there is one unique cell. The subject knows that he has encountered that unique cell if he tries to write in the hole that is in fact the unique cell and his pencil goes through the paper, through a hole in the backing plate. His task, in the Standard Game, is to locate that unique cell. He has one opportunity to designate the unique cell, and if he is correct, he is paid a prespecified amount of money, say 20 cents; if he is wrong, nothing unpleasant happens, he merely goes on to the next problem, losing the opportunity to make any money on this problem. However, before he designates the unique cell, he may, if he wishes, buy an opportunity to "look" into (i.e., write in) one cell, say, for 2 cents. If he buys that opportunity and "looks" into a cell and that cell is in fact the unique cell, then he is way ahead because he can then designate the unique cell with complete confidence. If not, he still gains something. Now there are only 15 possibilities left instead of 16. Furthermore, at this point he may, if he wishes, buy another observation, and another and another and so on, until he has observed as many times as he wishes and decides to quit and designate. He may take as many as 15 observations. If he does that, he is sure where the unique cell is. He never takes the 16th observation. On the other hand, at 2 cents a look, 15 looks cost him 30 cents and this game pays only 20 cents if he is right, so he is a sure loser.

It occurred to us that the Standard Game is closely related, strategically, to another game which we call the Inverse Game. You obtain the Inverse Game simply by changing the signs of the payoff. That is, there is a unique cell, and if you find it, you are fined a specified amount of money. If, however, you look into a cell and do not find the unique cell, then you are paid a specified, but much smaller amount of money. You simply change the signs of the payoffs. We ran our subjects on both the Standard and the Inverse games.

A strategy in either one of these games can always be thought of in the form, "I will look, at most, at N cells, and if I have not found the unique cell, I'll stop looking." In the Standard Game they will thereafter guess, and in the Inverse Game they will thereafter refuse and decide to go on to the next game. The location of unique cells is randomly distributed so that what cell the subject looks in is immaterial. The only thing of interest to us is how many cells he looks into for a specified payoff and a specified cost per look in the Standard Game, or for a specified cost and specified payoff for looks in the Inverse Game. There will be associated with each strategy, that is, with each number of cells that the subject is willing to look into, an expected value that is a measure of how attractive that strategy is. For the Standard Game, the expected value of strategy N is given by the expression

$$EV_N = \frac{N}{16}(R) - \sum_{n=1}^{n=N-1}\left[\left(\frac{17-n}{16}\right) \times (C_n)\right]$$

where R is the payoff for locating the unique cell
N is the number of cells observed
n is the nth observation
and C_n is the cost of the nth observation

The expected value for the Inverse Game is given by

$$EV_N = \frac{-N}{16}C + \sum_{n=1}^{n=N-1}\left[\left(\frac{15-n}{16}\right) \times R_n\right]$$

where C is the cost for locating the unique cell
and R_n is the payoff for the nth observation.

These equations define ideal performance, since ideal performance in the game would consist of choosing that value of N for which this function is maximized, given the specified cost and payoff conditions. The capital *N* is simply the name that we chose for the subject's strategy and is the total number of cells that the subject contemplates investigating, including the free look in the Standard Game and the compulsory first look in the Inverse Game. We compel the first look in the Inverse Game so the subject can't just refuse to play at all. Incidentally, in the Inverse Game, we don't pay them for the first look. We pay only if he does not find the unique cell in subsequent looks.

Ten subjects participated in eleven sessions. In each session each subject performed 48 tasks. All problems within a given session had the same cost and payoff structure. The cost and payoff structure was carefully explained to the subjects ahead of time, and they always had in front of them a sheet of paper on which the cost and payoffs that applied to that day were spelled out in great detail. The subjects were randomly chosen undergraduates. Real money was at stake, and it was quite a lucrative game from the subject's point of view.

Table 1 shows the cost and payoff structures. There are two classes of games. At the top are those games in which the cost or payoff for finding the unique cell is fixed, and the cost or the payoff for each voluntary look is fixed and does not change as a function of the number of looks. In the other class of games, shown in the second half of the table, the cost or payoff for each look does change as a function of the number of looks. Subjects do about equally well in the more complicated task and in the simpler task. A distinction is made between Standard Mirror, and Fitted Inverse conditions. Mirror and Fitted Inverse refer to ways of making up the payoffs for the Inverse task, and it turns out that the difference was not very substantial. Furthermore, we made some errors that further blurred the distinction. It is best to ignore that distinction, and to think of Mirror and Fitted Inverse together, as an Inverse task.

Figure 2 shows a lot more about the task. In each panel, the upper function shows the expected value for each of the *N* strategies, that is, for each number of cells at which the subject might contemplate looking. The first column shows all standard conditions; the next two columns are all inverse conditions. There are several classes of functions. In session 1, for example, two good strategies are never to buy any information at all or else to keep buying information until you find the unique cell. In session 4 the optimal strategy is to buy 8 items of information and then stop. In session 6 the optimal strategy is never to buy information at all; and in session 9 the best strategy is always to buy information until you find the unique cell. Similar functions are shown for the Inverse case. In session 5 you ought to buy 8 cells and stop. In session 10 you should never buy any information at all. There is not, for the Inverse games, a case in which you ought always to buy as much information as you can.

The lower curves in each panel are the data of the experiment. These are the distributions of subject strategies, i.e., strategies actually chosen. How the subject's strategies were identified requires explanation. Some subject strategies are easy to identify. If the subject does not find the unique cell, the number of cells he looked at before he stopped is his strategy by our definition of strategy. But if he does find the unique cell, then he is through and all you really know about that particular problem is that on that problem his strategy was to look at at least as many cells as he did look at, maybe more. If finding or not finding the unique cell were not correlated with strategy, there would be no serious problem. Unfortunately, it is. If the subject looks at only a few cells, not acquiring much information, then he will seldom find the unique cell. If, on the other hand, he acquires lots of information, then he will almost always find the unique cell.

The key to a satisfactory solution was the recognition that it is necessary to infer the subject's strategy on the trials where he did find the unique cell from his strategy on the trials where he did not. The only rule for inference that could possibly be used is that for problems with the same cost and payoff structure,

Table 1

Costs and Payoffs

Tasks for which Cell Cost or Payoff Was Fixed

Session No.	1	2	6	7	8
Task Type	Standard	Mirror	Standard	Mirror	Fitted Inverse
Payoff or Cost for Finding Unique Cell	+20¢	-20¢	+12¢	-12¢	-12¢
Cost or Payoff for Each Voluntary Look	-2¢	+2¢	-2¢	+2¢	+1¢

Tasks for which Cell Cost or Payoff Varied

Task No.		3	4	5	9	10	11
Task Type		Fit.Inv.	Stand.	Fit.Inv.	Stand.	Mir.	Fit.Inv.
Payoff or Cost for Finding Unique Cell		-20¢	+32¢	-32¢	+40¢	-40¢	-40¢
Cost or Payoff for Look no. (n):	1	+3¢	-1¢	+8¢	-2¢	+2¢	+9¢
	2	0	1	3	2	2	3
	3	2	1	4	3	3	3
	4	1	1	4	3	3	3
	5	2	1	4	3	3	5
	6	1	2	5	3	3	4
	7	3	2	6	4	4	5
	8	2	2	6	4	4	5
	9	4	7	2	5	5	7
	10	4	8	3	5	5	6
	11	6	10	3	6	6	9
	12	7	12	3	7	7	10
	13	11	16	5	9	9	15
	14	17	24	8	12	12	22
	15		48		18		

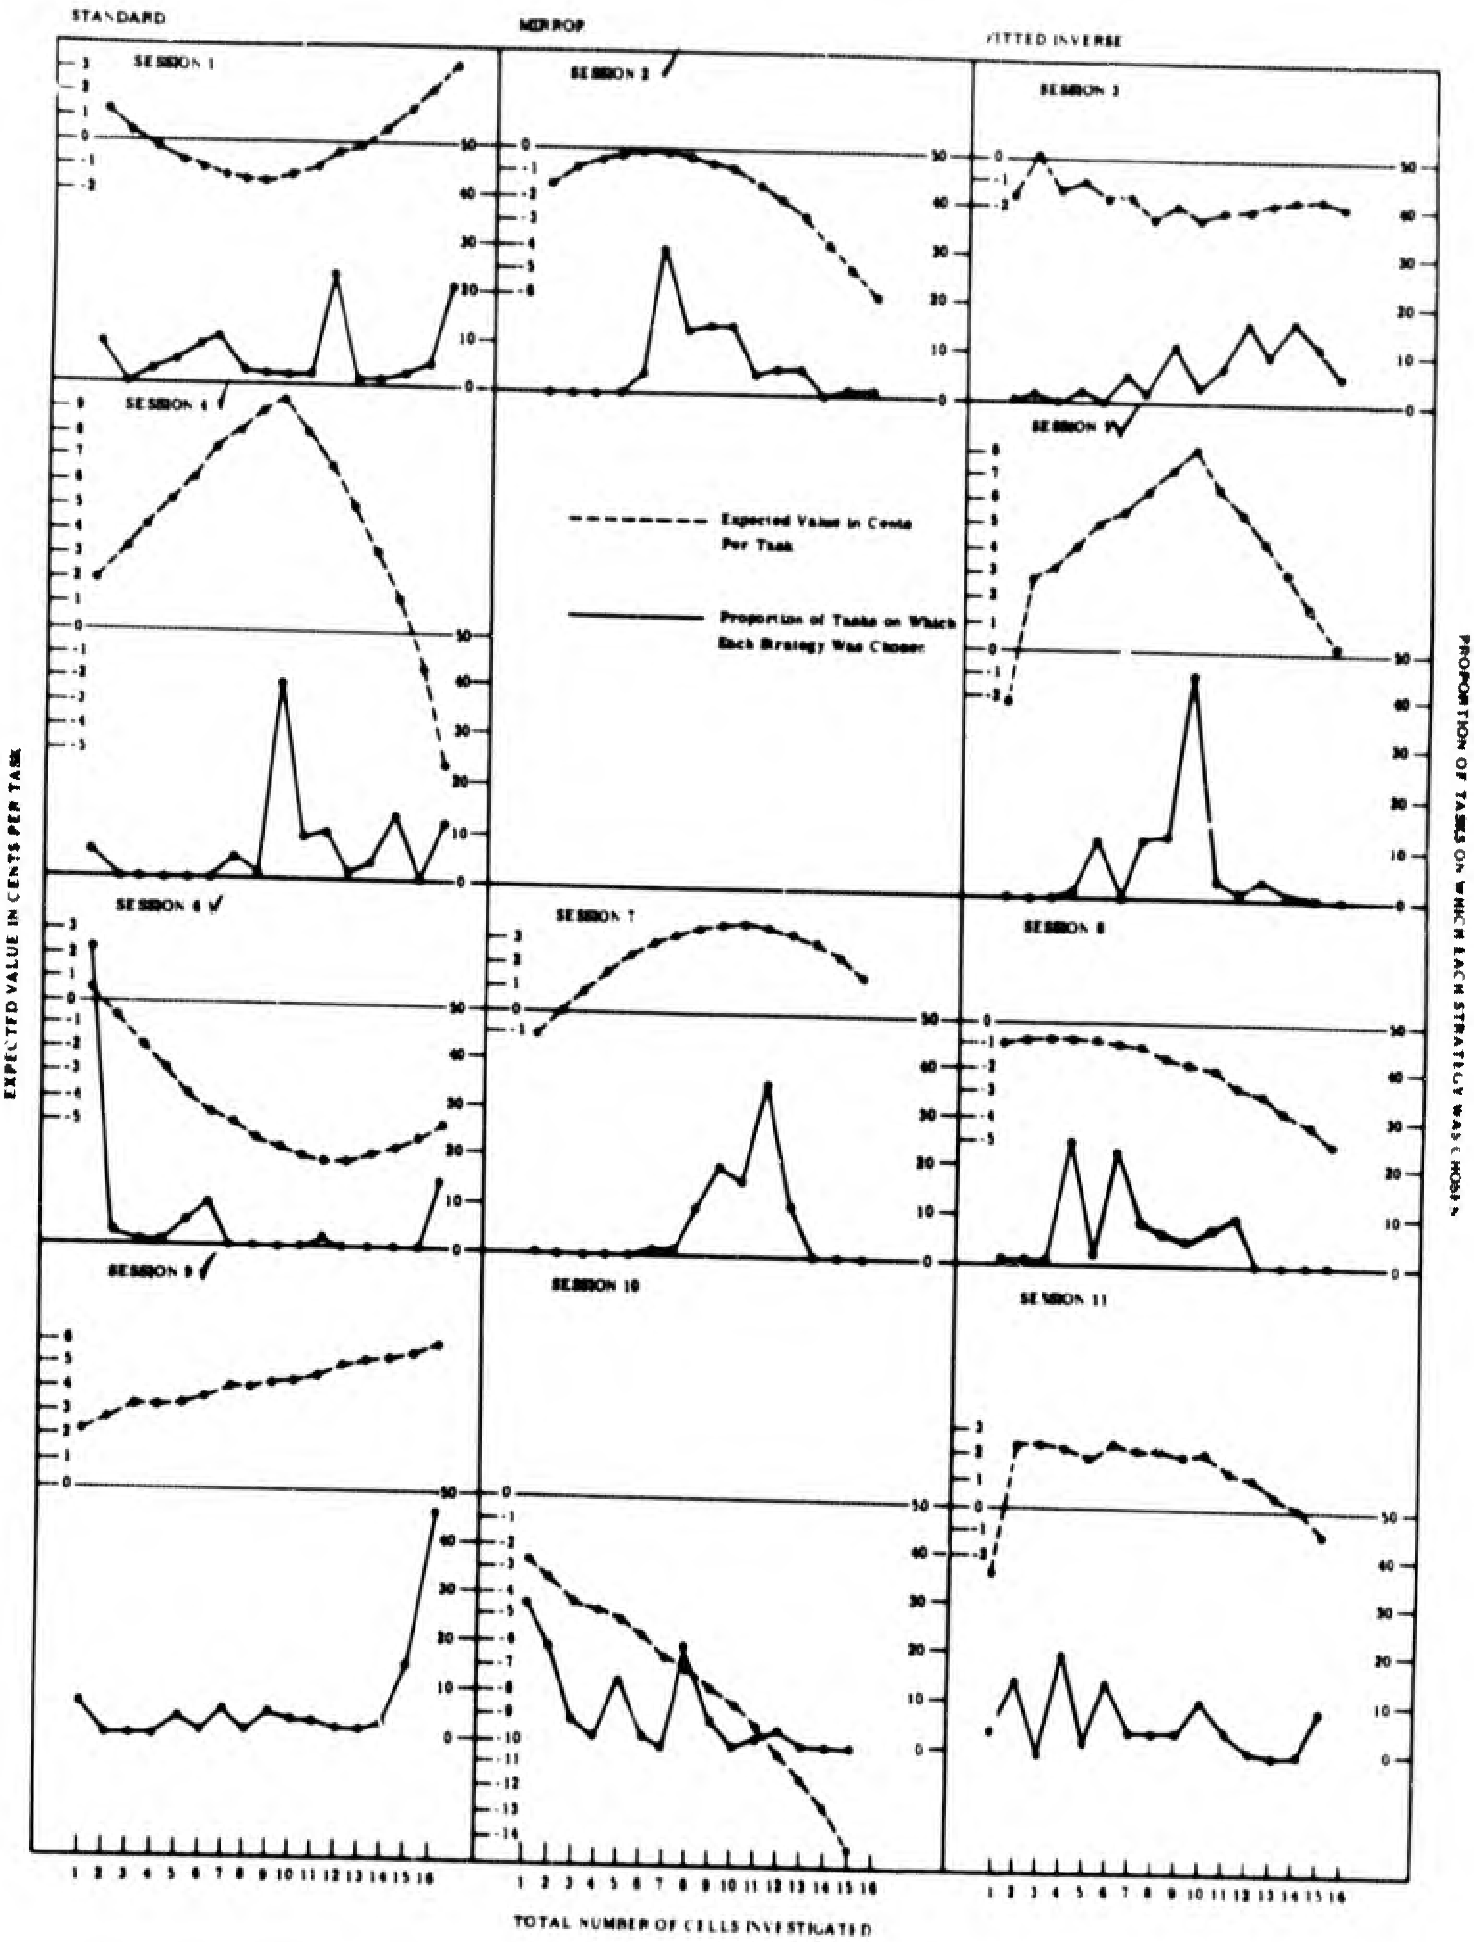

Figure 2. Expected Value of Each Strategy and Proportion of Choices of Each Strategy for Each Session. The Left-Hand Ordinate Is the Expected Value in Cents. The Right-Hand Ordinate Is the Proportion of Choices of Each Strategy

strategy changes as infrequently as the data permits you to believe.

With that idea in mind, we formulated an algorithm on how to perform extrapolations. The first rule of the algorithm says that in each task in which the unique cell is not found, the strategy score is equal to the number of the cells examined. The second rule says that in each task in which the unique cell was found, strategy score is equal to or greater than the total number of cells examined. Rule 3 is the crucial one. It says that within the restrictions of Rules 1 and 2, strategy scores are chosen so as to minimize the number of transitions from one strategy score to another within a sequence of 48 tasks, all having the same costs and payoffs. Now, that does most of the work but it doesn't determine all strategy scores, so we make two more rules. Rule 4 says that within the restrictions of Rules 1, 2, and 3, when a transition from one strategy score to another is necessary, it is placed as early in the sequence of 48 tasks as is possible. Rule 5 says that within the restrictions of 1, 2, 3, and 4, strategy scores are chosen so as to minimize strategy score. Strictly speaking, Rule 4 is not necessary. We could have done without it, but its virtue is that Rule 5 has a downward bias built into it while Rule 4 does not, and if you have Rule 4 in, then Rule 5 almost never operates and so you get rid of that bias.

One possible defect of this way of defining strategy is that Rule 3, which says that you should assume that as few changes in strategy occur as the data permit you to believe, may have the effect of reducing variance in strategy scores as compared with subjects' actual strategies. If that is the case, it ought to do so differentially, for high and low strategies. Rule 3 almost never applied for low strategies, because it only applies when the subject finds the unique cell, whereas it usually applied for high strategies when the subject almost always does find the unique cell. In fact, the correlation between the variance of strategy score and the value of strategy score is .002 and the scatter plot looks just as random as a .002 correlation suggests, so there is no effect of lowering variance at higher strategy scores as would be the case if Rule 3 was really reducing the variance compared with what it should have been. The lower curves of Figure 2, then, are in each case the distributions of subject's strategies over the 16 possibilities, summed over trials within a 48 trial block.

The first finding of importance is the remarkably high degree of optimality in some of the panels of Figure 2. For example, in session 6, the optimal strategy is never to buy any information, and well over 50% of strategies chosen were that exactly that optimal strategy. Similarly, in sessions 4 and 5, the optimal strategy is to buy about 8 items of information and that is just what subjects did. In session 9, the optimal strategy is to buy until you find the unique cell and that is just what they did, nearly 50% of the time. Sessions 2 and 7 show pretty good performance also. For sessions 10 and 11, performance was terrible. I think it's fair to say that on the whole these data exhibit quite good performance.

Figure 3 gives a more direct measure of performance. Figure 3 shows the percentage of strategies whose expected values deviate from the expected value of the optimal strategy by more than so many cents. Thirty-five percent of all choices were within 2 cents of the optimal strategy. This is good performance by almost any criterion.

Table 2 shows another measure of performance. Here, for each of the tasks, is the observed deviation summed over subjects compared with the deviation that would have been obtained by choosing randomly from among the strategies available. For all sessions except session 3, the observed deviation is substantially less. Some sessions are better than others. In general, performance was better on the Inverse tasks than it was on the Standard tasks, although there is clearly substantial inter-task variation. I think this is consistent with other findings in the decision literature which indicate that people take a far more realistic and calculating attitude when facing the possibility of large losses than they do when facing the possibility of large gains.

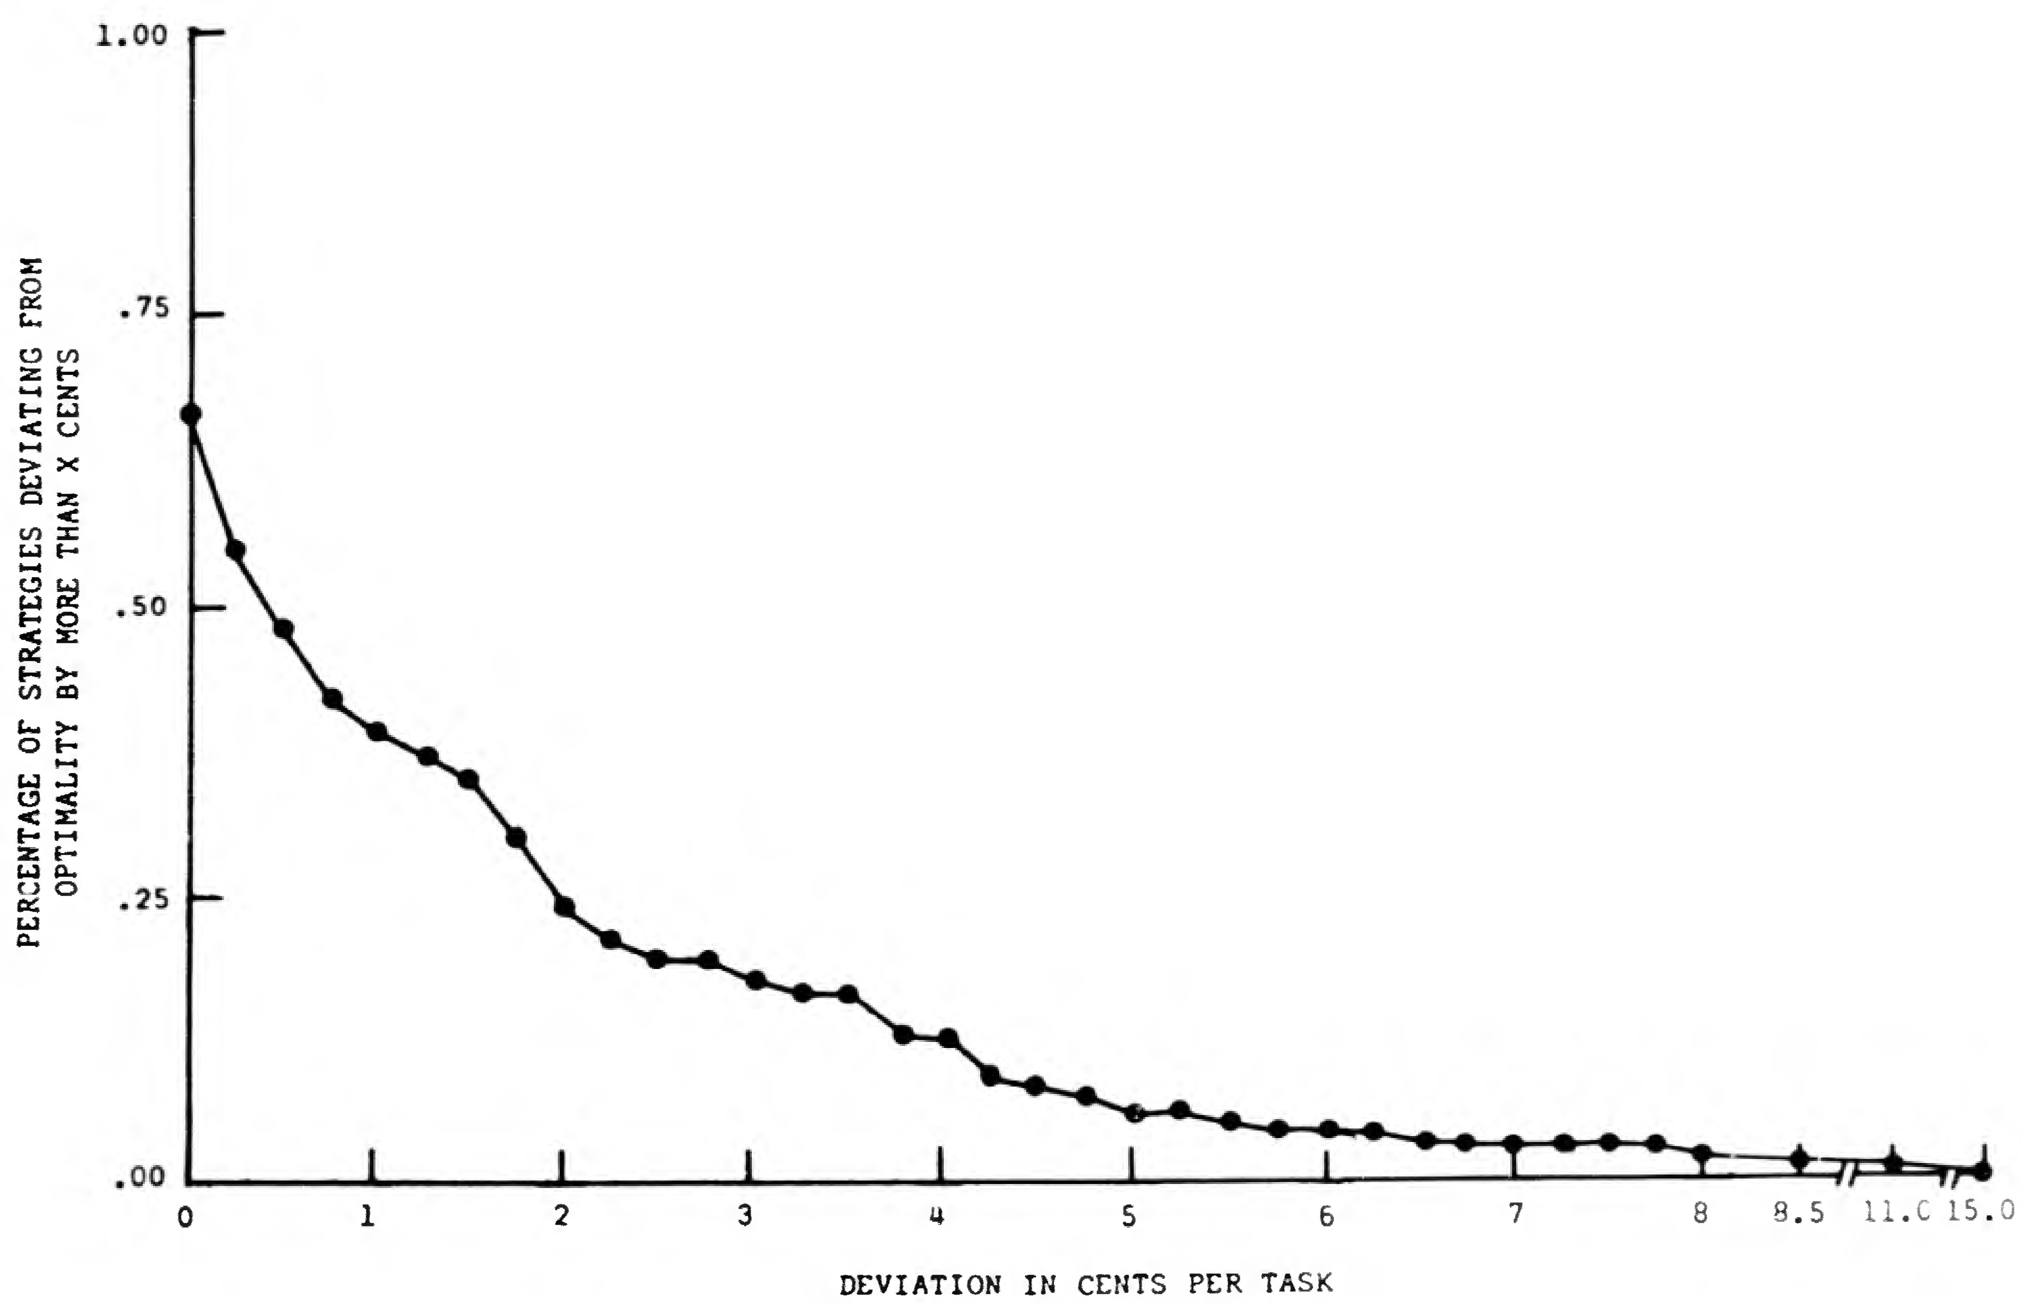

Figure 3. Cumulative Proximity to Optimality as a Function of Size of Deviation in Cents

Table 2

Mean Deviation from Optimality in Each Session

Session No.	Task Type	Observed Deviation in Cents per Task	Random Choice Deviation
9	Standard	0.76¢	1.72¢
6	Standard	1.86	5.00
1	Standard	2.67	3.12
4	Standard	4.15	5.23
7	Mirror	0.13	1.29
2	Mirror	0.69	1.54
10	Mirror	2.07	4.87
8	Fitted Inverse	0.51	1.21
11	Fitted Inverse	0.97	1.32
5	Fitted Inverse	1.18	3.84
3	Fitted Inverse	1.86	1.66

Table 3 shows deviations from optimality by subject. All subjects do better than the random choice strategy would do, and some subjects do very substantially better than others. Notice that this table is in cents per task, which means the deviations are fairly small. For some tasks there are two possible ways to deviate from optimality. Either the subject can be too cautious, which means that in the Standard Game he takes too many looks, and in the Inverse Game that he takes too few, or he can be too incautious, which means in the Standard Game he takes too many. These possibilities exist for sessions 2, 4, 5, and 7. For these four sessions, we calculated for each subject a cautiousness ratio, the ratio of the number of cents lost by being too cautious to the total number of cents lost. Subjects who are cautious on Standard Games are also cautious on Inverse Games, and subjects who are incautious in the Standard Game are incautious on the Inverse Games. As you see, the cautiousness ratios show a large range, from a high of .98 to a low of .12. The numbers are bimodally distributed; five subjects are cautious and five incautious.

We can look separately at how the cautious and incautious subjects perform in two classes of tasks, one the neutral tasks, being those tasks in which it is possible either to be cautious or to be incautious; the other tasks favoring cautiousness, that is those tasks where caution was the optimal strategy. There were not enough tasks favoring incautiousness for us to be able to look at that.

Table 4 is an analysis of variance of deviations from optimality. The blocks main effect is quite weak, and means that subjects do learn a little within each session. The important point is the substantial sessions by subject interaction. Table 5 shows the deviations from optimality for the two kinds of tasks. On the neutral tasks the incatuious subjects do better than the cautious subjects. Both do better than a random choice strategy. This, incidentally, reflects the fact that

Table 3

Mean Deviation and Cautiousness Ratio for Each S

Subject No.	Mean Deviation From Optimality in ¢/task	Cautiousness Ratio
1	0.50¢	0.
2	0.90	0.96
3	1.34	0.91
4	1.45	0.85
5	1.47	0.37
6	1.62	0.17
7	1.77	0.91
8	1.78	0.47
9	2.03	0.12
10	2.47	0.98
Mean	1.53	0.57
Random Choice Strategy	2.82	0.52

incautious subjects are in some absolute sense better decision makers. It has often been asserted, and I fully believe it, that in the college population from which these subjects are drawn, that undue cautiousness is much more common than the reverse. Consequently, in general, you would expect more realistic subjects would also appear the less cautious. However, on tasks favoring cautiousness, it is nice, though not surprising, to discover that cautious subjects do better than the incautious.

In this very complicated task, in which I don't think anyone could do much of a job figuring out the optimal strategy in his head, subjects perform rather well. Strategies, if they were not optimal, were very nearly optimal. There are some tasks in which it is pretty obvious how the subject managed to do it. The finding about cautiousness versus incautiousness indicates that when subjects do deviate from optimality, these deviations are often orderly and consistent. There are consistent individual differences in approaches to decision making, perhaps this experiment taps one of them. If it does, and if it could be shown to tap it reliably, tasks of this sort might be useful selection devices.

This task is far from ideal as a simulation of real life information-seeking. For one thing, in real life, information is not equally likely to be wherever you might look, and for another thing, in real life, information, when you get it, is not guaranteed to be valid. Perhaps most important of all, in real life, information has intellectual content; different items of information hang together to make a picture. That is not the case with this experiment and couldn't be the case with one like this. That is obviously the heart of the problem of information processing.

Table 4

Analysis of Variance of Deviations from Optimality in Dollars (not in Cents)

Source of Variation	df	Sum of Squares	Mean Square	F
Sessions	10	6461	646.14	3.81***
Blocks	3	61	20.48	3.36*
Subjects	9	1437	163.70	NS
Sessions x Blocks	30	238	7.92	NS
Sessions x Ss	90	15280	169.78	27.81***
Blocks x Ss	27	107	3.98	NS
Sessions x Blocks x Ss	270	1648	6.10	9.49***
Within Cells	4840	3115	0.64	
Total	5279	28384		

* $p < 0.05$

***$p < 0.001$

Table 5

Deviations from Optimal Strategy in ¢/task for Cautious and Incautious Ss on Neutral and Caution-Favoring Tasks

	Neutral Tasks (Sessions 2, 4, 5, 7)	Caution-Favoring Tasks (Sessions 6, 8, 9, 10)
Cautious Ss	2.19	0.91
Incautious Ss	1.01	1.76
Random-Choice Strategy	2.98	3.21

Bayesian tools give you a handle on some aspects of it.

At any rate, even in a task as complicated as this, it seems to be possible to ask the question: What were the subjects doing? and to give the answer: They were doing the best they could, and that best was pretty good.

REFERENCES

1. Becker, G. M., "Sequential Decision Making: Wald's Model and Estimates of Parameters," Journal of Experimental Psychology, Vol 55, pp 628-636, 1958.

2. Bellman, R., Dynamic Programming, Princeton University Press, Princeton, 1957.

3. Edwards, W., "The Theory of Decision Making," Psychological Bulletin, Vol 51, pp 380-417, 1954.

4. Edwards, W., "Behavioral Decision Theory," in Annual Review of Psychology, Vol 12, pp 473-498, 1961.

5. Edwards, W., "Dynamic Decision Theory and Probabilistic Information Processing," Human Factors, Vol 4, pp 59-73, 1962.

6. Howard, R. A., Dynamic Programming and Markov Processes, Wiley, New York, 1960.

7. Irwin, F. W., and W. A. S. Smith, "Value, Cost and Information as Determiners of Decision," Journal of Experimental Psychology, Vol 54, pp 229-232, 1957.

8. Kullback, S., Information Theory and Statistics, Wiley, New York, 1959.

9. Neimark, E. D., "Information Gathering in Diagnostic Problem Solving: A Preliminary Report," Psychological Record, Vol 11, pp 243-248, 1961.

10. Pruitt, D. G., "Informational Requirements in Making Decisions," American Journal of Psychology, Vol 74, pp 433-439, 1961.

11. Raiffa, H., and R. Schlaifer, Applied Statistical Decision Theory, Division of Research, Harvard Business School, Boston, 1961.

12. Savage, L. J., The Foundations of Statistics, Wiley, New York, 1954.

13. Schlaifer, R., Probability and Statistics for Business Decisions, McGraw-Hill, New York, 1959.

14. Schlaifer, R., Introduction to Statistics for Business Decisions, McGraw-Hill, New York, 1961.

15. Swets, J. A., W. P. Tanner, Jr., and R. G. Birdsall, "Decision Processes in Perception." Psychological Review, Vol 68, pp 301-340, 1961.

BLANK PAGE

DISCUSSION OF EDWARDS' PAPER

SHUFORD (Discussant). My comments are mainly questions, not with what you are saying, but that arise in connection with a research idea. The first thing that strikes me is that we are discovering more and more that subjects do pretty well. Maybe they maximize expected utility, particularly as we get into more complicated situations. On the surface, these results disagree with the results of betting experiments. I think this needs some discussion, and I'm not sure what is happening. I would like to suggest one possibility, that in betting experiments, most of them have different bets of the same utility. In other words, when you use objective probabilities (which are what you think they are, I guess), and actual monetary values, the expected utilities of the bets are all the same. When you use these two functions, which can be thought of as some function of monetary value, and subjective probability, which also might be close to objective probability, to write the subjectively expected utility of different bets, these bets are very close in terms of subjectively expected utility. Decision theory, in singling out the two major variables of utility and subjective probability to explain behavior, minimizes the difference due to utility and subjective probability, so to speak, gets down into the "noise level" of the subject's behavior and not surprisingly other factors come to determine the choices of the subjects. This is one possible explanation. I'm sure there are others, but it puzzles me that as we get into more complicated situations, subjects do better than the very basic tests of the axioms of the decision theory.

Another thing I'd like to hear discussed is the interpretation of cautious and incautious behavior. One explanation, I would guess, is that the cautious and the incautious subjects have different utility functions. The cautious subjects are winning money, and winning a certain sum means more to them, or has a greater utility for them, than not losing or paying out a certain sum, and just the opposite for the incautious subject. Maybe this would account for this cautious-incautious behavior observed here. Now, if this is so, this implies a more limited (although still interesting) general personality explanation of cautious and incautious behavior. This suggests that when you're dealing with money, this happens, but that in non-monetary situations where the cost of information is in terms of thinking or waiting, time delay, something like that, the same cautiousness and incautiousness may not show up. On the other hand, it may be that cautiousness and incautiousness are general things in the sense that the subject is using a generalized strategy, an approximation that he has learned to adopt to all sorts of behavior. This generalized strategy for the cautious subject means that the subject tries to get all the information he can and goes too far. The incautious subject is more inclined to take risk. This is quite general whether or not monetary or other types of payoffs are used.

The third question is this. I think most of us believe that as we move to more and more complicated situations, that the deviation between ideal and actual performance is likely to get bigger. I would like to suggest that perhaps by more complicated behavior we mean behavior which requires more complicated mathematical techniques to describe them. While we all seem to feel that the deviation between ideal and actual performance is likely to increase, I would like to suggest that in actual fact, this may not be happening.

This is a funny phenomenon. Take the computer engineer in contrast to the psychologist. The psychologist's point of view of the psychologist is that a human being does not run a maze as well as a rat; he does pretty poorly in various psychophysical and other experiments; he is not very much of an optimal individual. We tend to run down the individual. But when you consider the computer engineer trying to build complex electronic systems, you see a different picture. The computer engineers wish that they could build a system which did halfway as well as the human being. With one approach they are simply trying to find out what techniques human subjects are using and try to get the computer to do the same thing. Another approach is Bionics in which the circuits themselves are based upon neurons, and are designed to function roughly like the brain. On the one hand, we see psychologists who have the impression that human performance is not optimal, pretty inefficient. On the other, we have the computer engineers who are impressed with human performance. In psychology, it comes as a discovery that as we get to more complicated tasks, that humans do better and better, and seem to do more poorly in simple tests than in the complex tests. Well, these are some of the things that puzzle me and that I would hope to be able to discuss.

EDWARDS. I don't think I can provide the answer to any of these questions, but I can at least say a word about them. Emir asked about the inconsistency between the simple choices of betting experiments, where the impression has been widespread that rationality does prevail, and the more complicated situation in which it does, but he then answered his own question and I'm in full agreement with his answer.

It is perhaps appropriate to add that the very simple principle of maximizing expected value is capable of accounting for at least 95% of all real choices ever observed, but we knew that before we ever went into the laboratory and so we don't study those choices. If you offer a person a choice between a 50-50 chance to win a dime with no loss otherwise and a 40-60 chance to win $10, he is certainly going to take the 40-60 chance to win the $10, and I'm not about to run an experiment just to verify that belief. I think the general aura, "Well, man isn't so darn rational after all," that came out of the whole static bets literature is deceptive. I guess I'm a contributor to that particular kind of deceptiveness. At any rate, I now believe that in static as well as dynamic situations men are quite rational and the only reason why we didn't realize it before is because we designed experiments which made it very hard for the fact to exhibit itself.

Emir also raised the question of whether cautiousness versus incautiousness should be thought of in terms of the forms of utility functions. Alternatively, cautiousness versus incautiousness might be associated with subjective probability functions. To my mind this is more attractive than the utility interpretation. In fact, I speculate that there is no profit whatever to be gained from thinking about utility functions for money that deviate from linearity within the range of amounts of money that we are accustomed to dealing with daily. If there is one thing that we have learned more thoroughly than anything else by the time we get old enough to be college sophomores and experimental subjects, it is the value of money, or at any rate of small amounts of money. For that reason, explanations of the results of gambling experiments based on the belief that our utility functions are curvilinear rather than linear seem unattractive. I prefer to look primarily at subjective probability functions rather than utility functions to explain gambling experiment results. It certainly is entirely possible that cautiousness versus incautiousness is exactly a subjective probability sort of dimension. It would be extremely easy to write subjective probability functions to explain the data I presented. On the other hand, I don't know how much help that would be in pursuing the practical task of measuring cautiousness as a trait of people, because to find general enough subjective probability functions is difficult. I have just about given up on it.

ATKINSON. I think that Dr. Edwards' introductory remarks today indicate that research on decision making has reached a degree of maturity that is very encouraging. When he pointed out that this research has to be viewed in the context of the more general trends and fluctuations of behavior as the culmination of past experience and information-seeking, and so on, and that decision processes have to be viewed in a larger context, he said, in effect, that decision making research has gotten to the point where it can begin to look around and can profit by what experimental psychologists now demand of research and have been interested in for years. For example, I think I see a clear analogy between all of the research on animal behavior which has attempted to relate decisions at a choice point of information-seeking and to a confirming or disconfirming process during a series of learning trials. I want to call attention to this because I think it means that research on decision process is now branching out into the study of learning processes, information-seeking processes, concept development, and personality differences.

Now, let me direct a comment to the more puzzling question Dr. Shuford raised. He called attention to places in the decision literature where the optimal strategy isn't performed. I think the question that someone from outside of decision theory has to raise is this: Do the theorists in the decision making field intend to have the concepts of probability and utility embrace all of the various influences that determine the immediate action of the person? If they do, then clearly, everything that occurs in an experiment of this sort has to reduce finally to probability and utility. If they don't, then in a sense they are saying that they don't believe that the principle of probability times utility, is to be the simple principle of how actions are determined, but rather that decision making theory will have to fall into a context of a broader and more general theory of behavior which does account for these other variables. As for the question of whether we need more complex mathematical concepts, or whether this simple one is useful, I think my position would be, at this point, that this simple conception will turn out to be useful. The implication here is not that we need more complicated mathematics; rather, we need a more complicated and insightful orientation in the approach to the experiment. I would raise this question: Is the utility of the outcomes for the cautious and incautious subjects to be linked solely to the monetary price in this experiment, or do we imagine that the subjective satisfaction to be attained has many different sources? I think decision making theorists propose that the concept of utility should be used to map any kind of satisfaction, and in my own research on achievement motivations I have taken a clue from this. So the question I would raise for Dr. Edwards would be: What other qualities can be imagined to attend to the winning or not winning of the monetary prize in this experiment? Can the cautious and incautious differences that are observed be reduced to some people feeling more ashamed or more happy or more surprised when they win a certain amount of money under certain conditions, that is, with higher or lower probabilities of winning it? I'd like to hear him comment on the more general question: Does the decision theorist intend that the concept of utility should embrace all different sorts of satisfactions that might enter in? If his answer is yes, then I think the question is how we sensitize these experiments that manipulate only money as the major experimental source of utility to these differences.

EDWARDS. I guess the only reply to Jack's comment that he is glad to see decision theory rejoining psychology, catching up with what the psychologists have seen for 20 years as the important problems, is to bow our heads and ruefully say, "Yes, that's right." Decision theory is very late in coming to a serious concern with the sequential aspects of behavior which have been the heart of learning theory.

Now for your other comment about utility. It should be possible to order the different levels of any source of value, that ordering of levels of value should be transitive, and that there probably should be quite a number of steps on it. None of these are axiomatic properties of utilities; these are what people who think this way about values find it comfortable to think about utility. In thinking about any decision problem, different sources of value must be brought into relationship with one another, and must be examined for their joint over-all effect. Hopefully, that over-all effect will also be ordered and transitive. Utility is a name for the problem of conflicting value systems. It is also a name for the belief that the problem is a soluble one.

You raised the specific question whether things like achievement motive and fear of failure can be incorporated into the concept of utility. Not only can but should, I think; but to say that is to say nothing more than, "Let's try." It doesn't say we know how to do it.

LANZETTA. I am curious whether in your functions showing the ideal strategies, if it is true that where you have a maximum that you run into a transition point where the subject actually is investing more money than he can get back? Would the simple notion that he does not take more information than what he can possibly gain account for an awful lot?

EDWARDS. There are two panels out of the eleven for which that notion is adequate; for the others it is not.

FLOOR (K. V. Wilson). I wonder if it might be profitable to look at your results about cautiousness in terms of differential weighting of gains and losses, assuming linearity of utility with money, with a point of discontinuity at zero. It may well be that some of the cautious subjects have attached a good deal of negative utility to losses. I also wonder if cautiousness may not be a sort of desire for some kind of cognitive simplicity, as for example, in the Robinson tracking experiment. Are you possibly getting this second kind of cautiousness: some subjects simply assume that there is a stationary process generating the light probabilities, and the best strategy is to average over-all the trials, versus another kind of subject who says, "Well, there's some kind of probability the probabilities will change. If I notice a big discrepancy, I should ignore the earlier information and act on the recent information."

EDWARDS. There are no representatives of the first population; everybody is in the second population.

FLOOR (Wilson). With respect to the combining of values, or the combining of various kinds of utility, I wonder if here also there may not be some differences in cognitive complexity. Some subjects make a lexicographic ordering, saying, "I must avoid risk, period," while others are willing to take a little bit of risk in order to gain values along other dimensions; or possibly even some subjects are willing to take less of a greater value in order to satisfy a great many subordinate values. I think here we can see some very interesting examples of different cognitive styles in weighting decision making situations.

TODA. Dr. Edwards stressed the importance of looking upon sequences of decisions rather than separately. I would further stress the hierarchical structure of the decision. Sometimes a decision is concerned with a very objective, real situation. But there are other kinds of decisions such as, for example, how to make decisions, decisions about decisions. I think perhaps cautiousness-incautiousness has something to do with this: perhaps the subjects are evaluating their decisions in terms of success and failure, the cost of making errors, the cost of failures, and so forth. Of course, there are many idiosyncracies, but this might be a possibility.

ZINK. Ward, in another report of part of this study, you mentioned that during pre-testing, you discovered that the initial amount of money that the subject had was rather important. Did you notice any variations in the strategies of the subjects as a function of the total amount of money that they had in the bank at a given time?

EDWARDS. The answer to that question is no, I did not notice any such effect. This experiment was not, however, ideally designed for examining that question. However, as it happens, Sarah Lichtenstein has just finished her Ph.D. thesis using a static choice situation, addressed to the over-all problem of variance preferences, in which she did an extremely thorough and detailed job of asking exactly that question. She looked at it in a number of ways and I'm delighted to report that she found just about nothing. She had designed her experiment in the tradition of static business experiments to try to minimize the effect of current financial status, and she was very successful in doing just that. On the other hand, the thing I found in my pre-test, is that it does create a discontinuity if the subject uses all his poker chips and has to buy some more. The main thing I had to do was to provide him with a sufficiently large supply so this would not occur. I had earlier underestimated the amount of potential gain or loss during the course of an hour in this experiment, and it was necessary to discard some pre-test subjects who might otherwise have gone on in the main experiment.

FLOOR (M. J. Warrick). Dr. Edwards, would you explain more fully just what is the information-seeking act in your experiment, and what is the decision?

EDWARDS. The information-seeking action is easy enough to say: it is the process of looking into a cell before deciding to stop and say, "This is the unique cell." The decision is somewhat fictional, really, because it is the final designation of the unique cell. Of course, in the inverse game there isn't any decision, and you might think of this as being, in a sense, an information-avoiding task, rather than an information-seeking task. Of course, that is wrong. It is a unique cell avoiding task; in a sense you are acquiring information. Perhaps you are getting at the fact that the subject can't learn anything from one task that he can carry to the next; that is perfectly true.

FLOOR (Warrick). Within any one 16-cell matrix, if the subject punches a hole and his pencil doesn't go through, how does this affect his subsequent behavior?

EDWARDS. It changes the probability he will find it on the next punch. The first time he has one chance in 16 of finding it, and on the second, 1 in 15, and so on.

FLOOR (Warrick). Since you give the subject the cost and payoff structure in advance, couldn't he do it all at once, not sequentially? Are there any sequential effects at all? It's sort of clear that he isn't gaining information; by his punch each time, he omits so many subsequent punches.

EDWARDS. I agree, except I don't think they do it that way, really. In the inverse task, the subject says, "Do I dare do another?" In the standard, he says, "Can I afford to do another?" I do think they probably make a sequence of decisions rather than make a single decision. Nevertheless, the over-all point you are driving at is completely correct, namely, this task is not a good simulation of real world information-seeking. Compared to the real world, information is sparse in this experiment and the decision to which it is relevant is, in a sense, pretty unimportant relative to the information-seeking process itself. I think that if one is going to do more experiments on information-seeking, both of those characteristics of the task ought to be replaced by characteristics which sound more realistic.

FLOOR. Could the nature of the task itself, that is, the punchboard, influence the subject's decision? It could be a boring kind of game; some people like crossword puzzles, others are bored stiff.

EDWARDS. I have never had a subject yet, in a game involving real transfers of money which may amount up to $15 or $20 over the course of an hour, who was bored. The punchboard is no more the game than the roulette wheel is the game of roulette. The game is the money.

BELIEF STATES, EVIDENCE, AND ACTION*

Thornton B. Roby
Tufts University

Although there are probably no clear lines of demarcation between decision making behavior and garden variety behavior, there are two features that seem to characterize decision making processes as ordinarily encountered. One of these is the presence of juncture points beyond which behavior or actions cannot be continuously adjusted to environmental feedback. Thus, decision making, as contrasted with primitive behavior, generally entails an action commitment that extends well into the future (of Roby, ref. 6).

A second characteristic feature of decision making is the fact that the reaction is to an inferred or constructed environmental situation rather than to immediately present stimulus cues. This aspect of the decision making process will be the central topic of the present paper.

The approach to be suggested got its initial impetus from an attempt to conceptualize the effective information of small task oriented groups or teams. A natural objective of this attempt, which is reflected in the present paper, is to find a mode of representation that is useful for the purpose of studying information seeking and processing in either individuals or groups and which remains valid, though perhaps not as directly cogent, if it is extended to statisticians or amoebae.

The chief purpose of this paper will be to develop a catalogue of the principal defining characteristics of information gathering and processing systems. The basic procedure for identifying these characteristics or properties will be to examine a comparatively abstract and simplified epistemic system -- namely, the sequence of probability estimates based on Bayesian inference from a successive sampling scheme.

The resulting set of properties will have a rather heterogeneous form. Some of them are purely descriptive whereas others are expressly normative -- that is, they specify "rational" reactions to evidence. Some of the properties to be identified will be either present or absent in particular cases, while others may

*Partial support for the theoretical development discussed in this paper, and for related empirical research was provided by the National Science Foundation under Grant # G-10947; by the Office of Naval Research under Contract #94 (15); and by Contract AF 19(628)-2450 with the Decision Sciences Laboratory, Electronic Systems Division, USAF.

be present in various degrees. Finally, some of the properties attach to the situation within which an organism is operating; others attach to the organism; still others describe the relations between the organism and the situation.

The ultimate intention of defining all these properties is, however, identical. It is to serve as the basis for a set of measures that can be used to characterize, and perhaps to evaluate, the observed behavior of persons or groups of persons in real task situations.

Turning to substantive matters, the central entity to be discussed is the *belief state* (*B-state*) that an organism (*O*) maintains with reference to some specified environment, and which is used as a basis for action selection. Typical environmental referents might be: a) the deployment of enemy forces in a military situation; b) the particular malfunction that renders a machine inoperative; and c) the current trend of certain issues in the stock market. The chief functional properties of *B-states* to be discussed are the way in which they may react to external evidence and the way in which they are reflected in action decisions or in further information seeking processes.

REPRESENTATION AND SEMANTIC PROPERTIES OF B-STATES

It is assumed that, generally speaking, ***B-states*** are referred to a real environment and that *O* uses the dimensional structure or category system of this environmental space upon which to construct the ***B-states***. This assumption risks the philosophical heresy of naive realism, but this is not an imminent danger at the present level of discourse. For the most part we are concerned with information seeking processes in which ultimate verifiability is at least theoretically possible.

A very simple sampling problem will help to give definiteness to further remarks. Suppose the physical situation to consist of three drawers, ***U, M, L,*** each of which is known to contain 2 tokens. In one drawer the tokens are both white; in one drawer there is one black and one white token; and in the third drawer they are both black. The problem is to determine which of the drawers is associated with each of these compositions. The space of interest is then the set of 6 permutations of 0,1,2 white tokens over the ***U, M, L*** drawers, respectively. Holding the latter in fixed order, we can conveniently describe the points of the space by:

$$\begin{aligned} E_1&: (0\ 1\ 2)\\ E_2&: (0\ 2\ 1)\\ E_3&: (1\ 0\ 2)\\ E_4&: (1\ 2\ 0)\\ E_5&: (2\ 0\ 1)\\ E_6&: (2\ 1\ 0) \end{aligned}$$

In this notational system, E_4, for example, is the state of the world in which there is one white token in the upper drawer; two in the middle drawer; and none in the lower drawer. The six ***E-states*** are mutually exclusive and, in view of the information given, exhaustive.

Modality. Assuming a fixed finite substrate environment, consisting of distinct environmental points, there are two rather different forms in which ***B-states*** may be expressed. In the ***categorical*** form, the ***B-state*** is defined in terms of the subset of *E-states* that are deemed possible at the time of specification. Thus every element of the *B-state* space is a member of the power set of the *E-state* space.

Examples based on the three-drawer problem would be $\{E_1\}$, $\{E_1, E_3\}$, or $\{E_1, E_2, E_3, E_4\}$. The first means that only E_1 is accepted as possible; the second means that either E_1 or E_3 may be the case; and the last one admits the possibility of the first four *E-states*. These *B-states* make no assumption about the relative probabilities of the component *E-states* -- for example, $\{E_1, E_3\}$ does not imply that equal likelihood is assigned to these two *E-states*, but only assumes that those two (and no others) are possible. Of course, only one *B-state* of the kind shown above may be held at any instant.

A second mode of expression (or representation) would assign a distinct probability to each *E-state*. This form will be referred to as the *vector* mode as it is described by a probability vector. Examples, based on the three-drawer problem, are B_1 = (.5, .3, .2, .0, .0, .0); B_2 = (1.0, .0, .0, .0, .0, .0); and B_3 = (.16, .16, .16, .16, .16, .16). In each case the probabilities sum to unity. B_1, for example, means that *O* assigns 50% of the total probability to E_1, 30% to E_2, and 20% to E_3. B_2 corresponds to certainty that E_1 is the correct *E-state*; and B_3 corresponds to complete uncertainty.

In practice the choice between modes may depend upon several factors. It may be determined by the nature of the substrate environment: for example, a relatively homogeneous environment with many *E-states* seems to lend itself best to the categorical formulation. Again, the choice of modality may be determined by the nature of the available evidence: the categorical mode is perhaps better suited to situations in which complete or exhaustive induction is possible. Finally, the choice may depend upon the kind of "game" that *O* thinks that he's playing; the categorical mode is safer in some ways if the environment or the evidence may be controlled by an antagonistic intelligence.

Since one of our objectives is to analogize from the situation in which rather strong Bayesian inferences are possible -- that is, from sampling situations -- we shall deal primarily with the probability vector mode.

Resolution. An obvious measure to apply to a probability vector of the type $B_1 = 1/4(1, 0, 1, 0, 0, 2)$ is an informational measure of entropy or uncertainty, in this case 1.5 bits. This might be compared with the baseline measure for a uniform distribution $\log_2 6 \approx 2.5$ bits, to give a measure of the negentropy or resolution of the *B-state*, in this case 1.0 bits.

An alternative index of resolution, more in keeping with the rest of the present development, is to identify the resolution with the variance of the probability components of the *B-state*. This index takes the value 0 for the uniform distribution and is maximal for a *B-state* assigning certainty to a single component. For the illustrative *B-state*, the index of resolution (using the estimate $\sum x_i^2/(n-1)$) is .25 whereas it is 1.0 for the *B-state* with a single unit entry.

It might be remarked that this index of resolution is highly correlated with the information theory index and it is closely related to the measure of information first proposed by R.A. Fisher (see Reza, ref. 5). Evidently the resolution condition is closely related to the market place idea of confidence in a judgment. It does not have any direct relation to either the logical consistency or the validity of a belief. The two properties that follow bear on these latter aspects.

Consistency and Compactness. The essential notion of internal consistency of a *B-state* that we want to incorporate is that the most probable alternatives envisaged by the *B-state* are mutually compatible. This seems to imply in turn that the information or structure reflected in the *B-state* is congruent with the primary environmental dimensions on which distinctions are being made.

To illustrate, the assertion that a certain public figure is either a saint or a scoundrel seems at first blush to be an internally inconsistent *B-state*. However, this impression is largely due to the assumption that our informant's judgment is made primarily on moral grounds. It might turn out, on the other hand, that the judgment was based on the great persuasiveness of the public figure, so that the moral judgment was secondary, and a comparatively minor readjustment of the *B-state* would be required for final resolution.

Again, in the three-drawer problem, it seems more reasonable to assign probability .5 to E_1 and E_2 respectively than to assign probability .5 to E_1 and E_6. The former divides the probability between *E-states* that are intuitively "close" to each other in the environmental space, while the latter divides the probability between highly discrepant *E-states*. Here too, though, it is possible to conceive of situations in which the E_1, E_6 *B-state* might be a natural one. For example, E_1 and E_6 are the only two states in which the middle drawer contains 1 white and 1 black token.

In order to translate this into numerical terms it is noted first that, although the 6 *E-states* of the illustrative problem are logically independent (in the sense that one and only one obtains at any time), they are substantively interrelated. That is, there is a relation between any two states of comparative similarity or proximity. Perhaps the simplest measure of similarity in this case is the vector product of the *E-states* when they are expressed in numerical form (e.g., $E_1, E_2 = (0, 1, 2)\,(0, 2, 1) = 4$). This measure, however, is essentially the same (up to a linear transformation) as the more familiar Pearson correlation coefficient, and the latter will be used here.

The first step in the proposed 'compactness' measure is therefore to compute the matrix of intercorrelation between *E-states*.* For the present example, this is:

$$Q = \begin{pmatrix} 1.0 & .5 & .5 & -.5 & -.5 & -1.0 \\ .5 & 1.0 & -.5 & .5 & -1.0 & -.5 \\ .5 & -.5 & 1.0 & -1.0 & .5 & -.5 \\ -.5 & .5 & -1.0 & 1.0 & -.5 & .5 \\ -.5 & -1.0 & .5 & -.5 & 1.0 & .5 \\ -1.0 & -.5 & -.5 & .5 & .5 & 1.0 \end{pmatrix}$$

---

*If this seems like an arbitrary procedure, note that the same table of intercorrelations is obtained if, instead of correlating the numerical components of the *E-state* description, one correlates the components associated with the *E-states* in the clue operators introduced below. This latter operation is the fundamental one: the *E-states* are related insofar as they imply similar clues.

In order to use this as a measure of compactness, Q is then used to set up a quadratic function on the *B-state* of the form $B\ Q\ B^T$ where B^T is a column vector, transposed from B. For example, if B = (.5, .5, 0, 0, 0, 0) the resulting value is .75; if B is (.5, 0, 0, 0, 0, .5) the compactness index is 0. A completely resolved *B-state* (probability 1.0 associated with any *E-state*) has a compactness index of 1.0; and the "neutral" *B-state* (equal components throughout) has the compactness index of 0.0.

This index of compactness appears to agree with intuitive notions of "consistency" in this situation. It has the additional virtue, though, of tieing in closely with the other two "semantic" properties considered in this section. First, note that it is a sort of generalization of the "variance" measure of resolution suggested above. If we set all off-diagonal elements in Q equal to 0, then the resulting quadratic form consists precisely of the sum of squares of the components of B. Thus the property of resolution is identical with compactness in an *E-state* space in which all *E-states* are strictly independent. Incidentally, this is necessarily true of a two-state space. We shall return later to the relation between these properties and that of validity.

Verdicality and Validity. The simplest index of the correctness of a *B-state* would be the degree of probability ascribed to that *E-state* which in fact obtains. This index does not, however, seem adequate for several reasons. First, it does not give credit for assignment of probability to *E-states* that are similar to the correct *E-state* though not identical. Thus, it would seem to underestimate the veridicality of *B-states* that are comparatively detailed and sophisticated. Perhaps more serious, it may yield the same value for a *B-state* that is relatively unresolved or fuzzy, and a *B-state* that is highly aberrant.

In order to take these considerations into account, it is necessary, as with the compactness notion, to use some concept of proximity. In fact, we can use the same Q matrix, only now it is postmultiplied by a vector B^*, which contains a single unit entry corresponding to the true *E-state*. As numerical illustrations, let the true *E-state* be E_2 and the *B-state* be (.5, .5, 0, 0, 0, 0). Then B^* is $(0, 1.0, 0, 0, 0, 0)^T$ and $B\ Q\ B^* = .75$. If, on the other hand the *B-state* is (0, .5, 0, 0, .5, 0), then $B\ Q\ B^* = 0$ -- that is, a *B-state* with 0 compactness, like the latter, cannot have positive validity under any circumstances.

QUERY PROGRAMS AND CLUES

This section will be concerned with the genesis and modification of *B-states*. This investigation is limited strictly to modifications which are due to external inputs or evidence. It does not deal with changes in *B-states* that may be a function of time alone. The most important aspect of this process of *B-state* modification is the interpretation of evidence, but this implies some auxiliary mechanism for obtaining the evidence in the first place. To avoid the double meaning of the word "observation", the act of seeking or receiving information inputs will be referred to as a ***query***, and the result of a query will be referred to as a ***clue***.

In the three-drawer problem described above, a query would be a draw with replacement, from any of the three drawers. The resulting clue would of course be the token, black or white, that happened to be drawn.

The query may entail a very active process of information seeking or it may entail only passive attending. It appears that the important characteristic is the predictability or controllability of the queries that are applied to the environment. Thus a query program may be *imposed* or *elective*. If it is imposed, it may be fixed or random: if it is elective, it may be elected on a "bulk" basis, or it may be subject to continuous readjustment.

For the three drawer example, a query program would be any sequence of draws from the upper, middle and lower drawers in specified order. An imposed program would be one over which *O* had no control. A fixed, imposed, program would be one, such as continuous cycling over the three drawers, which was not subject to choice by *O* nor affected by other contingencies. A random imposed program would be one in which the drawer from which the sample was drawn was determined by a chance process of any kind. A bulk elective program would be one in which *O* selected a preferred sequence of queries which was then followed without deviation or further adjustment by *O*. In a readjustable elective program, *O* can select queries one at a time, basing his choice on previous outcomes.

A second very general aspect of query programs is the extent to which queries interact with the substrate *E-state*. Several distinctly different cases may be described: first, where the query -- whatever its outcome -- leaves the basic *E-state* unchanged; second, where the query is associated with specific *E-state* changes; and, third, where the query affects the *E-state* in an indeterminate fashion.

Sampling with replacement illustrates the first case, and sampling without replacement is a typical example of the second case. A rather unrealistic example of the third case would be one in which replacement of successive samples would be determined by a chance process of which the outcome was unknown to *O*. The third case is actually the most prevalent one, however, in many real life epistemic processes -- e.g., testing a subject in a psychological study affects the subject to an unknown degree.

A final aspect of the query program that will here be considered is that of the *dependence* or *independence* of distinct queries. Two queries may be regarded as dependent if any outcome of one query affects the probability of any outcome of the following query. This relationship of dependence is presumably symmetric in most cases of interest.

For the three-drawer problem, each of the three queries is dependent on the other two. If a white token is drawn from the upper drawer it decreases the probability that a white token will subsequently be drawn from either of the other two drawers. Typically, but not necessarily, a query is self-dependent; any outcome within the query increases the probability that the outcome will occur again. The exception applies to intrusive queries such as sampling without replacement.

<u>An Operator Description of Clues</u>. In general, we consider the available queries in an epistemic process as families of potential clues, each clue having a defined effect on the existing *B-state*. It is convenient to represent these effects in operational notation, so that the results of a query, say α, may be the clues α_1 or α_2, ..., and the consequence of any of these clues α_i, is a change in the existing *B-state* -- for example, $B_i^* = \alpha_2 B_i$. Further comments can be most conveniently related to the example introduced earlier.

For the three-drawer problem, we have three possible queries that may be described as *U*, *M*, or *L*, depending on whether a single token is sampled from the top, middle, or bottom drawer. Each of these queries may have one of **two** outcomes *B* or *W*. Thus, the six potential operators on the existing *B-state* are U_B, U_W, M_B, M_W, L_B, L_W.

It is convenient, and as we shall see, justifiable to represent these operators as diagonal matrix multipliers. The entries along the diagonal are, in this case, the conditional probabilities of obtaining the outcome *B* or *W* if each of the respective *E-states* obtains. The resulting clue operators for the three-drawer problem represented as diagonal matrices are:

$$U_W = \begin{bmatrix} 0 &&&&& \\ & 0 &&&& \\ && 1/2 &&& \\ &&& 1/2 && \\ &&&& 1 & \\ &&&&& 1 \end{bmatrix} \qquad U_B = \begin{bmatrix} 1 &&&&& \\ & 1 &&&& \\ && 1/2 &&& \\ &&& 1/2 && \\ &&&& 0 & \\ &&&&& 0 \end{bmatrix}$$

$$M_W = \begin{bmatrix} 1/2 &&&&& \\ & 1 &&&& \\ && 0 &&& \\ &&& 1 && \\ &&&& 0 & \\ &&&&& 1/2 \end{bmatrix} \qquad M_B = \begin{bmatrix} 1/2 &&&&& \\ & 0 &&&& \\ && 1 &&& \\ &&& 0 && \\ &&&& 1 & \\ &&&&& 1/2 \end{bmatrix}$$

$$L_W = \begin{bmatrix} 1 &&&&& \\ & 1/2 &&&& \\ && 1 &&& \\ &&& 0 && \\ &&&& 1/2 & \\ &&&&& 0 \end{bmatrix} \qquad L_B = \begin{bmatrix} 0 &&&&& \\ & 1/2 &&&& \\ && 0 &&& \\ &&& 1 && \\ &&&& 1/2 & \\ &&&&& 1 \end{bmatrix}$$

Although applications of these operators to *B-states* does not immediately result in legitimate probability vectors, the vectors that do result can be set in direct equivalence with the legitimate probability vectors. All that is required is a scalar correction.

To illustrate, suppose that *O* begins with the *B-state* 1/6 (1, 1, 1, 1, 1, 1), that is, a uniform probability for all six *E-states*. He tests the top drawer and draws a black token. The result, using conventional matrix multiplication, is

$$1/6\ (1,\ 1,\ 1,\ 1,\ 1,\ 1) \begin{bmatrix} 1 &&&&& \\ & 1 &&&& \\ && 1/2 &&& \\ &&& 1/2 && \\ &&&& 0 & \\ &&&&& 0 \end{bmatrix} = 1/6\ (1,\ 1,\ 1/2,\ 1/2,\ 0,\ 0)$$

The sum of the entries in the resulting vector is 3/6 and, adjusting for this, we obtain the legitimate probability vector 1/6 (2, 2, 1, 1, 0, 0). These components are the same as would be obtained individually by a conventional application of Bayes' theorem. For example,

$$P(E_3|U_B) = \frac{P(U_B|E_3)\ P(E_3)}{\Sigma_j\ P(U_B|E_j)\ P(E_j)} = \frac{1/12}{1/2} = 1/6$$

The denominator term in this case is identical with the adjustment made in the probability vector. Substantively, it reflects the probability of the obtained clue in light of the preexisting *B-state*.

Balanced Clue Effects. Before considering what can be done with these matrices, the question arises as to where the entries would come from in the general case. Although, for strictly Bayesian processes, these coefficients are definitely prescribed by the appropriate conditional probabilities there is no reason to assume in the general case that the O knows what the exact probability will be of various clues, as a function of the substrate environment. We assume, however, that he does make some estimate to this effect and it is then desirable to impose a condition that relates this estimate to the operative effect of the clues. The suggested property of *balance* is based on the fact that the operators associated with a single query constitute a family and that the a priori probability of each clue in this family appearing on any occasion is defined by the existing *B-state*. Any clue resulting from a query will result in a change in that *B-state* and the less probable the clue, the greater the ensuing change. The query is balanced if the expected algebraic change in the *B-state* vanishes over all possible clues.

This says that while O must expect some change in the *B-state* following a query, his original *B-state* is the weighted center for all potential changes: he is cheating if he expects a clue to change his *B-state* in any particular direction.

For the three-drawer problem, O begins, say, with a uniform probability distribution over the six states. For query U, he may receive U_W or U_B. U_W leads to the *B-state* 1/6 (0, 0, 1, 1, 2, 2); the average of this and the *B-state* 1/6 (2, 2, 1, 1, 0, 0) which would follow U_B, is again the uniform distribution. If, however, O stops testing the upper drawer after receiving a black token and tests the second drawer, then the probability of receiving M_W given B_1 = 1/6 (2, 2, 1, 1, 0, 0), symbolized by $P(M_W|B_1)$, is 2/3 and $P(M_B|B_1)$ is 1/3. The *B-states* that would be produced by these operators are respectively,

$$B_1'|M_W = 1/4\ (1,\ 2,\ 0,\ 1,\ 0,\ 0)$$

$$\text{and } B_1'|M_B = 1/4\ (2,\ 0,\ 2,\ 0,\ 0,\ 0)$$

The combination 2/3 $B_1'|M_W$ + 1/3 $B_1'|M_B$ is again B_1 itself.

Associativity (Path Independence) and Commutativity. Operating on B_1 = 1/6 (2, 2, 1, 1, 0, 0) with the clue operator M_W yields the vector $B_1'\ M_W$ = 1/4 (1, 2, 0, 1, 0, 0). But B_1 is already the result of the applying U_B to B_0, the uniform distribution vector. The question arises as to whether these operations can be combined into a single step -- that is, whether $(B_0)\ U_B\ M_W = (B_0\ U_B)\ M_W = B_1\ M_W$.

For the Bayesian sampling problem this is indeed the case. In fact

$$U_B\ M_W = \begin{bmatrix} 1/2 & & & & & \\ & 1 & & & & \\ & & 0 & & & \\ & & & 1 & & \\ & & & & 0 & \\ & & & & & 1/2 \end{bmatrix} \begin{bmatrix} 1 & & & & & \\ & 1 & & & & \\ & & 1/2 & & & \\ & & & 1/2 & & \\ & & & & 0 & \\ & & & & & 0 \end{bmatrix} \begin{bmatrix} 1/2 & & & & & \\ & 1 & & & & \\ & & 0 & & & \\ & & & 1/2 & & \\ & & & & 0 & \\ & & & & & 0 \end{bmatrix}$$

and this latter matrix, multiplied by the uniform probability vector, clearly leads to the vector 1/4 (1, 2, 0, 1, 0, 0) which was obtained in a less direct way above.

A query program may be described as *associative* if the result of applying a sequence of clue operators $\alpha 1, \alpha 2, \alpha 3, \ldots \alpha n$ to an initial *B-state*, B_0, is the same as the result of applying the terminal sequence $\alpha_{k+1}, \ldots, \alpha_n$ to the vector, B_k, that results from applying the initial subsequence $\alpha_1, \ldots, \alpha_k$ to B_0.

It should be observed that this formulation of the associativity condition also implies "independence of path". That is, it suggests that, if B_k is the effective *B-state* before the terminal subsequence of clues is obtained, the final *B-state* does not depend on how B_k was originally arrived at.

From the definition of the operators for the sampling problem, it can further be seen that the order in which a sequence of operators occurs is immaterial. All diagonal matrices are automatically commutative if they can be multiplied at all. It can be ascertained directly that this condition does in fact hold for Bayesian sampling of this type. Accordingly, an epistemic process may be defined as *commutative* if the effect of a set of clue operators on any initial *B-state* is independent of the order of operation.

At this point it may be well to reemphasize that the properties described above, and those to follow, are not assumed to characterize flesh and blood epistemic processes in general. In fact the most interesting problems, from a psychological standpoint, might be related to specific departures from these essentially normative properties. For example there is evidence that everyday epistemic processes are not commutative at all -- that people may be unduly influenced by early information on a subject or by recent information. The psychologically fertile problem concerns the conditions under which either of these conflicting departures from the normative process may occur.

Asymptotic B-states. Given the properties of associativity and commutativity, it is legitimate and convenient to pool the effects of a sequence of clue operators into a single ***composite operator***. Thus if 10 successive U queries should produce 6 white and 4 black draws, this can be represented by the operator $U_w^6 U_b^4$. The components of this composite operator are of course calculated by termwise multiplication of the components of the individual operator.

This section will consider briefly the expected limiting or asymptotic *B-states* obtained after repeated applications of such composite operators. The central question is of course whether this limiting *B-state* will be the unit probability vector that assigns probability 1 to the *E-state* that in fact obtains.

The answer to this question depends upon two aspects of the reciprocal relation between queries and *B-states*. First, a favorable answer (i.e., the possibility of ultimate selection of the correct *E-state*) requires that the available set of clues be sufficiently discriminating to distinguish between the correct *E-state* and all other *E-states*. If E_i is the true *E-state*, this means that, for some composite clue, the probability component corresponding to E_i must be the unique maximal element.

In the three-drawer example, no single query has sufficient discriminating power to lead to an unequivocal asymptotic *B-state*. Thus the query U can only result in one of the three *B-states*, $(0, 0, 0, 0, p_5, p_6)$ $(0, 0, p_3, p_4, 0, 0)$ or $(p_1, p_2, 0, 0, 0, 0)$, in which the values of the non-zero components are determined by the initial state. On the other hand, combinations of two or more queries result in composite operators that discriminate among all six *E-states*.

The availability of a discriminating set of queries is not sufficient, however, if the initial *B-state* is not *open* to evidence with respect to the correct *E-state*. The *B-state* would be considered closed in these terms if the relevant *E-state* was not represented or, equivalently, was represented with zero probability.

To see what happens when the *B-state* is closed for the correct *E-state* consider the fate of the *B-state* vector 1/4 (0, 0, 1, 1, 1, 1) under repeated application of the query *U* if E_1 is in fact correct. The typical outcome under these conditions is of course a repeated black draw, represented by U_b. This results in the asymptotic *B-state* 1/2 (0, 0, 1, 1, 0, 0) which is, in this case, incorrect.

As a guard against this anomalous result, it is proposed to define the *hypercharacteristic B-state* corresponding to any composite operator as that closed *B-state*, B_y, (containing one or more zero probabilities) such that every neighboring open *B-state*, $B_y + \epsilon$ will return to the same limit B_y under repeated application of the composite operator. Notice, for example, that if a small probability vector, ϵ, is added to the false asymptotic value above -- e.g., (ϵ, 0, 1-ϵ, 1, 0, 0) -- the resulting limit vector under repeated application of U_B, will be (1, 0, 0, 0, 0, 0). In this case, a similar shift occurs if the second zero is perturbed, but this is because of insufficient resolving power of the single query .

Thus a necessary condition for convergence of an epistemic process to an asymptotic *B-state* that is correct, is that it converge to a unique hyper-characteristic vector. Presumably, this criterion would be applicable too for more complex epistemic systems than here considered.

One would like to extend this to accommodate "creative" epistemic processes, in which the actual dimensions of description are changed. This should occur if it turns out that obtained clues do not correspond to *O*'s a priori understanding of the *E-space*: in the example above, if R_W, M_W, L_W, was received, *O* would know that there was something wrong with his original understanding of the contents of the three drawers.

It is proposed to define an epistemic process as *serendipitous* if there exists a "reservation" space of finite probability attaching to a non-explicit set of outcomes.

The analogy here is to a mixed probability distribution which has finite probability for several events and a continuous probability over an infinity of other events. Strictly speaking these other states, considered a priori, have zero probability but if a sufficiently compelling clue occurs (that is, a clue that can only be due to certain previously unanticipated *E-states*) it reorganizes the space of potential *B-states*.

In coin flipping, for example, we may imagine that *O* attaches a probability (p(1-ρ), q(1-ρ), ρ) to the three events, heads, tails or unanticipated, where ρ is the probability, for example, that the coin lands on edge, stays in the air, or lands with some totally unforeseen obverse facing upward, ad infinitum. This can be most readily illustrated by a special numerical example. Suppose it is required to determine how many white tokens there are in an urn, given that there is exactly one black token, by drawing tokens one at a time with replacement. The a priori information is that there may be "0, 1, or more" with equal probability.

The explicit *B-state* vector has the form 1/3 (1, 1, 1). The "reservation space" refers to all integers from 2 to ∞. The total reservation probability is 1/3 but this cannot be assigned to specific elements of the reservation space in any definite way.

In this case, it is apparent that the *B-state* space will be transformed by successive clues in such a way as to insure final convergence on the correct *E-state*. Upon the first draw of a white token -- if it occurs -- the point E_0 (no white tokens) is annihilated. Then if there are, say, 1065 white tokens, the black token will appear every 1066th trial on the average. *E-states* with finite numbers of white tokens will be conjectured after the first black token appears, and the correct value will ultimately be accepted.

PRAGMATIC ASPECTS

The aim of this section will be to classify action choices and to examine their relations with the existing *B-state*. Two aspects of action choices will be stressed. The first is the "sharpness" of the required choice; and the second is the completeness or latitude in the set of available action alternatives. Both of these factors are assumed to have a bearing on the pragmatic value of particular *B-states*.

It is assumed that an action choice is represented by a set of actions, generically A_k, each of which may be applied under any *E-state* condition. For each *E-state* condition E_j, the action has a specified utility $u(j,k)$ which is known to O. Thus, to evaluate any action in the light of a particular *B-state*, B_i, O uses the vector product of B_i and the column of utility scores of A_k for corresponding *E-states*. He then selects the action having a maximal row-by-column product with his current *B-state*.

Action Vector Norms. This condition may be introduced by noting that real life decision situations fall rather naturally into two distinct categories. In some situations, the indicated action is selected by choosing the most likely *E-state* and going all-out for that choice, in other situations there are "hedging" or broad spectrum responses such that the choice of an optimal action can be tempered by various degrees of uncertainty.

A familiar military example of the former situation is the so-called "calculated risk". For example, if we must select a date, April 19th or April 26th, for an invasion, it is not ordinarily prudent to launch half an invasion on the 19th and half an invasion of the 26th. It is generally better to go all out in one direction, even if it may be somewhat the wrong direction. An example of the latter situation would be deployment of infantry troops along a defense line. If the precise point of attack is not known, it might be decided to deploy troops along a line in rough proportion to the expected probability of attack in various sectors. The critical difference between these situations is that: in the former case, the optimal action is based on maximum likelihood or modal probability; in the latter case, these are actions which are not optimal for any actual state of the world but are optimal for certain belief states.

In order to describe these various payoff conditions in parametric form, it will be useful to consider a two-dimensional space, such as that in figure 1. The axes of this space are the possible *E-states* of the system, E_1 and E_2 and the interior of the space is used to represent both *B-states* and action utilities. In figure 1, each point on the diagonal line from (0, 1) to (1, 0) is a possible *B-state* and each point in the square is a possible action-utility vector. Thus, the point B_1 = (.3, .7) indicates a subjective probability .3 that E_1 is the actual state and .7 that E_2 is the actual state. The point A_1 = (.2, .5) indicates an action leading to a payoff .2 at E_1 and .5 at E_2. As already noted, the utility ascribed to this action according to the indicated *B-state* will be B_1.
A_1 = .3 x .2 + .7 x .5 = .41. Graphically, this is the projection of B_i on A_1.

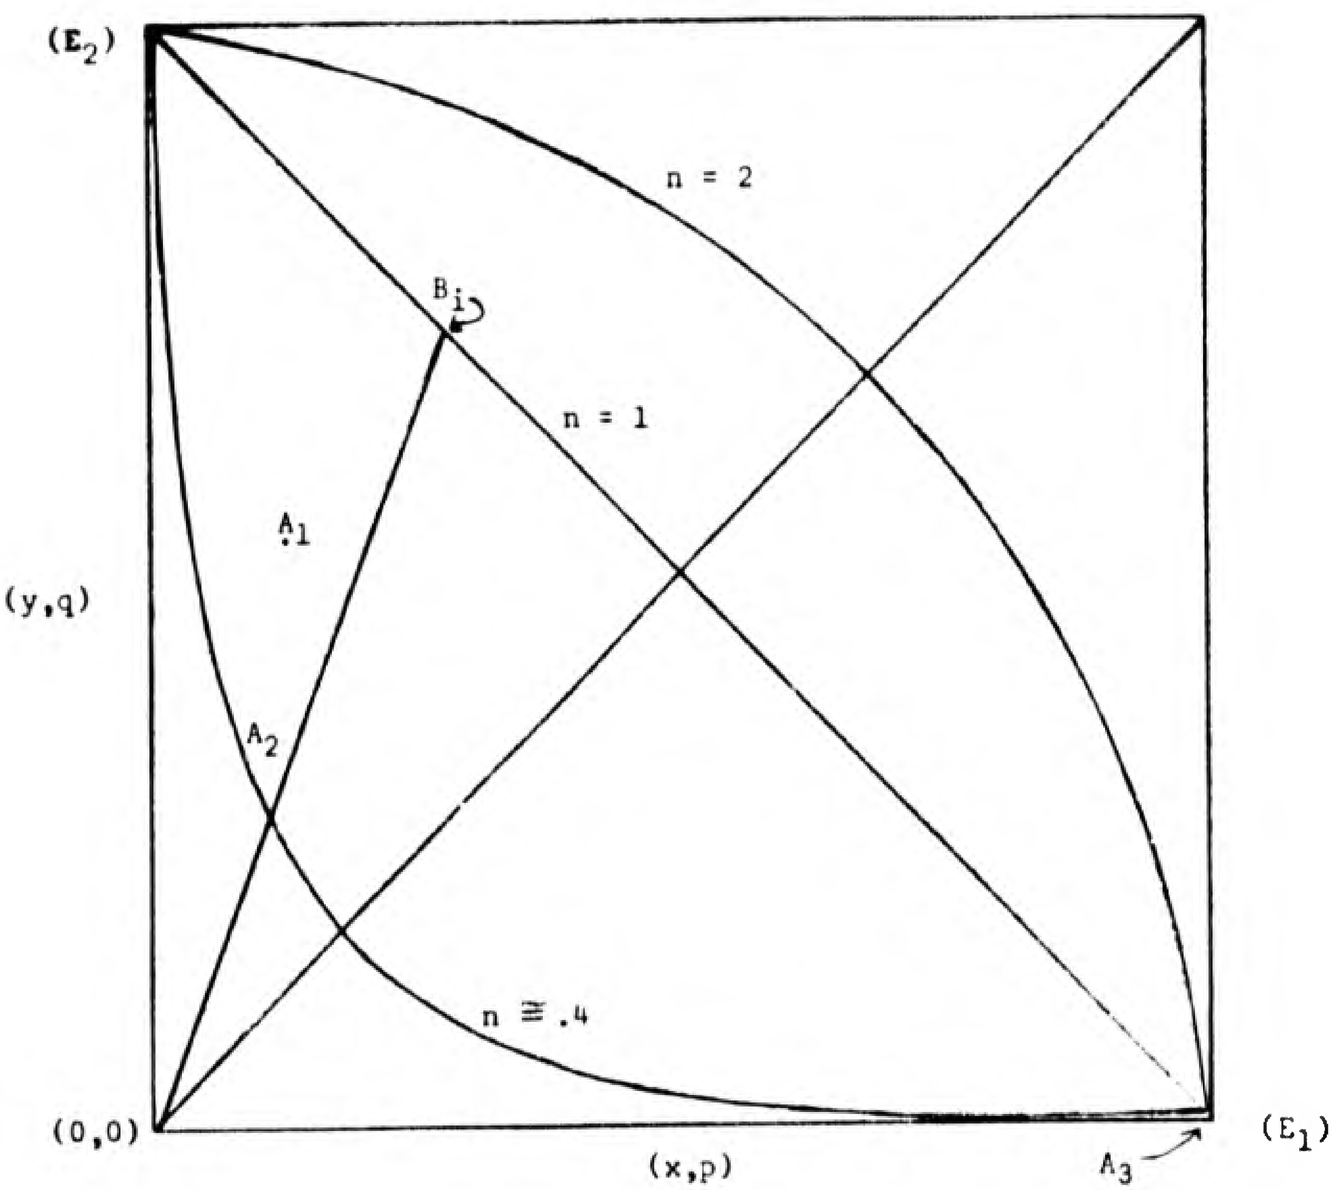

Figure 1. The space of action utility vectors with special cases of curves $x^n + y^n = 1$

We shall consider choices made among action elements lying on continuous, well-behaved, curves running from E_2 to E_1...that is, roughly paralleling the set of possible *B-states*. The chosen set of curves is of the form $(x^n + y^n) = 1$, $(n > 0)$, where n is defined as the ***hedging*** parameter of the curve. In the limit, this set of curves generates the entire inside of the square. Families of actions of this kind might be generated by mixtures of the response actions which are successful at the precise states E_1 and E_2, or the curves might represent continuously graduated responses.

If is convenient to divide these curves initially into two categories; the ***concave*** curves $(0 < n < 1)$ lying between the origin and the diagonal line $x + y = 1$ and the ***convex*** curves $(n > 1)$ lying between $x + y = 1$ and the right upper corner of the square.

The concave curves lying inside the diagonal correspond to "calculated risk" situations. For these curves, compromise actions have a lower mean value than those which are closer to one of the axes. Thus, in this region one would select "way out" actions rather than hedged actions even in the absence of definite information as to a *B-state* favoring E_1 or E_2. In fact as the curves approach the origin it may become expedient to select actions that are contraindicated by the *B-state*. In figure 1, for example, the action $A_3 = (1,0)$ is superior to $A_2 = (.1, .3)$ although the *B-state* (.3, .7) assigns a higher probability to E_2 than E_1.

This assumes, of course, that the choice is between these two actions only -- for example, the action $(0, 1)$ is not available. This latter aspect of the problem is discussed below.

Perhaps the most striking feature of decision making when the action utility functions are of this concave type is that the relation between *B-states* and indicated optimal actions is discontinuous. In general the optimal action is the one closest to the appropriate action associated with the most probable *E-state*. As the *B-state* changes, this optimum remains the same until it shifts abruptly to some other extreme action if another *E-state* becomes more probable.

For the convex curves lying between $x + y = 1$ and the vertex $(1, 1)$ the compromise actions have greater mean utility than more extreme or committal actions and will be selected in the absence of definite information. In the extreme curve of this set the action $(1, 1)$ is optimal under all conditions of information. The interior elements of this upper diagonal, then, correspond to what is ordinarily connoted by "broad spectrum" choices.

In order to study the general behavior of the optima in this region, consider the function

$$u\,[1, k] = p\,x + q\,y$$

where $q = 1-p$

and $y = (1-x^n)^{1/n}$

whence $u(p,x) = p\,x + (1-p)(1-x^n)^{1/n}$

Differentiating with respect to x,

$$\frac{du}{dx} = p-(1-p)\ \frac{1}{n}\ (1-x)\ \frac{1^{-1}}{n}\ (nx^{n-1})$$

$$= p-q\ \frac{x^{n-1}y}{1-x^n}$$

$$= p-q\left(\frac{x}{y}\right)^{n-1}$$

Thus, du/dx is 0, and the function is at its maximum value, when $\frac{p}{q} = \left(\frac{x}{y}\right)^{n-1}$

The most interesting curve in this region is the quarter circle defined by $(x^2 + y^2) = 1$. For this curve, the above expression for the optimal action becomes $x/y = p/q$: that is, the optimal action lies on the projection of the *B-state* vector onto the action vector curve. In a sense, "probability matching" becomes the rational procedure if the payoff functions are as described by this curve.

Between the diagonal and this curve the optimal action choice (if a full range of choice is available) lies between the projection of the *B-state* vector on the utility curve and the *E-state* axis on the same side of the 45° line. For action utility curves beyond the quarter circle, the optimal action lies between the projection of the *B-state* vector and the 45° line.

The special properties of the curve $x^2 + y^2 = 1$ carry over directly to higher dimensional spaces. Specifically, consider a set of action vectors such that *O*'s action choice is tantamount to the selection of a vector $(p_1, p_2, p_3 \ldots p_m)$ with $\Sigma p = 1$. Let the payoff, if the true *E-state* is E_j, be defined by $p_j/\sqrt{\Sigma p_1^2}$. Then *O*'s best choice, if his *B-states* are realistic, will always be the action vector that is equivalent to his current *B-state*.* Among other things, this would seem to be a very useful relation for the direct measurement of *B-states*.

<u>Cogency</u>. A second aspect of the choice situation is the range of options. Ordinarily, *O* will not have a free choice among actions of various hedging norms or even among all the actions in a particular family. Rather, he will have to select an action from the interior of the rectangle in figure 1 (or its generalization to a higher dimensional convex polyhedron). This restriction may be imposed by inherent conditions in the environment or by limitations in *O*'s own action capability. If the choice is restricted it may lie within a single curve (that is, action vectors with homogeneous norm) or it may be across curves. These two situations lead to rather different problems.

---

*If he selects this "probability matching" action on all occasions when his belief state is B_i, it can be readily seen that his expected payoff will be $(B_i \cdot B_i)^{\frac{1}{2}}$ -- the Euclidean length of the probability vector. For any other vector B_k, his expected payoff will be $(B_i \cdot B_k)/(B_k \cdot B_k)^{\frac{1}{2}}$. If this were greater than the payoff for B_i, we should have $(B_i \cdot B_k)^2 > (B_i \cdot B_i)\ (B_k \cdot B_k)$ which violates the Cauchy-Schwartz inequality (see Bellman, ref. 2).

If the choice is made within a curve, then the relevance of a vector depends rather directly on its resolution and on the spread along the curve of the action alternatives. Obviously the information becomes more useful if the *B-state* is highly resolved and the action vectors have highly discrepant loadings or components on the *E-state* axes.

On the other hand, if the choices are heterogeneous over curves, it is possible that the mean value of one action vector over all *E-states* may exceed the mean value of other action vectors to an extent that makes the information almost irrelevant. The extreme case of this would be a choice among action vectors all lying on the same ray from the origin but with different hedging parameters. For such a set of alternatives no amount of information makes any difference since action vectors with the larger parameter are superior under all conditions of information.

A CATALOGUE OF PROPERTIES

To summarize the principal points of the preceding sections, they are now listed as a set of 14 properties. The first four describe the *B-states* themselves; the next three relate to the query program; the next set of five properties describe the effect of clues on *B-states;* and the final set of two properties describe the pragmatic context within which the *B-states* are used. The present definitions and indices are tentative: either theoretical or empirical investigation may lead to extensive revision of this list.

B-State Properties

Modality. A *B-state* based on a finite environmental space is in categorical form if it asserts the possible presence of a particular subset of *E-states* and denies the possibility of remaining *E-states*. The *B-state* is in vector form if it assigns a distinct probability to each *E-state* in the space.

Resolution. The resolution of a *B-state* in vector form is the ratio of the variance of its components to the maximum variance possible for a *B-state* of that dimension.

Validity. The validity of a *B-state* is the weighted proximity of its probability components to the true *E-state*.

Compactness. The compactness of a *B-state* is the weighted proximity of its components to each other.

Query Properties

Query control. A query program is elective if *O* may control the order of query application, otherwise it is imposed. An elective program may be chosen in advance, on a bulk basis, or modified as a result of received clues. An imposed program may be fixed or random.

E-state-Query interaction. A query is intrusive if it modifies the *E-state* upon which it operates.

Query dependence. Two queries are independent if no possible outcome of either query changes the probability of any outcome of the other.

Clue Operator Properties

Balance. The potential effect of a query is balanced if the implied probability of various clues and the consequences of those clues if they occur are so related that the expected change in *B-state* vanishes.

Path independence. An epistemic system is path independent if the effect of applying any clue operator to any *B-state* is unrelated to the earlier history of that *B-state*.

Commutativity. An epistemic system is commutative if the net effect of any combination of clue operators is independent of their order of application.

Openness. The open space of a *B-state* is the set of dimensions (defined by *E-states*) that are attributed non-zero probability.

Resolving range. The resolving range of a set of clues is the set of distinct hypercharacteristic vectors that can be associated with the set of clues.

Pragmatic Properties

Hedging norm. The hedging norm of an action A_k is the parameter n such that the action utilities, u_{ik}, obtained if A_k is used and E_i is the true *E-state* will satisfy the equation $\sum_i u_{ik}^n = 1$.

Cogency. If a clue modifies B_k to B_k' then the cogency of the clue is the expected utility of the action A_k', selected and evaluated on the basis of B_k', less the expected utility of A_k, selected on the basis of B_k but evaluated by B_k'.

DISCUSSION

The set of properties introduced in the foregoing sections has been intended as representative rather than exhaustive. Even so, it has not been possible to devote as much space to any of these conditions as either rigor or expository clarity might require. Instead of trying to discuss any of these conditions in greater detail, however, we shall now attempt to put them in clearer focus by examining the relationships among various areas.

Figure 2 presents the major constructs that have been discussed, together with a schematic picture of the linkages between them. Three types of entities are represented: boxes represent the fluid states describing the environment and the organism's belief; ovals enclose the sets of conditions that are taken as fixed or given at any time -- the query program, the operator dynamics, and the utility scheme; diamonds represent the events that occur, including queries, clues, actions and payoffs.

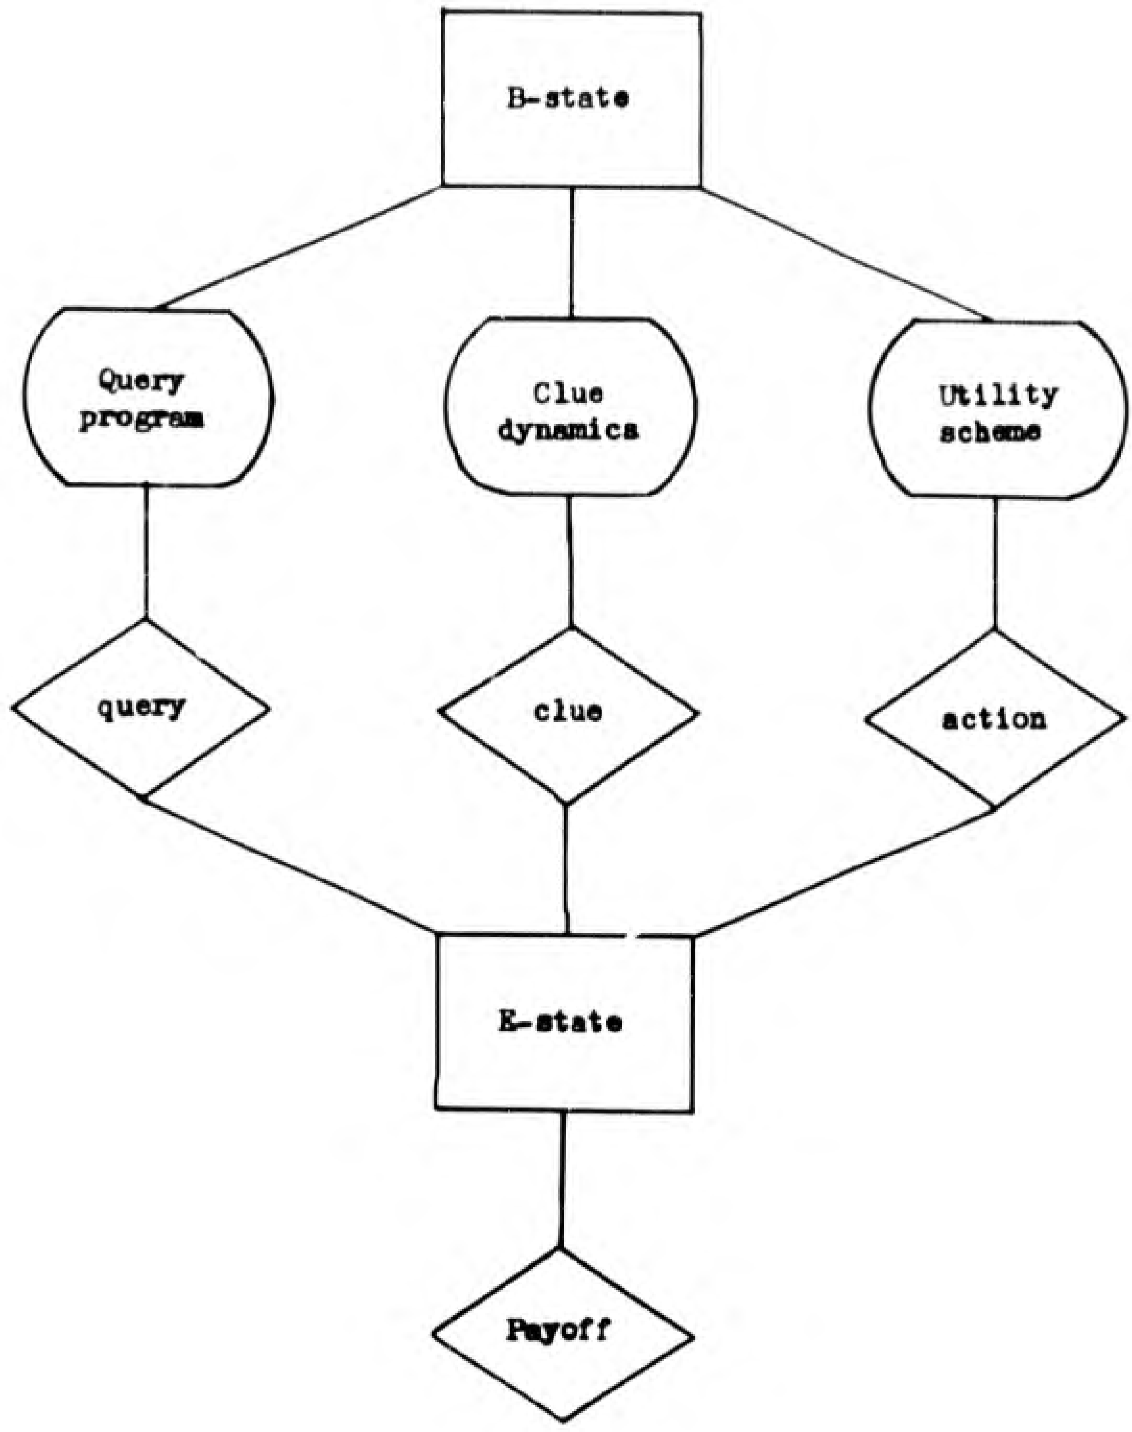

Figure 2. A schematic flow diagram of decision making with an associated epistemic process

A sort of flow description of the overall process is as follows. First, the existing *B-state* acts through the query program to select a query (if selection is possible). The query acts on the current *E-state* to produce a clue. The clue, in turn acts through the rules described under operator dynamics to effect a change in the *B-state*. Finally, the *B-state*, acting through a utility scheme, selects an action which, in conjunction with the current *E-state*, yields some payoff or income.

Considering the overall process in this way, several rather general conclusions emerge. First, it appears misleading to identify the value of a *B-state* with its effectiveness for specific action choices. As the present schema emphasizes, the *B-state* enters into three sets of interactions and it may have to be evaluated rather differently for each of these.

The primary purpose of *B-states* is of course to guide the choice of action at a particular juncture, and for this purpose some set of conditions such as those outlined in the section on pragmatic aspects may determine its value. In addition, however, the *B-state* is used to guide the selection of queries, and its value for this purpose may be quite different. As an obvious example, a highly resolved but aberrant *B-state* may be as useful as a fuzzy *B-state* for selecting actions; however, it is much less likely to lead to a query program that guarantees its own refinement.

The third function of a *B-state* is as a context for interpreting incoming clues. It is easy to show, however, that for this purpose one may need a much more detailed picture of the world than is required for individual action choices. As a homey example, most of us carry around in our heads enough information to know whether we can safely write a check for a certain size. In order to maintain a valid *B-state* on this matter over a long period of time, however, we need a more precise bookkeeping procedure.

In brief, then, we conclude that the value of *B-states* can be assessed only by reference to the entire decision-making process. For what it is worth, this can be more formally stated in a summary equation

$$E_u = \Sigma_i \, P\,(E_i)\, P\,(u)|E_i$$

$$= \Sigma_i \, P\,(E_i)\, \Sigma_k \, P\,(A_k|E_i)\, u(i,k)$$

$$= \Sigma_i \, P\,(E_i)\, \Sigma_j \, P\,(B_j|E_i)\, \Sigma_k \, P\,(A_k|B_j)\, u(i,k)$$

In this expression, $P\,(E_i)$ may be a constant or may be affected by query interaction; $P\,(B_j|E_i)$ is a function of the epistemic process; $P\,(A_k|B_j)$ will be 0 or 1 for particular choices, but will depend on the available set of actions over all choice junctures; and finally $u(i,k)$ is assumed fixed.

A second summary observation is that this whole process is extremely cumbersome and tenuous as compared with the more direct mapping of actions on external cues. If the payoff is meager from a process such as that outlined here, *O* cannot be certain whether the fault lies with his epistemic process or with his utility scheme. Even if he has good reason to suspect his epistemic process, he cannot be sure whether he is receiving poor evidence or is not interpreting the evidence properly.

If in spite of this awkwardness some such procedure as described herein does characterize rational behavior, two compensating advantages may be adduced. First, the indirect procedure permits the economical compression of what would otherwise be many-many functional relations. That is, instead of relating many clues to many action choices, *O* can relate the clues to a comparatively few *B-states* or constructed environments (or at least to comparatively few dimensions of these) and he can then apply the *B-states* to a variety of action choices.

A second, and probably decisive, advantage is that behavioral procedures mediated through a constructed environment lend themselves to symbolic manipulations in a way that is not possible with direct stimulus-response adjustment. Without belaboring this point it is quite clear that the use of ***B-states*** permits the application to behavior of a fund of vicarious experience -- either hypothetical or borrowed from the adventures of other organisms.

REFERENCES

1. Ashby, W. R., An Introduction to Cybernetics, Chapman & Hall, Ltd., London, 1956.
2. Bellman, R., Introduction to Matrix Analysis, McGraw-Hill, New York, 1960.
3. Hardy, G. H., J. E. Littlewood, and G. Polya, Inequalities, Cambridge University Press, New York, 1934.
4. Marshak, J., Remarks on the Economics of Information, Cowles Foundation for Research in Economics, New Haven, Connecticut, 1960.
5. Reza, F. M., An Introduction to Information Theory, McGraw-Hill, New York, 1961.
6. Roby, T. B., "Commitment," Behavioral Science, Vol 5, pp 253-264, 1960.

DISCUSSION OF ROBY'S PAPER

TODA (Discussant). My general reaction is quite favorable to Prof. Roby's excellent schematic and operational representation of the mechanisms involving belief states, evidence and action. But, just a statement of general agreement will not stimulate interesting discussions. So, I will first make a critical comment on his general conceptual framework, and then make a pair of suggestions which may be of some technical interest.

According to Prof. Roby's flow diagram, a circular process, *B-state--query--clue--B-state*, takes place independently of each individual decision problem and its accompanying utility system. This is certainly a valuable point to make, that is, that the belief system of a person is maintained and refined not necessarily expecting a certain, specific decision problem. On the other hand, however, this complete detachment of information-seeking and processing cycle from the information utilization process is, I think, an over-simplification. For example, a person A would wish to refine his *B-state* about a person B, if the person A anticipates a possibility of having an interaction with the person B, even though A is not very specific about the kind of interaction that would take place. Further, A would want to have higher resolution for the *B-state* when the class of decision problems he is anticipating is generally more important--important in terms of the utilities concerned, than when he is anticipating a class of unimportant interactions.

In general, I believe that how much and what kind of resolution satisfies a person and terminates the circular process of information-seeking and processing, can be decided only when a relevant utility system is given. So, this information-seeking and processing cycle is often restarted when a specific decision problem is explicitly given, even if the person has somehow anticipated the problem and prepared for it, since the amount of resolution might not have reached the necessary level to effectively handle utilities concerned. So, what I want to ask Prof. Roby is to add a backward arrow from "*utility scheme*" to "*B-state*" in his diagram.

One more point I would like to mention about this issue is the concept of the cost of information inquiry. Without introducing a kind of cost-concept into the scheme, it seems to me impossible to determine when a person is satisfied by a resolution which is not complete.

My next point is concerned with the use of matrix representation for his clue operator. I think that nothing much is gained by using matrices to express Bayesian inferences as far as the kind of matrices is restricted within the class of diagonal matrices. What impressed me in this matrix representation, however, was the vast possibility in the use of non-diagonal matrices. A clue matrix involving non-zero non-diagonal elements, in a sense, reorganizes the *B-state* vector. Take, for example, the saint-or-scoundrel situation. What this person will do allows for double interpretation. If the person is a saint, his behavior means one thing, and if the person is a scoundrel, it means something else. So, the *B-state* space concerning the person should consist of at least two subspaces, the saint-subspace and the scoundrel-subspace, and each item of information about the person should be entered into both of the two subspaces as partial evidence.

Now, suppose that a clue is given which eliminates every possibility that the person is a saint. Then the matrix representing this clue should have non-diagonal elements which move the evidence stored in the saint-subspace into the scoundrel-subspace. Of course, reorganization of the *B-state* space through a non-diagonal matrix is more effective when more elaborate reorganization than the above example is required. Obviously, conditional evidence like that in the above example are very popular in military situations, and the possibility of the kind of operations might be worth investigating.

The last technical point I would like to mention is concerned with the pragmatic aspects of Prof. Roby's scheme. There he has pointed out that the organism maximizes expected utility of the form

$$u = px + qy,$$

where x and y are payoff functions satisfying the condition

$$x^2 + y^2 = 1,$$

by choosing such action that makes the ratio x/y equal to the ratio p/q. Prof. Roby points out that this is probability matching behavior and this type of payoff function is useful in measuring *B-states*, or, in other words, useful in measuring subjective probabilities. I agree. I just want to add a couple of comments which might make the issue conceptually clearer.

First, suppose that the organism's choice is made by picking a number P from the interval $(0, 1)$. Then the payoff couple $x(P)$ and $y(P)$ are defined as functions of the choice, P, in such a way that

$$x(P) = P/\sqrt{P^2 + Q^2}$$

$$y(P) = Q/\sqrt{P^2 + Q^2}$$

where $Q = 1 - P$.

Obviously, this payoff couple satisfies the above condition, and the expected utility u is maximized by choosing such P that is equal to p. The choice certainly deserves the name of probability matching.

Secondly, the type of utility considered by Prof. Roby is not the traditional utility measured by an interval scale. The payoff ratio x/y is meaningless unless utility has an absolute origin. I think that a ratio-scale of utility is more realistic than an interval-scale of utility. But, the point is, if there is zero utility, there should be negative utilities. And, if x and y are both negative in Prof. Roby's equations, it is easily proved that the probability matching behavior gives the <u>minimum</u> expected utility.

Thirdly, there is a payoff which makes the probability matching behavior optimal when the payoffs are both negative. The payoff pair is given as

$$x(P) = \log P$$

$$y(P) = \log Q$$

With this payoff, expected utility u is expressed as

$$u = p \log P + q \log Q$$

and this is maximized by choosing $P = p$ so that the maximum is

$$Max_p\, u = p \log p + q \log q$$

This looks like an information theory concept; I don't know however, if it goes any deeper than that.

ROBY. I too have been uncomfortable about the separation of the information acquiring loop from the pragmatic loop and I welcome your suggestion on this matter.

As to the comment on the diagonal matrices, I was assuming that one would be concerned with matrices in general. However for Bayesian processes, the appropriate matrix is apparently a diagonal matrix and that was why I was emphasizing those. I would like to turn the discussion over to the floor at this point. I'm afraid any further discussion of the various possible utility schemes would get rather involved.

EDWARDS. I have two questions to ask, both resulting from my lack of clarity about some of the deeper concepts that were being talked about. I was trying to hunt around for some familiar landmarks among the large set of words that I had not encountered before. I certainly recognized some of them, "Bayes' Theorem," and "maximization of expected utility;" but it was not clear to me how many other ideas were built into this schema that I had not recognized. It would be helpful if I could be given some kind of clue that would enable me to determine how much of this follows from Bayes' theorem, and how much is something different.

I also noted that in both Professor Toda's comments and in Dr. Roby's paper there were some statements about the circumstances under which probability matching can be thought of as optimal behavior. An average can never be as large as the largest of the things being averaged; and since an expected value is an average, the implication would be that it is never possible in a game against nature, for a mixed strategy to be also optimal. That is, as a matter of fact, a very important theorem of decision theory, and yet somehow they seemed to be saying that probability matching, which is mixed strategy, could in some sense be optimal in the context of this model. Could you tell me what it is that I missed?

ROBY. I think you are being Socratic, but I will respond like a straight man anyway. I have indeed used Bayes' theorem, but I have not gone over entirely to being Bayesian. What I wanted to do was to use the structure of the Bayesian type of inference as an analogy for information acquiring processes in general. I also used, to a certain extent, some notions from quantum mechanics which at the level I used them are probably fairly familiar, e.g., the notion that you can multiply vectors by matrices and get characteristic states that mean something. As to terminology, I tried to use the most common sense terms that I could for the processes that I wanted to describe. I can't say how much of it is new, in part because I am probably not as familiar with this area in an intimate sense as many of the other people here. A lot of this work is just incidental to work directed toward formalizing small group processes.

The other point, on probability matching, is unclear to me. Are you saying it is an impossibility in the real world that a mixed response could be more useful in the blind than a pure response?

TODA. I don't think this is a mixed response. The subject should choose a fixed behavior; in this case the response is made by picking a number from a continuum, and if the interval is $(0, 1)$ then the subject should choose the number that corresponds to the probability. It is not mixed behavior.

EDWARDS. Then this is a model for deciding on the likeliness of an event, and not a model of how to choose among various things that you might do in order to exploit that opinion point to maximize some quantity.

ROBY. There is no reason why these can't be isomorphic.

EDWARDS. I am saying that they can't be isomorphic if by isomorphic you mean that the probability of choice ought to be equivalent to your opinion of the probability that one of these events will in fact obtain. If, for example, in a standard probability learning experiment, you are attempting to infer the relative frequency of the more frequent event, and discover that it is .7, then there is no sense in which would be optimal for you to predict that event 70 per cent of the time.

ROBY. Oh, no, this is what we do. We ask "How often is this event going to occur?" We give a subject a payoff which is tantamount to the probability that he in fact assigns to it. If he says it will occur 70 per cent of the time, then his payoff when the event does occur, is

$$\frac{.7}{\sqrt{.7^2 + .3^2}}$$

If it comes up tails it is

$$\frac{.3}{\sqrt{.7^2 + .3^2}}$$

This is an artificial kind of payoff, though I don't think it is any more artificial than others, and I think as the term "broad spectrum" implies, we do have such situations in the real world. For example, we give people drugs, because it is cheaper to make good drugs than it is to teach diagnosis. This is a kind of intermediate response, and you base it on the fact that you do have limited knowledge. When you have this kind of response, it pays off not only to know something about the world, but to know how much you know.

ZINK. Early in the paper you stated that the person is cheating if he expects a query to change his *B-state* in any particular direction; but later you say one function of the *B-state* is to guide the selection of queries. Are you allowing him to cheat anyway?

ROBY. What I mean is that if he has predetermined that, if some particular clue occurred, it is going to move his *B-state* so much in a given direction, and it is fairly probable that that clue will occur, then I say he is cheating, and that the query program is balanced only if all things considered, his net position should be the same after the query. Now, at the same time, there is an interaction, in that he can on the basis of his *B-state* select a query for which the clues are most likely to cut the *B-state* in some particular direction. He can, to that extent, use the *B-state* to guide him.

LANZETTA. Is there any attempt in this model to deal with the problem of conditions under which the organism ought to elect to make queries, and the conditions under which he ought to terminate querying and select an act in terms of the present belief state? You made the point that it may be either imposed or elective.

ROBY. I didn't go into that, but I think that one could develop that from within this general framework.

FLOOR. In your empirical work have you given some attention to the conditions under which the subject would adopt a probabilistic point of view towards the phenomenon in question versus when they adopt a more deterministic attitude? I would also suggest that you might have some interesting research possibilities suggested by the flow diagram, in connection with things like bridge bidding, when you are attempting to communicate to your partner, but conceal from your opponents; that is, you want to reserve a certain amount of ambiguity as to what you are communicating.

ROBY. We have done some work on this matter of trying to find out when people adopt a probabilistic or when they adopt a point type of representation. This work has not been very successful yet. I don't have an experimental situation that I like as far as sorting people out. I suspect that there are individual differences in addition to situational differences as to which of these modes people tend to adopt.

PREDECISIONAL PROCESSES RELATED TO

PSYCHOPHYSICAL JUDGMENT

mir H. Shuford
University of North Carolina

For the past several years some of us at the Psychometric Laboratory have been engaged in interpreting human behavior in terms of Bayesian decision theory (refs. 11, 13, 18). Our approach is similar in spirit to that of Swets, Tanner, and Birdsall (ref. 14) in the area of signal detection and to that of Masanao Toda (refs. 15, 16) in the area of probability learning.

There are a number of reasons for our attempting to use decision theory as a starting point for a theory of human behavior. The first and most obvious reason is that a decision-theoretic analysis of a task yields an upper limit on the performance possible in the task as presently structured. This information is useful in many applications. For example, if one is designing man-machine systems, the optimum decision-theoretic strategies can be compared for different systems and the system yielding the largest average expected utility selected for further development. On the other hand, it may turn out that the decision-theoretic analyses indicate that none of the proposed systems is capable of yielding an adequate level of performance. In this case the design of a new system should be attempted and the decision-theoretic formulation can prove to be an indispensable tool in guiding this design. The upper limit of performance yielded by a decision-theoretic analysis is also relevant to evaluating the performances of human operators and the adequacy of training programs.

The second reason for using decision theory as a descriptive theory of human behavior also derives from the fact that it is an optimality model or a model which represents *efficient* behavior. Consider developing a descriptive model for a certain segment of behavior. A very large number of models are mathematically possible which would adequately describe this limited area of behavior. Suppose we choose one such model, adequate for the particular data at hand. Now consider incorporating new segments or areas of behavior into this model. If we had chosen a model which yields *less efficient* behavior than that empirically found in the new areas, we would have to scrap our inefficient model and start over again as has happened many times in the history of psychology. On the other hand, if we had started with an efficient model, it usually proves to be an easy task to modify this model to correspond to the apparent relative inefficiency of behavior by introducing certain constraints such as a limited and imperfect memory, information processing errors, and the "cost of thinking". The validity of these modifications can be empirically determined.

There are a number of possible decision-theoretic formulations. They all share in common the notion of a payoff matrix in which: (a) the rows represent the alternative courses of action available to the decision maker; (b) the columns represent either the alternative courses of action available to the opponent as in game theory (refs. 1, 17), or the possible states of nature as in decision making under uncertainty (refs. 1, 3); and (c) the entries in the cells of the matrix represent the utilities associated with the possible outcomes. The utilities are measured on a psychological continuum defined up to a linear transformation (refs. 1, 12, 17).

The decision-theoretic formulations differ, however, with respect to the criterion for an optimal strategy (ref. 1). It often has been assumed that the *minimax strategy*, i.e., that course of action which results in maximum utility assuming the worst possible outcome, is reasonable at least for the case in which one is playing a strictly competitive game against a rational opponent. This view has recently been attacked by, among many others, Schelling (ref. 5) and Wiener (ref. 18). Schelling essentially argues against the existence of strictly competitive games in diplomatic, political, economic, military, and other decision making contexts and gives many examples of partly competitive, partly cooperative games. Wiener's criticism is deeper in the sense that he argues that even in a competitive situation one should take account of the past behavior of the opponent along with all other possibly relevant information. This is a policy which is in general incompatible with a minimax strategy, but which is, in fact, an intuitive restatement of a *Bayes strategy*, i.e., use all available information to assign probabilities to the opponent's actions or to the states of nature, use these probabilities to compute the expected utility for each alternative course of action, and then select that course of action which has the largest expected utility.

It seems apparent that Bayesian decision theory is a much more promising candidate for a descriptive theory of human behavior than is a theory based on the minimax principle. It is more powerful and more general because it includes an additional psychological variable, subjective probability. Also, with the development of theories of subjective probability (refs. 8, 15, 16, 19) which can account for the modification of probabilities on the basis of past experience and new information, we have a descriptive decision-theoretic model which is dynamic in the sense of accounting for changes in behavior over time. Further, the model describes and, in fact, can be used to define adaptive behavior.

Since 1959 with the support of the Air Force Office of Scientific Research we have been developing and testing Bayesian models for individual decision making. Most of our work has been in the area of psychophysical judgment which seems particularly well suited for testing our models. We conceive the subject's task in a psychophysical experiment to be a problem in decision making under uncertainty in which the subject responds on each trial of the experiment in such a way as to maximize his average payoff during the course of the experiment. Two sources of information are available to the subject: (a) the information obtained from observing the particular stimulus presented on that trial; and (b) the background information as to the probabilities with which the different stimulus values might be presented. This background information is provided by the subject's experience before participating in the experiment, by the instructions, and by the context of the experiment, and during the course of the experiment by the stimuli previously observed by the subject. This background or prior information is available at the beginning of the trial. The subject observes the stimulus and obtains imperfect information as to the value of the stimulus on this particular trial. The subject, in effect, combines this specific information with his prior information and obtains a posterior probability distribution representing the subjective probability of the stimulus value given the observation. These posterior probabilities determine the expected utilities of the different response categories and, thus, the response made by the subject. These types of models use Bayes' theorem in the computation of the posterior probabilities and in many cases are formally analogous to the models used in Bayesian statistical inference (ref. 3, 6, 7, 10).

Now, let me suggest a definition of predecisional processes. The decision process itself might be taken to be the computation of expected utilities from the previously given probabilities and utilities and the selection of a course of action. In other words, the decision process starts once the probabilities and utilities have been determined. If this is so, then everything else, including the structuring of the payoff matrix in terms of relevant courses of action and states of nature, can be classified as predecisional processes relative to a particular decision process.

I think that his classification and the resulting emphasis on predecisional processes has some heuristic merit, particularly if we keep in mind the fact that the distinction applies to our theories and not necessarily to the way a subject might view the situation. That is, in many situations which might be analyzed in decision-theoretic terms it is quite unlikely that the subject views his task as one involving a number of predecisional processes leading to a solution of the decision problem itself.

The value of the distinction between predecisional processes and the decision itself lies in the implied criticism of research on decision making behavior. My own research has tended to concentrate too much on the decision process and to ignore the subject's activities leading up to the final decision. This procedure poses a special danger when decision theory is used as a descriptive theory of behavior. The great generality of decision theory and the concomitant advantages have been referred to earlier. Some of this generality is due to the fact that a decision-theoretic model has many, many parameters. Restle (ref. 4) has pointed out that every probability and utility in a decision-theoretic formulation is, in effect a parameter that needs to be assigned some value. The situation is, in reality, more extreme than this in that the courses of action and states of nature must also be determined. This aspect of decision theory has led to the formation of two "camps", the believers and the non-believers in the rationality of human behavior. Both camps use decision theory to prove their assertions. The believers in the rationality of behavior will manipulate the parameters of decision theory until it fits any observed behavior. The other camp, which might better be called believers in the irrationality of behavior, will fix the parameters of decision theory in such a manner that it does not fit any observed behavior. This is a rather sterile sort of activity and it may turn out that an increased emphasis on predecisional processes will serve to reduce its popularity.

To be more explicit, the study of predecisional processes should lead to a better understanding of how a subject interprets and structures a decision task and to theories of subjective probability and of utility. It is probable that different theories will be required for different areas of behavior. Part of the value of such theories resides in the fact that they can effect a great reduction in the number of free parameters involved in the decision-theoretic model. I would like to illustrate this in some detail by discussing some examples of predecisional processes in the area of psychophysical judgment. Some of the discussion will refer to experimental data; some will be based on speculation.

Let us return to Bayesian decision theory and note that it contains a predecisional process. A distinctive characteristic of this theory is the use of Bayes' theorem to modify probabilities in the light of new information. This process can be demonstrated in a psychophysical judgment task.

Notice the 16 x 16 random matrix in the upper half of figure 1. The proportion of *1*'s in this matrix is .25. We could study, and indeed have studied (ref. 9), the scale of proportion by having subjects estimate the percentage of *1*'s in such matrices. If this procedure is followed, we are faced with the problem of having to estimate a parameter of the model, the sample size, which represents the amount of information extracted from the matrix by a subject (refs. 11, 13). Thus, we would have to fit the model to the data of each subject. This can be avoided by showing the subject only the small 3 x 3 matrix shown in the bottom half of figure 1. This small matrix has been obtained by placing a mask over the larger matrix and thus is a random sample of nine elements from the larger matrix. The subject's task is to estimate the percentage of *1*'s in the large matrix using only the information provided by the the small matrix. The courses of action available to the subject are the 101 percentages from 0 to 100. The relevant states of nature are the 257 possible proportions of *1*'s in the large matrix containing 256 elements. It is helpful to use continuous rather than discrete mathematics for such a payoff matrix. For example, the subject's probabilities of the states of nature may be approximated by a continuous *beta distribution* defined over the interval $(0, 1)$, which corresponds to the proportion of *1*'s in the large matrix. A beta distribution is completely specified by two parameters, a and b (ref. 3). Thus by using this approximation we have reduced the number of parameters needed to describe a subject's probabilities from the original 256 down to 2.

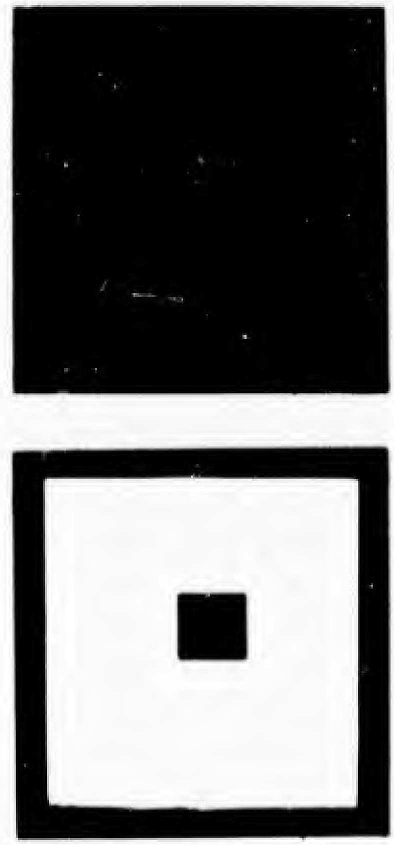

Figure 1. Random Matrix of 0's and 1's

Figure 2 shows four different beta distributions. The distribution shown at the upper left might represent the subject's prior distribution at the beginning of the experiment. The distribution shown at the upper right represents the distribution of proportions actually used in the experiment. This distribution also might be taken to represent the subject's prior distribution after having experienced a large number of trials in the experiment since the subject is told the actual proportion of *1*'s in the large matrix at the end of each trial.

Suppose we have the subject *predict* the percentage of *1*'s in the large matrix at the beginning of each trial and pay the subject 50 cents for each prediction which is within one percentage point of the actual percentage. The courses of action and the states of nature are the same as in the estimation task. The probabilities of the states are specified in terms of the two parameters of a beta distribution. How many parameters do we need to specify the subject's utilities? Notice that there are two things that can happen to the subject in the experiment. His prediction is correct and he is paid 50 cents or his prediction is in error and he receives no payment. Therefore, we might assume that two different utilities are involved in this task. However, we do not need to know the numerical

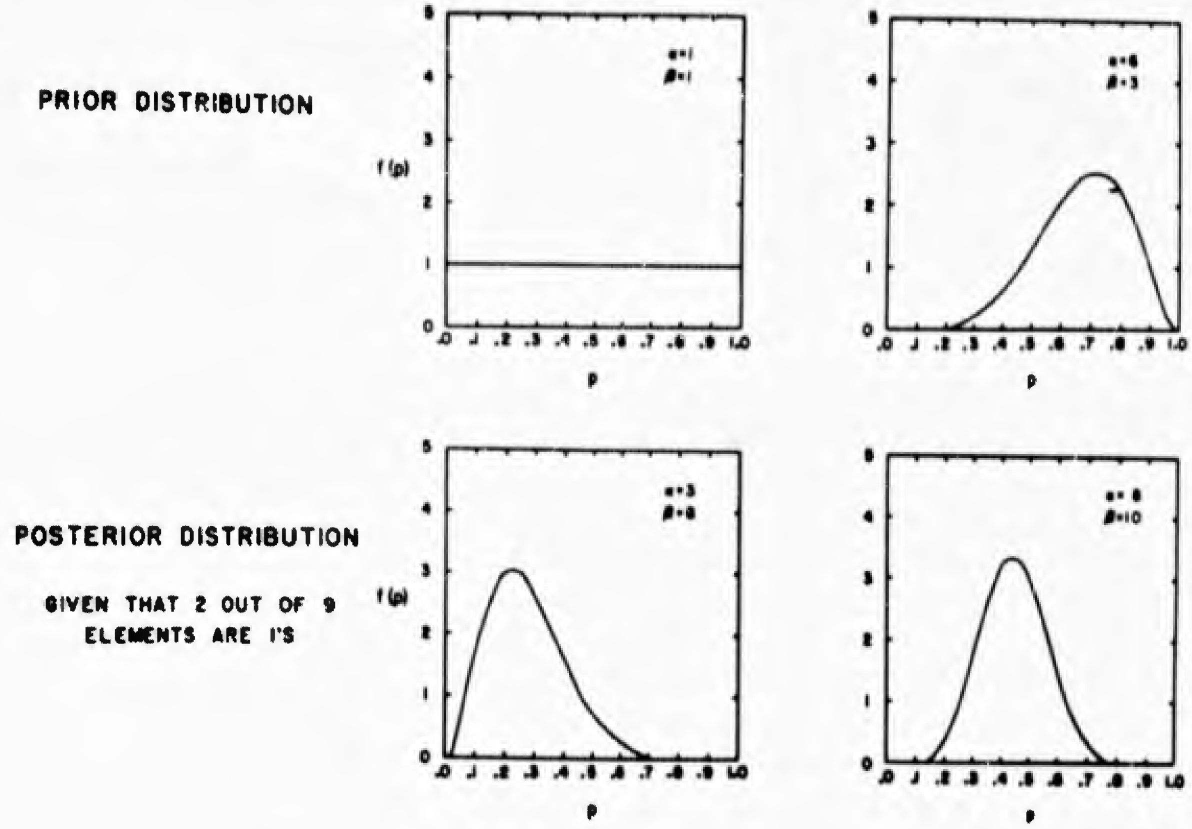

Figure 2. Four Different Beta Distributions

values of these utilities in order to obtain a prediction from the Bayes model. All we need to know is that the utility of being correct and receiving 50 cents is greater than the utility of being wrong and receiving no payment. This is so because the Bayes strategy in this decision task is for the subject to select as his prediction the percentage corresponding to the most probable proportion, i.e., the mode of his probability distribution. If the distribution at the upper left describes the subject's probabilities, any percentage is as good as any other since there is no one most likely proportion. However, if the distribution at the upper right describes the subject's probabilities, the best prediction is 71 per cent corresponding to the mode of this beta distribution. This strategy is valid for all cases in which the subject wishes to maximize the expected number of correct predictions and thus does not depend upon how much the subject wishes to do this. The exact numerical values of the two utilities are irrelevant in terms of this Bayes model.

In fact, the Bayes model suggests that the subject's behavior is relatively insensitive to the particular payoff scheme used in this decision task. For example, suppose we pay the subject an amount of money proportionate to the square of the difference between his predicted percentage and the actual percentage. In this case the Bayes strategy is for the subject to select as his prediction the percentage corresponding to the mean of his probability distribution. If the distribution at the upper left describes the subject's probabilities, the best prediction is 50 per cent corresponding to the mean of .50. If the distribution at the upper right describes the subject's probabilities, the best prediction is 67 per cent. For another example, suppose we pay the subject an amount of money proportionate to the absolute value of the difference between his predicted percentage and the actual percentage. In this case the Bayes strategy is for the subject to select as his prediction the percentage corresponding to the median of his probability distribution. The median, of course, lies between the mode and the mean of these distributions. Therefore, to the extent that the difference between the mode and the mean of the subject's probability distribution is small, the difference in the subject's behavior induced by these three different payoff schemes are small.

So far the main thing accomplished by the Bayes model is a reduction in the number of parameters required by a decision-theoretic formulation. We are left with the two parameters describing the subject's probability distribution and with an assumption about the subject's utilities which the model suggests is not too critical. This model can do one more thing for us. Suppose we also have the subject *estimate* the percentage of *1*'s in the large matrix using only the information provided by the small matrix as mentioned above. The subject is paid 50 cents for each estimate which is within one percentage point of the actual percentage. It should be apparent that this *estimation* task is formally analogous to the *prediction* task described above and that the same strategies apply. The major difference between the two tasks is that more information is available in the estimation task than in the prediction task, i.e., the subject can make use of the information provided by the small matrix. This additional information leads to a new probability distribution for the estimation task. The Bayes model provides an explicit process for modifying the prediction or *prior distribution* into an estimation or *posterior distribution* in order to take account of new information.

Suppose that the beta distribution shown at the upper left represents the subjects prior distribution and that the subject observes a small matrix in which two out of the nine elements are *1*'s. This new information leads to the beta distribution shown at the lower left, which represents the subject's posterior distribution, i.e., the probability of a proportion given *all* available information. The subject's estimate of the percentage of *1*'s in the large matrix will be 22 per cent, the mode of this distribution. Suppose that the distribution shown at the upper right represents the subject's prior distribution. Then, after observing the small matrix his posterior distribution is described by the beta distribution shown at the lower right. In this case the subject's estimate will be 44 per cent. It is clear that the prior distribution can have a significant effect on the subject's judgment.

This example shows how the Bayes model can be used to relate different decision tasks. This is accomplished by writing the probabilities for one decision task in terms of the probabilities for the other decision task and the additional information available for the new decision task. In the case of the prediction and estimation tasks, the probabilities for the prediction task are described by two parameters, a and b, while the probabilities for the estimation are described in terms of these same two parameters plus two more parameters, the size of the sample of elements in the small matrix and the number of *1*'s in the small matrix. It should be noted that values for the latter pair of parameters can be obtained from a knowledge of the physical environment: they do not have to be estimated for each subject. Thus, we can obtain a model for a new decision task without having to estimate any new parameters for the subject. The significance of this procedure can be made more dramatic by considering a sequence of decision tasks, each differing from the previous one only by the addition of new information about the actual proportion. The Bayes model relates all of these tasks using only the two parameters, a and b, and "objective" information about the environment. This is a major advantage of the predecisional process contributed by the Bayes model. The process describes the modification of probabilities to take account of new information about the states of nature. In this context the Bayes model describes a learning process in the sense that more information yields improved performance.

Now, let us see if the behavior of the subject actually corresponds to that predicted by the models. On each trial of the experiment, the subject predicts the percentage of *1*'s in the large matrix; the subject sees the small matrix and estimates the percentage of *1*'s in the large matrix; the subject is told the actual percentage and is paid 50 cents for correct predictions and estimates. There are 150 trials in the experiment and the distribution of proportions for these 150 trials corresponds to the negatively skewed beta distribution shown in the upper right of figure 2.

If we fix the values of the two parameters of the prior distribtuion, we can use the Bayes model to obtain theoretical values for the subject's predictions and for the subject's estimates. If we assume that the subject's prior distribution is not changing too rapidly, we can average the data over 15 trials in order to compare the subject's mean prediction and his mean estimates with the theoretical values given by the Bayes model.

It seems reasonable to expect that the subject might begin the experiment with a prior distribution similar to the rectangular distribution shown in the upper left of figure 2. In this case the theoretical value for the mean prediction is 50 per cent. The mean of the subject's predictions for the first 15 trials of the experiment actually turns out to be 54 per cent. The theoretical function relating mean estimate to the number of *1*'s in the sample shown in the small matrix is a straight line given by the equation

$$e = 100(r + a - 1)/(n + a + b - 2)$$

where r is the number of *1*'s in the sample
n is the number of elements in the sample
and a and b are the parameters of the prior distribution.

The rectangular prior distribution has parameters, $a = 1$ and $b = 1$, so in this case the mean estimate equals the percentage of *1*'s in the small matrix. Figure 3 shows the mean estimates of the subject plotted as a function of the number of *1*'s in the sample for the first 15 trials of the experiment. The straight line is the theoretical function obtained from the Bayes model. Notice that some of the data points are based upon only one observation.

Let us take a look now at the subject's behavior during the last 15 trials of the experiment. It seems reasonable to expect that by this time the subject might have learned the stimulus distribution. If this is so, his prior distribution can be represented by the negatively skewed beta distribution with parameters, $a = 6$ and $b = 3$, shown at the upper right in figure 2. The theoretical value for the mean prediction is 71 per cent. The mean of the subject's predictions for the last 15 trials of the experiment actually turns out to be 63 per cent.

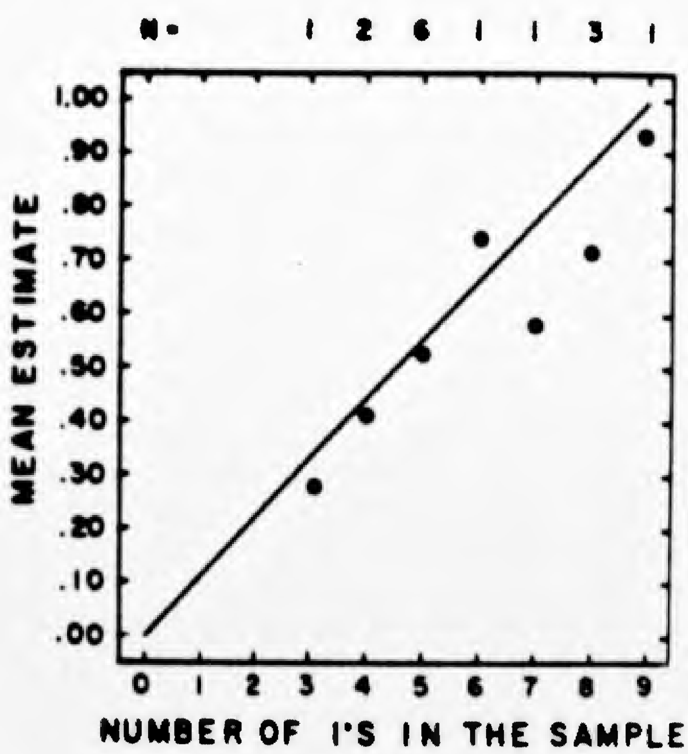

Figure 3. Mean Estimates for First 15 Trials (Average Subject)

Figure 4 shows the mean estimates of the subject for the last 15 trials of the experiment. Notice that in this case the theoretical function has a slope less than one and an intercept greater than zero. Thus, the Bayes model yields an effect very similar to the central tendency or regression effect frequently observed in psychophysical experiments.

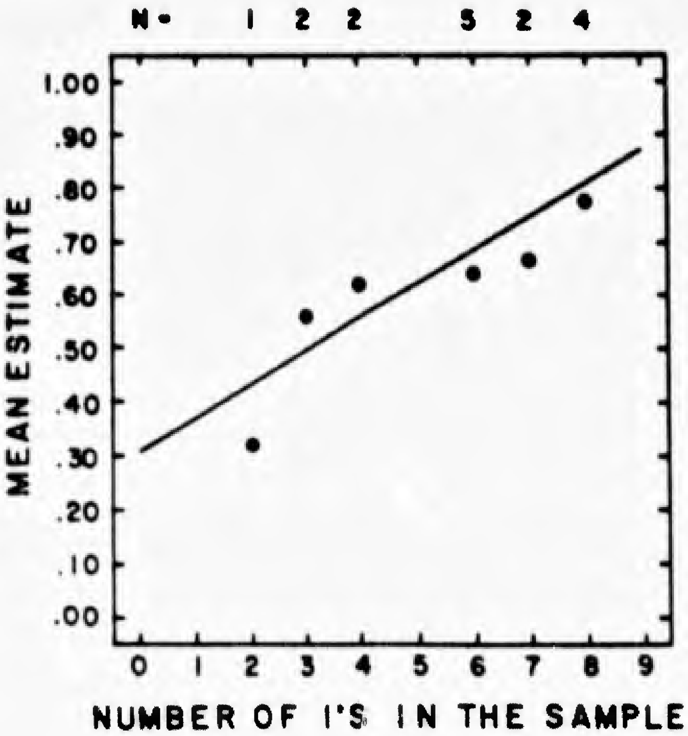

Figure 4. Mean Estimates for Last 15 Trials (Average Subject)

The subject's behavior appears to have changed somewhat from the beginning to the end of the experiment. The course of this change throughout the experiment can be studied by estimating the parameters of the subject's prior distribution for each of the 10 15-trial blocks. These results are shown in figure 5. Since I find it more natural to think in terms of the mean and variance of a distribution, I have inserted the estimates of the parameters, a and b, into the equation

$$m = \frac{a}{a + b}$$

for the mean, and into the equation

$$v = \frac{ab}{(a + b)^2 (a + b + 1)}$$

for the variance of a beta distribution. These estimated or inferred means and variances are plotted for each 15-trial block. The dashed lines represent the mean and variance of the stimulus distribution and might be taken as theoretical asymptotic values for these learning curves. The learning curve for the subject's predictions is similar to the one shown for the inferred means. It might be noted that this learning of the stimulus distribution is a learning-to-learn type of process relative to the "learning" occurring between the subject's prediction and his estimate.

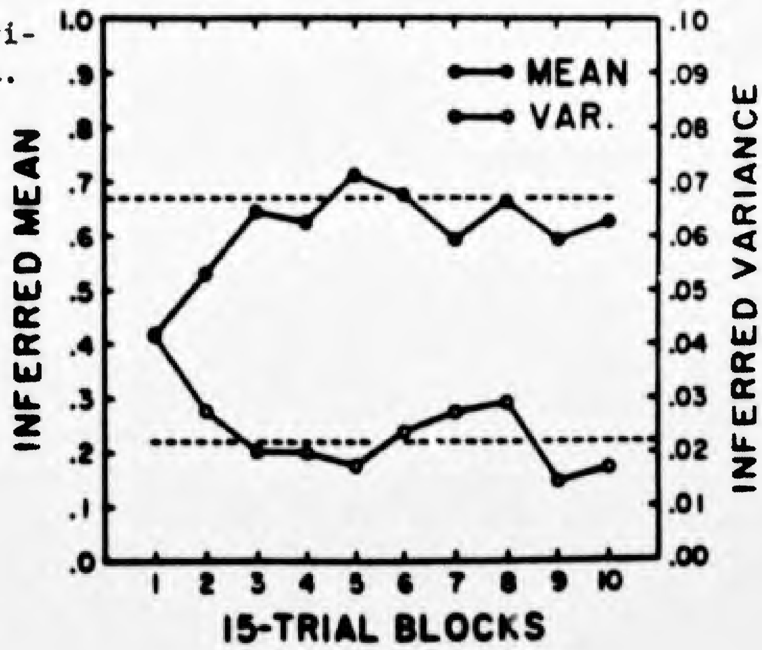

Figure 5. Estimated Parameters of the Subject's Prior Distribution (Average Subject)

At this point I feel compelled to admit that there were 21 other subjects in this experiment. Since we have a digital computer which makes data analysis very inexpensive, I could spend the next several days describing the data analyses of the other subjects. I hope that I might be excused from this task by saying that the behavior of some of the subjects looked somewhat better in terms of the Bayes model and the learning of the stimulus distribution while the behavior of some of the subjects looked very much worse, especially with respect to the learning of the stimulus distribution. An example of the latter type of subject is shown in figure 6. The behavior of this subject was quite variable and there is no evidence that he made use of the information provided by the stimulus distribution. In fact, it is possible to distinguish two groups of subjects, those who make use of the information provided by the stimulus distribution and improve their performance during the course of the experiment and those who do not make use of this information and show no improvement in performance. Further experimental sessions having different stimulus distributions a learning-to-learn type of effect occurs in that more and more subjects come to make use of the information provided by the stimulus distribution (ref. 19).

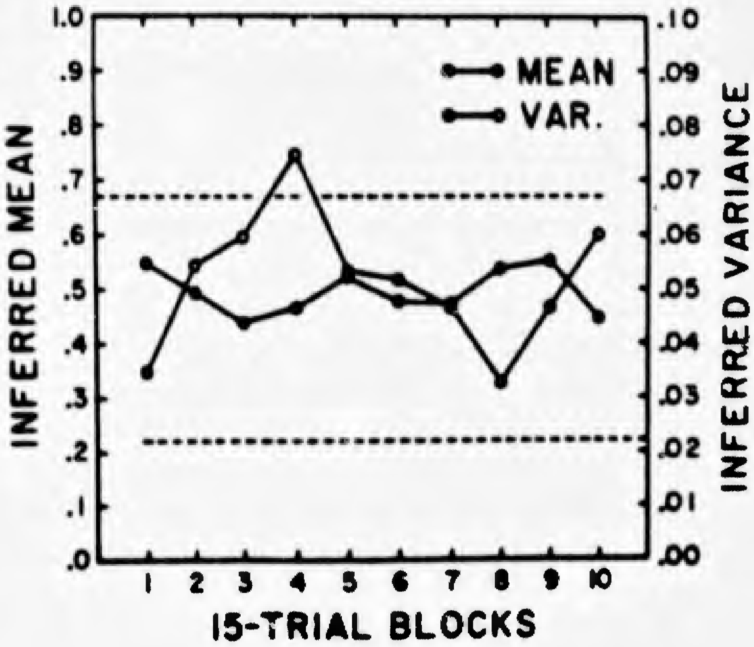

Figure 6. Estimated Parameters of the Subject's Prior Distribution ("Poor" Subject)

The learning of the stimulus distribution which has been demonstrated by this and other experiments (ref. 13, 19) is a second type of predecisional process. In contrast to the first type which deals with the effect of the addition of new information on the decision task and which is an integral part of the Bayes model, this second type of predecisional process deals with the accumulation of information at a higher level and thus lies outside of the elementary Bayes model. There is some reason to expect that it may be possible to develop a Bayes model for the learning of a stimulus distribution (ref. 19). Such a model would probably require that two "learning" parameters be derived for the subject. This implies that a total of four individual parameters would be needed to describe the behavior of a subject throughout the course of a prediction and estimation experiment such as that previously described.

I find it useful to distinguish between these two types of predecisional processes in this way. It is possible to look upon each trial as an "experiment" in which the state of nature is fixed while the subject gains more information about this state. Thus, the first type of predecisional process deals with the course of learning associated with the accumulation of information about this particular state. If each trial is interpreted as an experiment then the sequence of trials corresponds to a sequence in which the state of nature may vary from experiment to experiment. Thus, the second type of predecisional process deals with the course of learning associated with the accumulation of information about how the state varies from experiment to experiment.

I would like now to discuss a third type of predecisional process; one which lies outside of the elementary Bayes model but which is conveniently studied with the help of a Bayes model.

Consider having a subject judge which of two random matrices contains the larger number of *1*'s. The subject is paid a fixed amount for each correct judgment. Presumably the subject looks at one matrix, looks at the other matrix, and then compares his two percepts in order to decide which matrix contains the larger number of *1*'s.

Assume that the average percepts are described by the Bayes model previously described. Assume further that the information that the subject gains from looking at a matrix is equivalent to that yielded by a sample of n elements in the matrix. If the experimental situation is such that the subject can classify the stimuli into two or more sets then it is possible for the subject to have different prior probabilities for the different sets of stimuli. For example, if the distribution of the proportion of *1*'s for the matrices presented on the left in an experiment has a smaller variance than the distribution for the matrices presented on the right, the subject might be expected to notice this during the course of the experiment and to take account of this information in his giving his judgments. If this is the case, the subject's percept for a matrix on the left is, on the average, regressed more toward the mean of its prior distribution than is his percept for a matrix on the right.

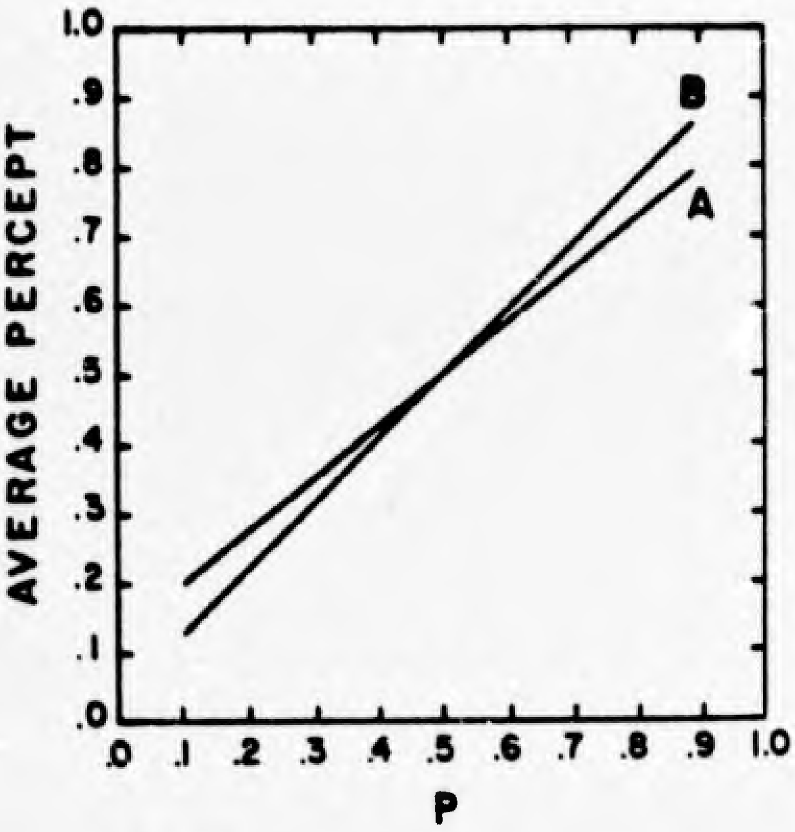

Figure 7. Two Possible Relations between Average Percept and Proportion of 1's. A: Prior Distribution with the Smaller Variance; B: Prior Distribution with the Larger Variance

Figure 7 shows two possible relations between average percept and the proportion of *1*'s in a matrix. The two lines have been computed from the Bayes model assuming that both prior distributions have a mean of .50. The line with the smaller slope is for the prior distribution with the smaller variance and can be taken to represent the subject's percepts for a matrix on the left while the other line represents the subject's percepts for a matrix on the right.

Suppose that the subject is shown two matrices which have the same proportion of *1*'s. If this proportion is *greater* than .50, the subject tends to perceive the matrix on the left as containing a smaller number of *1*'s than the matrix on the right. If this proportion is *less* than .50, the subject tends to perceive the matrix on the left as containing a *larger* number of *1*'s than the matrix on the right. The amount of this average bias increases with the distance between the mean and the actual proportion of *1*'s. When the proportion of *1*'s is equal to the mean of the prior distributions there is no regression and, thus, no average bias.

Now, consider changing the experiment so that the two matrices are presented successively rather than simultaneously. This, of course, is the traditional context for the study of time-order errors. The subject can, and probably will, distinguish two sets of stimuli, i.e., the matrices appearing first and the matrices appearing second. If the distribution of "first" stimuli is different from the distribution of "second" stimuli, we have the case described above and the resulting biases would be called time-order errors. However, in a well-controlled experiment using a completely counterbalanced design the two stimulus distributions are identical and would probably be so perceived by the subject. So, let us assume that the prior distribution for a "first" matrix is the same as the prior distribution for a "second" matrix and see if a Bayesian interpretation can still yield time-order errors.

There is some reason to expect that a subject is less certain about an event which has occurred at a remote time than about an event which has just occurred. This decrease in certainty with time might be considered a third type of predecisional process dealing with the value of information as determined by the time elapsed since receiving the information. In contrast to the first and second type of predecisional processes this third type is, in a sense, a forgetting process rather than a learning process.

This process suggests that a subject may be less certain about the proportion in the matrix shown first than about the proportion in the matrix shown second since his perception of the first matrix occurs at a time more remote from the time of decision. In terms of the Bayes model the sample size represents the certainty of a percept. Therefore the sample size, n_1, associated with a matrix shown first might be smaller than the sample size, n_2, associated with a matrix shown second. If this is the case, figure 7 can be used to represent this effect. Now, the line with the smaller slope is for the smaller sample size and can be taken to represent the subject's percepts for a matrix shown first while the other line represents the subject's percepts for a matrix shown second. If the two matrices have the same proportion which is *below* the mean of the prior distributions, the subject tends to judge the first matrix as having a *larger* number of *1*'s than the second matrix. If the proportion is *above* the mean of the prior distributions, the subject tends to judge the first matrix as having a *smaller* number of *1*'s than the second matrix. If the proportion is equal to the mean of the prior distributions, there is no time-order error. In general, this interpretation predicts a reversal in the direction of the time-order error as proportion is varied from zero to one.

Notice that this interpretation is not the same as the usual fading stimulus trace theory of time-order error. According to the traditional theory the subject's percept of the proportion of *1*'s decays toward zero with time. Since the first matrix contains the same proportion of *1*'s as the second matrix, the two percepts start decaying from the same initial value. However, at the time of decision the percept of the first matrix has been decaying for a longer period of time and thus has a smaller value than the percept for the second matrix. Therefore, the subject judges the first matrix as containing a smaller number of *1*'s than the second matrix. This effect is independent of p and depends only upon the time intervals involved. Thus, there can be no reversal in the direction of the time-order error as predicted by the Bayesian interpretation.

It should be noted that this Bayesian interpretation does contain some elements of a fading stimulus trace theory in that the percept can be considered to be "fading" with time. However, this type of fading leads to increased uncertainty on the part of the subject as to the proportion of *1*'s in the matrix which is interpreted in the model as a decrease in sample size.

The Bayes model can be used to derive theoretical values for the percentage of time the subject selects the first matrix over the second matrix. The model requires four parameters. Two of the parameters, a and b, describe the subject's prior distributions. The other two parameters, n_1 and n_2 represent the sample sizes for the first and second matrices respectively. Figure 8 shows four theoretical functions for the relation between choice percentage and proportion when both first and second matrices contain the same proportion of *1*'s. In computing these functions it has been assumed that the subject's prior distribution is the same for both first and second matrices and that it is rectangular with parameter $a = b = 1$. The sample for the second matrix has been held constant at 24 while the sample size for the first matrix has been given the values 0, 2, 12, 24 to obtain the four functions. There is, of course, no time-order error when $n_1 = n_2$ but notice that when n_1 is less than n_2 the direction of the time-order error reverses as p increases from zero to one and that there is no time-order error at $p = .50$, the mean of the prior distributions. Since our choice of values for n_1 and n_2 are somewhat arbitrary, it should be noted that larger values yield smaller time-order errors and smaller values yield larger time-order errors.

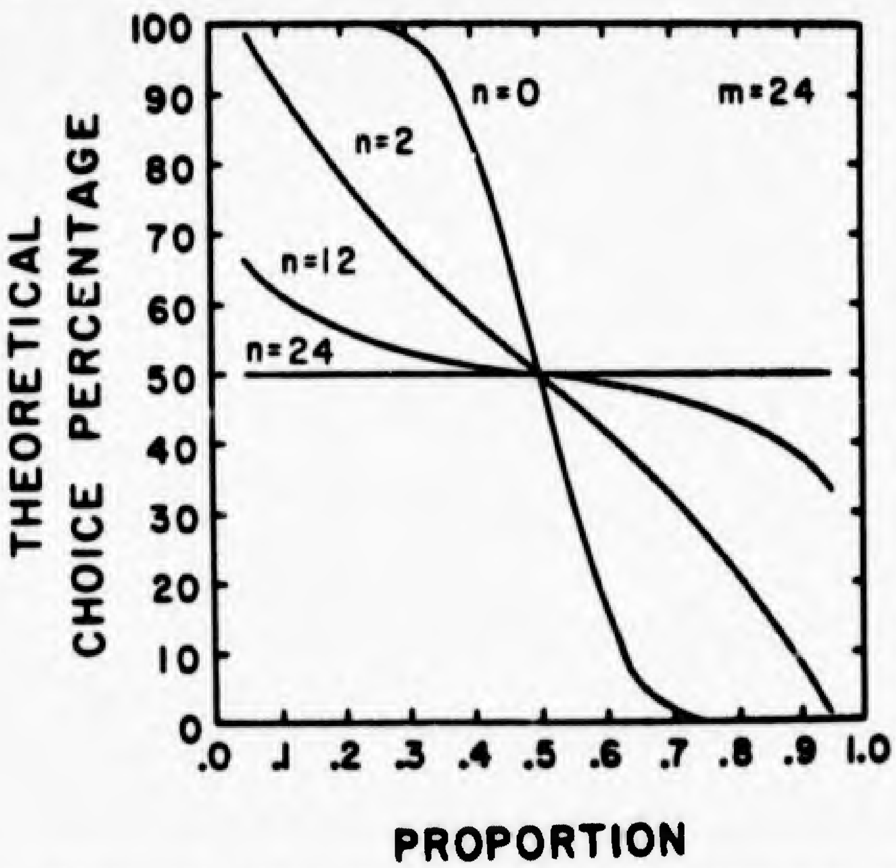

Figure 8. Four Theoretical Functions between Choice Percentage and Proportion of 1's. Same Proportion for Both Matrices; Prior Distributions Assumed Identical

Unfortunately, we cannot use these functions to make quantitative predictions about the behavior of a subject since we do not have a method for determining the values of the four parameters for a given subject prior to the experimental session. In other words, we do not have a quantitative model of this third type of predecisional process. The two parameters of the subject's prior distribution pose no special difficulty since we can use an approximately rectangular stimulus distribution and be reasonably certain that the subject's prior distribution is close to a rectangular distribution. On the other hand, n_1 and n_2 present great difficulties. The sample size for the second matrix might be assumed to remain constant regardless of the interstimulus interval. However, because sample sizes vary dramatically from subject to subject (refs. 11, 13) a prior estimate must be obtained for each subject. Even if this were done, we would still be left with the problem of obtaining a prior estimate for the sample size of the first matrix which, according to our hypothesis, depends upon the interstimulus interval.

Though we cannot predict the choice percentages of a given subject we can use the Bayes model and the assumption that n_1 decreases as the interstimulus interval is increased to make some qualitative predictions about the behavior of a subject in this choice situation. For example, the horizontal line obtained when $n_1 = n_2$ can be taken to represent the behavior of a subject when the interstimulus interval is zero. This, of course, is impossible so that the horizontal line really represents an extreme limit on the shape of the relation between choice percentage and proportion. The function for $n_1 = 12$ might represent a relatively short interstimulus interval; the function for $n_1 = 2$, a long interstimulus interval. The function for $n_1 = 0$ represents another extreme limit in that the subject retains no information about the proportion of 1's in the first matrix. In this case the subject compares his percept of the second matrix with the mean of the prior distribution.

Let us look at the data for two subjects as shown in figure 9. Each graph shows the percentage of times the subject chose the first matrix as containing the larger number of 1's plotted as a function of the proportion of 1's. Therefore, these graphs can be compared with the theoretical functions given in figure 8. The three graphs for each subject represent three different interstimulus intervals, 5, 2, and 1/2 seconds. The data for the 5- and 2-second intervals were obtained during the same experimental sessions. Each data point is based on 30 choices. The data for the 1/2 second interval were obtained during later experimental sessions. Each of these data points is based on 80 observations.

Notice that the behavior of subject RW is relatively insensitive to the duration of the interstimulus interval. Though time-order errors are present and in the predicted direction, there is not much difference between the graphs. The behavior of subject AM is at the other extreme in terms of sensitivity to the duration of the interstimulus interval. The 5-seconds interval yields a very strong time-order effect suggesting that much of the information about the first matrix has been lost. The 2-seconds interval yields almost no time-order error suggesting that approximately equal amounts of information have been retained for the two matrices. The 1/2-second interval yields a time-order effect in the reverse direction suggesting that more information is retained for the first matrix than for the second. This result was unexpected. It may be that an after image of the first matrix interfered with the perception of the second matrix. At any rate, this phenomenon is consistent with the Bayes model but not with the hypothesis that retained information necessarily decreases as the interstimulus interval is increased.

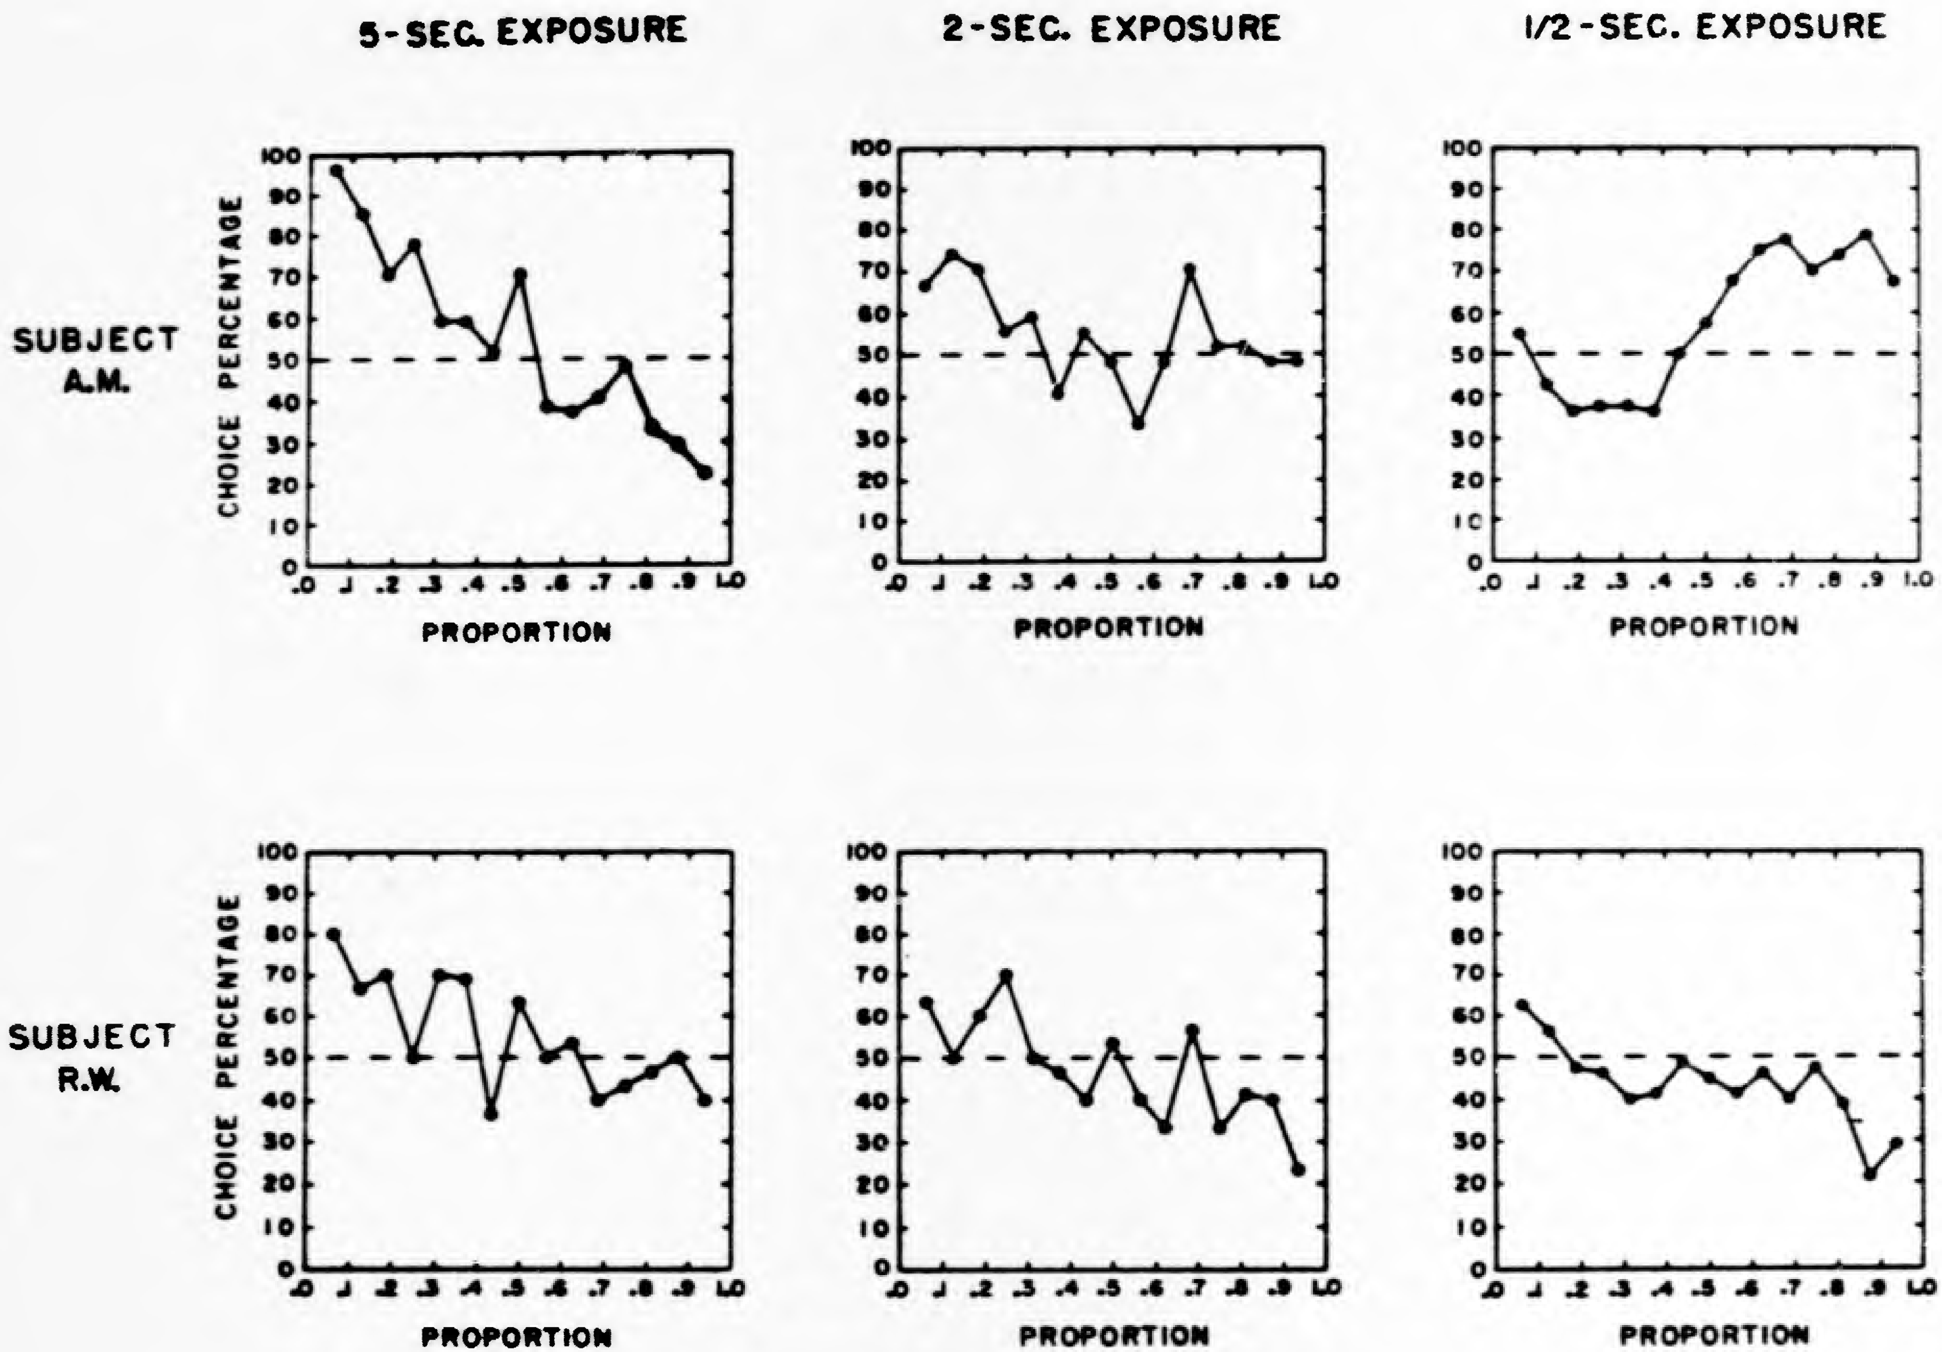

Figure 9. Choice Data for Two Subjects

I was the other subject in this experiment. My behavior was moderately sensitive to the duration of the interstimulus interval with the amount of time-order error decreasing to about zero for the 1/2-second interval.

The data of all three subjects yield curves which are not too different from the theoretical curves derived from the Bayes model. All of the empirical curves either indicate a reversal in the direction of time-order error near the middle of the stimulus distribution or else indicate no time-order error. The most frequently observed difference between the empirical and theoretical curves is that some of the empirical curves do not turn down or up near the ends to approach a choice percentage of zero or 100. Analyses of other data obtained from these subjects indicate that sample size is not the same for all proportions, but instead, is a *U-shaped* function of proportion. This finding agrees quite well with my experience as subject in the experiment. Because of the regularity of the location of the elements in the matrices, one sees or takes in a larger area of those matrices which contain almost all *1*'s or almost all *0*'s. In fact, if one observes a matrix containing all *1*'s or a matrix containing all *0*'s, it is immediately apparent that the matrix contains all *1*'s or no *1*'s. If larger sample sizes for the proportions near zero and one are used in the Bayes model, the theoretical curves take on an appearance similar to some of the empirical curves.

It is apparent that it is a difficult task to obtain quantitative predictions for the behavior of a subject in an experiment dealing with time-order and related errors. However, a Bayes model appears capable of making qualitative predictions and yielding some insight into the processes involved in these experiments.

Now, I would like to leave behind the discussion of experimental results and enter the realm of speculation. First, I will suggest a decision-theoretic model for response probabilities. Second, I will attempt to relate the subjective feeling of confidence to a predecisional process. Finally, I will comment about the use of computers in studying predecisional processes.

Most decision-theoretic models are deterministic in the sense that once the expected utilities of the courses of action have been determined the subject is expected to respond with that action which has the largest expected utility. If the identical choice situation were repeated indefinitely, he continues to respond with probability one with the same optimal course of action each time.

Our Bayes models have, up till now and in this sense, been non-probabilistic. When a subject is allowed to look at only a specified sample of, say, 9 elements from the large matrix he can count the number of *1*'s. In this case we have arbitrarily assumed a probability process in order to perform statistical analyses of the data. (This aspect of the situation is somewhat improved when the subject is allowed to look at the complete matrix in that we can assume that he observes a sample of the 256 elements in the matrix and that this sampling process induces a probability distribution over his responses.)

The rather arbitrary choice of a response distribution is quite unsatisfactory. On the other hand, the implication that the subject "knows" the posterior distribution and, thus, "knows" the expected utilities associated with each course of action and, as a consequence, responds with probability one is a rather unpalatable characteristic for a descriptive model of behavior. Further, some real benefits might result from deriving an explicit decision-theoretic model for response probabilities. Let us examine one such model.

In a simple two-alternative decision problem such as a paired comparison judgment task in which the utilities remain essentially constant from trial to trial, a very simple Bayesian decision rule can be derived. This same rule applies on every trial of the experiment. What is the rule? The utilities are used to determine a constant, c. The rule can be stated as-if the mean of the posterior distribution is greater than c, choose the first alternative; if the mean of the posterior distribution is less than c, choose the second alternative (ref. 10). This rule always selects the alternative that has the largest expected utility. Since this expectation is taken with respect to the posterior distribution which may change from trial to trial, the rule may yield different alternatives on different trials. However, this is still not a response probability model in the sense intended.

This rule implies that the subject has perfect knowledge of a parameter of the posterior distribution, the mean of this distribution. A possibly more realistic assumption is that the subject has available a process which, in effect, estimates the mean of the posterior distribution. The average of a random sample of k observations from the posterior distribution serves as an estimate of the true mean. This estimate is inserted directly into the decision rule, i.e., if the average is greater than c, the subject chooses the first alternative; if the average is less than c, the subject chooses the second alternative.

This averaging process seems to be a reasonable assumption--at least psychologists are quite prone to assume the existence of averaging mechanisms. The random sample generator with probabilities modified by previous experience, stimulus uncertainty, and other factors which determine posterior probabilities, might be thought of as an extension of a traditional learning theory in which response probabilities are modified of the basis of past experience or as deriving from a field-theoretic neurophysiological explanation of behavior.

Notice that the modified decision rule based on the use of an estimate of the mean of the posterior distribution does not always choose the optimal alternative. If the first alternative is best, i.e., the true mean of the posterior distribution is greater than c, the modified decision rule will choose this optimal alternative with probability, π where π is the probability that the average of k samples from the posterior distribution is greater than c. The other, non-optimal, alternative is chosen with probability $1 - \pi$. Therefore, the average expected utility to the subject lies between the larger expected utility of the optimal alternative and the smaller expected utility of the non-optimal alternative. The larger π becomes, the closer the average expected utility yielded by the decision process approaches the expected utility of the optimal alternative. Obviously, the subject can attain the maximum expected utility only if $\pi = 1$ and this can be achieved in general only by taking a sample of infinite size from the posterior distribution.

Of course the subject does not take an infinite sample and the most natural way of accounting for this is in terms of the cost of sampling or, more generally, the cost of thinking. The introduction of this assumed sampling process and the notion of the cost of thinking leads to a basic reformulation and expansion of the decision problem. The decision task is no longer just the choice between the first alternative and the second alternative, but involves an implied choice as to the number of observations taken from the posterior distribution. After the sample is taken, the decision rule is applied to choose between the first and second alternatives of the basic decision task which I will call the terminal decision (ref. 3). Inherent in this formulation is the fact that average expected utility of the terminal decision is a monotonic increasing function of k, although each successive increase will be smaller and smaller because the number of observations, k, has a diminishing effect on the probability of choosing the optimal act. In contrast to this effect, the cost of thinking, which must be subtracted from the average expected utility of the terminal decision, increases as a linear function of k. Thus, there is usually a sample size which is optimal for the given posterior distribution.

Notice that the choice of k is determined completely by two factors, the average expected utility of the terminal decision and the cost of thinking. The cost of thinking is a function only of sample size: it is independent of the posterior distribution and of the other utilities involved in the decision task. On the other hand, the average expected utility of the terminal decision is determined by a number of factors. The most obvious determinant is the size of payoffs utilized in the experiment. The larger these payoffs the larger is the average expected utility and thus, the larger is the optimal k. A less obvious determinant is the variance of the posterior distribution. When this variance is large, the probability of choosing the optimal alternative tends to be near $1/2$ and thus can be increased more easily by additional sampling from the posterior distribution. Therefore, the larger the variance of the posterior distribution, the larger is the optimal k.

In giving an empirical interpretation to this phenomenon, it might be expected that both the probability of a correct response and reaction time will increase with the number of observations from the posterior distribution. Therefore, the following predictions can be derived for a paired comparisons experiment with the random matrices.

1. Increasing the monetary payoffs given a subject for a correct judgment increases both the probability of a correct judgment and reaction time for a given pair of proportions.

2. Since the variance of the posterior distribution tends to be larger for proportions near .50 than for more extreme proportions, reaction time will be longer for pairs of matrices with proportions near .50.

3. Since the variance of the posterior distribution is partially determined by the variance of the prior or stimulus distribution, reaction time for a given pair of proportions will be longer when a wide range of stimuli is employed.

4. Learning and extinction-like effects will occur as the subject learns a stimulus distribution which differs from his initial prior distribution.

This response probability model can also be used to obtain other predictions, e.g., increasing the number of response categories and the monetary payoffs serves to increase the amount of transmitted information (ref. 2). All of these predictions await experimental confirmation.

It has been implicitly assumed in the response probability model that the subject has available some information concerning the variance of his posterior distribution. It would be inconsistent to argue that the subject "knows" this variance since I have maintained that it is unreasonable to assume that the subject "knows" the mean of the posterior distribution. The subject can only observe the average of k random samples drawn from this distribution. It might be assumed that the subject uses these k samples to estimate the variance in the same sense that he estimates the mean of his posterior distribution. Such a process might or might not exist. At any rate, the subjective correlate to the variance of the subject's probability distribution, whether prior or posterior, is almost certainly a feeling of confidence.

I have observed this feeling of confidence myself while participating as a subject in the time-order error experiment. When the proportions in the two matrices are near .50 there often occurs a feeling of low confidence which results in an almost irresistable impulse to suspend judgment and to try to obtain additional information. When the two proportions are near zero or one but are not identical there occurs a feeling of high confidence, tendency to respond quickly and vigorously, and a wish that one could back up his judgment with a $100 bet. In these cases it is clear that confidence reflects the smallness of the variance of my posterior distribution.

Let us distinguish between the basic terminal decision problem, e.g., judging which matrix contains the larger number of *1*'s, deciding whether or not to accept the offer of x dollars for your used station wagon, or deciding whether to classify a blip on the radar scope as a hostile or friendly aircraft, and the expanded decision problem where the choices are between deciding *now* to make a terminal decision based on the presently available information and *waiting* to obtain more information before making the terminal decision. This expanded decision problem might be considered a predecisional process relative to the terminal decision.

In general, the larger the variance of the probability distribution in the initial terminal decision problem the higher is the expectation of gain by holding off the terminal decision and trying to obtain more information. This additional information serves to decrease the variance of the probability distribution of the terminal decision problem (ref. 3). If the cost of waiting and obtaining additional information is not too large, it is wise to hold off on making the terminal decision. On the other hand, if the variance of the probability distribution is quite small, there is nothing to be gained by obtaining further information unless the utilities involved in the terminal decision are so large that even a very small relative gain in expected utility can offset the cost of a very large amount of thinking, information seeking, research, etc. Therefore, the occurrence of a large variance, or equivalently great uncertainty as to the state of nature, should lead the decision-maker to delay his decision and to seek out more information or to reformulate the decision problem. It appears not unlikely that through extensive experience with the environment or through innate organization of the nervous system, man has come to encode uncertainty as to the state of nature as a feeling of confidence loaded with emotional overtones in order to serve as an affective warning to postpone action and to think.

It appears that the validity of this proposed relation between a feeling of confidence and the variance of a posterior distribution can be evaluated empirically by studying the nature of the relation between confidence, as measured by the technique developed in the area of speech intelligibility and communication, and the variance of a posterior distribution, as inferred from our Bayesian interpretation of psychophysical judgment.

This interpretation of confidence also suggests some other empirical phenomena. For example, an individual might be expected to have a generalized strategy of obtaining additional information until his feeling of low confidence is reduced to a certain, acceptable, level. Since low confidence is related to uncertainty as to the value of a state and is not directly related to the payoffs or to cost of information in a decision task, the amount of information that the individual attempts to obtain may be relatively independent of the expected gain associated with obtaining additional information. Thus, we may find in experiments dealing with the amount of information obtained before making a decision, that individuals are relatively insensitive to the payoffs of the terminal decision and to the cost of additional information. This should be especially evident in experiments in which the payoffs and costs are small or which deal with perceptual or other processes which are relatively stable because they are used in the everyday life of the individual. Due to their different life histories, some individuals may consistently gather too little information while others gather too much. It should be noted that such a generalized strategy which leads to non-optimal behavior in a specific experiment might be considered optimal in the larger sense of serving to reduce the cost of thinking to the individual. The strategy of obtaining information until the feeling of low confidence reaches an acceptable level as determined by previous experience with the environment can substitute for a careful analysis of the payoffs and costs involved in a decision task. If these utilities are small relative to the cost of thinking or if the generalized strategy serves as a close approximation to the strategy which is optimal for the particular decision task then the individual is not losing anything and, in fact, may gain by employing the generalized strategy. Of course there are many decision tasks where such a generalized strategy is clearly non-optimal in either a large or a small conceptualization of optimality.

In summary, I have touched upon five predecisional processes related to psychophysical judgment. The first two I have described as learning processes; one reflecting improvement in performance due to additional information about the value of a particular state, the other reflecting improvement in performance due to additional information about the distribution of a sequence of states. The third I have described as a forgetting process reflecting changes in the value of information with time. The fourth deals with the efficiency of utilization of existing information and yields a model for response probabilities. The fifth deals with the amount of information required for a decision as related to uncertainty about the value of a particular state.

Though I have discussed these processes primarily in the context of psychophysical experiments, I hope you will find that they are relevant to a much broader range of decision tasks.

REFERENCES

1. Luce, R. D., and H. Raiffa, Games and Decisions, Wiley, New York, N. Y., 1957.

2. Miller, G. A., "The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information," Psychological Review, Vol 63, pp 81-97, 1956.

3. Raiffa, H., and R. Schlaifer, Applied Statistical Decision Theory, Harvard Business School, Boston, Massachusetts, 1961

4. Restle, F., Psychology of Choice and Judgment, Wiley, New York, N. Y., 1961.

5. Schelling, T. C., The Strategy of Conflict, Harvard University Press, Cambridge, Massachusetts, 1960.

6. Schlaifer, R., Probability and Statistics for Business Decisions, McGraw-Hill, New York, N. Y., 1959.

7. Schlaifer, R., Introduction to Statistics for Business Decisions, McGraw-Hill, New York, N. Y., 1961.

8. Shuford, E. H., A Comparison of Subjective Probabilities for Elementary and Compound Events, Psychometric Laboratory Report Number 20, University of North Carolina, Chapel Hill, North Carolina, 1959.

9. Shuford, E. H., "Percentage Estimation of Proportion as a Function of Element Type, Exposure Time, and Task," Journal of Experimental Psychology, Vol 61, pp 430-436, 1961.

10. Shuford, E. H., Applications of Bayesian Procedures Based on Discrete Prior Distribution, Psychometric Laboratory Report Number 31, University of North Carolina, Chapel Hill, North Carolina, 1961.

11. Shuford, E. H., and W. J. Hall, A Decision Theory Approach to Psychophysics, Psychometric Laboratory Report Number 22, University of North Carolina, Chapel Hill, North Carolina, 1959.

12. Shuford, E. H., L. V. Jones, and R. D. Bock, "A Rational Origin Obtained by the Method of Contigent Paired Comparisons," Psychometrika, Vol 25, pp 343-356, 1960.

13. Shuford, E. H., and R. A. Wiesen, Bayes Estimation of Proportion: The Effect of Stimulus Distribution and Exposure Time, Psychometric Laboratory Report Number 23, University of North Carolina, Chapel Hill, North Carolina, 1959.

14. Swets, J. A., W. P. Tanner, Jr., and R. G. Birdsall, "Decision Processes in Perception," Psychological Review, Vol 68, pp 301-340, 1961.

15. Toda, M. Chapter in M. Toda, Y. Umeoka, and P. Suppes, (Editors), Models and Strategies in Psychology: Contributions by Japanese Mathematical Psychologists, Stanford University Press, Stanford, California, 1963.

16. Toda, M., "Subjective Inference Versus Objective Inference of Sequential Dependencies," Japanese Psychological Research, Vol 5, pp 15-37, 1958.

17. von Neuman, J. and O. Morgenstern, Theory of Games and Economic Behavior, Princeton University Press, Princeton, N. J., 1947.

18. Wiener, N., "Some Moral and Technical Consequences of Automation," Science, Vol 131, pp 1355 1358, 1960.

19. Wiesen, R. A., and E. H. Shuford, "Bayes Strategies as Adaptive Behavior," in E. E. Bernard, and M. R. Kare, (Editors), Biological Prototypes and Synthetic Systems, Vol 1, Plenum Press, New York, N. Y., 1962.

BLANK PAGE

DISCUSSION OF SHUFORD'S PAPER

ROBY (Discussant). My first comment is a seconding of your point on the virtues of a decision model. You said, "The second reason for using decision theory as a descriptive theory of human behavior derives from the fact that it is an optimality model or a model which represents efficient behavior." I think this is an easy point to overlook, and I would like to underline it because it's a very important point. There is a great deal of resistance to using a normative model of this kind as a descriptive model. I think people don't often consider the alternative of what kind of a base line model you can use, assuming that you are really interested in behavior which does represent departures from this. The alternative base line seems to be a model that assumes either that people are utterly inert except for what you predict from these other processes, or that they respond strictly at random. In the physical sciences we have a lot of examples of using an ideal model, which corresponds to a normative model, and cutting down on the range of filling in that you have to do in order to account for departures from the model.

My second comment is this. Probability has been around a long time, and it's been recognized for a long time that probability is very important in guiding human behavior. It's not much of a step from there to the notion that conditional probability is important, and that Bayes' theorem applies. What I am unclear about is what you mean by Bayesian theory with a capital B. If you want to apply Bayesian models, what are you excluding as alternatives? Is it simply an exclusion of the minimax type of assumption, or do you mean for it to have some more restrictive use? If it is an exclusion of the minimax assumption, let's say, how do you get to the Bayesian model as superior to the minimax assumption, when it does in fact apply for game theoretical type situations where you have a hostile environment?

SHUFORD. When I speak of Bayesian models, I mean that I am selecting from a class of models defined as those that layout states of nature, probabilities over the states of nature, alternative courses of action, and consequences of courses of action expressed in utilities. Bayes' Theorem is a natural thing to use in order to modify the probabilities of the states of nature with new information. Of course, this is a wide class of models. The La Placian criterion assumes equally probable states of nature; the minimax criterion assumes that the worst possible thing is going to happen. These are special cases of the Bayesian model.

LANZETTA. I was curious about the Bayes model also. Are you embracing, within the notion of Bayesian models, the explicit way in which probabilities are modified as a function of information, or simply suggesting that such probabilities are modified? If the latter, then of course this is not much different than looking at it as set, or prior hypotheses, and how these influence the weight people accord to any new information that they get.

You mentioned that you were being speculative about the relationship between uncertainty and search behavior, and possible emotional concomitants. I think these are not quite as speculative as you are implying. There is considerable data at both the physiological level and the behavioral level indicating that it is very likely that uncertainty is related to the arousal state. Dr. Berlyne has done a fair amount of research primarily on that point.

SHUFORD. It is useful to distinguish two purposes of a model. The decision theory model is unique in its normative aspect. The Bayes model is not normative. The Bayes model as a mathematical theory of probability says only that probability should be modified. As a descriptive model of how probabilities are modified, I suggest that we try Bayes' Theorem. It's an empirical matter whether people really do modify this way or not. Certainly, Bayes' Theorem doesn't do everything for us. There are some very important unsolved problems concerned with Bayesian statistical inference, that is, the question of how can you combine information from different experiments dealing with rather different aspects of one theory? How does this increment the probability of the theory? You don't have a normative model for that.

LANZETTA. In a gross way, I think, the Bayesian model is already in accord with a lot of data in the descriptive sense. We seem to know that when people have strong hypotheses they are less prone to modify those hypotheses on the basis of some new information. If they are weak hypotheses, they are much more prone to react to small bits of external data. This is exactly what you would predict from some sort of Bayesian model.

EDWARDS. I think everybody is in agreement with each other, and I want to comment on this interchange just so it doesn't look as though everyone who espouses Bayesian points is now in monolithic agreement with one another on all matters. In the portion of your remarks, Emir, which were concerned with the notion of probability of choice, that is, stochastic models generally, you were discussing the thought that you didn't really know what your posterior distribution was, that you somehow attempted to discover this by sampling the inside of your own head, or something of the sort. I don't know what a probability means to you, but to me it means an opinion, and Bayes' Theorem, in fact, the whole Bayesian approach to information processing is, for me, nothing other than a formal model for the revision of opinion in the light of evidence. So that to say you don't know your posterior distribution is to say that you don't know what your own opinions are. Well, people sometimes say that; they are willing to say, "What do I really think about whether Kennedy is going to win the next election?" In a sense, that is self interrogation about your own opinion, and surely self interrogation is an intelligent, meaningful procedure and does give you information which, in one sense of the word, you didn't have before. Yet at the same time, you suppose that it's meaningful to think about somehow having a posterior distribution and not knowing what it is, and therefore being uncertain about it, and having that uncertainty be describable by another distribution. This last step worries me. I don't understand what the posterior distribution you might currently be entertaining would mean. I can't phrase my problem more precisely than that because I don't think it's a precise problem; it's a problem in the fundamental conception.

SHUFORD. The subject, in effect, samples from a posterior distribution. Granted, this generates another distribution, but it is completely specified by the posterior distribution that is sampled.

EDWARDS. But what *is* the distribution that he is sampling?

SHUFORD. I think what is at issue is this: We can say that the subject knows his posterior distribution, and we can observe it or obtain some slight bit of information about his posterior distribution by having the subject make a choice. Now, when we observe it again it turns out to be different. We observe it again and it turns out to be different again, so we have a sequence of different probability distributions. Most people would say, "Well, there's some error coming in here, the subject really doesn't change his opinions that drastically from trial to trial." So you have a choice of the way you can build error into the model. You can arbitrarily say that this is normally distributed or some other kind of error, and try to get at the true posterior distribution, or you can try to use some kind of more suggestive model for this. This suggests some interesting things about reaction time and proportion of correct responses. But one way or another, you have the problem of error to contend with.

BERLYNE. One of the questions that is bound to come up, and ought to come up, in a symposium like this is, of course, the relationship between a normative theory of decision making and a descriptive theory. Today, already, we have had two arguments, if I understand them correctly, for taking a normative theory as a starting point for a descriptive theory. Earlier we had Dr. Edwards' assertion that 95 times out of a hundred, the way people actually make decisions isn't very far from the way they ought to make decisions. Now, we have Dr. Roby's rhetorical question, "If we don't use the normative theory as a base line, what can we use?" I'd like to suggest that there might be some dangers in this approach. First of all, why do you take a normative theory as a base line? It depends partly on the questions you want to ask of your descriptive theory. If you have people making decisions on your behalf, you might want to know, before anything else, how far their decision making deviates from optimality in any one direction because you want to know what you are going to get out of it or what you are going to lose by it. If, on the other hand, you are interested in improving decision making processes so that actual decision making can come to resemble optimal decision making more than it does, then I sense that there is an implicit questionable assumption here that the way to bring actual decision making nearer to optimal decision making is to begin by studying ways in which actual decision making differs from optimal decision making. This, I suggest, is questionable for a number of reasons. One thing is this. If you start doing experiments with optimal theories as bases, you are going to study the effects on actual decision-making of the variables that determine which decision is optimal. Now, it is likely, a priori, that these variables are among the most important variables that determine actual decision making, but this, in itself, is questionable before it has been demonstrated. Second, it is more than questionable, that these are the only factors that play an important part in determining actual decision making.

SHUFORD. I think we are in some agreement; I don't know anybody who would claim that probability or utility are the only variables that are important. We are trying to build descriptive models, but we aren't describing in detail all aspects of behavior. We are saying that these tend to be important aspects of behavior that we want to describe. There's a lot left over there as exemplified by the betting experiments.

LANZETTA. I think a more serious criticism that Dr. Berlyne raised is whether it is possible to discuss what you might even mean by an optimal model. We are all suggesting here that such a model must obviously reflect the information base that the subject has at the time he makes a decision. We haven't the faintest notion how to even describe that information base, or how to describe its modification as a function of current information. I don't know how you can talk about a rationality criterion other than with respect to the information base. The only other way is to go outside the subject altogether, say, with respect to some little god up here who has assigned utility and knows the probabilities and accepts this particular criterion of rationality. It is impossible to even define what you might mean by optimal model, it seems to me.

SHUFORD. That applies whenever anybody calls anybody else's behavior rational or irrational. You're acting as a god, using your utilities and probabilities with respect to the other person. Quite frankly, we do not have optimal decision theories for many complex situations. The basic concept is okay, but working out the model is something else again.

EDWARDS. This point could go too far in the other direction. To imply that we have no basis for speaking about optimal behavior is, I think, to carry permissiveness too far. I would be content to dispute as vigorously and as long as necessary to assert that, for example, intransitivities are non-optimal, and if somebody comes up with an intransitive preference structure and asserts that it is optimal, I will invite him to define his assertion and proceed to take away all the money he owns. If he prefers A to B, B to C, and C to A, and will pay me a little each time to take away one and give him the one he prefers, we can go around the circle indefinitely until he is bankrupt. I just don't think that is optimal.

LANZETTA. It is not optimal if I have to choose them simultaneously; but if I do it in succession my utilities might change. If I have had A, I may prefer not to have A again, therefore the next time around, I would get an intransitivity because my preference or utilities have changed.

EDWARDS. I think you are agreeing with me because you are saying that your preferences should not be intransitive unless your utilities change, which is to say you agree that intransitivities are undesirable. You don't want a stable permanent structure which includes intransitivities.

LANZETTA. I am implying that you can't evaluate what you mean by intransitivity, other than mathematically, unless you assume simultaneous choice. The term has no meaning for successive choices. I am never intransitive in that at any one moment I do prefer one to another, or I'm indifferent, period. But as soon as you have successive choices, then the concept of intransitivity, in itself, implies the notion of a stable utility scale. Otherwise, the term has no meaning, so I think I would turn your argument right around the other way.

BERLYNE (To Edwards). Is your point that intransitivity, as such, is undesirable, or is it that if somebody has an intransitive structure, you can always beat him in some kind of gambling game? Suppose you have the situation of an art dealer trying to buy a picture, and it so happens when he is shown Pictures A and B, he prefers Picture A, and so on. You can have intransitive preferences, because the preference depends on the pair that he is shown.

EDWARDS. My point is simply that any intransitive stable preference pattern makes you so vulnerable to anyone who would like to exploit you, that if you have one, you darn well better keep quiet about it. I don't think, in other words, that the two thoughts you were suggesting that I might be entertaining are independent of one another. The fact that it's undesirable in a gambling game is only one of many good reasons not to have intransitivities.

FLOOR. I think it might be interesting to look into what kinds of additional information we can give subjects in order to produce rationality. For example, in the instance of intransitive choices, if such intransitivities are pointed out to them, do they still remain, or is there some revision of the choice pattern? Or in making a decision to buy an automobile, if we point out that the lovely young lady in the evening gown does not come with every Ford, or if they prefer a Cadillac, that it costs a great deal to keep it up, do these choices still remain? It occurred to me when Dr. Edwards was talking earlier, that while not all of us are rational, most of us would like to be more rational, and that perhaps the rationality of human beings is a preference for being rational and that is given information to aid rationality they would attain such states. I'd say that those who argue for human irrationality would be well advised to show instances where irrationality is consciously and deliberately preferred.

BLANK PAGE

PRE- AND POSTDECISIONAL PROCESSES OF THE "FUNGUS-EATER"*

Masanao Toda
Harvard University

I will talk today on the pre- and postdecisional processes, not of humans, but of robots called fungus-eaters (see Toda, refs. 1, 2). Fungus-eaters are the robot miners working on a planet outside the Solar System to pick uranium ores, while sustaining themselves by eating wild fungi growing on the planet. This is obviously a science-fiction stuff, and there are two reasons why I like a science-fiction type stage setting for my investigation. First, the stage is used to run subjects in decision making games, just like one runs subjects in war-like setup, or stock market-like situation. But I like science-fiction type stage more than anything else, since there in Outer Space we are free. I can set up any kind of environment I like there, and fortunately, the present-day subjects are usually ready to accept any image of Outer Space with a certain sense of reality as far as it is described to them with a sufficient realism. The second reason is that, in science-fiction, we can not only design environment but also man himself in a form of robot. So, the both ends--organism and environment--being theoretically cleaned up, we would be able to derive, at least in principle, a rigorously formulated theory of behavior. You may wonder what is the use of this; it is a theory of robot's behavior and not of human behavior. To answer this question, just imagine you have a model of man. Any model of man is not man himself. And if the model is really good, then you should be able to design a robot according to the model.

At the present stage, however, the focus of my theoretical attention has not yet reached the point of aiming at a plausible model of man. Both of my fungus-eaters and the kind of environment I am assuming are so simple that there is little element of simulation. But my primary attention is on the clarification of some of the basic concepts in psychology. For example, the questions like "What is perception?" "What is learning?" and "What are the functional relationships between perception and learning?" have been discussed many times by many psychologists, but these issues have inevitably been trapped in the immense complexity of human being and the human environment. But there is a hope that these questions can be answered in the neatly defined small world of fungus-eaters.

*This work was done during the author's visit at the Center for Cognitive Studies, Harvard University, and was supported (in part) by a grant No. G-16486 from the National Science Foundation. Preparation of this report was supported by a research grant No. GS-114 from the National Science Foundation. The author is indebted to Dr. George A. Miller and Dr. Paul A. Kolers for their valuable suggestions.

Of course, the perception of fungus-eaters could not be the perception of human beings, but if we know how fungus-eater's perception is functionally related to the fungus-eater's decision making, there is a good hope that we use this as a heuristic model. The type of question I will raise today is: "What is decision making, after all?" As far as I know, this question has never been asked very seriously.

Let me begin with a simple set of assumptions about a single fungus-eater and its environment. The assumptions may be modified when necessity arises. The basic set of assumptions about the environment of the fungus-eater is as follows:

(1) Uranium ores distribute at random on the surface of the planet Taros.

(2) Fungi of identical size grow at random on the surface of the planet Taros.

(3) The land of Taros is entirely flat. There is nothing that disturbs fungus-eater's travel or object-detection.

The basic set of assumptions about the fungus-eater itself is as follows:

(4) The fungus-eater is so programmed as to maximize expected utility.

(5) The fungus-eater has a computer of a finite capacity, called the *brain-computer*, by which expected utilities are computed or estimated.

(6) Each uranium ore has a utility proportional to its weight. Utility is additive.

(7) Fungus has no utility.

What the assumptions (4) through (7) mean is that the fungus-eater is a uranium miner and nothing else. Its efficiency as a uranium miner solely depends upon the efficiency of the computer installed in it.

(8) Fungus is the only energy source for the fungus-eater. When the fungus-eater reaches a fungus, it picks up the fungus and stores it in the stomach.

(9) The fungus-eater can travel in any direction at a constant speed. It can also stand still. When it travels, the fungus storage is consumed at a constant rate. The distance traveled with the consumption of a single fungus is taken as the unit of distance, and is called one *fut* distance.

(10) The brain-computer has a number of routines, one of which can be run at a time. When any of the routines is run, a certain amount of fungus is consumed for the run. The amount of fungus consumption generally varies from routine to routine. When the brain-computer is run with any of the routines, we will say that the fungus-eater is thinking.

(11) The fungus-eater has a sensing device by which it can detect any uranium ore or fungus within a semi-circular area of radius a futs with the current position of the fungus-eater at the center. The area will be called the fungus-eater's *visual field*.

Before entering into the theory, I ought to give you some explanation as to why I do not give any utility to fungus even though the fungus is the only energy source for every activity of the fungus-eater. Also there is another question which is relevant to the same problem. Somebody suggested that I drop uranium from my fungus-eater scheme, and make the fungus-eater an entirely survival-oriented mechanism. And he added "It's more realistic." Is it? Is human being a completely survival-oriented organism? If so, nobody commits suicide, and further, there would be no hero in the battle, no political fanatic who does not spare his life for the glory of the ideology he puts his faith in, and there would be no tragic lovers like Romeo and Juliet. And if these were all true, we would have more peace in the world.

After all, I think that by having uranium in this fungus-eater scheme we obtain more descriptive power of the scheme when it is applied to any closed real human life situation. In a closed real life situation, there are means and ends, and also some objects which are means and ends simultaneously. Any objects which are ends are represented by uranium in the fungus-eater scheme, and objects which are means are represented by fungi. Objects which are both are represented by a certain combination of uranium and fungus. If one takes a combat situation as a closed situation, one may consider that the victory of the combat is the only uranium in the situation, and the occupation of a bridge, a house, a hill, etc. may be considered as fungi of different sizes at different locations.

But if one still insists on having a survival-oriented mechanism, although I think this is a foolish idea, one can get it without any revision of the above assumptions. One need further to assume only that uranium distributes everywhere. Then, maximizing expected utility (or, maximizing expected amount of uranium exploitation) is equivalent to maximizing the length of its life. So, this assumption of ubiquitous uranium makes a survival-oriented fungus-eater as a special case of generally production-oriented fungus-eater.

Now the question why fungus has no utility. The answer is simple: Because we don't need it. Even if fungus has no utility, the fungus-eater will prefer fungus to uranium when it is hungry, since saving life is ordinarily a better way of maximizing expected utility than collecting a small amount of uranium and killing itself by starvation. Now, imagine that a decision making theorist observed this kind of preference of the fungus-eater. Then he would conclude that the fungus has a greater utility than the uranium according to the traditional interpretation of the situation. And further, he would get the law of diminishing marginal utility of fungus, since it is obvious that the fungus-eater tends to become a uranium lover as it gets satiated.

The problem I want to discuss first is: "Decision making on what?" There seems to be no general agreement in decision theories on what the subject is deciding. Sometimes the alternatives are just a set of objects for which the subject has an option of taking one. Sometimes the alternatives are elaborately specified courses of actions. These views are both unsatisfactory to apply to the fungus-eater. The fungus-eater will never go for fungus if the fungus-eater just compares utilities of the objects being present and disregards the relations among the objects, since fungus has no utility. On the other hand, to take all the possible alternative courses of actions precisely specified in every minute detail is unrealistic for the following reasons: First, there is an infinite number of alternative courses for the fungus-eater's travel even if we confine the specification of the courses within the present visual field. And, furthermore, there seems to be no easy rule to reduce the whole set of alternative courses into a pleasantly small finite class.

So, this numerosity of alternative courses is one of the difficulties of having the set of spelled-out courses as the alternatives. But there is another reason for my objection, and this might be more important. Consider the following case illustrated in figure 1: There is a fungus at B. Beyond B, there are a couple of objects in the area C. For the sake of simplifying the discussion, we will only consider those courses connecting objects with straight lines, even though straight courses are not optimal under certain circumstances. By virtue of this illegitimate assumption, however, the infinite set of spelled-out courses is reduced to a finite set, since the number of objects in the visual field should be finite, and it is meaningless to spell out a course beyond the visual field. Now, in the situation like this the reduced set of spelled-out courses can further be reduced since any course that goes to the area C first and then comes back to B is obviously suboptimal. Then all the spelled-out courses that remain in this further reduced set share the characteristic that they all go to B first and then branch there. Now, the point is: Should the fungus-eater at A further continue the reduction until only one spelled-out course is left which maximizes expected utility at A? This is obviously absurd, because, first, the fungus-eater can start executing the optimal decision by traveling toward B even though the optimal course beyond B has not yet been specified; second, the optimal course beyond B should be decided at B where the fungus-eater has the better view beyond the area C than when it was at A.

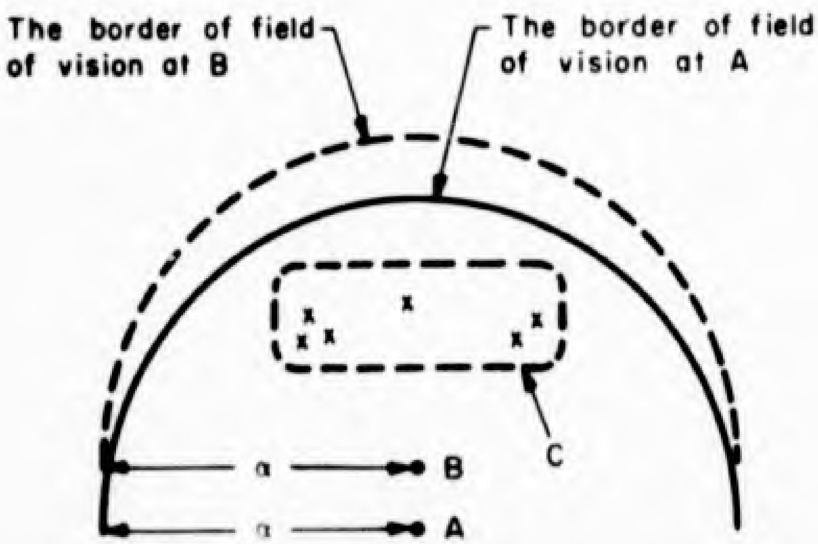

Figure 1. A decision situation for the Fungus-Eater at A. A fungus is at B, and other objects (fungi and/or uranium ores) in the area C.

Thus, probably redundant decision making activities can be avoided by delaying the part of decision which can be delayed, and the redundant activities should be avoided since any routine of thinking requires a finite amount of fungus consumption.

Now we can conclude: The alternative upon which the fungus-eater must decide at the choice point is not necessarily a fully spelled-out course, but usually is a class of courses which are temporarily equivalent, i.e., which overlap at the choice point. Let me call this the ***principle of the minimum course specification.***

However, here is one unpleasant aspect about this business of minimum course specification. Consider the following case illustrated in figure 2: The fungus-eater at the position **A** is fairly hungry, and should eat one fungus to reach the area **C**. At **B** is one fungus, and at **B'** is another fungus. The distance between **B** and **B'** is greater than one fut. So, here are two classes of courses between which the fungus-eater must decide: "Go to **B**" or "Go to **B'**." To make this decision, the fungus-eater must consider the objects in the area **C** and their positional relations to the positions **B** and **B'**, even though what the fungus-eater must precisely decide is whether to go to **B** or to **B'**. But, if the fungus-eater must take into account the objects beyond **B** and **B'**, what benefit can it get by making the course specification minimum? Since this question is touching upon the most essential part of the problem, we will consider it in a somewhat different context where the point will stand out more clearly.

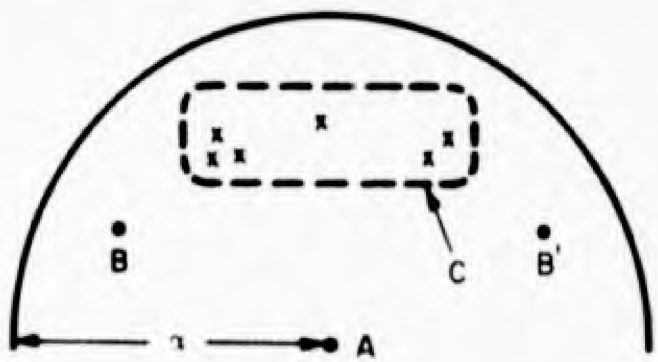

Figure 2. A decision situation for the Fungus-Eater at A. A fungus is at B, another at B', and other objects in the area C.

Consider a game of chess. What a player should precisely decide at each of his turns is only the immediate move. The decision on the next move can wait till the next turn comes up when he knows the opponent's response to his move. And at least theoretically, the decision on the next move made after he knows the opponent's response should not be worse than the decision made before he knows the opponent's response. However, to make a good first move, the player should somehow anticipate his next move and his third move and so on together with the opponent's possible responses. My interpretation of this situation is as follows. The first move, moving a certain knight to a certain position, say, forms a class of plays where a play means a completely spelled-out course of moves, those plays which share that movement of the knight to that position as their first moves. For this move to be optimal, the class should involve the optimal play, the play that maximizes expected utility under the light of the currently available set of information. (The expected utility in a chess game is simply equivalent to the subjective probability of winning.) Now, if the discovery of the optimal play and the discovery of the optimal class which involves the optimal play mean the same thing, this principle of minimum course specification is meaningless. But in general they mean different things--different in particular in the cost of thinking required to obtain the corresponding solutions. For example, the player may be able to use the method of elimination to find out the optimal class, e.g., when all the moves except one are obviously absurd. In this case, the optimal class is easy to find, but the discovery of the optimal play itself may not be as easy as that even at the end of the game. Therefore, the very thing that makes sense of this principle of minimum course specification is the economy of thinking and nothing else.

One might feel uneasy about the exclusion of the element of time in the context of decision making so far discussed. However, the element of time can be incorporated, at least as a first approximation, within the concept of the cost of thinking. We might plausibly assume that, the more elaborate the routine of thinking, the greater the cost, and the longer the time for the completion of the routine. Therefore, the fungus-eater has a propensity to avoid too long a pre-decisional thinking even without explicitly introducing the element of time into the context. Of course, you can make explicit the contribution of time to the cost of thinking if you like, just by assuming the existence of more than one competitive fungus-eater in the world. Whenever a fungus-eater detects another in its visual field, it should act quickly, since otherwise the other will devastate the land and the first one might be starved. This problem of competitive fungus-eaters is interesting, but I will not go into the problem now.

Now let us return to the principle of minimum course specification. Once a local course $A \rightarrow B$ is fixed, we may say that the decision is made. Once a decision is made, what should follow is the postdecisional process. The essential feature of the postdecisional process is the execution of the locally fixed course, which we will call the *local program.* The execution of a local program will be completed, of course, if nothing new happens during the period of the execution. But as the fungus-eater proceeds from A to B, its visual field will also move toward the same direction, and new objects may be found in the newly accrued part of the visual field. Such a new piece of information may or may not keep the local program optimal. But to find out which is the case, the fungus-eater must think. If it thinks and finds out that the local program under execution is still optimal it loses some amount of expected utility corresponding to the cost of thinking. But if it thinks and finds out that the information makes the local program suboptimal, the thinking might have been worthwhile. So, here arises a decision problem whether to think or not to think, and the trouble is that if the fungus-eater thinks to solve this problem of "whether to think," it might be a double waste of fungus. Take a military situation for example. Suppose that a surprise attack on an enemy fortress by a battalion is decided as a local program. As the battalion secretly approaches the fortress, new information will be acquired from time to time. Each bit of information may or may not make the surprise attack suboptimal. The set of reports from reconnoitering parties would reveal, if carefully analysed, that the enemy is or is not aware of the attack. But if the commander stops the battalion for fear of being trapped each time he gets a report, he will miss the optimal time for the attack and lose the battle.

But if he always disregards the reports, he might be trapped by the enemy. Probably, the only way out from this dilemma will be to give the commander, either through instruction or through training, a set of rules to perceive each report in terms of the amount of "problem" a report may imply.

In general, the fungus-eater should be provided with this kind of rules, the rules for perceiving a problem. Whenever a new object is detected by the fungus-eater, the rules are automatically applied to the percept, and the percept is classified either as problematic or not. This of course presupposses that the act of fungus-eater's perception includes the act of classification of the percept, and this seems to be a plausible assumption. Although the cost of perception may not be negligible, the fungus-eater must perceive anyway, and the assumption that the cost of perception is by and large cheaper than the cost of thinking will be realistic.

Now, if the perceptual classification process concludes with "no problem," the previous decision remains valid, and the execution of the local program is continued. That is, the consideration of the new information is postponed until the execution of the local program is completed. But, if the perceptual classification ends up with a "problem," the execution of the local program is suspended, and a thinking routine is started to take care of the new information. This is then to be regarded as the end of the postdecisional process and the beginning of a predecisional process.

To simplify the discussion, assume that perception consists of classifying the information into one of n categories, where n is a constant. Then the rules of problem-perception may be represented by a binary function which assigns a value, either "problem" or "no-problem" to each of the n categories. Now, this problem perceiving function can hardly be fixed once and for all for every possible situation, since as I have mentioned previously, the meaning of a fungus at a certain location varies depending upon how many other fungi are in the field, how hungry the fungus-eater is, and so forth. Therefore this function must be determined afresh at each decision making and must constitute a part of the local program.

It would be silly, however, if the fungus-eater uses a thinking routine precisely to specify this problem-perception function at the time of decision making, since most of the n categories would never be used during the execution of the local program. Usually, but not necessarily always, the specification of this function may automatically be taken care of through the process of predecisional thinking. First, there would exist a natural relevance-ordering among the percept-categories with respect to the local program to be executed. Any object which is newly detected in the general direction of the present travel gives rise to little problem, since it usually just reinforces the optimality of the present local program. On the other hand, an object found in a direction fairly far from the current direction usually reduces the optimality of the present program, and therefore, there is a high chance that this means a real problem. Therefore, the local program itself induces a sort of importance gradient upon the percept-categories. What is further to be done for the specification of the function is to determine the cut-off points to this natural order of percept importance separately for fungus and uranium.

Instead of discussing the cut-off points, however, we will consider the notion of thresholds to problem perception. The absolute threshold to problem perception concerning fungus is defined as the ratio of the number of fungus-percepts having the value "no-problem" to the total number of fungus-percepts. Of course, as this ratio increases, it is less probable for a fungus-percept to be perceived as a problem. The absolute threshold to problem perception for uranium is similarly defined. Also, the relative threshold for fungus is defined as the ratio of the absolute threshold for fungus to the absolute threshold for uranium. This relative threshold for fungus will automatically be determined by how hungry the fungus-eater is, since, obviously, fungus provides a more serious problem than uranium does when the fungus-eater is starving. On the other hand, the mean threshold, the mean of the two threshold values for fungus and uranium, is partly determined by the degree of optimality of the local program employed. The degree of optimality of a local program will be defined as the difference in expected utility between the optimal local program and the next to optimal local program. If this difference is large, there is relatively little chance that the optimal local program turns into suboptimal by new information. We may say that the fungus-eater is confident in his local program if its degree of optimality is great. Then, the more confident the fungus-eater is in his local program, the higher the mean threshold to problem perception during the execution of the program. Another determiner of the mean threshold is the cost of thinking the fungus-eater has paid to obtain the local program under execution. If the predecisional process required to obtain the present program was long and elaborate, a reconsideration of the present program might also require similar elaboration, and therefore, it should be avoided as much as possible.

After all, the problem-perception function set to work during the period of postdecisional process is usually determined as a by-product of the predecisional thinking process, although it is possible that some part of the function is precisely worked out by thinking.

To recapitulate: One of the major characteristics of the postdecisional process is the neglect of information. How much and in what way the fungus-eater neglects information while it is executing a local program depends upon the given situation. When the expected utility is more sensitive to fungus than to uranium, stronger attention is focussed upon fungus than upon uranium. When the fungus-eater has a strong confidence in the optimality of the local program under execution, new information is apt to be neglected. When the thinking routine which the fungus-eater must go into by perceiving a problem is expectedly expensive, the fungus-eater will generally avoid perceiving a problem.

Now let me consider the postdecisional processes of important decisions in real life situations. To be more specific, let us consider a marriage decision. Suppose that Mr. A has decided to marry Miss B. Of course, what he has precisely decided upon should be a local program which leads him to the matrimonial ceremony. Before making this decision, however, there should have been a long and elaborate predecisional process, and he would not want to go into the same process again and again. One may point out that Mr. A might have enjoyed his predecisional process. He might have. I concede up to that point. But to repeat the same thinking process again after he made the decision would not be so happy. Just take this for granted. Yet, Mr. A should go into this repeated vacillation and indecision whenever he perceives a problem after he gave her the engagement ring, since he will willy-nilly receive new information about Miss B while he is proceeding toward the ceremony. So this factor of expensive thinking, if taken alone, will increase the mean threshold to problem perception. (Of course there is no ambiguity about relative threshold. His attention must be focussed upon the things concerning Miss B; otherwise he is not taking the marriage seriously, and therefore this is not the case of an important decision problem.) Now, suppose that there is Miss C who has been another candidate for his spouse, and Mr. A is not very sure about the optimality of marrying Miss B. This will then be a very strong factor that decreases the threshold to problem perception.

There is one more thing to be taken into account here. So far in my discussion of the fungus-eater I have extremely simplified the issue concerning subjective probability. I assume that the fungus-eater knows every object in the visual field with complete certainty, and there is no way of knowing about objects outside of the visual field. But in real life problems, subjective probabilities play a more essential role, and they are particularly susceptible to newly acquired information. The important decision problems one may be concerned about in real life are usually characterized by a large variance of utilities of things concerned, and therefore, even a slight change in subjective probability will greatly affect the expected utility. Therefore, even how much Mr. A may hate vacillation and indecision and the accompanying expensive thinking, he could hardly get away from them until he finally is dragged to the church for the ceremony.

Now let me proceed to the following problem: What is the difference between the predecisional process and the postdecisional process? Postdecisional process begins at the moment of decision making and ends with the completion of the execution of a local program or with a problem perception. (It is also possible to include the completion of the execution of a local program as a special case of the problem perception.) Then predecisional process begins with problem perception and ends with decision making. In the context of fungus-eater's world, the alternation of these two processes is usually enough to describe the behavior of the fungus-eater. But if we want to describe the real life behavior process, this simple-minded sequential description is hardly satisfactory, but we should consider a hierarchical structure of these two processes. That is, as a rule, a predecisional process in a higher context will consist of alternating sequence of pre- and postdecisional processes in a lower context, and each of these lower processes will again consist of the two processes in further lower context. For example, consider Mr. A's predecisional process for his marriage decision. This predecisional process may start with a lower level predecisional process for deciding how to collect necessary information about Miss B. When this process is completed, i.e., the local program for the informational inquiry is worked out, then the local program will be executed as the postdecisional process in the lower context. Here, the point is that this postdecisional process is still a part of the predecisional process to the marriage decision. Suppose that this decision about information inquiry is to ask Mrs. D who is a friend of Miss B for information. But this postdecisional process may again be initiated with another predecisional process concerning how to get in touch with Mrs. D, and so on.

Here, then, we run into a difficulty: If both the predecisional process and the postdecisional process consist of these two kinds of processes in lower levels, then what is the characteristic difference between these two processes? The difference is usually obvious in the lowest level, or in an ordinary fungus-eater context. The fungus-eater does not usually move during the period of predecisional process, but it is moving during the period of postdecisional process. The same is true for a child who is permitted to take one candy among a couple of them. While he is thinking, comparing, and asking questions about the taste of each candy, it is predecisional process that is taking place. When he stretches his hand and picks up one candy, that is postdecisional process. But once one gets into higher level processes, there seems to be no easy and clear-cut way of formally defining pre- and postdecisional processes, although the distinction may intuitively be clear. It would be a useless attempt to define them as the processes preceding and succeeding a certain decision making moment, since there seems to be no other way of defining the decision making than as the moment when a predecisional process turns into a postdecisional process. So, it is hard to avoid the whole set of definitions becoming circular.

Consider an example. Imagine a big boss of a company. He has made up his mind to take a certain policy for the company, and he has enough power in hand to make such a decision completely by himself. But the boss pretends himself to be democratic, so that he convenes an executive meeting formally to discuss the policy. The executives may then argue pro and con to the policy, until they all know that the boss's decision has already been made. After that moment, the opposers to the policy will quickly change their minds, and the meeting will be closed with a unanimous agreement upon the policy and with a clean conscience of the boss about being a democratic president.

The problem is: How could the executives know that the boss's decision was already made? In other words, how could they know that what the boss was engaging in at the meeting was already a postdecisional process, even though he was making a pretense of a predecisional process? I think the only clue is the raised threshold to problem perception set up in the boss's mind, which may betray itself in the way the boss attends to the discussion.

Now, further suppose that there is a young executive who is green enough not to notice the boss's mind. He is bright, and by seeing clearly a grave drawback in the policy, he makes a powerful argument. Then this might make the boss suspend the decision, and let the boss think over again the adequacy of the policy. Now, what does this suspension of the decision mean? Ostensibly, there is no apparent change in what the boss is doing, since he is doing nothing yet. But what happened in his mind will be described this way. Some new information penetrates into his mind breaking the high threshold to problem perception characteristic to a postdecisional process. So the boss now sees a problem, and at that moment, the total set of his mind changes in such a way that, first, the threshold to problem perception is lowered and second, the thinking routine for reconsideration of the policy is started. In short, the postdecisional process in the boss's mind is terminated and a predecisional process is restarted.

I hesitate to use only this informational aspect, that is, the raised and the lowered threshold to problem-perception, for the definition of decision making and decision suspension respectively. Still, I believe that thev are the best parameters for the identification of pre- and postdecisional processes.

In general, in case of an important decision, the whole decision will be made step by step. First, a fairly large class of optimally-looking local programs will be selected, and a fairly high threshold to problem perception is set up which screens out information relevant only to the discarded local programs. Then the class of local programs will be further narrowed down, and the corresponding thresholds to problem perception are established. The process will be repeated until only one local program is left out. This moment would certainly deserve the name of the moment of decision making if the execution of the local program immediately follows. But sometimes, the local program specifies a certain future time for the starting date of execution. In this case, there remains some doubt if we may include in the postdecisional process the period running between the moment of the final selection of the local program and the moment when the execution is triggered, since the threshold to problem perception is usually further raised when the execution is actually started.

In case of important decisions, and this is one thing that makes the decision important, the moment of the triggering execution is often the moment of almost no return. That is, it is often very hard to stop the execution once it is started. The major reason for the difficulty of stopping the execution is, I think, the inertia involved in the social processes. Once the president of a country declares war, a war is started. And once a war is started, even the president who started it cannot stop it easily. Sometimes, people who are bad at establishing high threshold to problem perception intentionally use this inertia to raise their low thresholds. One may tell his family, his friends, anybody else he meets, that he will stop smoking. By this he is trying to increase the inertia, since he knows that he will soon come to doubt the optimality of his decision.

Anyway, it is easy to simulate this business of inertia in the context of the fungus-eater. Just give it a large physical inertia. Then once it starts moving, it is hard to stop itself without wasting a large amount of fungus.

REFERENCES

1. Toda, M., "The Design of a Fungus-Eater: A Model of Human Behavior in an Unsophisticated Environment," Behavioral Science, Vol 7, pp 164-183, 1962.

2. Toda, M., "The Optimal Strategies in Some Simple Fungus-Eater Games," Report #1, Division of Mathematical Psychology, Institute for Research, Pennsylvania State University, State College, Pennsylvania, March, 1963.

DISCUSSION OF TODA'S PAPER

EDWARDS (Discussant). Let me get the inevitable and trivial comment out of the way first: the fungus-eater doesn't really look very much like a human being. Everyone will agree with me there; I can't think of anything less relevant to consider, so we can forget that and go on from there.

What is a fungus-eater for? The fungus-eater serves the purpose of systematizing and integrating intuitions. If we want to think about behavior or, the consequences of an intellectual trait or property, or to examine the joint effects of several interacting traits or properties, we can put them together into the fungus-eater to see how they fit together. In other words, you might think of a fungus-eater as a sort of a projective test for scientists, into which they project their own prejudices and preconceptions without perhaps as much as immediate threat as if they attempted to project these ideas into the behavior of a subject in the laboratory. So I look at the fungus-eater as an inkblot and attempt to report to you what I see there. I hope you will recognize these are my perceptions, and it is both obvious and appropriate that your perceptions would be very different.

To my mind the most important single thing that I see in thinking about the fungus-eater is that at least as much care, attention, and importance must be associated with the model of the fungus-eater's environment as with the fungus-eater himself. The fungus-eater and his environment interact in a way such that the behavior of the fungus eater is strongly influenced by the nature of the model for the environment, and it would be almost impossible to think of a fungus-eater with exactly the sort of characteristics which Professor Toda has just discussed in a radically different environment. Now I think this is just about as true of any model I have heard of. I think that all of the models for behavior that we live with are also models of the environment in which the behavior takes place, fully as much as they are models of the behavior itself. I think we tend to ignore this fact sometimes and it is very useful to be reminded of it in some dramatic way as the fungus-eater does.

One reason why we tend to ignore this fact is that we very often tend to think of models for behavior in the context of the experiment. We think of the environment represented by the experiment as being the environment in which the model for behavior makes sense. Well, that is entirely true, and of course, what that means is that it is also entirely true that the experiment itself is a model of the environment, both in that it creates an environment and in picking out what is important and what isn't; by manipulating what is conceivably important, you specify what you think is important about the environment. You are creating a model and you can expect that your opinion about how the theory ought to go is just as strongly a function of that model of the environment as they are of the processes in the organism.

The second thing that I see in looking at the fungus-eater is that the informational requirements for survival, or rationality of whatever you conceive of the goal of the fungus-eater, are jointly determined by the need for information, that is, by the complexity of the environment, and by the computational or information exploiting capabilities of the organism. I find myself often forgetting the latter point.

It seems to me that one of the things one can state about statistical learning models is that they seem to work a lot better for the simple sorts of tasks than for the more complex ones, and perhaps that's because they have very small informational requirements and so they meet well with informationally very simple environments.

If I were going to design a competitive brand of fungus-eater to send to Taros to compete with Dr. Toda's design, I would want to have him a little bit different: as I said it is a projective test. The crucial difference that I would have in my fungus-eater is that I would like him to have a great deal more computational capability. I would like him, in fact, to have enough computational ability so that he wouldn't have to worry whether he had enough. I would also like him to have a capability which the fungus-eater doesn't have, namely, the capability of being bored. It seems to me that many of the properties which Professor Toda was attempting to explain in terms of the cost of thinking are equally explainable, perhaps even better explainable, in terms of the undesirability of boredom. Once you have made a decision, you don't want to think it over again. Not because it is expensive or unpleasant to think, but rather because it is boring to think the same thing twice. Therefore you don't want to go into it again unless there is some good reason to suppose you won't be thinking the same thing twice. So that the point of this, is that you get much the same sort of behavior in the very simple environment that has been proposed for the fungus-eater with either of these two assumptions about what is costly about the thinking process. However, a fungus-eater with more intellectual capabilities would be capable of surviving, doing better I would think, in competition with a fungus-eater having less intellectual capabilities. He would, of course, still make mistakes because he wouldn't be unlimitedly bright, but more importantly because in his reluctance to rethink problems and therefore become bored, he would fail to rethink something he should have "rethunk" on the basis of new information that came in; but he wouldn't make mistakes so easily the first time.

I'd like to discuss further the point about the model for an environment. It seems to me that a natural strategy for complicating the fungus-eater's environment would be to put statistical inhomogeniety in the distributions of fungi and uranium. This, of course, would raise serious problems of memory capacity and computing capacity. Memory capacity is particularly important here because if you have this inhomogeniety, you have to have careful records in order to discover what they are so that you can exploit them, and if they are complicated then this record keeping may have to be very extensive. Also you have to be able to extrapolate; this complicates the computing capability, and it seems to me that maybe these complications are exactly the sorts of complications that would make the fungus-eater's behavior more like the guide through the mountain pass.

I have two very minor technical objections about the fungus-eater. First of all, in connection with figure 1, it seems to me quite clear that it is always true that if you are going to go to B first, then the optimal strategy is not to decide what you are going to do after you get to B until after you get there. Furthermore this is rigorously true--this would be true even in the case where there is a greater uranium deposit in C than B, because for all you know there might be an even greater deposit just outside of the field of vision at A. So, in view of the fact that that has some probability greater than zero, the expected value of suspending judgment until you get to B has to be greater than the expected value of any program which says for sure that you are going to go somewhere else after leaving B. I really think, in other words, that maybe traditional decision theory does pretty well incompassing this fact that you don't want to think about an infinite set of courses of action.

My other technical objection is that I am very puzzled about how the fungus-eater is going to go about deciding whether new information does or does not constitute a problem. It seems to me that in order to decide that, he has to at some time or another do all the computing that that new information makes relevant. He has to figure out what the new information implies in terms of his strategy. If he doesn't do it after the new information comes, he has got to do it before the new information comes. If he does it before, he must do it not only for the new information he gets, but for all possible sets of new information that he could have gotten, so that the goal of attempting to incorporate thresholds of relevance, so to speak, into the local program is going to take a great deal of computing capability, probably a great deal more computing capability than if he just proceeded fat, dumb and happy until the new information actually came along and then computed it on the basis of it as it came.

You said quite a lot about the fungus-eater but very little about the fungus. So I started to look up some information about fungi and my wife called my attention to the existence of a class of fungi called "slime molds". These have the interesting characteristic in that they just sit there and vegitate as long as conditions are good, but if conditions get bad enough, they draw together into an amoeba-like ball and proceed to move away into some other location where times are less hard. Well now, if you have fungi like that on Taros, this would place some sort of premium on deceptive tactics on the part of the fungus-eater. It would be up to him to see that he always appeared well fed and that as he sneaked up on that fungus, it looked as though really what he was after was uranium.

TODA. I think it is true that every theory, as well as the fungus-eater, is a kind of projective test; and we will certainly agree that the biggest advantage of this kind of approach is that it makes it easier to integrate intuitions.

Ward suggested that the fungus-eater be given the capability of being bored; I prefer to think of him as having an incapability of being bored. The problem is how to specify the cost of thinking. If we carefully analyze what the cost of thinking really means, this might be many factors. Time is one element of the cost of thinking, boredom may be another, and there are probably many other things. As a first approximation, I think it is best to deal with a single parameter at first; and once things are clear, then we can divide the cost of thinking into various different variables.

Let me comment on your more technical points. I certainly agree, that in the case when uranium is by C that the decision must be made at B rather then at A about the courses extending from beyond point B. Sometimes human beings use a kind of negative control technique. In this case I am not assuming any error in traveling; sometimes, however, a fungus-eater may want to travel to B but will find out that he is off course a certain amount. If the fungus-eater uses a sort of negative technique to reduce the distance toward the goal, it is a useful technique to reduce the cost of thinking, and in such a case the fungus-eater may find another uranium deposit, and skip the last one, but because the probability of this is small I think the optimal course will be to go to B, and in that case, having this point as a goal and using a negative technique to control the direction of movement is, I think, useful.

About the problem perception function. If the set of courses that the fungus-eater has obtained as a solution is not far from another object that he detects, this will not be much of a problem; but if the fungus-eater detected some object which is very far from the set of courses, this certainly will mean a very big problem, so there is a very easy approximation technique to evaluate the amount of problems which new information may contain. This approximation technique, that the fungus-eater might use, might be cheaper than using rigorous calculations to find out each time what the optimal course is each time he finds a new object.

ROBY. It seems to me that there is some danger in this principle; the continuity sort of consideration leads you to freeze him right on point A. That is, if he is at A, then how does he know that the family of possible paths will take him all the way to B? Doesn't any incremental step from A raise the possibility that there will be a new optimal direction other than to B?

TODA. The problem perception function takes care of that. If the new information is classified as no problem, then he just keeps going in the same direction. So, if the threshold is very high, then the fungus-eater just neglects every object until he reaches B. If the fungus-eater is not hungry, and unless the fungus-eater has found a very huge amount of uranium in another direction, then only in that case does the fungus-eater stop and reconsider. If the uranium that the fungus-eater has found is small, he would postpone precise calculations until he reached B.

ROBY. Does this determine the length of the steps he can go before he starts to reconsider?

TODA. It depends on the distribution of the uranium and fungus.

ROBY. From A to B. Does the problem threshold tell him he can go precisely to point B before it is going to be time to start recalculating? What determines that distance that he goes?

TODA. Well, in this case I am assuming there is no object between A and B. Even if the fungus-eater had found many objects beyond it would not change optimality. This is the advantage of considering only discrete distributions of objects. If things were distributed continuously that could be a problem.

LANZETTA. I am curious about this too; are you trying to say that if B is far out on the periphery, that the problem threshold ought to be higher as opposed to when it is one fungus unit distance away? At A he has just one unit available; he has just consumed one fungus to move one more unit. Shouldn't that threshold vary as a function of what the distance to B would be?

TODA. I'm not sure. Suppose that there is only one fungus in the visual field, and the distance was very great; the fungus-eater would start moving in that direction because there is no other way. But if the fungus-eater found something else more attractive somewhere else, then he should change his course. Why he goes one way is because there is evidently no other way.

LANZETTA. Well, he might want to go that way because that's as far as he can go. Let's take the worst assumption, that the first fungus is at a distance that is the maximum amount of energy he has available. The problem threshold for that path ought to be very much higher than if the fungus is much closer, so that as he approaches B if he sees say five fungi out here which he could reach, he ought to change his path now and take the five rather than the one, if the energy for it is available.

TODA. The threshold for problem perception is a function of fungus storage.

BERLYNE. This model is reminiscent of one that Herbert Simon did some years ago, a hypothetical rat. He didn't have uranium with a kind of insatiable utility to it, however, but rather his prime concern was keeping alive by getting enough food. You said you introduce uranium, and didn't make survival the end-all, simply because human beings don't live for survival only. But there is a big step between that and uranium, because human beings may not spend all their time trying to avoid the calamity of death, but they may, conceivably, spend all their time trying to avoid some other calamity. They may, in practice, find that all their time is spent on avoiding one unpleasant thing if not another. This raises some of the most fundamental problems of all with regard to the relations between the fungus-eater and the human being.

TODA. Yes, there is a huge discrepancy between the fungus-eater and the human being. But there is little discrepancy between my model and Simon's model. The Simon model, purely a survival oriented mechanism, can be attained as a special case of the fungus-eater. If you assume that uranium is distributed everywhere, since at each moment the fungus-eater can collect a fixed amount of uranium, that means the longer he survives the more uranium he can obtain. Also, by introducing uranium into this scheme, I think I can get more descriptive power, although still not enough to describe human beings.

BERLYNE. You have allowed for two possible motivational functions, either or both of which may exist in human beings; the sort that is designed to avoid something unpleasant and therefore is satiable, and the kind that is designed to get something pleasant, which is unsatiable. This may be a merit.

ROBY. Isn't there a point of richness of fungus at which if you assume that he can eat as much as he wants and store it, that is, a point at which he pursues fungus only, that he can get into a situation in which he can assure himself infinite survival? Under these conditions he can ignore uranium completely and just pick it up if it is on the fungus' path.

TODA. The fungus should be rather sparsely distributed, otherwise, as you pointed out, the fungus-eater can survive forever and there is no problem. The expected lifetime must be finite, otherwise it is meaningless.

FLOOR. You noted the difficulty in distinguishing pre- and post-decisional processes, and then gave an answer, which in your model is that after a decision the range of information to which the organism is open is suddenly restricted; between decisions he sees far and the periods between these are such that his perception function has a threshold which prevents him from seeing anything except the overwhelming. This suggests a sort of cross-breed between the wagon master and the fungus-eater. If you make the situation three dimensional, and supply the fungus-eater with a ladder, and it costs so many "fungs per rung" to climb the ladder, you have a more human-like creature in the sense that he can either increase his search for information or not. It is costly to search for information: if he climbs the ladder he sees two alpha futs away, but it costs him beta fungs to get up the ladder; but if he doesn't climb the ladder he doesn't see two alpha futs away. Now you could say the predecisional process is the time when he is climbing the ladder, that is, he is maximally open to information.

TODA. Yes, I agree, and I did essentially the same thing, but I wanted to deal with perception. I assumed that there is a different degree of sensitivity of sensing devices and also introduced a cost of perception. With these notions we can deal with the same problem.

ROBY. Does he respond to just the semicircle at any one time, or does he have a memory?

TODA. This is another simplification: he should have a memory, yes.

ATKINSON. Can the fungus-eater begin the predecisional process because there has been a problem percept and continue execution of a previous decision or must he stop when the memory moves into the predecisional phase? Can he act and think at the same time?

TODA. Well, that's how you assume the fungus-eater to be. As far as my set of assumptions are concerned, there is no element of time. So in that case the fungus-eater should stop and reconsider the problem when he perceives it. I think it is much better if time actually means nothing. If time means something, then thinking while moving should be much better than stopping and thinking. I think this is a very interesting problem, but I haven't paid too much attention to it yet.

EDWARDS. It just occurred to me that if the fungus-eater is going to have a very limited memory, he needs to have a strategy whereby he does a thorough job of eating all the little fungi as well as the big ones as he moves along, because if he can't remember too well where he has been, he can't afford to get lured back by a few small pieces of fungi or uranium into the fungus and uranium poor areas that he has already been through. He will have to do a much more thorough clean up job. So you have a tradeoff here between how serendipitious he can afford to be about going after the big payoff and how much record keeping he does.

FLOOR. Going back to the possible inhomogeniety of the distributions, suggested by Dr. Edwards, it seems to me that the value in information per rung in the ladder depends upon the existence of homogeniety. If he has a homogeneous environment, he is not likely to benefit much by climbing the ladder. On the other hand, if there are big mountains to be seen over behind the little hills, then the value of climbing the ladder and looking for big game is thereby increased.

LANZETTA. I wanted to ask a question about the strategy of this approach. I am a little unclear as to your thinking about this, whether you started with some notion of what would be reasonable assumptions about the nature of the environment and then proceeded to try to postulate the minimum axioms, assumptions about the organisms in order to account for or to be able to predict survival or efficiency in such an environment, or whether the procedure was to postulate an organism with certain properties and then proceed to construct an environment which may confront this organism. I am really trying to get at your own thought processes, as to what is the best strategy in building such models: starting with environment, starting with organism, or trying to do both?

TODA. I prefer to do both. This is a very unrealistic and simple environment; the simple environment tells us something about the organization of fungi. On the other hand, if we stick to this kind of simple environment I'd never get to a model of man. At some moment you have to start from the human being himself and find some place to cross between the two lines of approach.

LANZETTA. I'll admit there is a certain amount of confusion as to exactly where that can ever lead. I could visualize either starting with some very realistic assumptions about minimum mechanisms of the organism in the light of what we know, and then proceed to ask the question of how much can I tax this organism, given these mechanisms; i.e., how complicated can I make this environment, and still assure some efficient adaptation of this environment, and ask what sort of minimum assumptions do I have to make about the organism. But to do both, I'm not quite sure where you could ever end up that way other than to always have an organism that meets the environment that you have postulated, and neither one really furthering the understanding of the other.

TODA. My inclination is to start with an abstract, simple environment. But to put some complications into this environment, I need some intuitions to make it more realistic, more close to the human environment. The other direction just serves as a guide as to how to proceed from this simple environment. I just want to first do some rigorous thinking about the simple environment.

FLOOR. You mentioned earlier that you have given some consideration to including another fungus-eater. It occurred to me that if the fungus-eaters clean up all the fungi at a given location, that they could very easily work themselves into a prisoner's dilemma situation. That is, if two fungus-eaters get together to exchange information about where there is no fungus, there would be a greater net mutual gain if they both told each other the truth; yet the fungus-eater that lied first would get all the fungus and all the uranium that remains. Have you built anything into your fungus-eater to deal with conflicts of this kind?

TODA. Each fungus-eater is programmed so as to maximize its own expected utility, so it should get into the prisoner's dilemma; I'm not sure if the fungus-eaters would eventually overcome this dilemma and begin to cooperate with each other.

BLANK PAGE

ATTENTION, CURIOSITY AND DECISION*

D. E. Berlyne

University of Toronto

My task is to discuss the bearing on predecisional processes of recent work on attention, exploratory behavior and related topics. This work has been carried on mainly within the current of contemporary psychology known as "behavior theory," a current whose strategy differs from those that have predominated in decision theory so far.

Most existing decision theory is the fruit of two approaches. Some of it, e.g., Wald's theory of statistical decision functions (ref. 30), consist of highly abstract models of the decision processes, fitting a wide range of situations but severely limited in the aspects of those situations that they can reflect. They do not pretend to take account of all the factors that can influence decision-making, and they do not aim to answer all the questions that could reasonably be asked about the decision process. Many of them, being normative models, are concerned solely with one question, namely, which decision is optimal.

The other prevalent approach, exemplified by some computer-simulation programs, produces powerful and detailed models for extremely specific and circumscribed sets of situations, frequently chosen for their tractability rather than for their inherent interest. The hope of their originators, which must remain a hope until it has been fulfilled, is that they can be adapted, through progressive addition of parameters, to encompass an ever expanding range of situations and a fuller and fuller collection of input and output variables.

Behavior theory, in contrast to these approaches, aspires after a conceptual scheme that will subsume as many known aspects as possible of as many psychological phenomena as possible and, above all, make manifest their interrelations. Its unmistakable drawback is its painfully gradual advance towards quantitative exactness. It is not inimical to the other approaches but must, on the contrary, aim eventually to merge with them. Meanwhile, it can interact with them and, at times, exert a corrective influence.

Behavior theory adopts the view that complex processes are best approached through concepts and principles derived from the study of simpler processes, including those found in lower animals. This view is often facetiously confused with the view that there are no important differences between the simplest animal behavior and the most complex human activities and that investigation of the former will supply all the principles that are needed to account for the latter. It is, however, quite possible to avoid this untenable position and yet to believe that at least some of the principles that govern simple phenomena must have a wide applicability. This belief is not only plausible in itself but borne out by the history of other sciences. Quite apart from principles, the study of simple phenomena can always yield questions to guide the examination of more complex phenomena.

---

*The preparation of this paper was facilitated by a grant-in-aid for research in the behavioral sciences from the Ford Foundation.

Perhaps the most important questions about decision processes that the study of simple learning can prompt us to consider and that current decision theory tends to overlook are genetic questions. It is commonly recognized that a particular decision process is not an isolated event but occurs in a temporal context. Decision theorists have paid attention to the fact that a decision has repercussions on future events, that it may result in the necessity of further decision, and that it may bias subsequent decisions in particular directions. For one thing, one decision usually supplies information on which subsequent decisions can be based.

Less attention has been paid to the antecedents of a decision process. The decision theorist generally begins at the point where the decider is in possession of some information relevant to an outcome, a set of possible decision, and a determination to select one of them within a finite time. It is worthwhile going back a few stages and asking how he acquired all these. It is especially worthwhile to ask what conditions determine when and how energetically an organism will embark on the preliminaries of decision making, which means inquiring into motivational factors.

The behavior theorist will be inclined to carry the process of tracing back still further. He will be dissatisfied with a theory that describes, however fully, the way in which an organism makes decisions. He will want to know how that organism came to acquire this or that decision making procedure. This means asking questions about natural selection, heredity, maturation and, above all, learning. The behavior theorist does not feel that he has understood decision making until such questions have been answered, but the same questions must also be of concern to anybody who wishes to change decision making procedures so as to bring the actual decision nearer to the rational decision.

All this is by way of introduction to the point of view from which attention and exploratory behavior is to be considered and to some of the directions in which decision theory may benefit from a junction with behavior theory. We shall not be able to review all that behavior theory might contribute to the study of decision processes, and, in concentrating on the notion of attention and exploration, we shall, in fact, be taking up matters on which there is relatively little to be said with confidence so far, since their systematic investigation has scarcely begun.

INFORMATION-COLLECTION AND INFORMATION-REJECTION

In order to view the phenomena we are discussing in biological perspective, we must first relate attention and exploration to the requirements of adaptation.

Biological adaptation certainly fits Wald's description, (ref. 30) of a "statistical decision problem". There is a space of parameters characterizing states of the world that affect the consequences of what the organism does. The organism selects responses from its repertoire of behavior in accordance with stimulus events whose probability distributions vary with the world-state parameters, and any rule by which the organism selects its responses must entail a certain risk depending on the values that these parameters assume.

There may nevertheless be advantages in using a language that is somewhat more general and more closely in contact with present-day psychology, namely the language of information theory.

Most writers on psychological information theory have concentrated their attention on the transmission of information from external environment to behavior, just as accounts of biological adaptation have commonly stressed the necessity for responses to vary concomitantly with external stimuli. Some have even tended to assume tacitly that the degree of adaptation increases with the amount or proportion of sensory input information that is transmitted.

It may, however, be better to express things rather differently. We assume that the organism has an ensemble of response-classes, one of which is being *performed* at any time, so that each has its *probability of occurrence*. We assume, further, that there is an ensemble of response-classes, one of which is *optimal* at any one time, so that each has its *probability of being optimal*. Adaptation is then seen as a matter of maximizing the degree of correspondence or statistical association between *optimal responses* and *actual responses* or, in other words, the "informational permeability" or "information-theoretic effectiveness" i.e., the ratio of rate of information transmission, R, to input entropy, $H(X)$ (see ref. 22, p. 141) of the channel connecting optimal with actual responses.

This way of looking at things draws attention to two aspects of the process that have not always received the emphasis that they deserve.

The first is related to the fact that the amount of information about (i.e., transmitted from) the optimal response in the present stimulus situation may be grossly inadequate, although additional information, making possible a much greater approximation between actual and optimal responses, may be obtainable through procedures that are within the organism's capability. This fact is recognized in Wald's theory of sequential decision functions, since he allows the decider a choice between making a terminal decision and deciding to "continue experimentation" in order to assign values to additional variables of the "stochastic process." The organism is thus not obliged to content itself with the information provided by its present environment but may act to change the stimulus field so that uncertainty about the optimal response is reduced further.

The second point is complementary to this last one. Not only must the organism seek some of the information that it needs for the specification of the optimal response outside its present stimulus field; it must discard most of the information that its present field offers.

It must do this for at least two reasons. The first, which has often been discussed, is the limited channel capacity of animal organisms. Sherrington (ref. 29) compared the mammalian nervous system to a funnel, pointing out that afferent fibers are five times as numerous as efferent fibers, with the result that afferent processes must compete for control over the final common path. The advent since then of information measurement has shown the occlusion to be even more severe, since the influx of information through receptors is calculated to exceed the output of information through effectors by over 100 times (ref. 20).

There is, however, another reason for not transmitting most incoming sensory information, namely that it would serve no function to do so. Most of the information contained in the external stimulus field does not come from the optimal-response ensemble. It does not help to identify the optimal response. As far as the channel linking optimal and actual responses is concerned, it is noise. Selecting information to discard is thus no less vital an aspect of adaptation than taking in information that ought to be transmitted.

MECHANISMS OF INFORMATION-REJECTION AND INFORMATION-COLLECTION

In the higher animals, there are two sets of mechanisms devoted to the task of rejecting superfluous information.

The first comprises processes that come under the heading of "attention". These are processes that block information coming *from particular receptors*, and recent neurophysiological research has caught sight of several of them (see ref. 7, ch. 3). It has revealed the existence of corticifugal fibers that can evidently suppress activity in sensory pathways and prevent it from reaching the cerebral cortex. Such suppression has been demonstrated in all the principal sensory modalities, and it can apparently intervene at various levels in the peripheral nervous system and brain stem.

This filtering process selects among stimuli from different sensory modalities. There must, however, be other processes that exercise a more subtle selection among items coming along parallel paths from the same modality, and these must presumably be brought into play after the sensory areas of the cortex have been reached. A number of writers (e.g., refs. 1, 3, 23) have drawn on available anatomical knowledge to speculate how this might occur.

Attentive processes not only determine which stimuli shall exercise control over overt behavior. They also determine what items of incoming information shall be stored, since "attention in performance" is supplemented by "attention in learning" and "attention in remembering". It was shown in Pavlov's laboratory that, if two conditioned stimuli are presented in combination and then followed by an unconditioned stimulus, only one of them will later be found capable of eliciting a conditioned response, the other having been "overshadowed". And studies of remembering, as well as everyday experiences, show how "incidental remembering", i.e., remembering of events that were perceived but did not occupy a focus of concentration, occurs but is extremely limited in extent.

The second set of information-rejecting devices selects among properties of stimulus patterns and comprises what is usually discussed under the headings of "stimulus generalization and discrimination", "selective perception", and "concept formation". The word "attention" is sometimes used in connection with processes of this type, but it is more conducive to clarity to reserve it for processes that select among stimuli exciting different points on the sensory surfaces and to refer to these other processes as *"abstraction"*.

Abstraction places behavior under the control of stimuli possessing properties that prior learning or natural selection has shown to be correlated with properties of the optimal response. They thus involve responding alike to, and thus treating as equivalent, a large variety of stimulus patterns with some biologically or motivationally important property in common but distinguishable in other respects, which are ignored. The equivalence classes that are formed in this way will vary with the organism's motivational condition, but, since their members can always be differentiated, their formation implies a failure to transmit some of the information received by the sense organs. "Selective remembering" exercises an analogous selection among features of remembering events.

Attention and abstraction can only reject information, since, although they may be said, and often are said, to facilitate the reception of certain stimuli, they are confined to the information that reaches the sensory surfaces and can thus do no more than block some of it and permit the remainder to proceed further.

The processes that subserve information-collection, on the other hand, contribute to information-rejection at the same time. They can be divided into *exploratory* activities and *epistemic* activities.

The function of exploratory responses is either to expose the organism to stimulation from sources that were not previously represented in the stimulus field or to augment stimulation from sources that are already active. This will generally mean affording access to new information, although it is doubtful whether informational properties of stimulus patterns are the only ones that influence exploratory behavior. The introduction or enrichment of information coming from some sources will, however, inevitably mean the withdrawal or impoverishment of information coming from other sources. Exploratory responses can thus assist the filtering processes of attention in their information-rejecting function.

Exploratory behavior may take the form of (a) *receptor-adjusting responses*, consisting of physico-chemical changes in sense-organs which heighten sensitivity or postural changes which focus sense-organs on particular stimulus-objects; (b) *locomotor exploration*, which increases the influx of information from particular objects by moving the whole organism nearer to them; and (c) the residual category of *investigatory responses*, usually manipulatory, which effect changes in environmental objects such as will make new information accessible.

Epistemic responses are responses that promote the acquisition of new knowledge or, in other words, the storage of information in the form of symbolic structures. The principal epistemic activity, which must play a part even when other epistemic activities like *observation* and *consultation* are used, is *directed thinking*. Directed thinking adds to the information originating in the present external stimulus field by supplying information retained from past experiences. This information helps to specify the present optimal response because of the redundancies that link the present with the past.

Nevertheless, directed thinking is also a collection of information-discarding devices and thus forms yet another aid to attention and abstraction. It forms equivalences between distinguishable entities and can, in fact, be regarded as an elaborate and laborious means of arriving at secondary or mediated stimulus generalization, i.e., the association of a common overt response with a set of physically dissimilar stimulus patterns through attachment of a common symbolic label to all of them (ref. 8).

THE SELECTION OF INFORMATION TO REJECT

We must now consider the bearing of these phenomena on predicisional processes.

The first relatively neglected group of problems that arises centers on information-rejection. Decision theory regards the decision as a function of the "stochastic process", the data available to the decider, as a whole. In experiments on decision processes, the subject is generally exposed to a highly restricted set of data whose information content is well within his processing capacity. But it is evident that, in any real life situation, the decider's first task is to identify which items of available information will be useful and which can be safely ignored. There are many experiments, especially in the field of thinking, to show how a superabundance of information may retard the solution of a problem by leading the subject to try out fruitless lines of attack.

One way to discover which dimensions of the sample space can guide him towards the optimal response is to act on various combinations in turn and observe the consequences. This is more or less what animals do when reduced to blind trial-and-error behavior. It is exemplified most clearly by discrimination-learning situations in which an animal adopts one "hypothesis" after another, making his responses depend on a number of stimulus properties in succession.

This procedure is, however, costly in time and effort and often dangerous. A safer and speedier alternative is take selections of information in turn and anticipate their consequences of acting on them by directed thinking. However, as Newell, Shaw and Simon (ref. 25) have shown, this "implicit trial and error" cannot profitably be carried out in a random order but must try out methods in an order that reflects how often each has proved successful in the past. Even the trial and error of the cat or rat is directed first to those actions and environmental features that, in the light of prior experience, have the best chance of being turned to account.

The organism must, therefore, be biased towards stimulus items whose likelihood of conveying information about the optimal response is especially high. And these stimulus items can be identified only by *content* or by *source*.

As far as content is concerned, the principal function of discrimination learning is to subject behavior to the control of stimuli with properties that have been correlated with the presence of rewarding or punishing agents in the past. This means that the occurrence of objects emitting these kinds of stimulation will override other information in swaying the decision process.

There are clearly also innate factors that give stimuli with particular qualities a predominant influence. The nervous system has to be so constructed that it gives special precedence to stimuli typifying injurious events or opportunities for organic gratification. An electric shock and the smell of fish have been shown to occasion inhibition of afferent auditory processes in the cat (ref. 18). Ethologists have described many examples of kinds of stimulation that are correlated, but not directly connected, with biologically important conditions and apparently have the power to oust other stimulation from control over behavior, e.g., the coloring of an opposite-sexed individual of the same species during the mating season, or the shape of a predator.

Both learning and instinctual mechanisms help to make the behavior-selecting process impervious to information from biologically neutral or indifferent stimulus sources unless and until new learning gives them a vital significance. Yet, such mechanisms cannot suffice by themselves. At any one time, there are likely to be several stimuli present that have become discriminative cues for action of one sort or another, and some selection must be made among these. This is especially apt to be the case in human decision-making when the decider is faced with a mass of information that is all highly pertinent but equivocal in its implication.

Further selection must thus be effected according to source. Instrumental or operant conditioning may give attentive processes a lasting bias in favor of information emanating from a particular direction. At other times, a signal, which owes its power over attention to learning, may indicate the sensory channel or environmental location that is at the moment most prolific of clues to the optimal responses.

Broadbent (ref. 12) has shown how human subjects, hearing two voices simultaneously asking different questions, can be instructed to concentrate on questions preceded by a certain call-sign, and how a much higher proportion of these than of other questions will then be answered correctly. Reynolds (ref. 29) has shown how a pigeon may be trained to let either the shape of a figure or the color of the background determine whether it is to peck or not, depending on whether a yellow or a green signal light is on. Martsinovskaia (see ref. 21) has shown how, in children aged five or more, the hand with which a motor response is performed can be made, by verbal instruction, to depend on the coloring of one or another portion of a visual display.

Finally, attention is particularly vulnerable to capture by stimuli, regardless of their quality or origin, that have properties like *novelty* and *surprisingness*. Presumably, when stimuli differ from those that have been experienced before, and especially from those that have been experienced recently, or when they deviate from established sequences and from the expectations founded on them, they are likely to betoken events with exceptionally cogent claims on motor functions. Thus, Poulton (ref. 26) found that subjects who were to write down information received through two loudspeakers missed much less of what came from an intermittently active loudspeaker than of what came from a constantly busy one. And I myself (ref. 4) found subjects less likely to respond to a visual stimulus of a kind that had been present for some time than to one of a kind that had been newly introduced.

These, then, appear to be some of the chief ways in which the nervous system executes a preliminary sorting of the information that is at its disposal and picks out the items that are most likely to repay further processing. But after this has been done, the final selection will generally have to depend on the assaying of successive samples by thinking. If this does not lead to quick success, the organism may well have to revert to some of the information that failed to pass the preliminary sorting. And the resources of exploratory behavior will have to be called on if the information already available proves insufficient for the designation of a terminal decision or if it is recognized from the outset as insufficient.

EXPLORATORY AND EPISTEMIC BEHAVIOR

Both common knowledge and accumulating experimantal evidence attest that, while the strength and direction of exploratory and epistemic behavior are susceptible to the influence of many factors, both external and internal, they are preponderantly governed by one group of variables. These are the variables that I have called "collative" variables, since their values depend on collation of information from several sources, whether simultaneously or successively active (refs. 8, 11). Collative variables include novelty, surprisingness, complexity, ambiguity, puzzlingness, and change. They are commonly spoken of as properties of external stimulus patterns, although they are more accurately described as relations between properties of external stimulus patterns and states of the organism.

A spate of research (see ref. 9, chapters 3 and 4), most but not all of which has gone on in Russia, has shown that at least some kinds of exploratory behavior are accompanied by a wide assortment of physiological processes, involving almost every system of the body. The whole complex is known as the "orientation reaction" or "orientation reflex", terms introduced by Pavlov. It included changes that are now recognized as indices of increases "arousal", and the currently much discussed concept of "level of arousal", which has been identified by a number of writers with the older concept of "drive level". Earlier work had shown that the electroencephalographic, electrocutaneous and muscular signs of heightened arousal also appear when subjects are thinking out solutions to intellectual problems, which means that they are engaged in a form of epistemic behavior (ref. 7, chapter 11).

A further hint toward understanding the workings of exploratory and epistemic behavior comes from the argument (see ref. 6, 7, chapter 2) that collative variables entail conflict. By "conflict" is meant the presence of stimulus conditions associated with incompatible responses. The argument is more direct in the case of some collative variables, like surprisingness and ambiguity, and more speculative in the case of others, like novelty and complexity. It is, however, supported by a number of experiments, in which situations that clearly involved competition between mutually exclusive motor activities or between performance and inhibition of the same responses, give rise to exploratory behavior (ref. 7, chapter 4). There is, moreover, evidence that the intensity of the orientation reaction is influenced by conflict (ref. 7, chapter 7, and ref. 9).

A picture may thus be tentatively sketched out as follows. Some forms of exploratory behavior and most forms of epistemic behavior are prompted by a state of heightened drive or arousal that we call call "perceptual curiosity" or "epistemic curiosity", as the case may be. The use of the terms "arousal" and "drive" implies that: (1) there will be an increase in the indices of arousal, including heightened motor activity and expenditure of energy; (2) there will be an increased probability of exploratory or epistemic behavior aimed in particular directions; and (3) conditions that remove or attenuate the state in question will be able to reinforce instrumental responses.

It seems that these curiosity states are generated by conflict due to collative properties of external stimulus patterns in the case of perceptual curiosity and due to discrepant symbolic response-tendencies (beliefs, lines of thought) in the case of epistemic curiosity. It seems, furthermore, that these states can be alleviated through the reception of information furnished by sensory processes and by symbolic formulae respectively.

The strength of curiosity is assumed to vary directly with the degree of conflict, which is, in its turn, held to be greater (1) the more numerous the competing responses-tendencies; (2) the greater their total absolute strength; (3) the nearer they come to being equally strong; and, in situations where partial incompatibility may occur, (4) the greater the degree of incompatibility between them. There is some evidence for an effect of each of these on reported epistemic curiosity (refs. 5, 7, 10).

A close relation undoubtedly exists between this conception of conflict and the information-theoretic concept of "uncertainty" or "entropy". When a subject has a certain degree of uncertainty regarding some impending event, he must be holding in readiness behavior patterns appropriate to the alternative possibilities that he recognizes, and there must be some incompatibility between behavior patterns corresponding to different contingencies. He will thus have a certain degree of conflict, and an external observer will have a certain degree of uncertainty about which behavior pattern he will ultimately perform.

It will be noted moreover, that uncertainty, as the information theorist measures it, reflects two of our determinants of degree of conflict, namely number of competing response tendencies and their nearness to equal strength or equiprobability, but it does not reflect the contribution of their total absolute strength. We must therefore multiply uncertainty by some factor representing this "importance factor", i.e., how strongly established the habits are that are in competition, how large the gains and losses are that are at stake, how powerful the motivation associated with each of the contesting courses of action is, if we are to measure intensity of conflict.

This conceptualization is applicable to the kind of exploratory or epistemic behavior that results from partial information about a particular object or event. This is the kind of exploration that I have called "*specific*" exploration. It needs to bé, but is not always, distinguished from "*diversive*" exploration, which can be reinforced by stimulation or information from any souce that is "interesting" or "entertaining" enough. Diversive exploration is not preceded by any state that can be termed "curiosity". If there is any drive behind it, it must be of a sort that is best called "boredom", since the strength of diversive exploration has been found in at least some instances to increase with time spent in an impoverished or tedious environment (see ref. 11).

Nevertheless, the factors that make for specific and diversive exploration must interact, since collative properties apparently determine how welcome the sensory consequences of the latter will be and how effective in reinforcing the responses that produce them.

Another distinction that will be relevant to what follows is that between exploration or epistemic activity with predominantly *intrinsic* and *extrinsic* functions. In the intrinsic case, stimuli of concern are indifferent or biologically neutral ones, and the stimulation or information that exploration wrests from them is rewarding in itself and not dependent on any ensuing effect on non-neural tissues. When an extrinsic function is operative, what is sought is information or stimulation that will guide a subsequent response with rewarding or punishing consequences of a different nature. The information-gathering processes that are of interest to decision theory typify the extrinsic case.

This distinction, like the other one, is not absolute, although certain problems are easier to study in the intrinsic case. The possibility of practical application is bound to play some part in many, if not most, attempts to relieve "disinterested" curiosity. And, while some of the reward-value or reinforcing power of information sought for extrinsic reasons may be attributable to secondary reinforcement, i.e., to association with the benefits derived from a judicious terminal decision, the relief of uncertainty and hence of conflict must make a weighty contribution here also.

THE STATE OF THE DECIDER

We must now ask what new aspects of the decision making process are brought to our notice by these considerations. The most important of them is undoubtedly the influence of the decider's state on the collection of information preparatory to decision. It seems that the vigor and duration of information seeking activities, as well as the directions in which information is sought, depend primarily on the subject's degree of conflict and on the nature of his conflicting response-tendencies.

Now, current decision theory depicts a subject who is offered a choice among a set of decisions and gives them all equal and unbiased consideration until he has identified the one that available information favors. But where do the members of this set of decisions come from?

When a subject is contemplating a number of alternative actions, without yet having embarked on any of them, each of them must be represented inside him by a fractional response, i.e., a reduced or curtailed version of the chain of neural events that would otherwise lead to motor performance. And each of these fractional responses must be a product of learning, either directly through having been practised in comparable situations in the past or indirectly through inferential processes. Even when a set of alternative decisions is presented to a subject in verbal form by an external agency, the verbal formulae designating them would have no meaning for him unless they appealed to existing habit structures.

Now, when fractional responses corresponding to a set of mutually exclusive actions are occurring in a subject, he is undergoing a certain degree of conflict.

It was suggested earlier that degree of conflict is a function of uncertainty multiplied by an "importance" factor, i.e., the total strength of the competing responses. If the responses corresponding to alternative decisions are actually equally strong, uncertainty will be a maximum. Incoming information will then increase conflict only if it reveals a new possibility of decisions to be added to those already under consideration or if it increases to importance factor by augmenting the absolute strengths of all the competing responses to a comparable extent.

It must, however, be unlikely that newly acquired information will increase the absolute strengths of the competing responses equally. If this were to happen, we should ascribe a motivational rather than an informative function to the stimuli in question. And, in any case, if some of the competing responses are better adapted to external reality than others, as will generally be the case, newly acquired information can be expected to make some of these responses relatively stronger and the other relatively weaker. This means that, as long as the subject is limited to the original set of alternative decisions, incoming information can only reduce conflict in cases where competing responses are initially of equal strength.

In practice, however, when we have a set of fractional responses corresponding to alternative decisions, we can be reasonably sure that previous learning has given some more reinforcement and more generalized strength than others. Frequently, one response will be markedly stronger than any other, even though its advantage is not yet sufficient to tip the scales. If this is so, incoming information has a good chance of increasing uncertainty and thus exacerbating conflict, by strengthening one or more of the weaker responses while weakening the strong responses, thus bringing response strengths nearer to equality*. Conflict can then be minimized

---

*The measure of information that is relevant to curiosity and curiosity reduction is the amount of information transmitted by a given signal, $I_T(y_j)$. This is the amount by which uncertainty about the input is reduced by receipt of a particular output signal, so that

$$I_T(y_j) = H(X) - H(X|y_j) = -\sum_i p(x_i) Log_2 p(x_i) + \sum_i p(x_i|y_j) Log_2 p(x_i|y_j).$$

$I_T(y_j)$ can take on negative values, i.e., uncertainty can sometimes be increased by the receipt of a signal. Samson (ref. 27) has pointed out the fallaciousness of the belief that information necessarily decreases uncertainty, presenting a counter-example involving a system with three possible states. A negative quantity of transmitted information can, however, occur even where there are only two input possibilities, as witness the following channel:

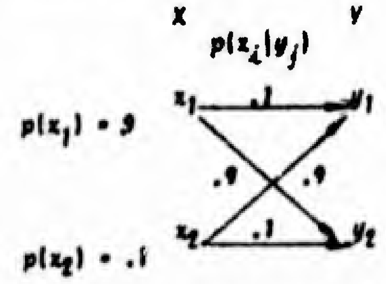

Here, if y_2 is received, uncertainty about the input falls from $H(X) = 0.47$ to $H(X|y_2) = 0.09$, so that $I_T(y_2) = +0.38$. If on the other hand, y_1 is received, uncertainty about the input rises from $H(X) = 0.47$ $H(X|y_1) = 1.00$, so that $I_T(y_1) = -0.53$.

The measure $I_T(y_j)$ is related to but distinct from a more widely used measure, "rate of information transmission (T)". T is the mean or expected amount by which input uncertainty is reduced on receipt of an input signal, i.e., $T = H(X) - H_y(X)$, where $H_y(X)$ (equivocation) = $-\sum_j p(y_j) \sum_i p(x_i|y_j) Log_2 p(x_i|y_j)$. T can take on only non-negative values; for example, in the channel depicted above, $T = 0.21$.

For the treatment of motivational aspects of human information-seeking processes, such as exploration and thinking, it may prove to be a more meaningful measure.

only if any information-gathering that occurs is directed selectively towards information that supports the response that is already uppermost and away from information that might support less favored alternatives.

Festinger (ref. 16) has spoken of "cognitive dissonance" in connection with situations where a subject has already committed himself to a decision and is thereafter inclined to avoid any information that might raise doubts about its correctness. Presumably, the same kind of mechanism can operate when the Rubicon has not yet been crossed but its distant bank is alluring.

To illustrate how a decision theory that neglects these factors may go astray, let us look at one of the classical chapters of bad decision making in military history, namely the Third Battle of Ypres in 1917. In Wolff's account of this episode (ref. 31), we read that "Haig, the British commander, was devoted to broad concepts in the Napoleonic vein that were no longer appropriate. He wrote that '...the role of Cavalry on the battlefield will always go on increasing...' He conceived of cavalry as the basic instrument of war, to which infantry and artillery were secondary. He resisted innovations, and in later years was to deprecate the airplane, the tank, and even the machine gun. He had come to epitomize what a renowned French cavalry officer had once said--that 'The British cavalry officer seems to be impressed by the conviction that he can dash or ride over anything; as if the art of war were precisely the same as that of fox-hunting.'"

In the cumbersome language of learning theory, Haig had evidently had responses leading to the engagement of cavalry strongly reinforced during his early training as a cavalry officer. These responses must have received subsequent reinforcement from the high prestige and aristocratic connections that the cavalry arm of the British Army enjoyed in his day, as well as a great deal of supplementary strength through stimulus generalization from experiences of riding to hounds.

We should not expect a man with that reinforcement history to be intensely curious, and inclined to wonder, about the usefulness of cavalry. He is not the sort of man that would eagerly sponsor commissions of inquiry, compilations of impartial opinion, or research projects, on the role of cavalry in modern warfare. Insofar as he sought information at all, he would presumably seek out arguments and associates that could be trusted to corroborate his predilections and encourage plans based on them and he would presumably avoid any conversation or thought or experience that might foster doubts.

PERSONALITY DIFFERENCES

Besides degree of conflict, we have to take account of several more lasting characteristics of the decider's state that must influence quite profoundly his manner of preparing for decision.

There must, for one thing, be sizeable differences between individuals in the extent to which they are disturbed by a given degree of conflict. Such differences in level of drive or arousal would mean differences in zeal for collecting information. Furthermore, and perhaps even more important, there are evidently differences in the level to which uncertainty and conflict must be reduced before different subjects are ready to make a terminal decision. Cohen and his associates (e.g., refs. 13, 14) have shown how, in various situations from shooting at the goal in association football to driving a bus through a narrow space, some individuals and even groups will act with greater risk of failure than others. This is so even when subject's estimate of their chances of success are reasonably accurate.

Then, there are differences in relative preference for various ways of handling conflict due to collative variables, exploratory and epistemic behavior being by no means the only devices that can reduce such conflict. A number of personality dimensions that have been widely discussed in recent years seem to reflect differences of this kind.

There are, for example, personalities characterized by high "intolerance of ambiguity" (ref. 19) by preference for "simplicity" (refs. 2, 15) or by a proclivity to "repression" (ref. 19). Persons with such traits apparently tend to ward off the effects of ambiguous, complex or otherwise troublesome stimulation by selective attention or by avoidance. Those who are low in "intolerance of ambiguity", who prefer "complexity", or who are classifiable as "intellectualisers" are, on the other hand, apparently willing to face such stimulation and to cope with it by means akin to exploratory and epistemic behavior.

This kind of behavior would seem to be more rational and adaptive. To ignore or escape from disquieting stimulation may allay conflict quickly, but its success is apt to be temporary, both because the eluded pattern of stimulation may recur and because it may continue to disturb through its representations in memory or thought. The other kind of reaction means enduring conflict or even intensifying it for a while until the situation responsible for it can be rendered innocuous once and for all. Nevertheless, a predilection for exploration and epistemic activity can be carried to maladaptive extremes, in the form of obsessive doubt, insistence on ascertaining pettifogging details, and an inability to act at all as long as the smallest residue of uncertainty remains.

OTHER FUNCTIONS OF STIMULUS INFORMATION

One of the most serious obstacles to rational decision making in real life situations is the fact that the information accruing from exploratory and epistemic activities can have functions other than the guidance of the terminal decision and the diminution of risk. In the kind of situation envisaged by Wald, the decision depends solely on data whose probability distributions determine which decision is optimal and how much would be lost by making any other decision. In real life, information-seeking activities may very well turn up data that incline the decision process in a particular direction and yet are independent of the parameters that govern outcomes of terminal decision. According to our information theoretic-formulation, the organism's limited transmitting capacity may be glutted with information that cannot be traced back to the optimal-response ensemble.

This is apt to happen because the consequences of exploratory and epistemic behavior can be rewarding for a variety of reasons and particular forms of exploratory and epistemic behavior can thus be reinforced in a variety of ways, not all of which are conducive to rational decision making.

The first intrusive function of stimulus information is its *autistic* function. Incoming stimulation that depicts or symbolizes a desired state of affairs may provide immediate gratification through secondary reinforcement or secondary-drive reduction. Psychologists of several schools have expatiated on the dangers of symbolic substitute satisfactions and especially on the danger that they may attract a subject away from the effort and uncertainty of rational action. More germane to decision making is the likelihood that autistically rewarding information will deflect a subject from the optimal decision by selectively strengthening a mistaken course of action toward which he is leaning. For an illustration, we may turn once again to Wolff's book on the Third Battle of Ypres.

"It naturally pleased Haig," he tells us, "to have carefully chosen and nicely cooked little tidbits of 'intelligence' about broken German divisions, heavy German casualties, and diminishing German morale served up to him...He beamed satisfaction and confidence. His great plan was prospering. The whole atmosphere of this secluded little community reeked of that sycophantic optimism which is the curse of autocratic power."

Haig's appetite for information of this kind and the rewards that his associates derived from catering to it must have helped materially to commit him to the bad decisions that he was continually on the verge of making.

The second function of stimulus information that can lead the decision process astray is its *intrinsically* satisfying function, in the sense discussed earlier. An exploratory or epistemic response may be reinforced and thus become resorted to more readily if it yields stimulus patterns that have "interesting", "entertaining" or "aesthetically pleasing" properties (thus turning the response producing them into a diversive response) or if it relieves "disinterested" or "idle" curiosity.

This latter case may seem paradoxical, since exposure to facts or to convincing symbolic representations of facts yielded by intrinsic exploratory or epistemic activities should, one might think, also help to steer practical decisions on to the right track. It has already been remarked that intrinsic and extrinsic forms of these activities cannot be sharply separated.

Nevertheless, the danger lies in the fact that the intrinsic forms cast their nets wider than their more extrinsic counterparts. A desire to have a puzzling point cleared up or to see whether one's hunches are correct might well prompt information-gathering operations that steal valuable time, resources and channel capacity from issues more vital to the choice of a decision.

Finally, we may note that this danger applies to the decision itself, since every decision has the accretion of information among its other consequences. The temptation to embark on a certain course of action "because it will be interesting to see what will happen" may incline the scales against rationality. Dr. R. B. Joyner, a colleague of mine at the University of Toronto, has obtained data illustrating this with subjects who had to guess which of two lights would come on in each of a succession of trials. Some of the subjects reported afterwards that they had been well aware that the best strategy was to guess the more probable event every time. But they had occasionally selected the other event for some reason such as "It was more fun that way."

That valuable information for the guidance of future decisions has been obtained is often pleaded in extenuation of wrong decisions once they have been made. Moorehead (ref. 24), writing of another exemplar of disastrous decision making from the First World War, states that "Gallipoli was a mine of information about the complexities of the modern war of manoeuvre, of the combined operation by land and sea and sky; and the correction of the errors made then was the basis of the victory of 1945." The prospect of having arguments of this kind to fall back on must not be overlooked as yet another source of detriment to the groundwork of decision.

REFERENCES

1. Anokhin, P. K., Vnutrennee Tormozhenie kak Problema Fiziologii (Internal Inhibition as a Physiological Problem), Medgiz, Moscow, USSR, 1958.

2. Barron, F., and G. S. Welsh, "Perception as a Possible Factor in Personality Style: Its Measurement by a Figure Preference Test," Journal of Psychology, Vol 33, pp 199-207, 1952.

3. Beritov, I. S., Nervnye Mekhanismy Povedeniia Vysshikh Pozvonochnykh Zhivotnykh (Neural Mechanisms of the Behavior of Higher Vertebrates), Academy of Sciences, Moscow, USSR, 1961.

4. Berlyne, D. E., "Attention to Change," British Journal of Psychology, Vol 42, pp 269-278, 1951.

5. Berlyne, D. E., "An Experimental Study of Human Curiosity," British Journal of Psychology, Vol 45, pp 256-265, 1954.

6. Berlyne, D. E., "Uncertainty and Conflict: A Point of Contact between Information-Theory and Behavior-Theory Concepts," Psychological Review, Vol 64, pp 329-339, 1957.

7. Berlyne, D. E., Conflict Arousal and Curiosity, McGraw-Hill, New York, 1960.

8. Berlyne, D. E., "Les Equivalences Psychologiques et les Notions Quantitatives," D. E. Berlyne and J. Piaget, Théorie du Comportement et Opérations (Etudes d'Epistemologie Génétique, XII), Presses Universitaires de France, Paris, France, 1960.

9. Berlyne, D. E., "Conflict and the Orientation Reflex," Journal of Experimental Psychology, Vol 62, pp 476-483, 1961.

10. Berlyne, D. E., "Uncertainty and Epistemic Curiosity," British Journal of Psychology, Vol 53, pp 27-34, 1962.

11. Berlyne, D. E., "Motivational Problems Raised by Exploratory and Epistemic Behavior," in S. Koch (Editor), Psychology--A Study of a Science, Volume 5, McGraw-Hill, New York, 1963.

12. Broadbent, D. E., "Listening to One of Two Synchronous Messages," Journal of Experimental Psychology, Vol 44, pp 51-55, 1952.

13. Cohen, J., "Subjective Probability," Scientific American, Vol 197, pp 128-138, 1957.

14. Cohen, J., and E. J. Dearnaley, "Skill and Judgment of Footballers in Attempting to Score Goals," British Journal of Psychology, Vol 53, pp 71-88, 1962.

15. Eysenck, H. J., Dimensions of Personality, Routledge and Kegan Paul, London, 1947.

16. Festinger, L., Theory of Cognitive Dissonance, Row Peterson, Evanston, Illinois, 1957.

17. Frenkel-Brunswik, E., "Intolerance of Ambiguity as an Emotional and Perceptual Variable," Journal of Personality, Vol 18, pp 108-143, 1949.

18. Hernández-Peón, R., H. Scherrer, and M. Jouvet, "Modification of Electrical Activity in Cochlear Nucleus during 'Attention' in Unanaesthetized Cats," Science, Vol 123, pp 331-332, 1956.

19. Lazarus, R. S., C. W. Eriksen, and C. P. Fonda, "Personality Dynamics and Auditory Perceptual Recognition," Journal of Personality, Vol 19, pp 471-482, 1951.

20. Luce, R. D. (Editor), Developments in Mathematical Psychology, Free Press, Glencoe, Illinois, 1960.

21. Luria, A. R., The Role of Speech in the Regulation of Normal and Abnormal Behaviour, Liverlight, New York, 1961.

22. Meyer-Eppler, W., Grundlagen und Anwendungen der Informationstheorie, Göttingen, Berlin, 1959.

23. Milner, P. M., "The Cell Assembly: Mark II," Psychological Review, Vol 64, pp 242-252, 1957.

24. Moorehead, A., Gallipoli, Hamish Hamilton, London, 1956.

25. Newell, A. J., C. Shaw, and H. A. Simon, "Elements of a Theory of Human Problem Solving," Psychological Review, Vol 65, pp 151-166, 1958.

26. Poulton, E. C., "Listening to Overlapping Calls," Journal of Experimental Psychology, Vol 52, pp 334-339, 1956.

27. Reynolds, G. S., "Attention in the Pigeon," Journal for the Experimental Analysis of Behavior, Vol 4, pp 203-208, 1961.

28. Samson, E., Information Theory: Questions and Uncertainties, Air Force Cambridge Research Center, Cambridge, Massachusetts, 1954.

29. Sherrington, C. S., Integrative Action of the Nervous System, Yale University Press, New Haven, Connecticut, 1906.

30. Wald, A., Statistical Decision Functions, Wiley, New York, 1950.

31. Wolff, L., Flanders Fields, Viking Press, New York, 1958.

BLANK PAGE

DISCUSSION OF BERLYNE'S PAPER

LANZETTA (Discussant). I think it is fairly clear that the approach Dr. Berlyne has presented here this morning is in sharp contrast to those we heard yesterday, and I am sure a lot of you would like to raise some questions about it, so I'll keep my remarks to a minimum. In a way, he anticipated a lot of things we might say about this approach in his introduction by first pointing out that, in fact, it is not incompatible with former normative models that have been previously suggested, and also by indicating that we have a long way to go, even within behavior theory, before some of the things that he was talking about receive firm empirical support. However, I think there were several things of interest in this approach that ought at least to be touched upon.

First of all, I think an issue that has not yet been touched upon in this symposium is the question "What is a decision?". At what point do we say the organism is making a decision? We all recognize the continuity between the selection of any act and what we call a decision situation, but it is very difficult to define it exactly. I think that the conflict notion, the notion of fractional anticipatory responses, suggests a possible continuum here; that is, some of these responses have been, in a sense, preprogrammed due to prior learning, so that, in fact, we have such an overwhelming response strength for a particular response that the organism does not phenomenally experience anything like hesitancy. There is no delay in the selection of the response. I am disinclined to call that a decision situation. This implies, of course, that there will be a continuum of decision situations, which vary in terms of the degree of conflict involved in the choice, and the uncertainty measure is one way of characterizing this dimension.

Dr. Berlyne also suggests some of the mechanisms of information collection and rejection, and their functional properties. Incidentally, I think the emphasis on rejection mechanisms is very important, especially since it is not often referred to in the literature of even the people who are most interested in dynamic decision processes. The assumption that the collection mechanism is the most important one, is, I think a useful re-emphasis, stressing as it does, that we ought also to consider the mechanisms for screening, filtering and rejecting information as it continuously impinges on the organism. These mechanisms, I think, are fairly interesting. I think what they imply is that, if in fact we have a fair amount of data already indicating what sort of mechanisms probably characterize the organism in terms of a normative model, we ought to use this information in postulating mechanisms. It seems unrealistic and - perhaps this is less of a problem than just poor strategy - possibly unfruitful to postulate mechanisms for either information collection or rejection which have no basis in existing literature. It would seem to me much more fruitful, and possibly leading to a much more rapid advancement, if the mechanisms we selected reflected what we do know about the organism on the basis of prior empirical work.

I think, in addition, that this approach suggests the conditions which affect elicitation of these mechanisms, which is again a point that most of the normative models do not deal with. Some do, but generally, the point of under what condition the organism does evoke information search responses is understressed. What are the conditions that maintain the search responses, and what are the variables that may determine future responses or truncate the process?

A fourth thing that I think his approach allows us to do, which most normative theories have not formally confronted, is to introduce, in a more natural way, emotional states into decision theory. They can be introduced into normative models, usually by postulates of negative utility for outcomes, and also by assuming some effect of either prior or present emotional disturbance states on the evaluation of the probabilities, that is, on subjective probabilities. There has been some work indicating that such effects do exist, but I think the postulation of a conflict mechanisms, or arousal states, allows one to introduce quite directly the notion of either prior emotional states or the current emotional state of the organism; it also allows one to predict, albeit a little bit, and precisely at the moment, what the effects of these emotional states of the organism might be on both information acquisition and processing activities.

I think it is probably apparent to all of us that in behavior theory generally there is certainly an understressing of the variables that have been talked about quite frequently in normative decision models, that is, such things as cost variables; that they can be introduced is typically understressed in behavior theory. I don't think there is any formal difficulty in bringing in these variables, but they were not explicitly considered in Dr. Berlyne's paper. There is some evidence, however, indicating that the cost of these processes themselves, the cost of acquiring information, and the costs associated with collecting and screening of information, are probably important in determining when these mechanisms get elicited and what the conditions are under which they are terminated. Another point that I think is not incompatible with behavior theory, but in Dr. Berlyne's presentation and in the form they have other than Dr. Roby's, is very much understressed, is the consideration of exactly what happens to this information when it is acquired. What is it modifying; what is the nature of this underlying cognitive structure; how do we describe it; what properties does it have; how is the information affected as a function of the properties of the information and properties of the organism? Here, again, there is a large amount of information from attitude and of opinion literature, which presumably is relevant, but as far as I know, has never been codified, and its implications for this problem have not been looked into. I don't know what the implications are; I do know, at least I feel, that we need to postulate something more than simply saying that the information coming in will modify the subjective probabilities. I don't think we can characterize the cognitive structure of the organism in quite that simple a form, that is, that he has a probability distribution of some sort and that this is the total content of his cognitive structure, in that the modification of it is simply a matter of introducing new information which directly operates on the structure and in certain determinate ways. In other words, I think another issue, one that presumably may be touched upon later, is this problem of the cognitive structure.

ATKINSON. Dr. Berlyne, you said that the degree of conflict is the antecedent of search behavior, and then you pointed out that sometimes the search might actually enhance the conflict and thus prolong the search. Now, is the main implication of this, for decision theory, that sometimes when the expected utility of one alternative exceeds that of another, but the difference isn't quite large, that the dominant tendency will, in fact, not be chosen and that the decision will not occur in this circumstance, but rather that a search will be undertaken? This seems to be rather a striking criticism of the assumption that the alternative having the highest expected utility will always be selected. This last point is the question. It appeared to me Dr. Berlyne said that sometimes when the strongest tendency is not sufficiently strong, relative to another tendency, instead of the act being performed, a decision being made, there will be a delay and a search undertaken which may, in fact, increase the degree of conflict, and change what the decision would have been in the initial state. One alternative had occurred with higher expected utility than the other and, in fact, was not chosen. I wanted to address that question both to Dr. Berlyne and also to some of the decision theorists here.

BERLYNE. Of course, I stress the role of conflict not because I think it is the only factor that plays a part, but because I suspect it plays quite a large role, and it has been neglected. Now, the idea is that, all things being equal, the motivation to gather information will be maximum when conflict is at a maximum, depending partly on the nature of the responses. As Lewin and Miller have shown, in an approach-approach conflict there may be a tendency to over-balance in the direction of one decision just by chance without gathering information; but in the case of more complicated kinds of conflict, such as approach - avoidance, in trying to escape from the conflict, the subject traps himself; the further he goes in one direction, the more he is pulled in the opposite direction. Now, one of the things that can happen in this situation is a search for new information. Additional information may reduce the degree of conflict by strengthening or weakening one or another of the competing response-tendencies, or it may introduce another possibility that crowds out the original alternatives. Of course, there are dangers in information-gathering: the collection of information is apt to be biased and consequently misleading; we know from the start that we may end up convinced that we ought not do the thing that we feel inclined toward at the moment. At any rate, the assumption is that conflict must be brought below a certain level - one response-tendency must exceed the others in strength by a certain amount - before the subject is ready to act.

BERLYNE. It is interesting to note that when you are in a situation in which you have to make a decision with not enough information, that is, you are uncertain as to which is the right decision, there is a tendency to supply spurious information. There is a good article by O. K. Moore with special reference to primitive fortune telling and devination devices such as tossing a coin or looking up your horoscope; these can provide you with stimuli to resolve your conflict, making you do one thing rather than another. This is spurious information, in that although it may be biologically adaptive, it doesn't inform you about the optimal response.

EDWARDS. This vacillation and the fact that information can hurt as well as help is a very clear property of decision theoretical models of the type we have been talking about. Consider a situation of straight choice-reaction time; suppose that you're doing a choice reaction-time experiment with two lights and two push buttons. You are to push the left button if the left light comes on, and to push the right button if the right light comes on; you're supposed to do this as quickly as possible, and there are well-defined cost and payoffs for being right or wrong, and the cost function is associated with time. If the left light turns on most of the time, but not all the time, you may be very willing to press the left button with no fear at all. If, on the other hand, you are quite willing to do that, and if in fact the right light was the one that turned on, then during the time in which you are discovering that it really is the right light, you must pass through a period in which you are more confused than you would have been if you hadn't made any observation at all. Nevertheless, under those circumstances, observing is indeed the strategy which will maximize expected value even though it may not maximize expected value on that single trial.

LANZETTA. All of this, it seems to me, implies a postulation of some threshold at which responses get elicited, either in the forced-choice reaction-time situation, or in the freer decision situations you're talking about; presumably there is some mechanism for initially eliciting the search and for terminating the search. Implicit in this is some notion of a threshold of uncertainty at which the organism is willing to make the decision, to commit himself to one choice or another. If not, continuing to search may increase the uncertainty rather than decrease it.

BERLYNE. Very often, you have to go to a stage of increased uncertainty in order to get to the goal of reduced uncertainty. This is true in many situations. This is one of the insights we get from Freud: the postponement of gratification. Very often, for example, we are hungry and we have to let things get worse before they can get better. We have to wait awhile more, suffer more hunger, before we can satisfy it in a more effective way. Perhaps there are better examples in the case of fear.

ROBY. I just want to enter a statement for the record. You used the phrase "even the cat". My experience with cats is that they have more intellectual curiosity than most of the college students I have come in contact with.

BERLYNE. Of course, perceptual curiosity. Possibly, epistemic curiosity. There is possibly even a cognitive element; it's too much to say intellectual, but at least a cognitive element in the simplest form of animal learning in light of some pieces of evidence.

SHUFORD. I have two comments and two questions. First, you gave an information theory definition of adaptation, and you said the point of it is that you want the distribution of response probabilities to correspond as closely as possible to the distribution of optimal response probabilities. As I understand that, this leaves out of consideration how optimal the response may be, or any definition of risk, or cost, or of utility. For example, it may be that one response is optimal 50% of the time, and another response is optimal, say 25% of the time; but it may be that the response which is optimal 50% of the time, is almost as good, in terms of expected utilities, as the response which is optimal 25% of the time, so there is very little to be gained by trying to match them. Things like this can happen. You don't take into account just how optimal each of these responses are.

Secondly, you say that each course of action must be represented with the organism by a fractional response. This leads into a problem which has concerned me, particularly since reading Von Neumann's book, The Computer and the Brain, which points out that the mathematics we deal with in theory may not correspond to the "mathematics" of the brain. I won't go into detail, but there is this problem of taking your models too seriously. Maybe the subject does think of alternative courses of action, but I think that it is dangerous to take all aspects of decision theory as actually being a process using addition, subtraction, multiplication and division within the individual.

Now, as to the two questions. First, you say that the point of behavior theory is to contain explanations of as many variables as possible. It seems that the program of behavior theory is to eventually - this is a long way off - completely describe all aspects of behavior, taking into account all variables, etc. Well, this is a very idealistic goal. But I wonder, if we ever do get to the point of being able to explain all behavior and taking into account all

BLANK PAGE

variables, whether the theory would be of much use because it would take so long to compute or come up with any predictions or anything from such a model? We might just as well look at the behavior of the operator or human being because the theory isn't serving much purpose. Maybe there's some consideration I'm leaving out. Also, you said that the examination of simple phenomena suggests the direction of study for more complex phenomena. Let me use an analogy. We have say, a desk calculator, and a Univac 1105. In the 1105 is a program for addition, multiplication, and subtraction and also the Newell, Shaw, and Simon program, which proves theorems contained in Whitehead and Russell. We put things into the desk calculator and see what we get out and measure the reaction time. We put the same things into the computer, and see what we get. It turns out that the computer is probably slower than the desk calculator. I don't see how studying the desk calculator would ever lead us to discovering this other aspect of the computer.

BERLYNE. First of all, the nature of behavior theory. I had to speak a little too briefly on some of these points, including the nature of behavior theory. I was trying to frame my ideal picture of behavior theory, the kind of behavior theory to which I would like to contribute. I was certainly not suggesting that all psychologists drop whatever they are doing and turn to this, but some people should do it. Behavior theory, as I conceive of it, is a sort of theory that attempts to encompass and to have something to say about everything that is known about behavior today, including everything known about the brain - and quite a lot is known now, especially from work done in the last ten years. When you do this, you have to do it at a cost.

The second question is certainly related to the first: the role of studying simple organisms. Now, first of all, I am not saying that you can learn everything you need to know about complex organisms from studying simple ones. A priori, if you study simple behavior, it will have something in common with complex behavior. Some of the features will apply to the more complex and some will not. It is, of course, difficult to know which, but at least this is giving you a starting point. You find out what principles govern behavior beyond conditioning, and then ask yourself which of these principles, if any, (and presumably there will be some) apply to complex decision making in a human situation. Even if you find some of these do not apply, this can be a valuable starting point to examine the differences and to explain the differences. At least, this is one useful strategy.

Another reason for this is very important if you are a behavior theorist. They are, of course, mostly psychologists the farthest from applied psychology, not because they are against applied psychology but because they believe somebody ought to be doing this thing which temperamentally they'd like to do. It means that we are guided by things - that we are guided by our own curiosity, if you like. Speaking for myself, I wouldn't be satisfied with a theory of decision making that didn't tell me how this was related to conditioning. I don't want to be told exactly; the cost-payoff conditions don't do it completely; the differences tell me that. I strongly suspect it is like the dog in some respects and unlike it in other respects, but I want to know; at least, tell me how it's related. Otherwise, I should say the psychologist hasn't done the job until he has answered these questions. That is my reason for bringing in this approach.

ROBY. I have a comment on the term "simple behavior". This has come up a couple of times. Yesterday you brought up the fact that complicated organisms don't seem to react well in simple behavior situations. The point I think I'd like to make is that probably, in complex organisms, we are not really getting simple behavior. What we have is simple-minded measurements of behavior that is just as complex as it is in more complex situations. I think the desk calculator is a fairly good analogy, and brings this point out.

BERLYNE. Yes; the point is very often made that the human brain, compared with the computer, can do many things that the computer cannot; but the computer can do the things it can do better than the brain. Many of the things our brain can do interfere with each other, and this is true of predecisional information gathering. These processes have, among other things, the function of helping to guide us to correct decisions. Unfortunately in some respects, and fortunately in other respects, they have other functions as well - for example, bolstering prejudice, removing dissonance, making you feel happier - and they get in the way of each other. While information gathering in order to find which decision ought to be adopted, we may find instead that we are being led in a direction interfering with the correct decision because these other functions are having deleterious effects.

FLOOR. I noted in Dr. Berlyne's talk quite a bit of emphasis on arousal as a kind of general property of emotion and motivation. Yet the utility concept is a relatively specific one where specific values are attached to outcomes. I also note in some psychophysiological writings that there is some emphasis on the specificity of the pattern of autonomic responses. The question I'm leading up to is this. Did you find, at least within a given individual, a general kind of arousal pattern, and secondly, at what point does it occur in, say, a gambling situation involving large positive payoffs if you win, and large negative payoffs if you lose? Would it occur at the point where the alternative bets are chosen, or would it tend to occur where the individual finds out whether he has won or lost?

BERLYNE. The answers to these questions are, of course, only speculative at this moment. There has been comparatively little done systematically on the individual aspects of arousal. As far as the gambling situation is concerned, it is expected that you get maximum arousal when you get maximum conflict. There is a bit of indirect evidence for this, at least from some of the studies made relatively early on arousal processes during thinking. It has often been found that you get maximum arousal when people are figuring out what to do with the problem, when it has just hit them and they are wondering what to do about it. In a later phase, when they have hit upon a course of action (maybe Dr. Toda's post decisional process), and it's just a matter of going through with it, then arousal seems to be reduced. In the gambling situation, I would suspect you would have much more arousal. When you have maximum uncertainty about what is going to happen, you might get a sort of rise again towards the end, just before you find out the answer, due to anticipation and other sources of emotionality. I think this would correspond to subjective experience as well. This would certainly be a most interesting and fruitful field for research.

LANZETTA. I'd like to comment that we do have some recent data from our laboratory which seems to indicate that at the point at which you present the problem, arousal is very high. The organism makes a query and arousal starts going down. Apparently at some point after the information has been digested, in some sense, it seems to jump up again and go down with the next query, in a sawtooth-like function; just at the point of decision it seems to shoot up and then drops off, and then another peak. We have some data indicating that this seems to track the GSR.

It seems to track what is happening in the decision situation, but as you all know, with the problems of GSR's, it doesn't allow you to feel very confident about this kind of result.

TODA. Dr. Berlyne mentioned that the onset and termination of information-seeking behavior should be related to the amount of uncertainty. This is not necessarily true because sometimes the person may realize that he cannot reduce uncertainty, and in that case, when he is certain that he cannot reduce uncertainty, then he will stop seeking further information and make the decision. So the termination of information seeking behavior cannot be described solely on the basis of uncertainty.

LANZETTA. I certainly would agree with that. I would tend to postulate that the threshold level is not the only factor terminating search, but that there is some expectation of reducing uncertainty by searching, and if the actual search does not produce uncertainty reduction commensurate with the expectation, this in itself is negatively reinforcing and could terminate the search process.

BERLYNE. With regard to the seeking of uncertainty reduction, sometimes we use rational language for convenience because it's shorter; but I, myself, am thinking in terms of reinforcement. It is not so much a matter of the organism looking for information that will reduce uncertainty, but working out what way is the best to get the information. In human beings that sometimes is the case, but it's more a matter that any response that happens in the past to have been followed by uncertainty reduction, and thus reinforced, is more likely to occur. Those are not quite the same things.

BLANK PAGE

PREDECISIONAL INFORMATION PROCESSES: SOME DETERMINANTS OF INFORMATION ACQUISITION PRIOR TO DECISION MAKING*

John T. Lanzetta and Joan Sieber**
University of Delaware

Empirical work in decision making has typically been performed within a normative theoretical framework and has generally focused on testing one or another of the assumptions or predictions of the theory. Thus, most studies in decision making have utilized experimental situations in which subjects are provided an "information base" in terms of which they make a choice; subjects are not permitted the option of delaying a choice until additional information is available or while actively seeking further information regarding alternatives, outcomes, or the contingent relationships between alternatives and outcomes, e.g., refs. 11, 14, 37.

A theory of decision making, however, must eventually consider the activities of acquiring and processing information that precede decisions; it cannot assume and thus leave unexplained the basis for the choice itself. The omnipotent rationality of economic man must be replaced by a concept of rationality which considers the capacities of the organism for assimilating and organizing information, and the information state of the organism at the time of decision (ref. 26).

Unfortunately, once the boundaries of decision studies are extended to include information acquisition and processing activities, a host of theoretical and empirical issues are raised. Under what conditions does an organism instigate search for information? What variables control the direction and redirection of search activity? How does the organism utilize the information acquired, i.e., how are tentative decisions generated and modified as new information is obtained? What variables control the termination of search and lead to a committment to decide? Although research from several disciplines has contributed to an understanding of such issues (e.g., refs. 8, 21, 27, 31, 41) we are, on the whole, still in the inchoate stages.

---

* This research was supported in whole or in part by the United States Air Force under Contract No. AF 33(616)-5845, monitored by the 6570th Aerospace Medical Research Laboratories, Aerospace Medical Division, Wright-Patterson Air Force Base, Ohio, and by the National Science Foundation Grant No. MH-03909-04.

** This is a revised version of the paper presented at the symposium by the senior author. Grateful acknowledgement is due Dr. Vera Kanareff for her active and invaluable collaboration; to Mrs. Jo Ann Davis for her assistance in conducting several of the studies and analyzing the results; to Dr. Louis Miller and Dr. James M. Driscoll for their stimulation and assistance.

The research to be reported in the present paper focuses on one of the simplest of the variables subsumed under information processing--the amount of information acquired prior to decision. Because of the paucity of empirical data and theory, the spirit of our inquiries was descriptive at inception, but theoretical notions suggested by early results have come to play a larger role in determining the selection of variables and experimental paradigms. The "theory" at this stage is a rather loosely formulated set of propositions borrowed in part from other contexts and in part representing efforts to account for our empirical data and serves primarily a heuristic function. In the discussion to follow empirical findings will be stressed and only brief attention will be direction to theoretical issues.

OVERVIEW

Our general approach has been to place adult humans in choice situations in which the alternatives and outcomes of choice are specified but with varying amounts of initial information relevant to the assignment of a probability distribution to the alternatives. Additional information useful in identifying the "best" alternative is available by reference to an external display. Such additional data is not costless, however: time, effort, and/or a monetary cost are levied for information items. Subjects have no control over the type of information available and interest is focused only on the amount of information acquired prior to making a decision.

Using this general procedure, we first examined the effects of variables such as the cost of information, time pressures, and level of aspiration on information acquisition behavior. Since large problem differences were observed in information acquisition and time allocation, an attempt was made to measure and identify those problem variables which are responsible for differences in amount of response uncertainty elicited in the decision maker. This led to the development of various measures of stimulus and response uncertainty and the exploration of their relationship to information acquisition and some other variables descriptive of predecision behavior. Information theory concepts proved of value in both the methodological and theoretical problems encountered here. The question of motivation remained unexplored: what were the mediating factors through which problem uncertainty led to information acquisition? Berlyne (ref. 6) has used the concept of physiological arousal to link the perception of complex or "uncertain" stimuli to manipulatory or exploratory behavior. Two studies which utilized measures of physiological arousal were performed in order to examine these hypothetical relationships between uncertainty, arousal and information search.

In the course of all of these studies, large and consistent individual differences in decision making strategies have been observed. Some organismic variables such as anxiety, conceptual structure, and academic achievement were found to account for some of these differences. A joint relationship of information processing to conceptual structure and uncertainty was empirically established, lending support to a theoretical model of information processing which utilizes concepts from theories of arousal and conceptual structure.

Thus, our research has been guided by concepts from such diverse sources as economics, physiology, and personality theory. The series of studies discussed here will follow somewhat the historical sequence of ideas described above.

INFORMATION ACQUISITION UNDER VARYING COST-PAYOFF SCHEDULES, LEVEL OF ASPIRATION, AND TIME PRESSURES

The earliest studies (Lanzetta & Kanareff, ref. 24) though descriptive in intent, focused primarily upon variables which normative models suggested were important: costs, payoffs, and utility of outcomes. An effort was made to control the expected utility of acquiring information since normative decision models suggest maximization of expected utility as a reasonable criterion of choice. Expected utility was manipulated by controlling the contingency between probability of payoff and amount of information acquisition, and by judicious selection of cost-payoff schedules. However, subjects were not informed of the probabilities of payoff or the contingency between probability of payoff and the number of information queries. From the subjects' point of view, they were faced with the necessity of making a difficult choice; the probability of being "correct" presumably improved with additional information not because of an experimenter controlled contingency but because the information was relevant to selecting an alternative.

In brief, the experimental procedure was as follows. The subject was presented with a series of decision problems. Each problem required a choice to be made among six possible alternatives and consisted of an information base (a statement of the problem, some facts or observations which might aid in selecting an alternative, and a statement of the six alternatives) plus five additional information items. The subject was presented the information base and provided the option of making an immediate decision or obtaining one of the additional information items. This same option was available for a maximum of five information items at which time the subject made a choice among the six alternatives. If correct, the subject received a chip redeemable for money at the end of the session. The number of queries made, the time delay between queries, and the subject's choice were automatically recorded.

Although the information obtained by making a query was phenomenally relevant to reducing uncertainty (as determined by pre-tests) it in fact had no bearing on the probability of obtaining a payoff. The probability of payoff was programmed by the experimenter and was a linear function of the number of information-seeking responses (queries) made by the subject. The schedule employed was as follows:

Number of Query	0	1	2	3	4	5
Probability of Payoff	0	.2	.4	.6	.8	1.0

Thus, a subject who made all five queries was assured of a payoff irrespective of his actual choice whereas one who made no queries would receive no payoff whatever his choice. Intermediate numbers of queries were associated with intermediate probabilities of payoff. Thus, information-seeking and not effective decision making or information processing was instrumental to obtaining a payoff.

The major independent variables in the first two studies were the information cost-payoff schedule, the level of aspiration for being correct, and time pressures. The conditions of information cost-payoff employed were equated for expected profit which varied with the number of queries in the following way:

Number of Query	Probability of Payoff	Expected Profit
0	0	$ 0
1	.2	.01
2	.4	.02
3	.6	.03
4	.8	.04
5	1.0	.05

In this situation,

where p = probability of payoff q = number of queries
$\$$ = amount of payoff c = cost of a query

$$\text{Expected Profit} = p\$-qc.$$

In the first study we examined two cost-payoff schedules, namely no cost for an information query with a $.05 payoff (0-5), and $.05 per query with $.30 payoff (5-30), and three levels of aspiration. The level of aspiration was manipulated by the presentation of fictitious group norms. The three levels were introduced by reporting to the subject before the beginning of the session that either 25, 50, or 75 percent of the problems were correctly solved by previous subjects. Two orders of problem presentation were used. The twenty-five problems were twice randomly ordered providing some check on sequence effects.

In addition to the "search" measure, we obtained records of time devoted to various aspects of the task. The total time per problem could be naturally divided into three periods: the time from the initial presentation of the problem to the first query (*problem time* or *P.T.*), the time from the first query to the last query (*total query time* or *Q.T.*), and the time from the last query to a decision (*decision time* or *D.T.*). It is assumed that *P.T.* reflects time spent in reading the problem and formulating an initial hypothesis as to the correct decision, *Q.T.* reflects time devoted to acquiring, assimilating, and organizing additional information, and *D.T.* reflects time spent in evaluating the alternatives and making a choice. The time measures, of course, can be considered only approximate indicators of the allocations of time to such functions since we had no control over subjects' cognitive activity. For example, subjects may very well have delayed reading new information until after the last query, resulting in a short *Q.T.* and long *D.T.* At any rate, the assumptions bear mainly on the interpretation of results; the measures themselves are unambiguous.

The overall level of information seeking was not very high even when information was costless, and there was an unexpectedly slow rate of acquisition of information-seeking, considering the instrumental value of a query, as shown in figure 1. The lower cost-payoff schedule elicited a greater number of queries under all levels of aspiration conditions, and there appeared to be some inhibition of search for the higher levels of aspiration.

In regard to the effect of the cost-payoff schedule, there was an increase in queries under the (0-5) schedule; however, the average number of information-seeking responses never exceeded 4 queries per trial. For the (5-30) schedule there was some indication of a decrease in queries over trials.

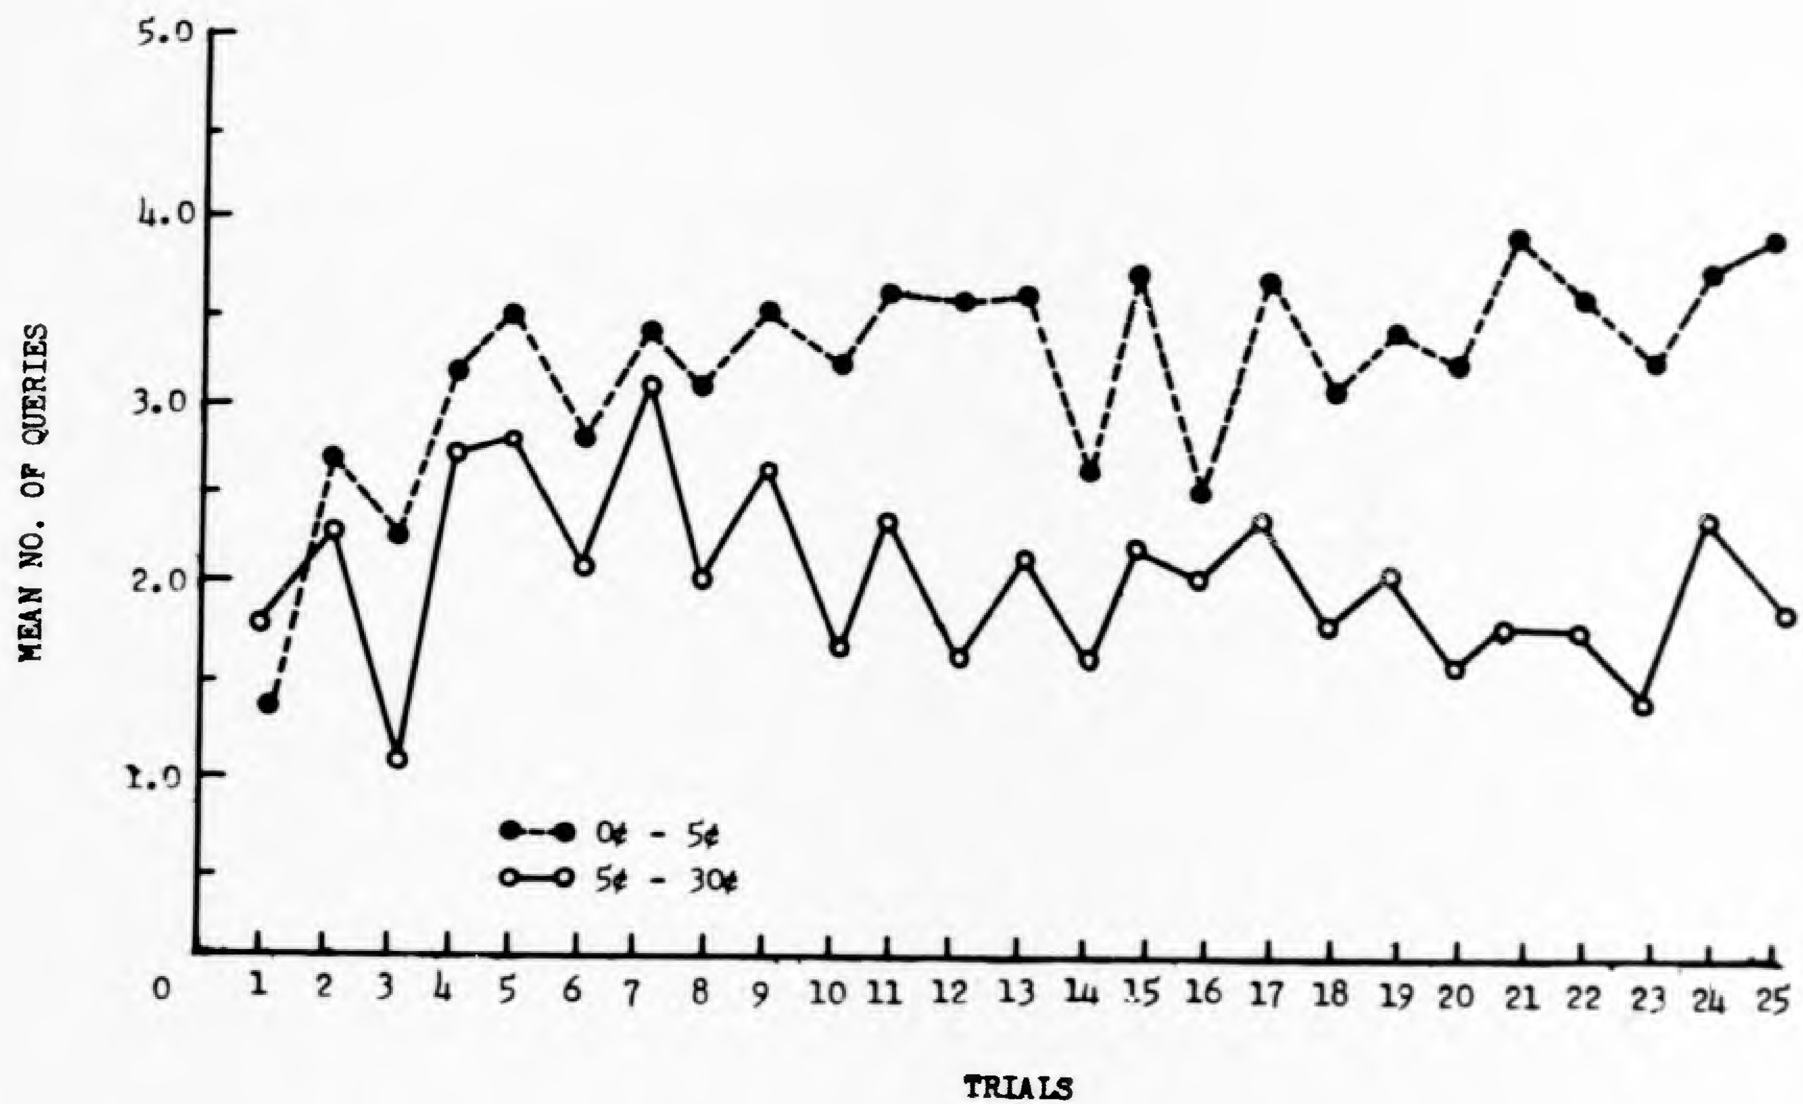

Figure 1. *Mean Number of Queries over Trials at 0¢ - 5¢ and 5¢ - 30¢ Payoff Levels*

Analyses of variance of the data on total time, problem time, average time per query, and decision time, indicated the following: Subjects within conditions and trials contribute significant variance for all time measures whereas none of the experimental variables contribute significant variance. Considering the difference between cost-payoff schedules in average number of queries the failure to find time differences between schedules was surprising. Apparently subjects took as long to handle a problem under (5-30) conditions as under (0-5) conditions even though they sought less information.

Although characterized by marked fluctuation, total time, average time per query and decision time decreased somewhat over trials. Some of this variation almost certainly reflects a "problem" effect and effects contributed by the interaction between problems and experimental conditions.

In the second study there were five conditions of information cost-payoff used: no cost per query - no payoff (0-0); no cost per query - $.05 payoff (0-5); $.01 per query - $.10 payoff (1-10); $.05 per query - $.30 payoff (5-30); and $.10 per query - $.55 payoff (10-35). For example, under the (5-30) schedule subjects "purchased" each additional information item at $.05 per item and won $.30 if they made a "correct" choice. In addition, two conditions of pacing were examined, one minute 30 seconds or two minutes 15 seconds. This refers to the total time available for each problem. The subjects were informed of the time constraints and given two practice problems to familiarize them with this pacing requirement.

The average number of queries per trial was consistently higher for the slow pace condition and relatively uniform for both conditions. On the other hand, there was great variability over trials and much overlap between conditions for the various cost-payoff schedules. Figure 2 presents the mean number of queries for the cost-payoff and pacing conditions. In general, there was a decrease in information acquisition under higher cost-payoff for both pacing conditions although the addition of any cost for information was much more pronounced for the slow pace than the fast pace condition.

As in the first study, there were no significant differences, for any of the time measures, attributable to the cost-payoff variable. However, as expected, there was less time spent in reading the problems, making queries, and in making a decision under the faster pace condition. The greatest change occurred in the average time spent per query, i.e., when time is limited subjects speed up information processing more than deliberation over the choice of an alternative. Also, subjects spent less time in all phases of the problem with increasing experience. Once again, the greatest change occurred in the average time per query--subjects spent about 25 percent less time per query on the last four trials than on the first four trials.

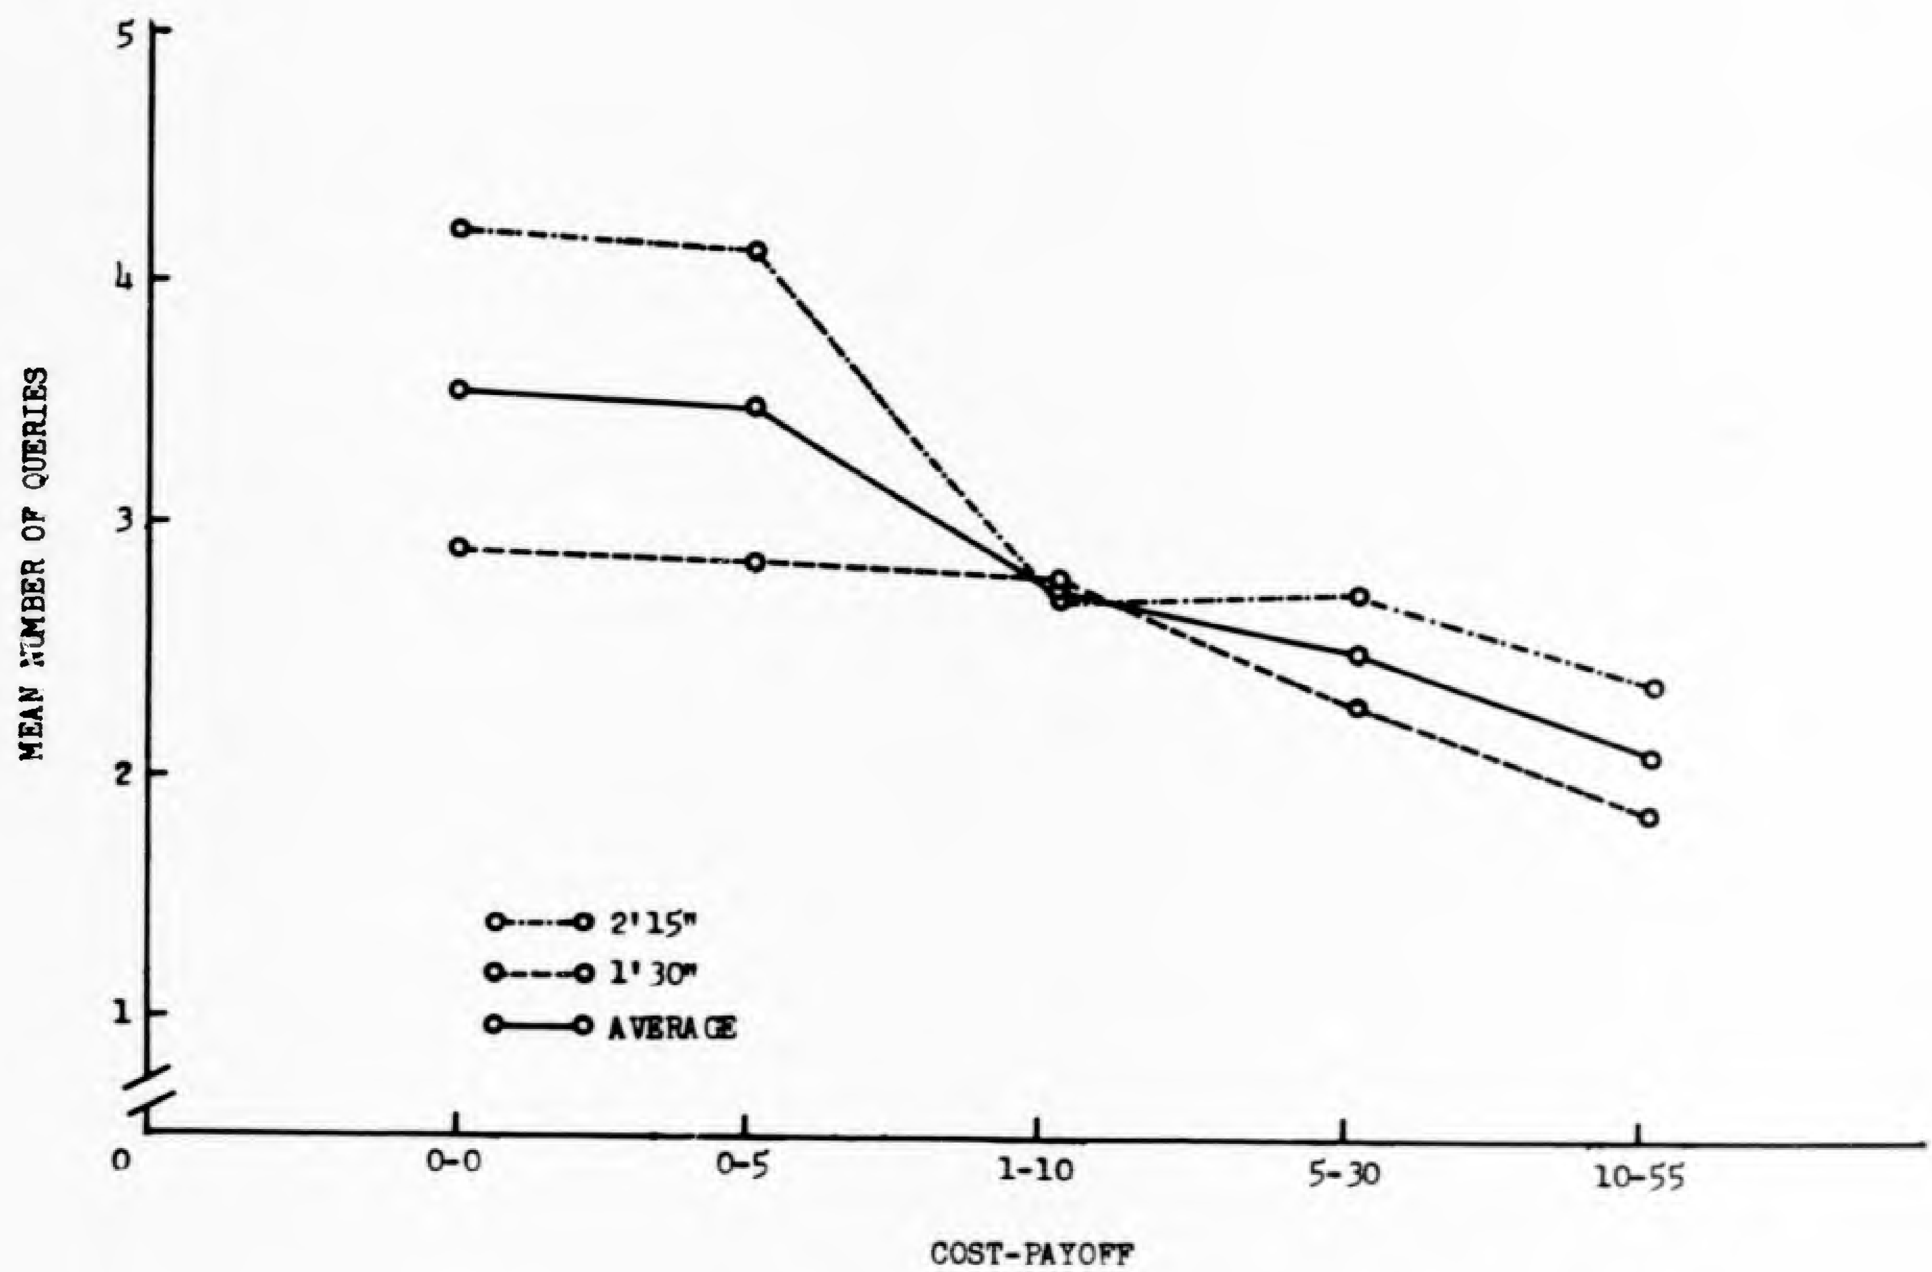

Figure 2. Mean Number of Queries over Five Cost Payoff Levels

In summary, the major results for the search and time measures were as follows:

(1) For none of the experimental conditions did a majority of subjects obtain all of the information available to them. This result is especially surprising for the zero cost conditions.

(2) The mean number of queries per trial decreased as the cost-payoff schedule increased. Subjects made their decisions on the basis of less information when information was costly even though, in terms of expected profit, the risks were the same for all cost conditions.*

(3) The mean number of queries per trial decreased as time pressures increased, as did the average time spent in making queries. In general, as time pressures increased, subjects sped up information processing more than deliberation over the choice of an alternative.

UNCERTAINTY, UNCERTAINTY REDUCTION, AND INFORMATION ACQUISITION

As was noted above there were pronounced differences in information acquisition and time allocation for the various problems used. The differences could result from a variety of factors intrinsic and extrinsic to the problems, e.g., length of problem, previous experience with certain types of material, extent of knowledge in specific areas covered by the problems. Such factors would presumably contribute to differences in the amount of perceived initial uncertainty concerning the best decision alternative. There is reason to expect that differences in uncertainty would result in differences in decision time (refs. 9, 10, 15) and possibly also in the amount of information seeking.

---

* In view of earlier findings by Edwards (ref. 14) and others with respect to probability preferences, it was considered possible that such differences between conditions might be due merely to specific preferences for certain cost-probability-payoff combinations over others. Accordingly, we conducted a study (ref. 30) in which three of the cost-payoff conditions were presented to 72 Ss in the form of four alternative gambling situations, in which the costs, probabilities of win, and payoffs were all clearly made known to the subject. The subject was required, on each of 80 trials, to choose one of these four alternatives to bet on.

In contrast to the findings reported above, choice behavior was not found to differ significantly between conditions, either in terms of an overall index of average choice behavior, or in terms of the specific distribution of choices among the four alternatives. It thus appears inappropriate to attribute the differences between conditions in mean number of queries observed in the above reported studies to preferences for certain cost-probability-payoff combinations. The decrease in information-seeking behavior which accompanied increase in monetary level of play remains to be accounted for.

Our first attempt to assess the role of uncertainty was based on post hoc analysis of the data from the second study. The problems used were scaled for uncertainty on an independent sample of 12 and 26 subjects respectively, who were presented the problems with and without the information available from queries. After reading each problem the subjects responded to several items assessing their choice of the correct alternative, their confidence in their choice, and their judgment about which alternatives could definitely be eliminated.

Two indices of uncertainty were developed. The *uncertainty measure* (U.M.) for each problem was determined from the following equation (ref. 3):

$$H = -\sum_{i=1}^{6} p(i) \log_2 p(i)$$

Three types of alternatives were distinguished and probabilities assigned as follows: $p(i_c) = \frac{1}{100} \times \textit{confidence score}$ for the alternative i chosen as correct; $p(i_d) = 0$ for the alternative(s) which were discarded; and $p(i_r) = \frac{1 - p(i_c)}{\textit{number of remaining alternatives}}$ for the remaining alternatives.

Thus $U.M. = - p(i_c) \log_2 p(i_c) + \sum_{i=1}^{5} p(i_r) \log_2 p(i_r)$

The *prominence measure* (P.M.) for each problem was based on the same equation, but with $p(i)$ estimated from the proportion of subjects choosing alternative i. It was assumed that the degree of initial uncertainty associated with a problem would be reflected in the distribution of choices of the most likely alternative: the lower the initial uncertainty the greater the likelihood that there would be a "prominent" choice. The average correlation between the uncertainty measure and the prominence measure for both samples is .245 which is not significant at the .05 level.

Based on the distribution of U.M. scores and on P.M. scores the problems were categorized as either high, medium, or low in degree of perceived initial uncertainty. One problem was discarded and cut-off points were selected so that eight problems were placed in each of the three categories. Separate analyses of variance for the number of information-seeking responses and total time for completing a problem were performed for the two bases of categorization.

Degree of initial uncertainty, as measured by U.M., was found to be positively associated with number of information-seeking responses and time per problem, but only the former difference was significant. An examination of the means for information-seeking responses for the three uncertainty categories suggests a slight curvilinear relationship: the largest number of queries for problems of medium uncertainty and the smallest number of queries for problems of low uncertainty. For the prominence measure P.M. there was also evident an association between uncertainty and the number of queries and time per problem, with only the latter difference significant. Again the greatest number of queries was obtained for problems of medium uncertainty but the lowest number of queries now occurred for problems of high uncertainty. Total time, as with the U.M. measure, increased monotonically with increasing uncertainty.

Degree of perceived initial uncertainty appeared to be curvilinearly related to the number of information-seeking responses and monotonically related to total time per problem, but did not interact with any of the experimental variables to a significant extent.

A third study explored further the relationships between problem uncertainty and information acquisition behavior. The earlier work suggested that information search was elicited by a response conflict engendered by response uncertainty: the greater the degree of uncertainty, the stronger the conflict and the stronger the instigation to search. However, the relationship obtained between information cost and amount of information acquisition also suggested that search was maintained only when the rate of uncertainty reduction exceeded an expected rate, the latter being a function of the cost of information. Thus subjects would initiate a search for information whenever a choice represented a response conflict, but they would maintain the search only so long as the amount of uncertainty reduction was commensurate with the cost of information, i.e., as long as information acquisition "paid off".

In this study subjects were presented with a concept attainment task, again with the option of delaying a choice until further data were gathered. Six groups of subjects were run under varying conditions of degree of initial uncertainty, rate of uncertainty reduction, and level of monetary play. The primary dependent variable, as before, was the number of information-seeking responses prior to decision.

The tasks were modeled after those used by Bruner, Goodnow and Austin (ref. 8) in research on concept-attainment. The task material consisted of an array of cards each containing a geometric pattern. The pattern on a card could vary in respect to four attributes: form, color, number of borders, and number of objects, each with three values.

Subjects were to identify a "concept" held by the experimenter, the concept being defined in terms of one value of each of two different attributes, e.g., red-square. The subjects were given an answer sheet for each problem consisting of groups of "concepts"; their task was to choose the group containing the correct concept. They could eliminate concepts by requesting cards showing *positive examples* of the concept, i.e., examples containing one *or* both of the attribute values of the correct concept. The elimination strategies were fully explained to subjects and several practice examples were used to insure that subjects completely understood the procedure.

The details of the construction of the problems to obtain varying rates of uncertainty reduction need not be elaborated upon here. Problems were constructed which varied with respect to level of initial uncertainty and rate of uncertainty reduction. Whatever the initial uncertainty and rate of uncertainty reduction, all problems had approximately the same terminal uncertainty level (.9 bits). That is, subjects could eliminate *almost all* of the alternatives by asking for all the available information, but could not reduce to zero the uncertainty regarding the correct alternative. Two levels of initial uncertainty, 3 and 2 bits, were obtained by varying the number of response categories from which the subject had to make a choice. Two rates of uncertainty reduction, approximately .2 and .4 bits per query, were obtained by varying the number of concepts eliminated by the particular card shown to the subject. In addition to the above variations, two cost-payoff schedules were employed, $0--$.05, and $.03--$.45.

In order to have a reasonable possible range of search responses it was not feasible to employ a high rate of uncertainty reduction with a low level of initial uncertainty. Thus three conditions of initial uncertainty and rate of uncertainty reduction were used: High initial uncertainty (*HIU* = 3 bits) and high rate of uncertainty reduction (*HUR* = .42 bits per query); *HIU* and low rate of uncertainty reduction (*LUR* = .2 bits); low initial uncertainty (*LIU* = 2 bits) and *LUR*. All three conditions were run under both cost-payoff schedules. The results, in summary, were as follows:

Subjects obtained more information when initial uncertainty was high. However, the expected greater amount of information acquisition with a higher rate of uncertainty reduction did not occur; in fact, the opposite effect was obtained. Also, in contrast to our earlier results there was little effect on search behavior of increasing the cost for information.

Since both the initial level of uncertainty and average rate of uncertainty reduction were controlled it was possible to determine the average residual uncertainty remaining at the time of decision. Table 1 presents the means of the residual uncertainty at decision for the various conditions. It appears that the uncertainty level at decision was relatively constant for all conditions. Residual

TABLE 1

MEANS FOR RESIDUAL UNCERTAINTY AT DECISION
FOR THE EXPERIMENTAL CONDITIONS

Cost	High Uncertainty High Uncertainty Reduction Rate	High Uncertainty Low Rate	Low Uncertainty Low Rate
High	1.15	1.61	1.23
Low	1.32	1.29	1.22

uncertainty was slightly greater for the higher level of initial uncertainty and slightly lower for the higher rate of uncertainty reduction, but neither of these differences were significant. Subjects varied their acquisition strategy depending on the initial conditions of uncertainty and the rate of uncertainty reduction so that the final level of uncertainty was relatively constant. It should be pointed out that this residual uncertainty level is higher than the terminal uncertainty level, i.e., the level of uncertainty remaining if <u>all</u> queries were made. This can be seen most clearly in figure 3.

This finding suggests that there exists a "commitment threshold" for decision: subjects acquire information until they reach this uncertainty threshold at which point they make a decision. There is also some suggestion that the "commitment threshold" is a function of the level of initial uncertainty and rate of uncertainty reduction although further research on the determinants of commitment thresholds is obviously required.

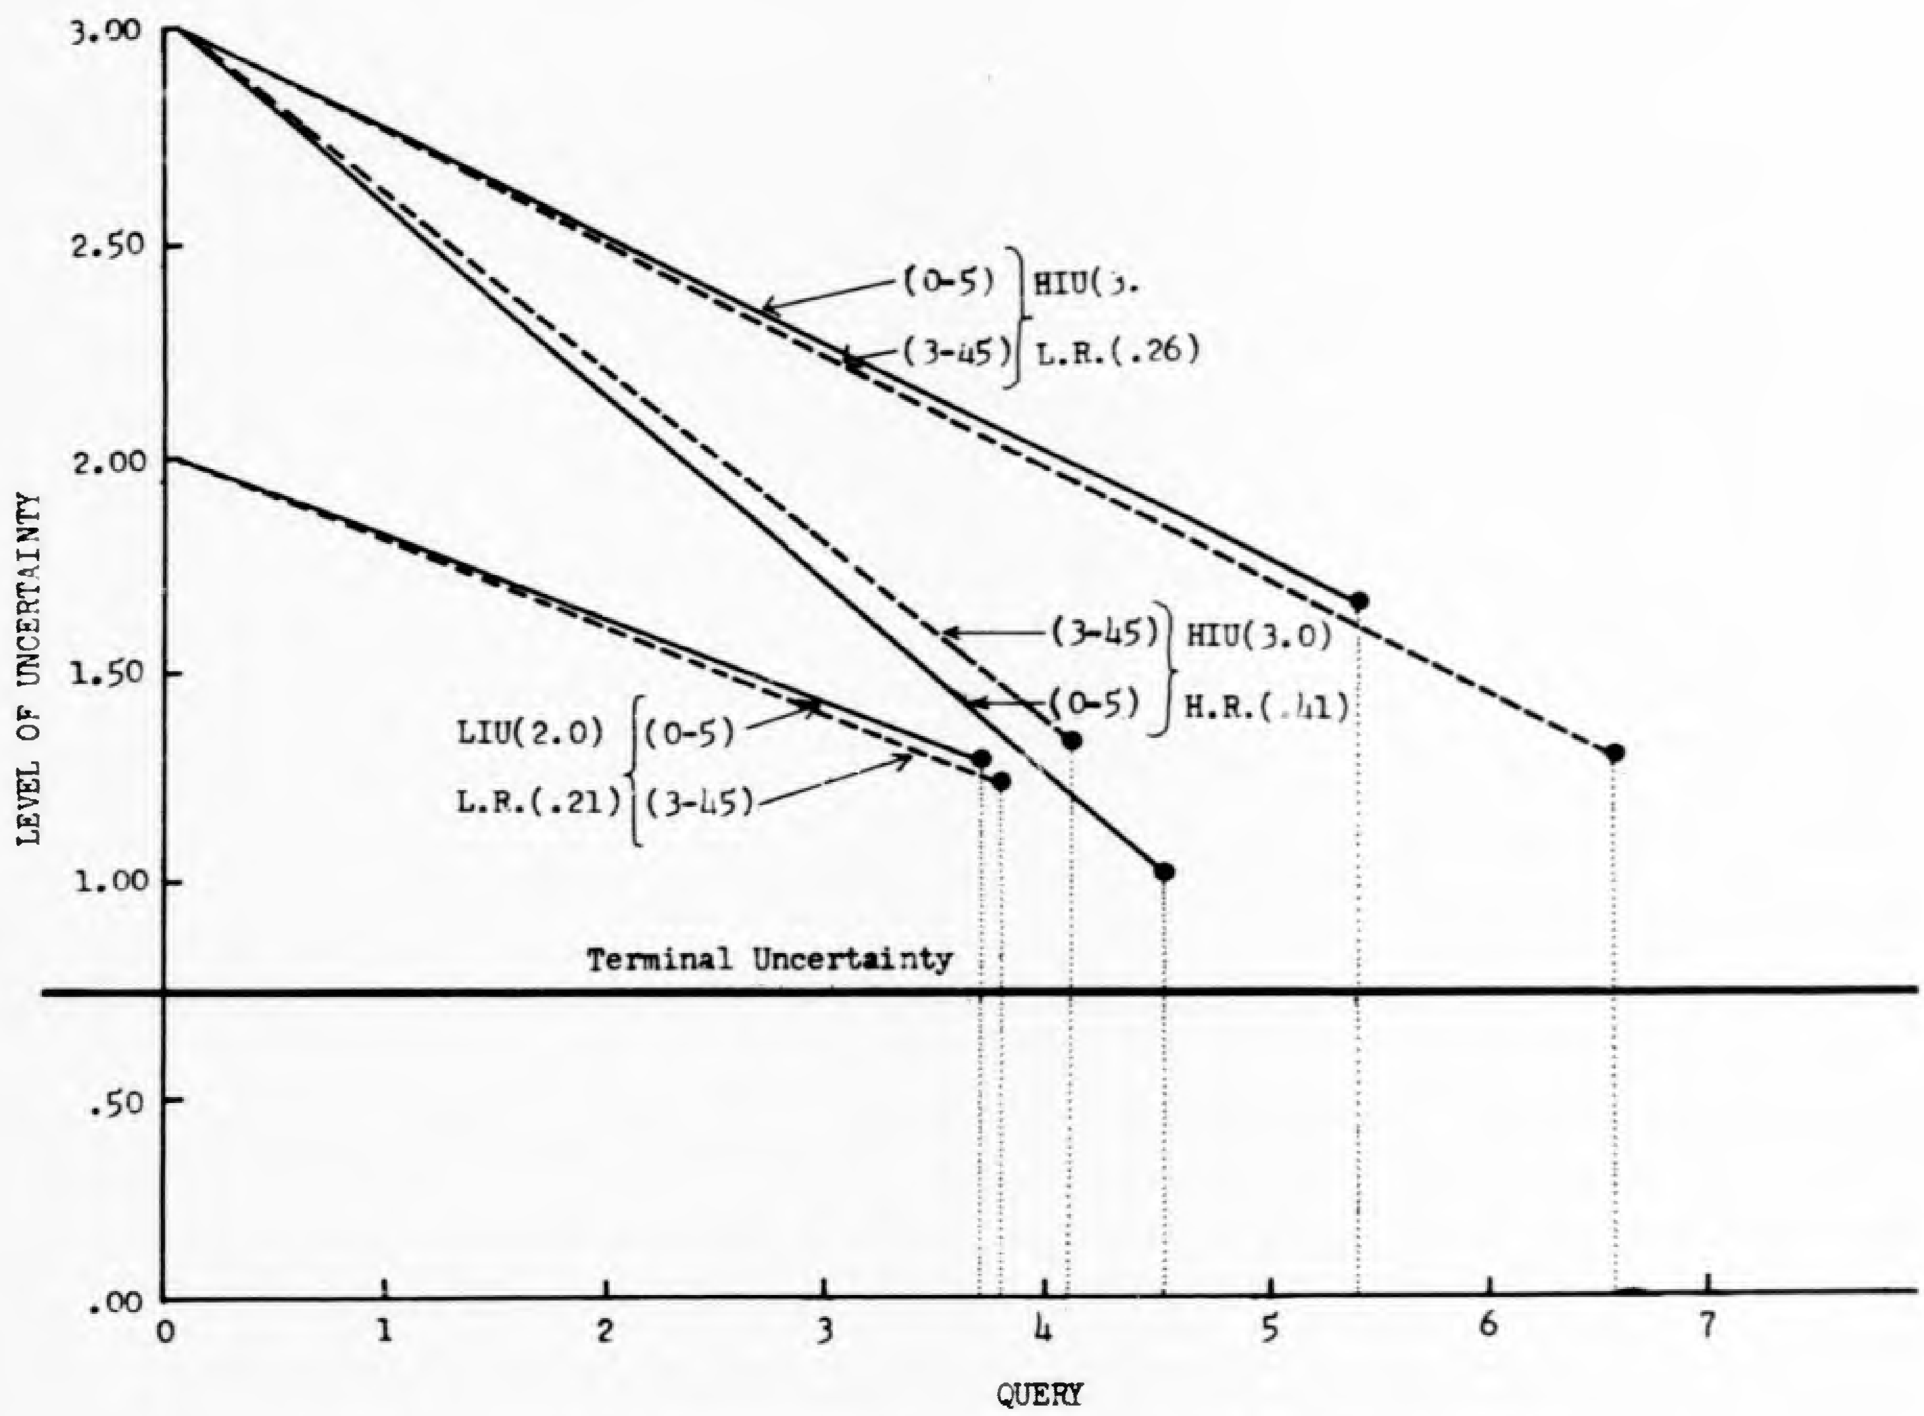

Figure 3. Level of Uncertainty as a Function of the Number of Queries

These obtained differences in search as a function of uncertainty were encouraging considering the crudity of the scaling techniques, but left many questions to be answered concerning the efficacy of our uncertainty measures. Two more studies were instigated to explore the relationships between stimulus uncertainty, phenomenal uncertainty and other aspects of predecision behavior such as information search and latency.

The first of these two studies was concerned primarily with the further development of our uncertainty scaling techniques and with ascertaining the extent to which our operational definition of uncertainty paralleled the subjective uncertainty of the decision maker.

The uncertainty measure was arrived at by presenting a problem to a group of subjects and allowing them to generate their own responses to the problem. The response distribution thus generated was converted into a probability distribution and substituted into the formula $H = -\Sigma\, p_i \log_2 p_i$ to yield the response uncertainty of the problem. Two assumptions which underlie this definition of uncertainty were tested in this study: (1) subjective uncertainty increases as a function of equiprobability of alternatives, and the number of alternatives; and (2) the distribution of response alternatives generated by a sample of culturally similar individuals is representative of the response hierarchy of any one individual within that population. Even if the second assumption is correct, there are two other related issues which should be raised about the efficacy of this measure: (1) are some subjects capable of entertaining more response laternatives than others, leading to the existence of individual differences in the degree of subjective uncertainty that may be elicited by a given problem; and (2) does an individual's response hierarchy actually resemble the group generated data, i.e., is the respondant really considering more than two or three words when response uncertainty is present? The information theory measure of uncertainty gives added weight to additional alternatives. For example, the uncertainty of a 51 alternative problem in which one alternative has a probability of .50 and the other 50 have probability .01 is 3.00 bits; the uncertainty of a two alternative decision where both have probability .50 is 1.00 bit. Intuitively, such a measure would not seem to reflect subjective uncertainty; responses with a .01 probability would probably not be considered unless a large reward accompanied their occurrence, as in a raffle.

The uncertainty scaling was carried out as follows: Words were chosen at random from the Lorge-Thorndike list of 500 most frequently used words. These words were photographed and slides of each were projected tachistoscopically on a screen. Six subjects attempted to identify the words and attached a confidence rating to each identification. The frequency distribution of subjects' responses was used to compute uncertainty.

We obtained a crude check on the assumptions underlying our measure of uncertainty, i.e., that subjective uncertainty increases with the degree of equiprobability of alternatives and with the number of alternatives, and that response distributions generated by a sample of subjects are representative of the individual response hierarchies of members of that population. This was accomplished by correlating uncertainty with the confidence scores obtained from subjects when deciding on stimulus identity. A correlation of +.85 was obtained between confidence ratings and our uncertainty measure. These data were also used to answer the question of whether extremely low-probability alternatives, which inflate the uncertainty measure, similarly serve to increase subjective uncertainty. To answer this question, uncertainty was recomputed after responses of extremely low probability had been excluded from the analysis. This measure of uncertainty did not correlate as highly with subjects confidence ratings (r = +.75). Thus, the measure of uncertainty which utilizes all of the group-generated data including the extremely low probability alternatives appeared to be the best descriptor of uncertainty.

Three more groups of 10 subjects each were then run. The 25 stimulus words were again shown tachistoscopically to these three groups. Experimental conditions were as follows:

Group 1 was sequentially presented with each stimulus problem and all of the response alternatives that had been generated for that problem in the previous study. Subjects were asked to determine which of these was the correct alternative.

Group 2 was given the stimulus problems and with each problem was given only two of the alternatives generated in the previous study. Subjects were asked to decide between the two alternatives which was the tachistoscopically presented word.

Group 3 identified the stimulus words under the same conditions as had obtained for subjects in the previous study, i.e., they had no alternatives from which to choose.

Subjects were instructed in the use of a magnitude estimation technique which they used to attach confidence ratings to each of their responses.

The mean uncertainty, mean confidence ratings, correlation coefficients between confidence and uncertainty and mean square of confidence ratings within words across subjects for the 25 words under each experimental condition are shown in table 2.

TABLE 2

RESPONSE UNCERTAINTY, CONFIDENCE RATINGS, CORRELATION COEFFICIENTS BETWEEN CONFIDENCE AND UNCERTAINTY, AND VARIANCE OF CONFIDENCE RATINGS ACROSS SUBJECTS WITHIN WORDS

	Entropy	Average Confidence	Correlation	Mean Square of Within Words Variance
Two alternative	.37	4.9	+.71	1.013
Multiple alternative	.78	6.5	+.75	7.981
No alternatives	1.44	5.9	+.71	17.237

These data provide more explicit answers to the questions raised above concerning the assumptions underlying our uncertainty measure.

As we expected, uncertainty increased with number of alternatives and was greatest in the no alternative condition. Given some advance notion of the identity of the word from the alternatives provided, subjects seemed to be more capable of judging the identity of the stimulus word; certain alternatives were never chosen in the two alternative and no alternative conditions although those alternatives had previously been generated by subjects who were responding to stimulus words without alternatives from which to choose. Three of the two alternative words which appeared on the two alternative group's test also appeared with only two alternatives on the multiple alternative group's test (since only two alternatives had been generated in the pilot study). Of these three words, two produced identical response frequency distributions for both groups of subjects, and one had one difference in 40 choices between the two groups, giving added support for the notion that individuals from the same culture have highly similar response hierarchies.

An unexpected finding, however, was that while confidence decreased from two alternative to multiple alternative conditions, the no alternative condition was responded to with an amount of confidence intermediate to the two alternative and multiple alternative conditions. Possibly, while a no alternative situation may elicit greater uncertainty when the subject is completely ignorant of the identity of the stimulus, after he has tentatively identified the word, he may then express a great deal of confidence in his decision in the absence of other reasonable alternatives to be considered. Personality factors may also operate to bring about a quick reduction of uncertainty in some subjects, especially under a no alternative condition. Subjects showed consistent individual differences in the magnitude of their confidence ratings, i.e., some subjects were consistently more confident than others. These differences (as shown by the mean square of confidence rating within words) were greatest in the multiple alternative and no alternative conditions, suggesting that under conditions in which only two alternatives are available, there may be little individual variability in perceived uncertainty, but, perhaps due to individual differences in the number of alternatives which can be entertained simultaneously, there exist, under conditions of objectively high uncertainty, great differences in subjective uncertainty.

While these data establish the psychological meaningfulness of our operational definition of uncertainty, the precise relationship between this measure and subjective uncertainty remained problematic. Furthermore, the distinction between stimulus uncertainty and response uncertainty had never been made in our research, though it seemed to merit attention. Consequently, a psychophysical experiment was performed to examine the relationship of stimulus uncertainty, and response uncertainty to subjective uncertainty (ref. 12). Most decision situations contain both stimulus and response uncertainty, i.e., the decision maker must both identify the nature of the situation (stimulus uncertainty) and decide what response to make to this situation (response uncertainty). These two sources of uncertainty were examined to see if they interact in the decision task to produce the overall uncertainty as experienced by the decision maker, i.e., subjective uncertainty. Subjective uncertainty was considered an intervening state mediating information acquisition and processing in the decision task and it was proposed that subjective uncertainty would increase with both stimulus and response uncertainty and that the amount of information acquisition in the decision task would be a function of subjective uncertainty.

In view of Miller's hypothesis (ref. 28) that there is a definite limitation on the number of unidimensional events that can be remembered and processed, and Broadbent's first principle (ref. 7) that ". . . .a nervous system acts to some extent as a single communication channel, so that it is meaningful to regard it as having a limited capacity," a ceiling effect was predicted for subjective uncertainty; all measures of uncertainty and measures dependent on uncertainty were hypothesized to reach an asymptote at about 3 bits of objective uncertainty.

Each subject was presented sixteen decision problems formed by the combination of one, two, three and four bit stimulus uncertainty conditions with one, two, three and four bit response uncertainty conditions. The stimuli were placed on slides so that the subject could reduce stimulus and response uncertainty by focusing a slide projector. This procedure permitted the subject to obtain successive one bit reductions in stimulus and response uncertainty so that he could eventually identify the stimulus pattern by calling out two, four, eight, or sixteen response labels assigned to the pattern in the projector. At the presentation of each problem, and after each information search response, subjects reported their feeling of uncertainty by the use of direct magnitude estimation. The experimenter recorded the number of information search responses made by each subject, their reports of subjective uncertainty, and the time of processing each one bit information item.

The results of the study indicated that subjective uncertainty increases to the three bit level of stimulus and response uncertainty, but further increases in subjective uncertainty; these results are shown in figure 4. This finding suggests that there may be a psychological correlate to Miller's proposed limit of the organism's ability to process information.

Figure 4 also shows that within the zero to three bit range of uncertainty, the rate of increase for subjective uncertainty was greater than the rate expected if increases in subjective uncertainty were proportionate to increases in objective uncertainty. This holds when subjective uncertainty estimates at the initial presentation of the problem are used. If, however, subjective uncertainty is scaled using the estimates of uncertainty obtained after one bit of information has been processed, subjective uncertainty increases proportionately with the objective uncertainty to the three bit level, as shown in figure 5.

This inflation of subjective uncertainty before information is acquired suggests a difference between making a decision without information and making a decision of the same uncertainty with information. Phenomenologically, one might speak of a change of state from that of uncertainty to some degree of certainty. One subject described this phenomenon as follows when his over-estimations were pointed out to him at the end of the experiment:

> "It's like on a multiple choice test, I think I would be more uncertain with four choices on a problem than I would be when there were eight and I knew four of them were not the right answer."

When stimulus uncertainty is separated from response uncertainty, the effects of response uncertainty on subjective uncertainty is seen to be predominantly between the zero and one bit level, i.e., between having to make a choice and not having to make a choice; this is shown in figure 6. Subjective uncertainty increases with response uncertainty between this zero and one bit level for the most part, but the rate of increase is a function of the amount of stimulus

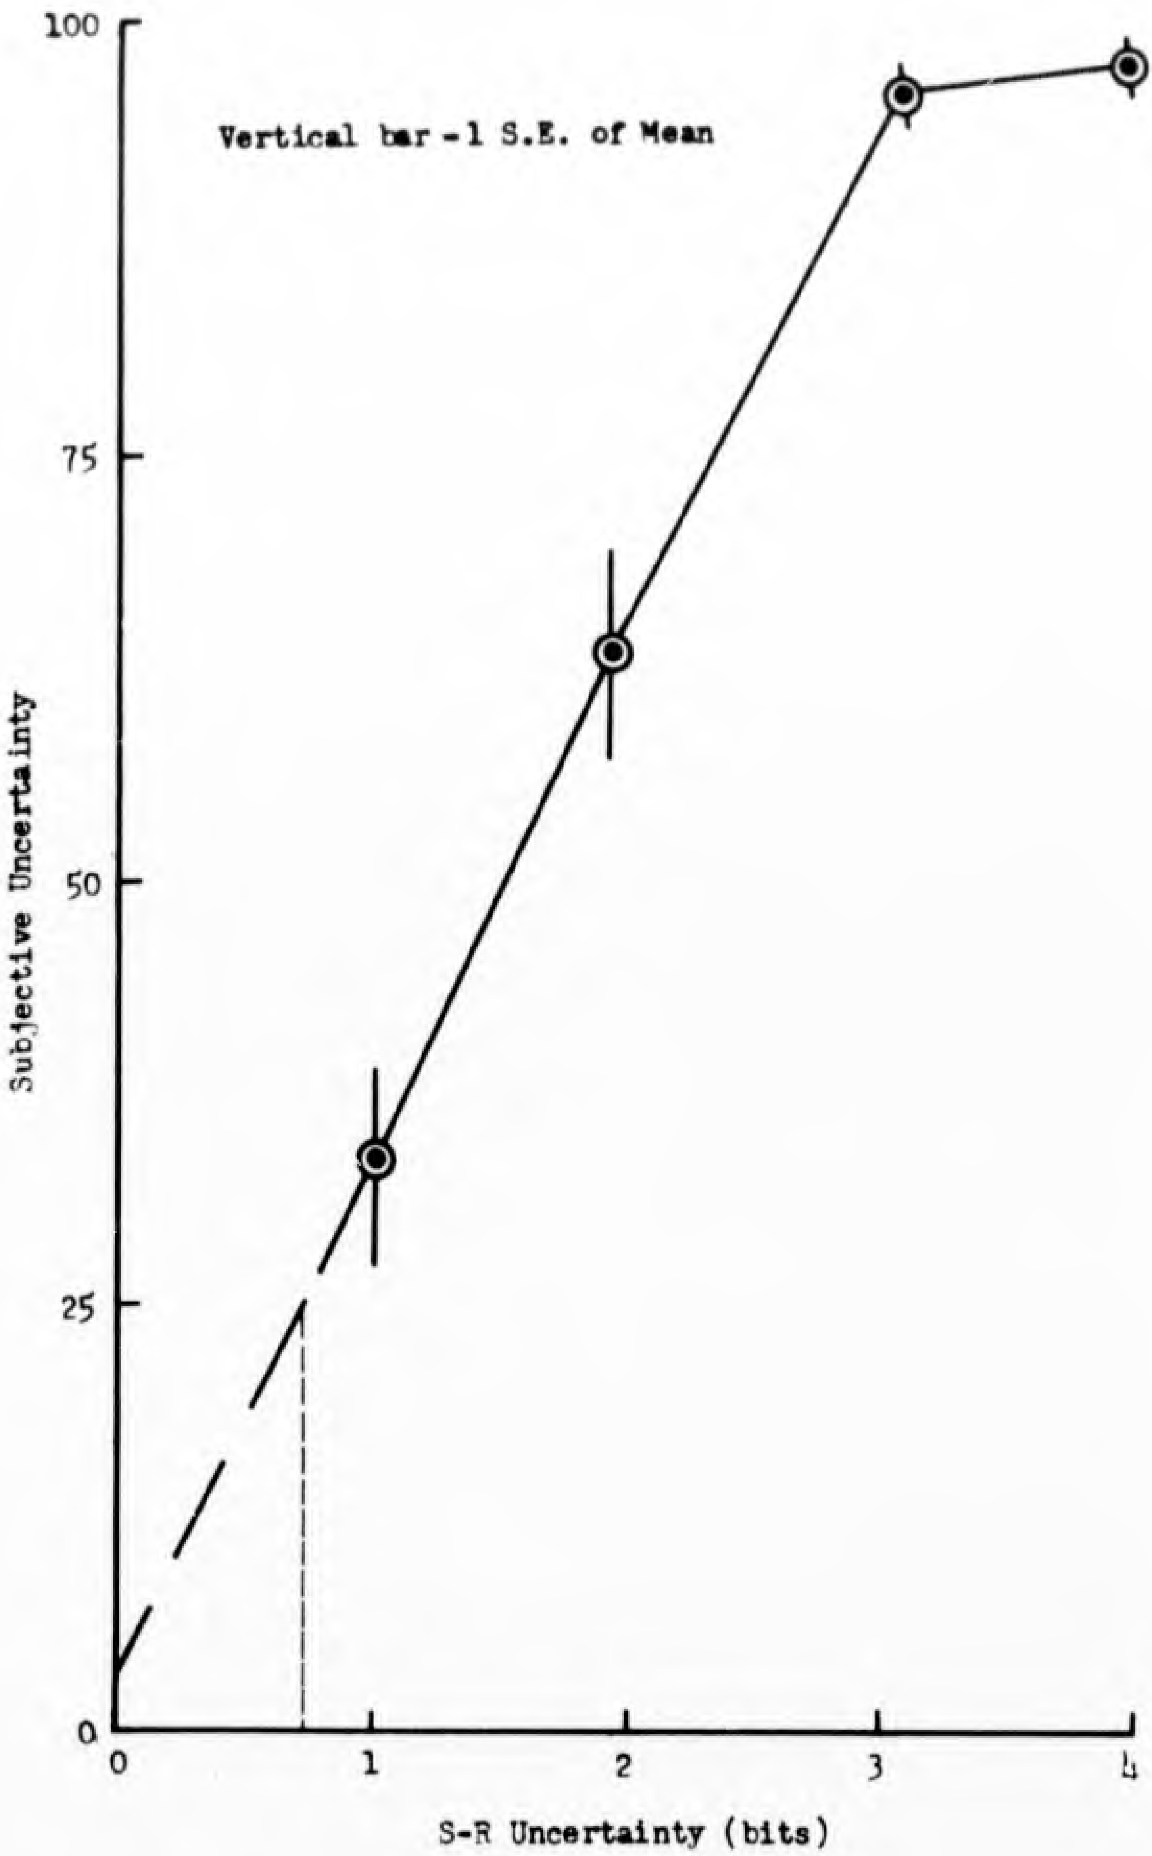

Figure 4. Mean Subjective Uncertainty across S-R Uncertainty

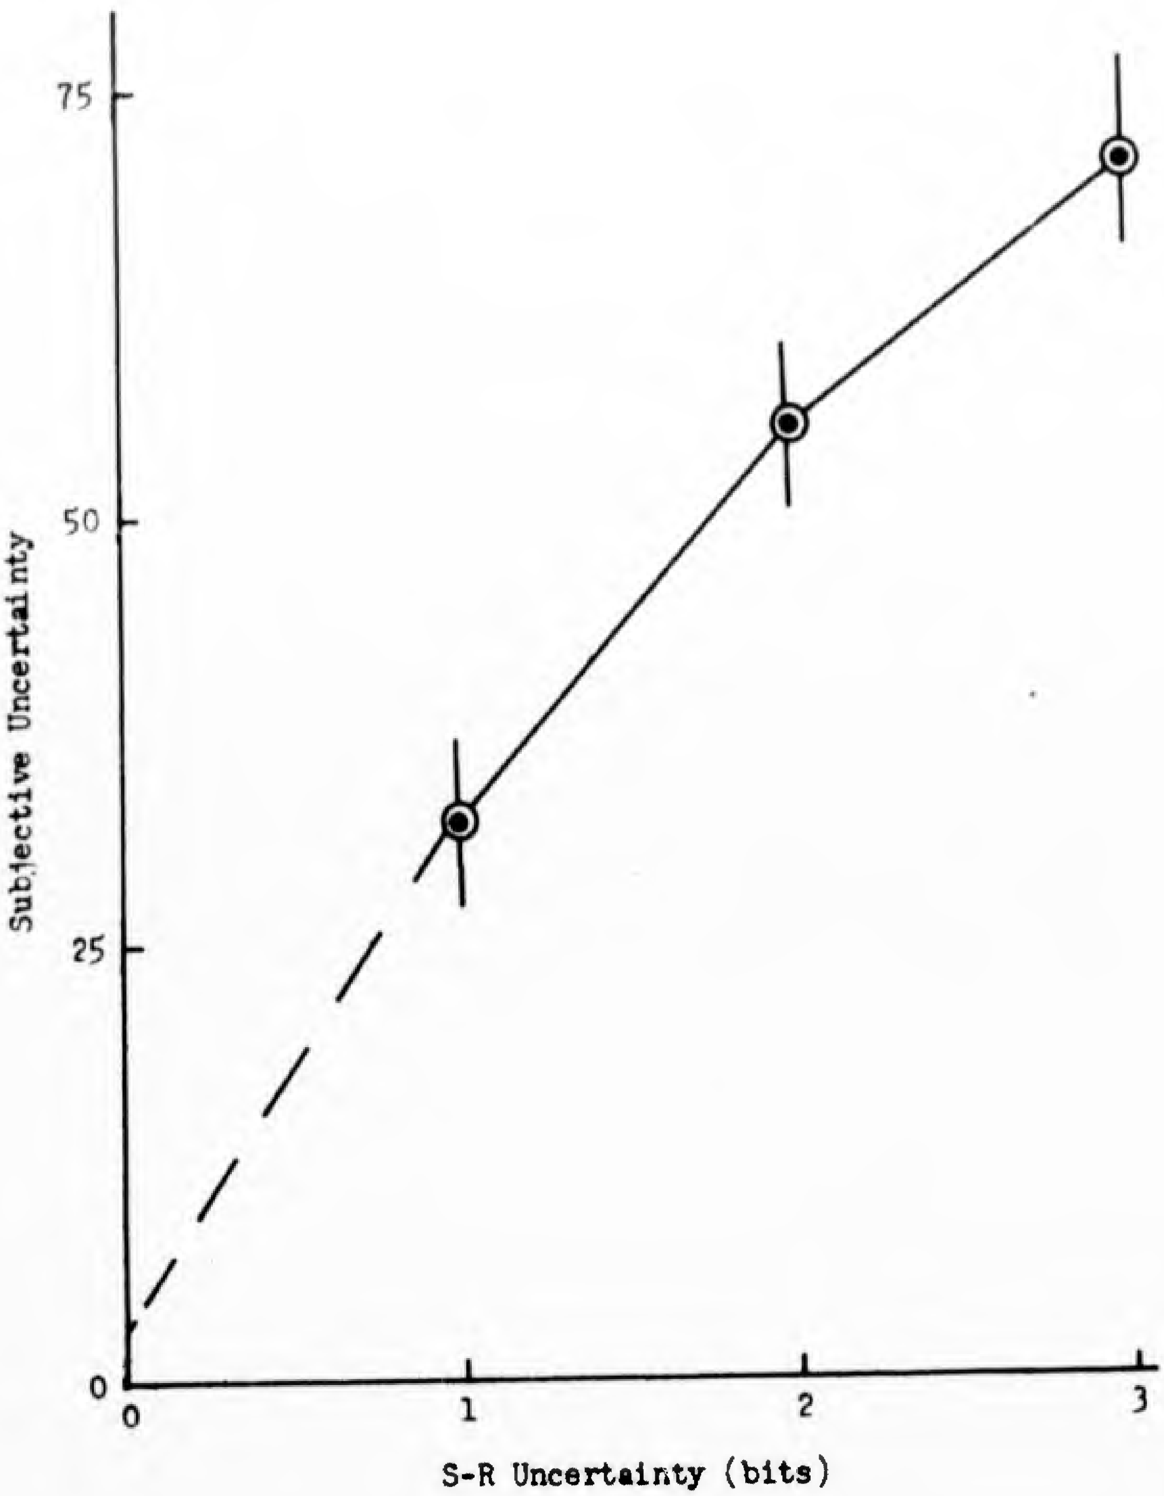

Figure 5. Subjective Uncertainty across S-R Uncertainty for Second Estimates

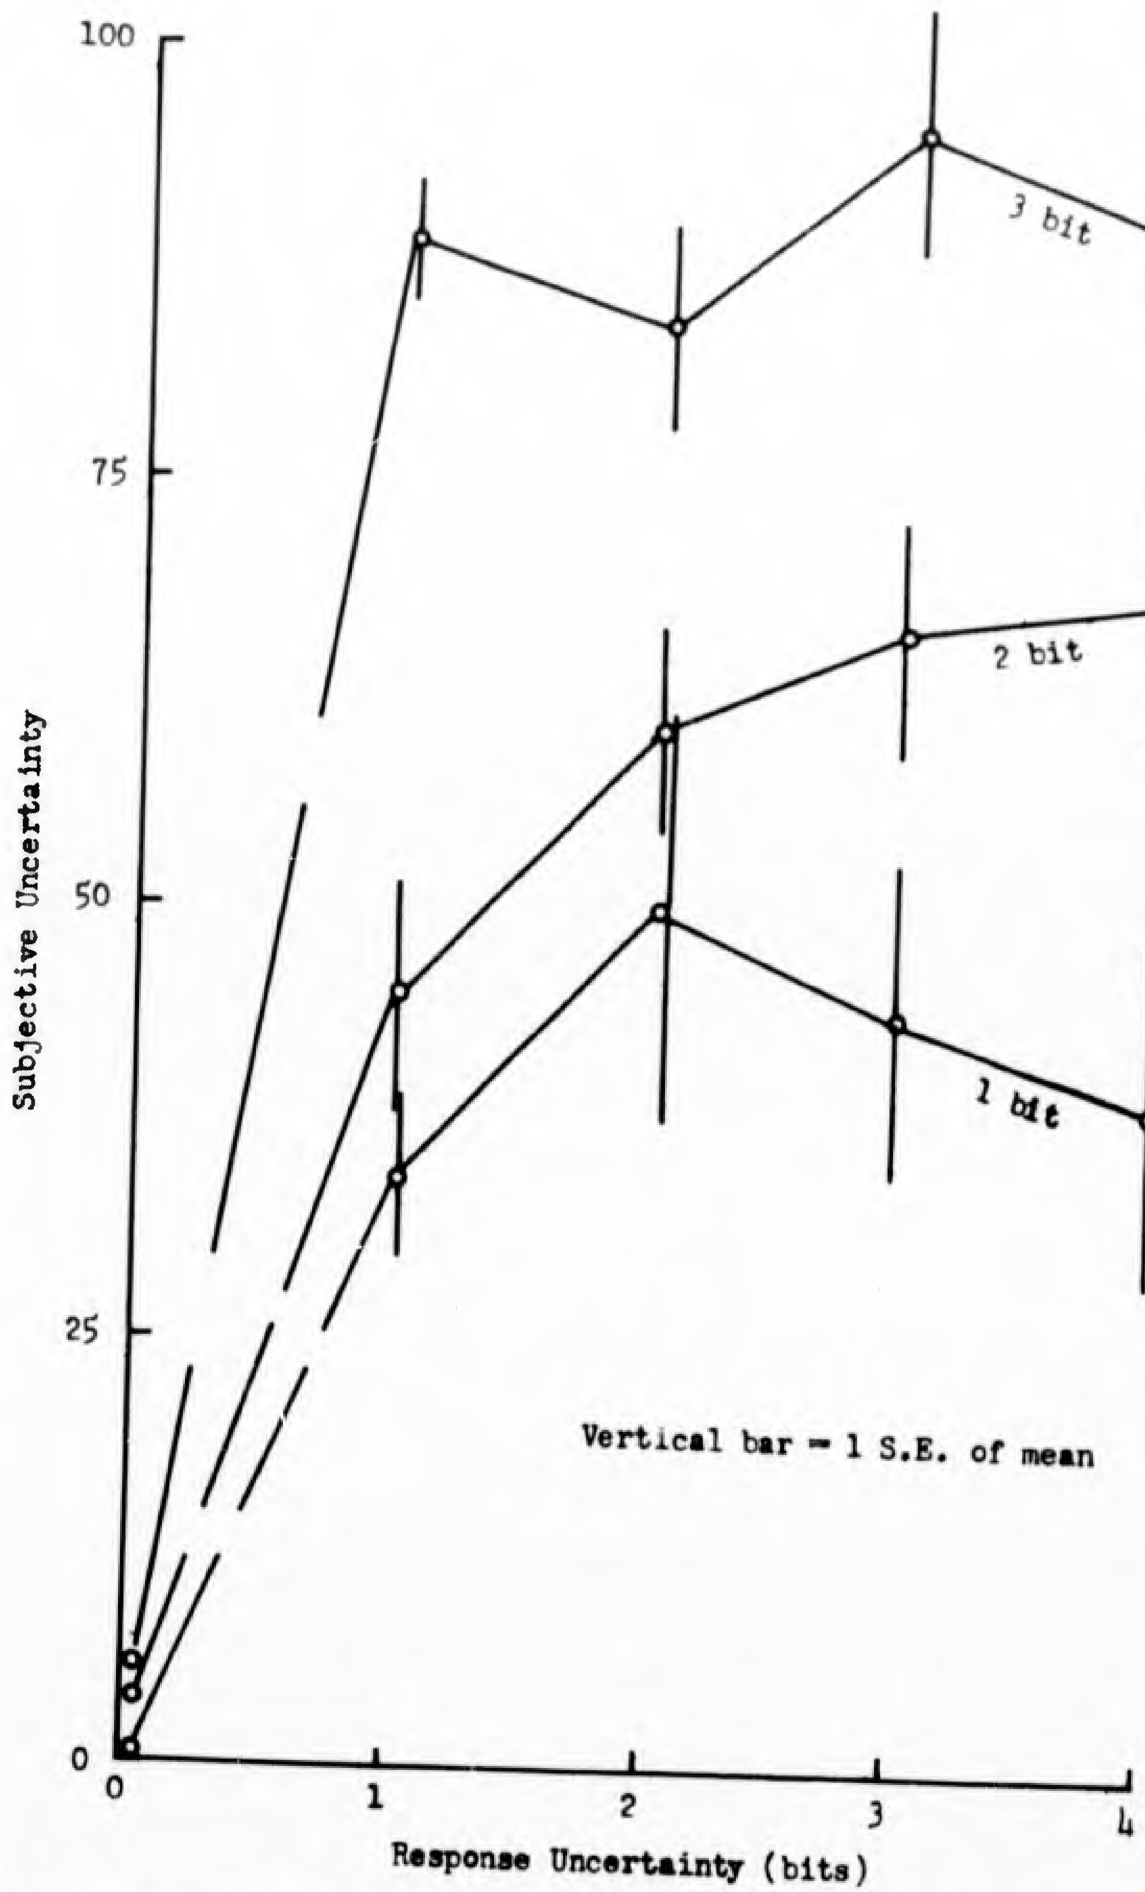

Figure 6. *Mean Subjective Uncertainty across Response Uncertainty for Four Stimulus Uncertainty Conditions*

uncertainty. When stimulus uncertainty is low, subjective uncertainty seems to increase gradually with increasing response uncertainty, but as stimulus uncertainty increases, this total increase for subjective uncertainty across response uncertainty occurs between the zero and one bit level.

Although the primary purpose of this study was to investigate the effects of different combinations of stimulus and response uncertainty on the uncertainty of the decision maker, some effort was made to relate both objective uncertainty and subjective uncertainty to information search behavior. It was found that when presented a decision problem, the decision maker proceeds immediately to engage in information search and that the strength of this tendency is a direct function of the uncertainty of the problem. The greater the problem uncertainty, the shorter the problem time and the faster the decision maker enters into the information acquisition phase of the problem. Once in the problem, the decision maker takes more information when uncertainty is high than when it is low.

Although uncertainty motivates this acquisition of information, the actual processing of information is more dependent on the structure of the decision task than on these motivational variables. Information is processed at a constant rate independent of the uncertainty level of the task and uncertainty is reduced at this rate until the 1 bit response uncertainty level is reached. Here, search depends on the amount of remaining stimulus uncertainty. If stimulus uncertainty is high, the decision maker will continue to search, but if little stimulus uncertainty remains, the decision maker prefers to make a decision with some response uncertainty remaining. There is evidence also that information search is more closely related to subjective uncertainty than to objective uncertainty, since a subject's estimate of his uncertainty correlates with the likelihood of his searching to certainty: the higher he reports his uncertainty, the more likely he is to search.

The failure of subjects to search to certainty in most of the conditions of the design is consistent with the earlier results on "commitment thresholds". In the present case, the subject seems to avoid complete certainty even when accuracy of identification has been emphasized by the instructions. This preferred level of uncertainty corresponds to the average subjective uncertainty of 26.5 (magnitude estimate on scale from zero to 100) judging from the estimates given by those persons who failed to reach at the 1 bit level. This is the uncertainty level that seems to be acceptable for making a decision, and its corresponding objective uncertainty value is .75 bits.

Although information search occurs in differential amounts as a function of uncertainty, the amount of time required of the decision maker to process a given amount of information is constant across different levels of uncertainty. In the present experiment, information was always given in 1 bit units which allowed the subject to eliminate half of the possible stimuli and responses at each search response. In some cases, reducing uncertainty by 1 bit required the elimination of eight individual stimuli and responses, but in other cases, it required the elimination of only one stimulus or response. Since the rate of processing was equal for these two cases, it would seem that the decision maker in this task was making a binary decision along one dimension of the stimulus patterns at a time.

CONFLICT, AROUSAL AND INFORMATION SEARCH

Our results clearly support the role of uncertainty in the instigation of information search and suggest that the rate of uncertainty reduction may also affect the number of search responses. We required some broad theoretical position, however, within which to interpret our results, and which might suggest the nature of the psychological mechanisms which mediate problem uncertainty and predecision behavior.

Berlyne (ref. 6) theorized that the uncertainty of a choice situation leads to conflict (competition among response tendencies) which, in turn, results in physiological arousal. The arousal motivates behavior instrumental in reduction of uncertainty and hence of arousal, e.g., perceptual and epistemic curiosity behavior. This is a curvilinear hypothesis; i.e., the effectiveness of the arousal reducing behavior is an inverted *U* function over increasing arousal.

Borrowing from and modifying Berlyne's ideas, we hypothesized that the physiological arousal engendered by the stimulus uncertainty of decision problems would motivate information search and processing responses instrumental to reducing the arousal and uncertainty. Since information-seeking presumably has a long and successful history in reducing uncertainty, the probability of information-seeking responses should increase with increases in conflict.

According to Berlyne's theory, conflict is assumed to increase with (1) the number of competing response tendencies elicited by the choice, (2) their nearness to equality of strength, and (3) their total absolute strength.

Accepting probability of occurrence of a response as an indicater of response strength, the first two variables are encompassed in the information theory expression for uncertainty, i.e., $H = -\Sigma\, p_i \log_2 p_i$, where p_i is the probability of occurrence of the ith response. The third variable, the sum of the absolute strengths of the competing responses (ΣE) is assumed by Berlyne to operate as a scaling factor so that degree of conflict (C) is equal to $\Sigma\, E \times H$. E may also be thought of as reflecting the importance of the choice, on the assumption that the more important the outcome of a choice, the stronger the elicited response tendencies. Degree of conflict then becomes a function of H and I, i.e., $C = H \times I$, where I may be defined as the average value of the outcomes of a choice.

To establish the validity of the position that arousal increases in proportion to uncertainty and serves as a motivator of information processing, two conditions should be met:

(1) Arousal should be demonstrated to be a concomitant of uncertainty in a decision situation and to co-vary with degree of uncertainty.

(2) Arousal should be independently manipulated to determine whether it is, in fact, an essential mediating variable, and not simply a concomitant of subjective uncertainty which serves no motivating function. If arousal does serve as a motivator, then different levels of arousal should produce differences in information acquisition and processing activity.

A study was performed (ref. 13) in which an effort was made to provide data on both these points by obtaining a continuous measure of "arousal" (as measured by GSR) in decision tasks of differing uncertainty, and by inducing different levels of arousal prior to placing the subject in these decision tasks.

Different levels of arousal were induced for each of three groups of subjects as follows:

In the arousal control condition, subjects were simply brought into the laboratory and GSR electrodes were attached to their right hand. They were told that the experimenter was calibrating it for use in a later study and were given 5 minutes to get accustomed to the situation before the decision making task was presented.

The medium and high prior arousal subjects were treated in a similar manner but, in addition, were given a task in pitch discrimination. These subjects were presented with two brief tones on each of 40 trials, and required to judge which of the two tones was higher in pitch. The required discriminations were impossible since the two tones on any one trial were nearly identical in pitch. This task lasted ten minutes for both groups. The medium prior arousal subjects received no feedback during the task, while the high arousal subjects were told they were wrong on a random 50% of the trials and were shocked on each "wrong" trial.

After this induction, subjects were seated before the decision making apparatus and instructed in its operation. They were told that a written problem and some choice alternatives would be presented on the screen before them. They were to read the problem and decide which of the alternatives was correct. If they wished, however, they could deposit a chip worth 1 cent in a slot, at which point an information item would appear on the screen. This could be done five times for each problem. At any point in this process, however, a response button appropriate to the choice of an alternative could be depressed. This caused another decision problem to appear on the screen and if the choice had been correct, a chip worth 10 cents was deposited on the table before the subjects. Each subject received the same series of ten problems, but the number of problem alternatives depended upon the uncertainty condition to which the subject was randomly assigned. A low, medium, or high uncertainty problem had two, four, or eight alternatives, respectively.

To reduce the effects of individual differences in problem solving ability, none of the alternatives were necessarily correct. Instead, the probability of "correctness" was programmed across the five queries and ranged from 0 to 1. The probability of reinforcement was 0 with no queries, .20 with one query, .40 with two queries, .60 with three queries, .80 with four queries, and 1.0 with five queries. Subjects were given 50 cents in chips at the beginning of the decision problem and allowed to keep everything they won minus the amount they were staked.

The GSR electrode that was attached to the subject at the beginning of the induction session remained there during the decision task. Three dependent measures were taken:

1. Skin conductance was recorded continuously during the decision task. From this record, measures were taken at the beginning of each problem in the decision situation. The log of the difference from the initial base line of the subject to the level of conductance at the beginning of a specific problem was used.

2. The number of information items taken by the subject for each problem also gave ten measures per subject, one for each problem.

3. Time recordings yielded data on total time, problem time, average time per query, and decision time per problem.

The data indicated that different uncertainty and prior arousal conditions did not elicit differential search behavior. The negative results for information search do not permit acceptance or rejection of the hypothesis that differences in the uncertainty of a decision problem result in differences in amount of information search and processing. Large individual differences in level of information search completely masked any effects of uncertainty. The only effect discernable is an apparent breakdown of information acquisition under conditions of high response uncertainty. It should be noted, though, that the amount of information search for the low and medium set of uncertainty problems was rather high (3.0 queries per problem). This high rate of responding suggests that this level of response uncertainty is a motivating condition for information search, even though differences in uncertainty within this range (.65 to 1.48 bits) made little or no difference in information-seeking.

A significant interaction between prior arousal, uncertainty and problem showed that there was a decrease over problems that was different for prior arousal and uncertainty groups. High and medium prior arousal combined with high and medium problem uncertainty produced long decision times initially which decreased slightly with problems. Low uncertainty combined with low and medium prior arousal resulted in consistently short decision times. Decision times for the high prior arousal, high problem uncertainty condition, however, begins low but tends to increase over problems.

GSR was not significantly different for problem uncertainty conditions despite a trend toward an increased arousal level with high uncertainty. There was an interaction between prior arousal and problem, as shown in figure 7 in which the

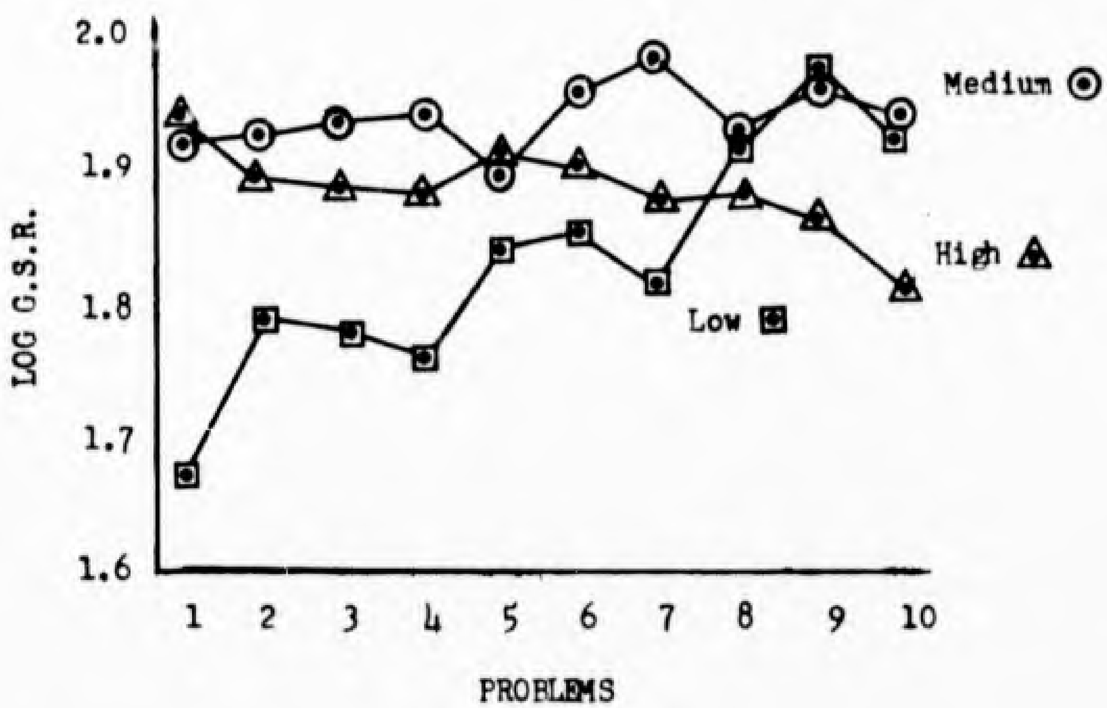

Figure 7. Mean GSR across Problems for the High, Medium, and Low Arousal Induction Conditions

arousal control group was initially lower than the medium and high arousal groups, but converges on these groups across problems, suggesting that a ceiling was quickly reached by all groups after they began the decision making task. In this respect, there is some correspondence between GSR trends over problems and the decision time trends. Decision times are high for the high and medium arousal groups which show high GSR. Both fail to increase with sequential problems. Furthermore, GSR and decision times in the low prior arousal groups begin low and the high uncertainty group increased as would be expected from the GSR trend, but low and medium uncertainty groups failed to respond to the increasing GSR level. This correspondence between GSR and decision time is close enough to support the assertion of a relationship between the two variables. The failure of the two lower uncertainty groups to increase with increases in GSR certainly suggests some limitations on this effect which must be explored before drawing more definite conclusions.

Every analysis of variance resulted in a significant between-subjects variance when compared to the intra-subject error term. In order to detect any relationship that might exist as a result of the wide range of individual reactions to the independent variables, correlations were performed between all dependent measures matched by subjects. Table 3 summarizes these intercorrelations. Similar patterns of correlations for search and time measures were obtained in the

TABLE 3

INTERCORRELATIONS OF DEPENDENT MEASURES

	Queries	Total Time	Problem Time	Av. Time Query	Decision Time
Queries		.56***	-.34**	-.15	.14
Total Time			.41***	.24*	.96***
Problem Time				.17	.33**
Av. Time/Query					.11
G.S.R.	.05				.31**

* $p < .05$

** $p < .01$

*** $p < .005$

first two studies reported in this paper (see tables 4 and 5) except that between search and decision time; the present correlatic. is positive instead of negative, but the reversal from a negative value is not significant.

TABLE 4

INTERCORRELATIONS BETWEEN NUMBER OF QUERIES AND TIME MEASURES

(Experiment I)

	S	T.T.	P.T.	Av.T/Q	D.T.
Search (S)	X	.46**	-.57**	.17	-.36*
Total Time (T.T.)		X	.40*	.73**	.34*
Problem Time (P.T.)			X	-.32	.70**
Average Time/Query (Av. T/Q)				X	-.22
Decision Time (D.T.)					X

* Significant @ .05
** Significant @ .01

TABLE 5

INTERCORRELATIONS BETWEEN NUMBER OF QUERIES AND TIME MEASURES

(Experiment II)

	S	T.T.	P.T.	Av.T/Q	D.T.
Search (S)	X	.33**	-.44**	.53**	-.52**
Total Time (T.T.)		X	.18	.27*	.11
Problem Time (P.T.)			X	-.69**	.58**
Average Time/Query (Av. T/Q)				X	-.46**
Decision Time (D.T.)					X

* Significant @ .05
** Significant @ .01

The high negative correlation between problem time and search indicated that those persons who search most spend less time reading and examining the problem material than those who search little. The probability of an information search response was inversely related to the latency of an information search response. As shown in table 3, there was no relationship between the level of arousal (GSR) of a subject and his information search, but decision time was positively related to a subject's arousal level. Apparently, if a subject was aroused by the situation, he also took a long time to make a decision.

In summary, uncertainty was shown to be an important parameter in the decision making situation, especially in determining the amount of time spent in information processing. Decision time was directly related to uncertainty as well as to amount of arousal prevalent during decision making. Prior arousal interacted with uncertainty to produce different rates of change in decision time across sequential problems. As always, individual differences in decision making were large, indicating different techniques in the method of handling problems.

This study gave no evidence in support of the main hypotheses derived from arousal theory, i.e., there was no evidence of a relationship between arousal level and uncertainty, and arousal level and "search" in a choice situation, and the independent manipulation of arousal level had no effects on search behavior. However, the weakness of the uncertainty manipulation and the GSR measurement employed, together with the few significant relationships observed between arousal, search and time measures prompted us to carry this line of research further and hopefully to resolve some of the ambiguities in previous results by (a) employing a wider range of uncertainty in the choices than had so far been used, (b) examining the effects of variation in the "importance" of the choice, and (c) obtaining continuous measures of GSR as an indicator of arousal level in order to determine whether changes in "arousal" are related to the independent variables of uncertainty and importance and to the dependent variable of search (ref. 23).

Assuming arousal to be a function of response conflict which in turn is a function of uncertainty (number of competing response alternatives) and importance (value of the choice outcomes) the following hypotheses were derived:

(1) Information search prior to decision will be a positive monotonic or curvilinear function of a degree of uncertainty and the importance of the choice. The form of the function will depend on the degree of conflict induced by the experimental manipulations.

(2) Arousal (GSR) will increase with uncertainty and degree of importance.

(3) Arousal will be positively correlated with the probability of an information search response.

The hypotheses were tested using a choice task in which uncertainty (number of alternatives) and importance (value of outcomes) were manipulated and the subject's GSR was recorded.

Subjects were presented 10 lights; on any trial 2 to 10 of these were illuminated. The subject was to predict which was the "pay-off" light. If correct, he received a pay-off. Before the choice, the subject could seek information by pressing a button eliminating half of the alternatives. Uncertainty was varied by the number of lights presented, i.e., 2, 4, 6, 8, or 10. Importance was varied by presenting, on half the trials, a signal indicating which of the two pay-off schedules was in effect. There were 60 trials per subject, six under each of five uncertainty levels and two importance levels.

Results were as follows. Search behavior increased monotonically with increasing uncertainty and was higher for the low than the high importance condition. The results on the importance variable are the converse of those predicted and not easily accounted for, even using the post hoc hypothesis that subjects were reacting to the greater cost of information under high importance, since the expected value of an information-seeking response was identical under the two conditions. The differences, though significant, are not very large, and since only two levels of importance were used, it seems best to reserve speculation until the findings are checked using additional levels of the variable and possibly other operational definitions of importance. The results for importance are not only opposite to prediction but there is no indication of an interaction between importance and uncertainty. Thus the assumptions that conflict is equal to the product of uncertainty and importance, and that search is a function of conflict, so defined, are not supported by the data.

As to the mediating role of arousal our results are ambiguous. If we accept the GSR as a reasonably sensitive indicator of transient arousal levels, then the results suggest arousal plays no important mediating role; arousal neither varies systematically with uncertainty or importance nor does it correlate with search. However, the differences in search between subjects classified as high or low in arousal on the basis of average GSR scores suggest that arousal operates as a generalized drive in that it increases the strength of responses, e.g., information-seeking, typically elicited in "uncertain" situations. Such differences in drive appear to be produced by exposure to a choice situation per se, and are not responsive to variation in the degree of uncertainty or importance of the choice.

INDIVIDUAL DIFFERENCES IN SEARCH STRATEGIES

As previously noted, large and consistent individual differences in pre-decision strategies were observed in all of the studies. Our analysis of these individual differences began with an examination of intercorrelation patterns found in our first studies. The observed differences suggested the role of certain conceptual structure variables which we have begun to systematically explore.

The pattern of individual differences found in the first two studies are revealed in the intercorrelations of the various time and search measures obtained. Tables 4 and 5 present these correlations between averages of the time and query measures across problems, by subjects, for the first and second studies respectively. Although there are differences in the pattern of obtained correlations, the similarities permit some license in generalizing across the experiments. The positive correlations between total time and the other indices are to be expected since total time is the sum of problem time, average time per query times the number of queries, and decision time. The significance of the correlations between average time per query, problem time, decision time, and number of queries is more suggestive of the nature of actual strategy differences: subjects who are slow and deliberate over queries as compared to those more rapid information gatherers make more queries, and tended to be faster in reading the problem and faster in making decisions.

The strategy differences suggested by these results are in accord with the experimenter's observations: some subjects, after reading the problem, seemed first to decide how many queries they should make. They then made this number of queries in rapid succession (short time per query), considered all of this information for a relatively long time (long decision time), and made their decision. The other strategy appeared to involve the subject's evaluation of additional information as he obtained it, i.e., those subjects whose predecision behavior was characterized by short problem time, long time per query, a large number of queries and short decision time seemed to take time to integrate all the information provided by a given query with all the information previously available and presumably on the basis of this, decided whether an additional query was necessary. Such careful consideration after each query would then require only a brief decision time after the subject had finished querying.

The differences in time apportionment patterns suggest that the subjects who spent little time per query and had a long decision time have a rigid prior assessment of what constitutes appropriate information acquisition behavior and adhere to this irrespective of the nature of the information they are able to obtain. The other strategy (much querying, long time per query, and short decision time) suggests that the decision maker has a relatively greater openness to the requirements of the immediate situation and a willingness to modify the assessment when new information suggests change. This difference brings certain personality variables to mind, e.g., the dogmatism dimension is concerned with the maintenance of rigid perceptions and subsequent failure to modify behavior in the light of new information. A similar variable, complexity of conceptual structure, is concerned primarily with the number of dimensions of a situation which persons consider when forming judgments and the flexibility and complexity with which they integrate diverse views and information.

The dogmatism variable was first explored by Pruitt (ref. 34) in a study of personality correlates of the amount of information acquisition before decision. The measure which we shall call a measure of dogmatism (but which Pruitt refers to as a measure of confidence) consisted of the subject stating whether he believed various statements were true or false, and stating how sure he was of his judgment. The statements were indeterminate ones such as "Hitler is still alive." The similarity between these and Rokeach's dogmatism scale items seems sufficiently great to permit the assumption that both tap the same personality variable.

Pruitt obtained a high negative correlation between amount of confidence or dogmatism and amount of information taken in sequential decision problems: the more dogmatic individuals consistently took less information. These results lend some support to the interpretation that subjects in the first two studies who took little information and spent very little time per query did so because they operated on the basis of dogmatic or preconceived notions of how a decision should be reached, failing to give careful consideration to the arguments for and against each alternative in terms of the currently available information and hence failing to acquire information in accordance with the actual demands of the decision problem.

Pruitt's data also indicated that anxiety level and academic achievement are related to search behavior. We replicated Pruitt's results as follows. Subjects were presented the task of guessing whether a red or green light was programmed to flash 60 percent of the time. Anxiety was assessed with a 17 item true-false scale adapted from the MMPI. Subjects' cumulative index was used as the measure of academic achievement. Our results were similar to Pruitt's: high anxiety subjects took less time in all phases of the decision task and high academic achievers took more information than low achievers. Various explanations are advanced for this relationship: high achievers may have stronger information acquisition habits; may be more cautious and less impetuous, generally; or may place a higher value on being "correct" than on maximizing profit.

Results similar to Pruitt's were obtained by Hoge (ref. 19) who studied the effects of strong prior attitudes regarding indeterminate statements upon information acquisition behavior. Sixty-four subjects were pre-tested on attitude toward Negroes. The twelve high scorers and twelve low scorers were used to test the hypothesis that strong prior attitudes predispose a person to acquire less information before making a decision on a relevant topic. Subjects were required to make decisions with the option of first acquiring additional information. These decisions were relevant to the attitudes tapped by the previously administered attitude scale. The data showed strong but nonsignificant trends between subjects who were very prejudiced toward Negores and subjects who were not. Highly prejudiced subjects took relatively less information and time to make decisions.

Hoge's data differ from Pruitt's in that his decision task tapped the same attitude area as did his attitude test, hence allow less latitude for generalization concerning the dogmatic personality. Results are consistent, however, with the conceptualization of the low searcher as one who has committed himself so rigidly to a given set of beliefs concerning the correctness of certain attitudes and strategies that he does not seek new information which could cause him to change his mind about either the appropriate decision or the strategy of reaching the decision.

While suggestive of personality variables responsible for the observed differences, Pruitt's and Hoge's studies were not performed within the context of a specific personality theory and hence do not lend themselves to very extensive theoretical analysis except on a post hoc basis. Harvey, Hunt and Schroder's theory (ref. 18) of conceptual structure was subsequently selected on the basis of its apparent relevance to our research and its power as a predictor of predecision behavior was empirically examined.

A conceptual structure is a set of cognitive mediating links which produces a relatively stable group of techniques by which one receives, processes and transmits information. The complexity of a conceptual structure is a function of (1) the number of dimensions along which stimuli are ordered, and (2) the complexity and number of different schemata with which the perceived dimensions of information are organized. Persons whose information processing is characterized by the use of few dimensions of information and simple integrating schemata are called concrete and those who typically process many dimensions of information and utilize complex integrative schemata are called abstract.

> "The greater the number of concepts and/or schemata simultaneously available for transforming a stimulus pattern, the greater the potential for these differentiated readings to be integrated; the more combinations of different integrations of the parts available, the more abstract the conceptual structure."

In theory, early training conditions establish such differences in conceptual structure. Concrete structures result when the training source: (1) determines the means and goals of the learner, (2) rewards the use of the prescribed means and attainment of prescribed goals, and (3) personally values the learner to the extent that he meets these criteria. Abstract structures are established when the training source: (1) allows the learner to select his own means and goals and encourages him to evaluate their acceptability in terms of feedback he gets from his environment, (2) rewards exploratory behavior (e.g., the learner attempts to understand causality, and to find new means and goals), and (3) values the learner as a person apart from the evaluation of his achievement.

The first attempt to relate conceptual structure to decision making behavior consisted of recalling as many of Driscoll and Lanzetta's subjects as could be located, administering the Asch Test of conceptual structure (ref. 36) to them and correlating their degree of abstractness with the previously obtained measures of predecision behavior.

Since abstract persons are characterized by more elaborate cognitive activity, it was hypothesized that abstractness would correlate positively with amount of information acquisition prior to decision, time taken to make each decision, and consistency of GSR measures. It was further hypothesized that mean GSR scores would be higher for abstract subjects than for concrete subjects. The correlations between abstractness and measures of predecision behavior are presented in table 6.

TABLE 6

CORRELATIONS BETWEEN ABSTRACTNESS AND SELECTED MEASURES OF PREDECISION BEHAVIOR

Total Decision Time	+.46*
Mean Time per Query	+.08
GSR Variance across Problems	-.20
Mean GSR	+.21
No. of Information Acquisition Responses	+.15
Variance of Search Responses across Problems	-.42

*p less than .05

Owing to the small number of subjects (12) and relatively low correlations obtained, the results were regarded as merely suggestive of areas of future investigation. The high correlation between total time and abstractness indicates that at some phase of the decision task, abstract subjects tend to deliberate longer. The negative correlation between abstractness and variability of GSR was interpreted as suggesting that concrete subjects experience differential arousal as they encounter or avoid ambiguity, and the positive correlation between mean GSR and abstractness may be explained as a reflection of the overall increased cognitive activity of abstract subjects. The low correlation between information acquisition and abstractness is difficult to interpret. Coupled with the negative correlation between search variance and abstractness, one may hypothesize that concrete subjects sometimes engaged in extensive search behavior with the expectation of finding a clear-cut answer which would reduce their uncertainty, whereas abstract subjects, sensing the lack of value of the additional information that could be purchased, may have preferred to save their money and rely on their judgment.

It was recognized that the low correlations obtained in this exploratory study could have been due, in part, to the poor obtained distribution of scores across the abstract-concrete continuum; of the 12 subjects recalled, most were concrete and no subject yielded an extremely abstract score. The data seemed sufficiently promising to prompt a more direct and powerful test of the notion that differences in conceptual structure are responsible for the large individual differences in predecision behavior under conditions of uncertainty. Accordingly, Sieber and Lanzetta (ref. 39) examined the effects of an individual's conceptual structure, problem uncertainty, and the importance of the decision on the amount of predecision information search, time per query, subjective uncertainty, and amount of new information generated. Predictions were derived from a model based on theories of arousal and conceptual structure.

Beginning with Berlyne's thesis that stimulus uncertainty elicits a host of possible responses, thereby eliciting conflict and response uncertainty, and that the amount of conflict elicited is a function of both response uncertainty and the "importance" of the problem, we further reasoned that stimulus variables alone may not account entirely for the amount of information processing; persons who perceive their environment via simple conceptual structures perceive less information, and thus may experience less uncertainty than persons who utilize complex structures. If perceived uncertainty is one of the determiners of physiological arousal, it would follow that arousal and information acquisition may also vary as a function of conceptual structure.

Whereas arousal theory analyzes cognitive phenomena in terms of the effects of perceived stimulus complexity on behavior, conceptual structure theory deals with cognitive determinants of the amount of complexity perceived. Differences in perceived arousal as a function of conceptual structure are predictable given the differences in training conditions assumed to underlie development of different conceptual structures. First, the arousal value of stimuli is partly dependent on their cue value to the perceiver. Having learned through reward of exploratory behavior to perceive more cues than concrete persons, abstract persons are usually more aroused by complex stimuli. Second, if abstract persons mediate information via a larger number of dimensions, then, as stimulus complexity increases level of arousal elicited by stimuli increases at a higher rate than for concrete persons.

Environmental and conceptual complexity are presumed to interact. Holding conceptual structure constant, there is an inverted U-shaped relationship between environmental complexity and complexity of information processing: stimulus-poor environments contain too little information to warrant complex processing, whereas very complex environments contain too much information (if all of the information is assimilated, it is too difficult to integrate via complex schemata, and if selectively assimilated, it may be impossible to integrate meaningfully).

As a consequence of this interaction, quantitatively different relationships between environmental complexity and information processing would hold for individuals who differ in complexity of conceptual structure. Figure 8 illustrates some of these differences: (1) information handling processes of concrete persons are less complex than abstract persons, holding stimulus complexity constant; (2) the differences in complexity widen as environmental complexity increases, up to some point at which persons experience an information overload and complexity is reduced. This widening of differences is due to the differential increases in both the number of variables recognized and number of ways in which variables are interrelated.

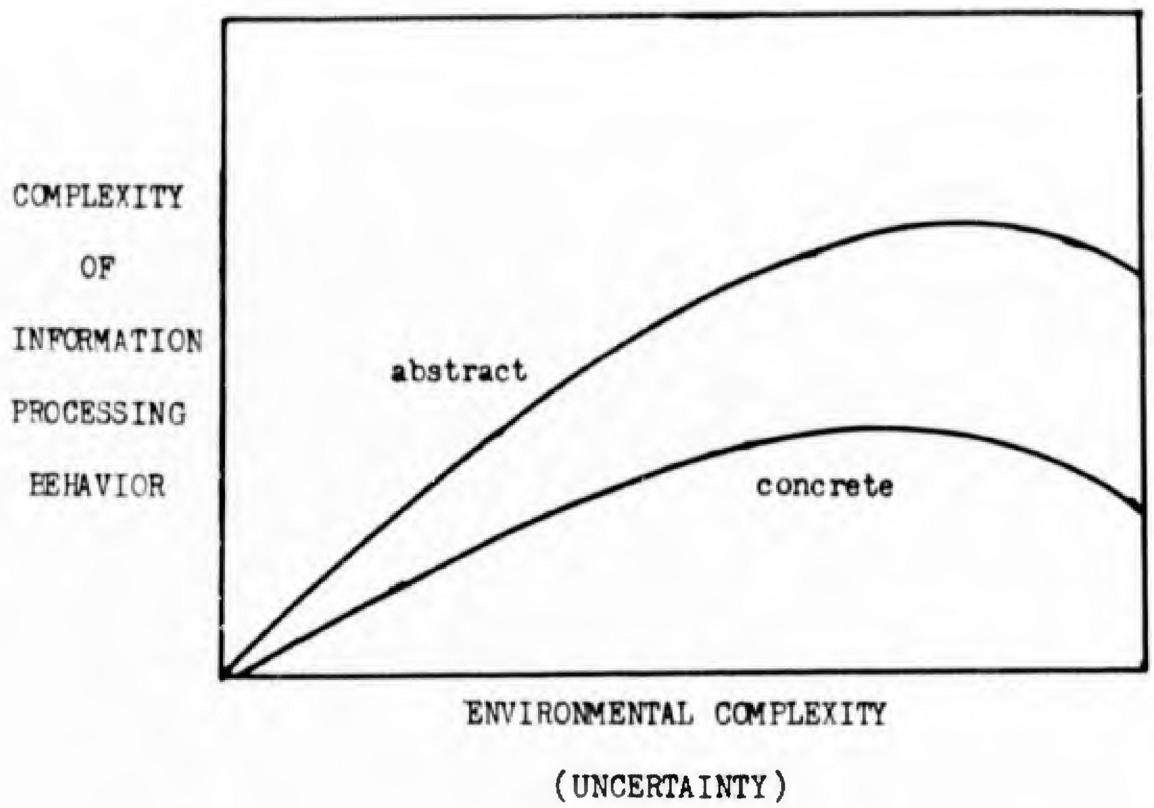

Figure 8. Complexity of Information Processing Behavior as a Function of Conceptual Structure and Uncertainty

In terms of these theoretical considerations it was predicted that the greater the uncertainty and importance of a decision the stronger the drive to acquire information in order to reduce uncertainty; the greater the abstractness of an individual's conceptual structure, the greater the complexity of his information acquisition and processing activity; and the greater the uncertainty of a choice the greater the differences between abstract and concrete individuals in information acquisition and processing activity.

The experimental procedure was as follows. The Sentence Completion Test of conceptual structure was administered to 332 persons. Of these, the 15 most abstract and the 15 most concrete were selected for this study. They were presented with 20 sequential decision problems of varying uncertainty. Problems consisted of identifying the objects depicted on slides which subjects presented to themselves tachistoscopically for successive durations of one-one-hundredth second each. They could look at each picture as often as they wished at this fixed duration of exposure before making a decision. Subjects' information search and decision responses were recorded automatically.

Problem uncertainty was manipulated by prior scaling of problem material, using the scaling technique described on page

The problems used ranged from 102 to 4.1 bits of uncertainty. Of the 20 slides used, 10 were solvable, i.e., could eventually be identified as some commonly known object. The other 10 were unsolvable, i.e., pictures in which the subject matter was either so unusual or so distorted that they were not identifiable. The unsolvable pictures were designed to present information much of which conflicted and was not amenable to a clear-cut decision. These two types of stimuli were used to explore the differences between behavior with problems having determinate and indeterminate solutions. It was reasoned that abstract persons might behave differently from concrete persons only in the types of indeterminate, unstructured environments which resemble the situations in which they developed their characteristic style of processing information, hence both solvable and unsolvable problem material was used to test this notion.

There were three levels of importance. In the low importance condition, subjects thought they were pilot subjects testing out some equipment. In the medium importance condition, subjects were told that they were participating in a decision making study designed to test problem effects. In the high importance condition, subjects were told that they were participating in a decision study designed to determine the relationship of intelligence grade point average and some personality factors to quality of decision making and were urged to take the experiment very seriously.

Exploring first those results which bear solely on the predictions derived from arousal theory, amount of information and time per query did increase monotonically with problem uncertainty as predicted. Contrary to prediction, however, these dependent variables were not monotonically related to importance: search was curvilinearly related and time per query unrelated to importance suggesting either that an extreme amount of conflict was induced in the high importance condition, resulting in breakdown of information processing, or that a norm of not searching was inadvertantly established via the high importance instructions. In any event, further exploration of importance manipulations would seem desirable before firm conclusions are drawn.

Concerning the nature of the joint effect of uncertainty and importance on decision behavior, if uncertainty and importance combine multiplicatively to produce conflict, as Berlyne suggests, the data should reveal any interaction between uncertainty and importance. The data yielded two interactions which suggest that such a multiplicative relationship may exist. In both cases, however, the interpretation of the interaction is complicated by the fact that information search was curvilinearly and not monotonically related to importance. In the uncertainty-importance interaction (figure 9) information search was facilitated

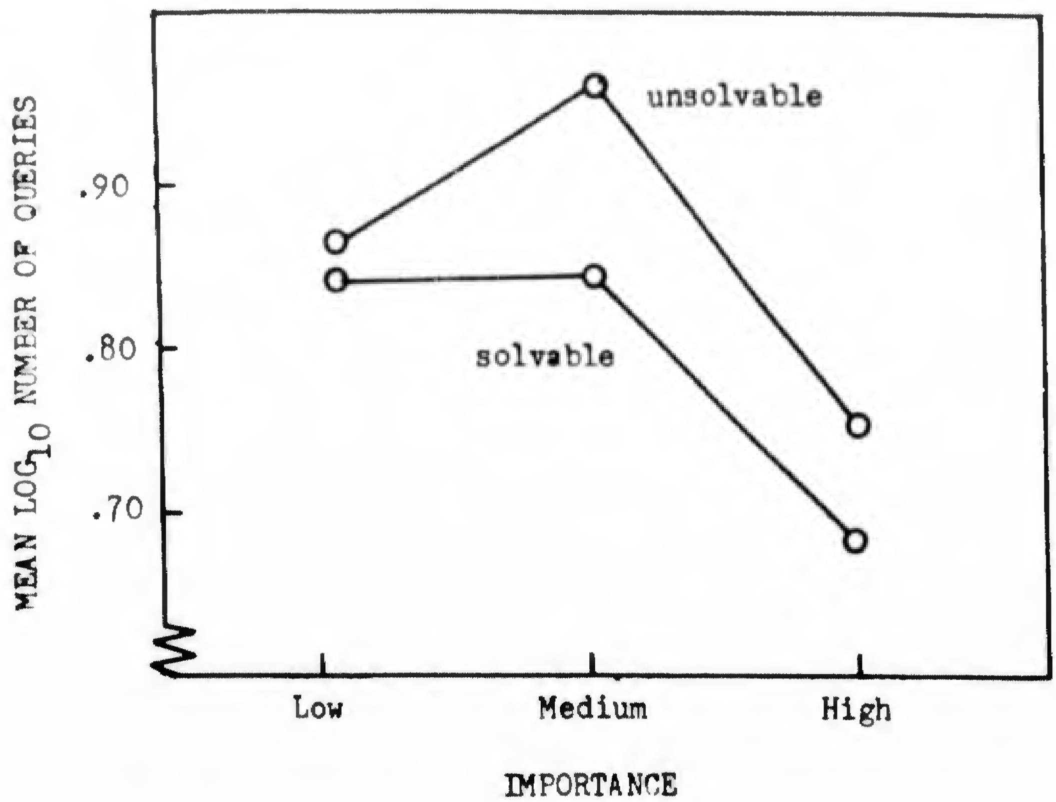

Figure 9. Mean Number of Queries Made Prior to Giving a Solution to Solvable and Unsolvable Problems under Conditions of High, Medium and Low Importance

by increase in importance under conditions of high uncertainty, but not under conditions of low uncertainty. However, the general form of the curve (i.e., the monotonic decrease in search with increase in importance under low uncertainty) cannot be accounted for by the theory. The solvability-importance interaction (figure 10) may also be interpreted within the arousal theory framework by assuming that non-solvable problems elicit greater uncertainty than solvable ones. Again, comparing the level of information search elicited by low and medium importance conditions, amount of search elicited seems to be a multiplicative function of uncertainty and importance.

Conceptual structure proved to be a strong factor in decision making behavior. The data are consistent with the theoretical conceptualization of differences between abstract and concrete decision processes, i.e., with the characterization of more abstract decision processes by a relatively greater openness to environmental ambiguity and information, leading to the examination of a larger number of choice alternatives, evaluation of these alternatives in terms of a wide variety of criteria, utilization of much information in order

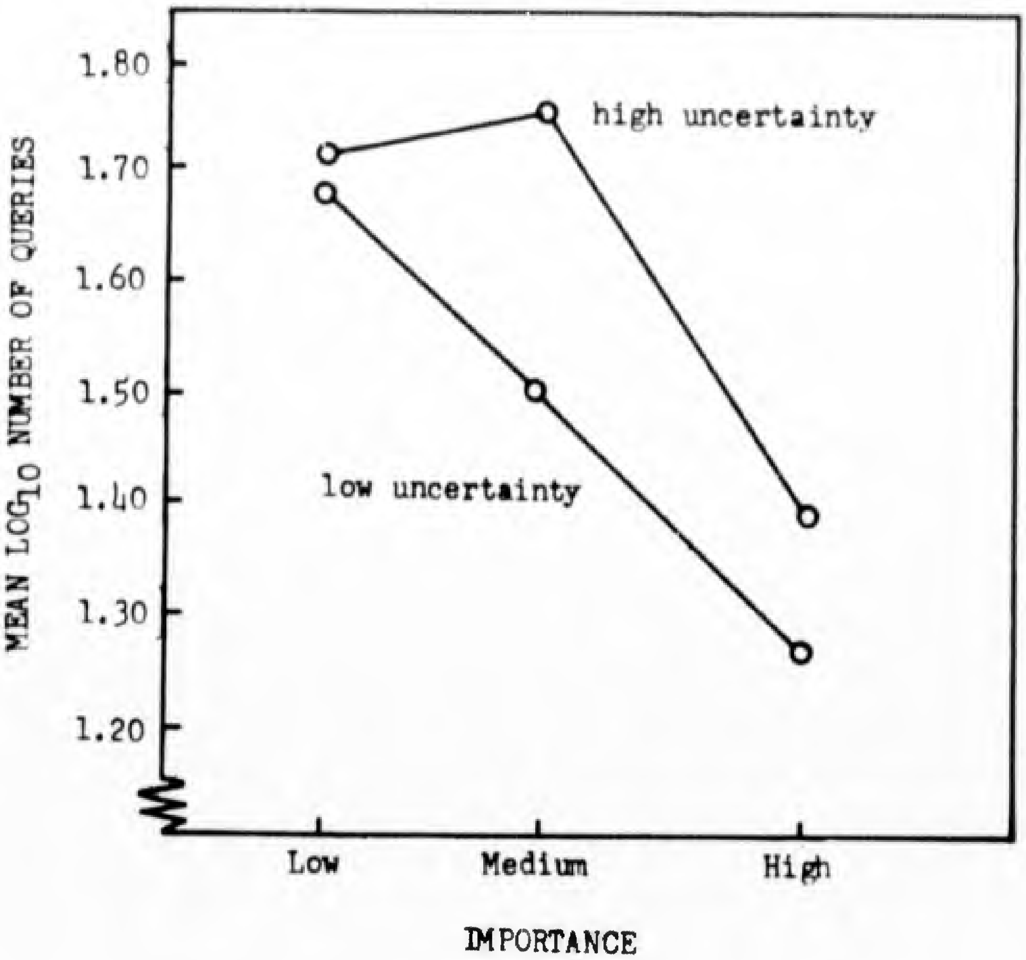

Figure 10. Mean Number of Queries for High and Low Uncertainty Problems over Three Levels of Importance

to effect the evaluation, and a tendency to remain cognizant of ambiguity and open to new information even after a decision has been reached. Persons who yielded more abstract scores on the Sentence Completion Test did, in fact, take consistently more information, suggesting that they are more aware of and inquisitive about uncertain situations. This impression was further supported by their increase in amount of search as problem uncertainty increased; such search increases did not occur for concrete persons (figure 11). These results seem consistent with the notion that abstract persons perceive more dimensions of information, form concepts based on more combinations of these dimensions, and habitually seek more information in order to have a more comprehensive map of their environment.

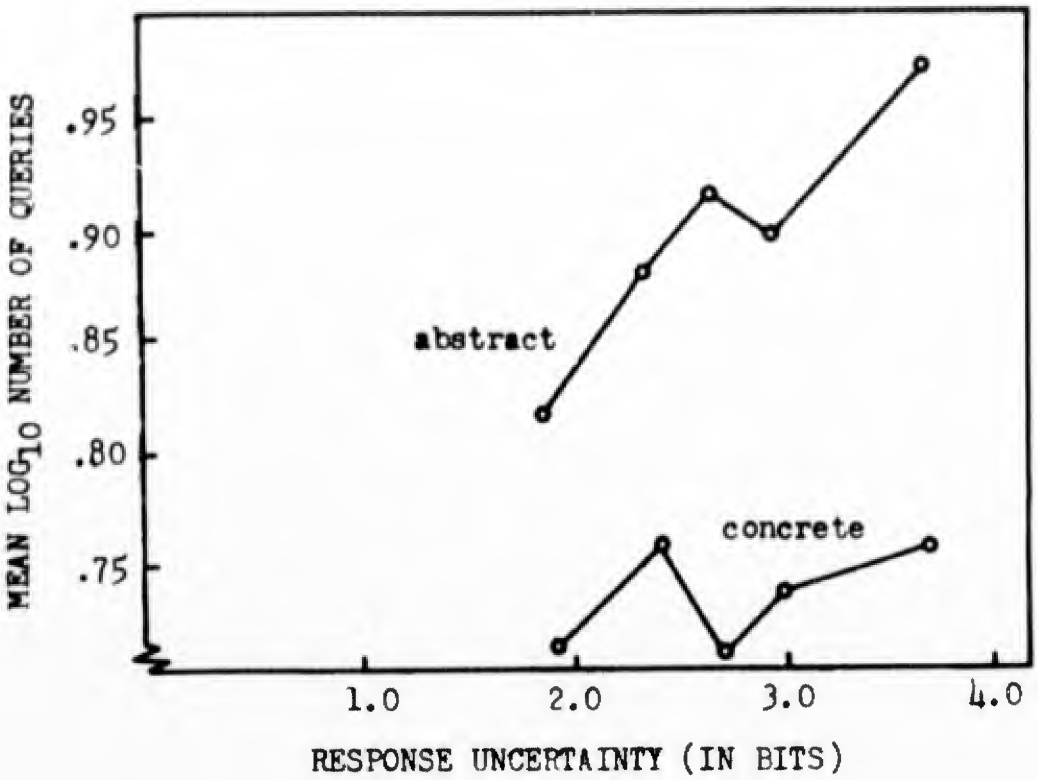

Figure 11. Mean Number of Queries of Abstract and Concrete Persons over Uncertainty

The verbal responses of subjects were also consistent with the theoretical conceptualization of abstract and concrete behavior. Abstract subjects made relatively more qualifying remarks with their decisions suggesting that they tend to entertain more complex and conflicting hypotheses about the problems they attempt to solve. In addition, abstract subjects made a greater number of qualifying statements under high importance than under medium importance, but concrete subjects did not show such a difference. Greater post-decisional uncertainty was probably experienced by subjects in the high than in the medium importance condition, since subjects took less information in the high importance condition before making a decision. However, only abstract subjects gave verbal indication of this increase in post-decisional uncertainty, suggesting that they are less prone to employ dissonance reduction techniques. This may be a consequence of their training: in theory, they have been rewarded primarily for remaining open to information in their environment, while concrete persons have learned the importance of "being right the first time".

Since abstract persons are presumed to consider more dimensions of information in their decisions, it was predicted that their decisions contain more information. Although the main effect of conceptual structure was insignificant, the interactions between uncertainty and conceptual structure gave more information about these differences. Abstract subjects in the medium importance condition consistently gave the greatest amount of information in their decisions, and abstract subjects in both importance conditions showed a tendency to increase information output as problem uncertainty increased.

The data suggest that there are certain limits on the extent to which one may generalize from theory regarding the nature of differences in information processing between abstract and concrete persons. The data on number of qualifying statements and amount of additional information given suggest that abstract persons' awareness of uncertainty is not always followed by information processing: under high importance, while abstract persons searched for less information, gave less information and showed less increase in information given as a function of uncertainty, relative to their behavior in the medium importance condition, they accompanied this with a greater number of qualifying remarks, suggesting that they were aware of the uncertainty which they did not explore and remove. So sharp a decrease in the information processing of abstract subjects under high importance is puzzling, however; perhaps a norm of not searching was inadvertently communicated in the high importance situation. Exploration of other intermediate levels of importance and different methods of manipulating importance seems warranted in order to discover precisely what types of environments produce rapid curtailment of search and information processing.

With regard to the relative effects of solvable and unsolvable problems on decision behavior, there is little to indicate that one type of problem consistently reveals more meaningful behavioral differences than another. The differences between behavior with solvable and unsolvable problems mirrored to a large extent the differences between behavior with problems of low and high uncertainty, respectively. Since abstract persons are more likely to cope with complex environments, the difference between abstract and concrete persons should be most salient with unsolvable problems. Conceptual structure did, in fact, account for more variance with the unsolvable than with the solvable problems; however, this was more than offset by the greater variance contributed by individual differences in the unsolvable problem condition. The nature of this additional source of individual variability may be elucidated by further research.

The data obtained thus far with regard to individual differences suggest a number of new research avenues. The inferred differences in cognitive dynamics used to arrive at a decision raise questions concerning the nature of prior training which resulted in these differences. Earlier, we discussed Harvey, Hunt and Schroder's theoretical specification of the types of environments which conduce to the development of abstract or concrete conceptual structure. These remain speculative in the absence of confirming research, however, and are in addition too general for theory testing purposes. Some specific questions need to be asked about the nature of abstract and concrete cognitions. For example, to what extent are they general within a given individual and to what extent are they topic specific? What is the specific content of an abstract and a concrete idea? Utilizing a matrix such as Abelson and Rosenberg's (ref. 2) within which to define the relevant dimensions of information and the relationship between dimensions, how would the matrices of abstract and concrete persons compare? One would expect a paucity of dimensions or a compartmentalization of dimensions with concrete individuals.

A lack of initial information, per se, may be responsible for concrete styles of decision making since it provides less potential for conflict; hence persons who are usually abstract may be expected to behave relatively concretely in new situations when they have little conflicting information. This suggests that the form of information input alone may dictate the style of decision making. The interaction between uncertainty and conceptual structure would support this contention, i.e., simple stimulus inputs did not conduce to abstract decision processes in either abstract or concrete persons. Presumably, habit and ability differences produce the greatest differences in predecision behavior under high uncertainty conditions. This raises the issue of differences in strategy habits.

Various techniques for attaining "cognitive balance" and hence reaching a decision in the face of conflict have been suggested by Abelson (ref. 1). These strategies vary with regard to the amount of information required and the nature of relationships which must exist between the relevant information. Through differential reinforcement, one strategy or another may become the dominant one for a given individual. The information and cognitive requirements of that strategy may then dictate the nature of predecision information acquisition behavior. More research needs to be directed at discovering the environmental factors responsible for the development of these strategy differences.

CONCLUSION

Several general issues directed the course of our investigations of information acquisition behavior. What are the conditions that elicit search? Under what conditions is a search response maintained? And, related to both, what variables determine the cessation or termination of information acquisition behavior?

One form of normative decision theory suggests a deceptively simple hypothesis to account for the instigation, maintenance, and cessation of search activity. In choice situations characterized by risk or uncertainty a decision maker will (or should) instigate a search for information when the expected gain of basing a decision on the additional information exceeds the cost of the search and/or information, and he should continue to search until this inequality reaches equality or reverses itself. When precise knowledge is available of the probabilities of outcomes associated with the choice of particular alternatives with and without additional information, a decision maker should base his search decision on this knowledge.

In sequential decision situations it is reasonable to assume a distinction between decisions leading to environmental consequences of value or profit, and decisions leading to the acquisition of information on which decisions of the former type may be based (ref. 22). Information of potential relevance to decision may be directly available to the decision maker, requiring a minimal "search" effort, or may require active acquisition behavior and be relatively costly in terms of time, money, or effort. In the latter case, the decision to seek information has implications for the profit that can ultimately be obtained since the "cost" of the information must enter into a calculation of profit. A "rational" decision maker would presumably be guided by the value of the information as well as its cost in choosing whether to invest in information acquisition or not.

But how should a value be assigned to information? Several alternatives are available; the one suggested here as a reasonable and simple assumption is that information has value to the extent that it increases the probability of choosing the alternative which yields the most favorable outcome. If the outcomes are being correct and receiving a prize or being incorrect and receiving nothing, then information has value to the extent that it increases the probability of choosing a correct alternative. In terms of this assumption the decision to acquire information may be treated in the same manner as the ultimate decision, in Coombs' terms as a "pay to play" decision situation. Some probability of achieving a payoff exists prior to obtaining any information and new information may or may not increase this probability. The expected "value of an inquiry" (ref. 27) is then assumed to be a direct function of the change in probability of choosing the correct alternative, the utility of the prize, and the utility of the cost of the information or,

$$E.V. = U(z)\,[Pa - Pb] - U(c)$$

where z = prize, in dollars

$U(z)$ = utility of the prize

Pa, Pb = probability of success after and before the new information, respectively

c = cost of information, in dollars

$U(c)$ = utility of the cost of the information

A "rational" rule would be to acquire the additional information when $U(z)\,[Pa - Pb] > U(c)$ and to make a decision in terms of the present information base when $U(z)\,[Pa - Pb] \leq U(c)$.

This result assumes that information is acquired in discrete "packets" and that a probability of payoff can be associated with each information level resulting from the acquisition of a new packet. In the usual case the probability of payoff will be an increasing monotonic function of the number of information packets since information generally assists in the selection of a correct alternative. The association of a probability of payoff with every information level allows the computation of an expected value for each level: $E.V. = P_q\,U(z) - {}_qU(c)$

where P_q = probability of payoff with q information packets

$U(z)$ = utility of payoff

$U(c)$ = utility of the cost of each information packet

q = number of information packets

A reasonable rule would be to acquire that amount of information which yields the maximum expected value. This rule would yield a result identical with that obtained by successive application of the criterion previously mentioned for all situations in which probability of payoff is a monotonic function of the number of information packets.

When the decision maker does not have precise knowledge of the probability of a successful outcome at each information level the "rationality" criterion is difficult to apply. One may reasonably assume, however, that under such circumstances subjects will expect the probability of successful choice to increase with the amount of information available at the time of decision and that this prior expectation will be reinforced, since in most decision situations "better" decisions are made when information is available. Thus subjects subjective probability of success, initially, would be a monotonic increasing function of the amount of information acquisition and, with experience, should approach the "true" probability of success. Of necessity, such subjective estimates of the relevant probabilities would have to be utilized in applying the "rational" criterion suggested.

The results on information-seeking responses (queries) do not support an expected value maximization assumption of the experimenter controlled objective probabilities of payoff are used since for all cost-payoff conditions, expected profit is greatest when the maximum number of queries are made.

However, a post-experimental questionnaire indicated that subjects' estimation of the cost-payoff function was quite at variance with the experimenter controlled objective probabilities. The subjects' subjective expected profit function for the $.05 condition remained monotonic increasing, but for the other cost-payoff combinations there was a transition from higher to lower expected profit occurring at progressively lower information levels.

TABLE 7

APPROXIMATE AVERAGE SUBJECTIVE PROBABILITY ESTIMATES AND EXPECTED PROFIT-QUERY FUNCTION BASED ON THESE ESTIMATES

Queries	Subj. Prob.	0-5	1-10	5-30	10-55
0	.20	$.01	$.02	$.06	$.11
1	.30	.15	.02	.04	.065
2	.40	.020	.02	.02	.02
3	.50	.025	.02	0	-.025
4	.70	.035	.03	.01	-.015
5	.75	.038	.025	-.025	-.087

In terms of these functions and a maximization of expected profit assumption, a decrease in search should be evident over the increasing cost-payoff combinations, a prediction clearly consonant with the obtained results of our early experiments. However, for none of the cost-payoff conditions did the average number of queries equal the information level of maximum subjective expected profit. Thus, in only a limited sense, is the subject's behavior consistent with the "rational" criterion proposed.

The results are encouraging and would support further refinement and development of the model if it had greater predictive and heuristic value for understanding the effects of other variables, e.g., uncertainty level and uncertainty reduction rate, and also if it provided a better conceptual apparatus for handling the problem of individual differences in search and time allocation. Within the framework of the theory, variables exert their effect on search by modifying either expectations, the utility of the cost of information, or the utility of payoffs. Individuals presumably differ in their estimation of probabilities, in the value they attach to the effort, time, or money involved in information acquisition, and in the value attached to the consequences of a decision. In addition, various situational factors, such as time pressure, would be expected to influence these variables. However, until more efficient procedures are devised for measuring such factors it is hard to take seriously the task of developing a comprehensive model encompassing the variables demonstrated to have an effect on search behavior.

The above considerations provide one basis for prediction in multistage decision situations which include a decision to acquire additional information or not. With some equivocation predictions can also be made from a drive reduction theory although not so precisely or easily since the application requires identification of the drive evoking stimuli, response habit hierarchies, and the conditions which serve to reduce drive and thus reinforce the behavior. With respect to cognitive processes there is considerable disagreement and confusion on all of these points.

Most "cognitive" theorists would probably agree with the statement that some form of dissonance, or incongruity, is a sufficient condition for eliciting search behavior either directly (ref. 16) or through the mediation of emotional distress or arousal (ref. 5). Dissonance and/or incongruity may arise from a discrepancy between a self concept and concrete perceptions of self-other interactions, an incompatibility between stored information and present inputs, or a discrepancy between stored cognitive elements such as beliefs.

Whatever its source, dissonance functions as a drive in eliciting responses that conduce to the mitigation of the dissonance. One possible course of action that may be instrumental in reducing an incongruity is information acquisition. For example, a discrepancy between our memory of an event and someone's report may be removed by checking an objective record, discordance between a conception of oneself as nonaggressive and an accusation of aggression may be reduced by checking the reactions of observers other than the accuser. Thus, information-seeking, being often instrumental in reducing incongruities, should have a high probability of evocation when incongruities occur. Some theorists insert an intervening event, emotional distress or arousal, between the evocation of incongruity and the elicitation of a search response. Arousal and emotional excitement are assumed to constitute the basic factor in a conception of generalized drive and additions to, or reorganizations of, information in memory is postulated as drive reducing.

A similar line of reasoning may be followed when the stimulus conditions elicit response conflict and/or uncertainty. Presumably "uncertain" decision situations evoke response conflict since the choice of a "best" alternative is equivocal. In such situations the subjective probability of a desired outcome occurring if, say, alternative A is selected is one determinant of the strength of the tendency to select A, and similarly for alternatives B, C,..., N. If one assumes that this anticipatory response strength directly determines the probability of the response then an information measure of response uncertainty can be derived from the decision makers subjective probability distribution relating alternatives to outcomes. More important, one would expect such a measure to be directly related to the degree of induced conflict and thus to various indicators of emotional distress and arousal. There is evidence to support the contention that arousal is a concomitant of conflict (ref. 32) and uncertainty (refs. 6, 33, 38) and that such arousal functions as generalized drive.

Thus, increases in response uncertainty should result in increases in arousal which constitute a basic determinant of generalized drive level. The strength of responses previously instrumental in reducing uncertainty should then be increased, as reflected in higher probabilities of response evocation. Since information-seeking presumably has a long and successful history in reducing uncertainty the probability of information-seeking responses should increase with increases in uncertainty. If one assumes that increases in disjunctive reaction time reflect, in part, more "search" activity then the large number of studies demonstrating a positive relationship between uncertainty and disjunctive reaction time would support the assumption that increasing uncertainty results in increasing search behavior (e.g., refs. 3, 20).

Our results on information acquisition are in accord with the assumption that increases in uncertainty result in increases in search behavior. The relationship between uncertainty and search appears to be negatively accelerating reaching an asymptote at about 3 bits of uncertainty. A similar asymptotic level is obtained for a plot of subjective uncertainty vs. objective uncertainty suggesting that subjective uncertainty is the important motivational variable. The results on "arousal" though somewhat ambiguous, do not support the assumption that arousal plays an important mediating role.

The non-linear relationship between information acquisition and subjective uncertainty can be accounted for within the "drive" framework by assuming that other responses, incompatible with search, are also strengthened as uncertainty and thus, motivation arousal increases. For example, it is conceivable that "internal" information search and symbolic information processing acts are also "aroused" by uncertainty and may, in fact, be dominant. When faced with uncertainty the organism may first search in memory for information which can provide a basis for choice and attempt to evaluate and integrate such data to assess its implication for the problem at hand. Only when these processes are completed and only if they fail to produce data of relevance, will the organism seek to acquire further information. Even then, since new information is rarely coded in a form suitable for direct application to a specific problem, these symbolic processing activities may again be evoked. Thus, at any stage in the sequential decision process the organism may be faced with a conflict between choosing on the basis of present information, seeking new information from the environment, and conducting a more intensive and extensive memory search and processing of recently acquired or stored data.

The development of a model to predict the transient dominance of one or another of these responses is obviously beyond the scope of this paper, but one can suggest some of the variables that may be important. The amount of relevant information in immediate memory, the similarity between the present situation and others experienced, the difficulty of understanding and synthesizing new information available, ability to process data and search memory, all should effect the relative priority the organism places on acquisition of additional data versus the processing of stored data. Thus, for example, one would predict less time devoted to information acquisition the greater the amount of information currently available, with the termination of information seeking being a function of the difficulty of processing the information.

Indirectly implied in the above formulation is the assumption that information search would not occur under conditions of certainty. As Berlyne (ref. 5) and others (refs. 17, 25) have noted, however, there is some evidence that organisms tend to maintain a rather specific amount of "arousal tonus" and thus by inference, would prefer situations of some uncertainty. This suggests that information-seeking would not be initiated at zero uncertainty but at that level of uncertainty which produces arousal greater than a preferred level of "arousal tonus". The probability of search should then increase with increasing discrepancy between a preferred uncertainty level and the level of uncertainty induced by the choice task and should terminate when the preferred uncertainty level is reached. Our results on residual uncertainty suggest the existence of an "uncertainty-commitment" threshold for decision. This may be a manifestation of an attempt on the part of our subjects to maintain a preferred level of "arousal tonus".

Our data on "personality" differences in search behavior support and extend the proposition that subjective uncertainty is a strong determinant of search activity: the complexity of predecision information processing increases as a function of problem uncertainty and of complexity of conceptual structure. These two variables interact so that the complexity of abstract persons' information processing increases more in proportion to increases in stimulus uncertainty than concrete persons'. Such data were obtained for all three indices of complexity utilized: amount of information search, number of qualifying remarks accompanying decision, and amount of additional information generated and given with the decision.

The data suggest that the determining factor responsible for the individual differences is the degree of perceived uncertainty; abstract persons perceive more uncertainty in decisions than concrete persons. Abstract persons apparently encode a greater number of dimensions of information from most stimuli, their integrative schemata within which they organize the information are more complex and more numerous, and they have a greater tendency to seek multiple interpretations of situations, hence find themselves faced with a greater number of alternatives than concrete persons.

If, as we have previously theorized, subjective uncertainty acts as a drive state which may be reduced by information, then these differences in conceptual structure which cause differences in perceived uncertainty result in different degrees of motivation to engage in information processing.

Undoubtedly, differences in drive are not the only determinants of individual differences in predecision behavior. Differences in conceptual structure would also result both in differences in perception of the problem base and of the ways in which additional information is seen to relate to the problem. The value of information is thus subjectively determined: it depends upon the way in which it adds to the existing knowledge, which in turn depends upon the conceptual organization of the individual; once information is acquired, the ways in which it is integrated with existing information would depend upon the way in which that information is organized.

So far, we have avoided confronting the issue of "what is the reinforcement" for information search. Two possibilities exist, the occurrence of a desired outcome, e.g., a payoff, or subjective uncertainty reduction.

In the studies reported the two factors are somewhat confounded as they probably are in most complex decision tasks. However, in the procedure employed the probability of payoff increased more rapidly than subjective uncertainty reduction (table 7) and, thus, the assumption that uncertainty reduction serves as a reinforcer could account for the generally low level of information acquisition by our subjects. The assumption could not in itself, however, account for the differences between cost-payoff conditions in information-seeking.

If the expectation of uncertainty reduction were a function of the cost of information, i.e., the higher the cost the greater the expectation of uncertainty reduction, and if a positive discrepancy between obtained and expected uncertainty reduction served as the reinforcer for search, the cost-payoff results could be explained. The higher the cost of information the greater the expectation of uncertainty reduction and the greater the likelihood that the obtained reduction will be less than the expected. Reinforcement for information search would then be less likely the higher the cost of information and the asymptotic level of queries would be lower the greater the information cost.

It is apparent that the collection of concepts and propositions discussed constitute the bare beginnings of a theory of "search" behavior. The "theory" is neither an *S-R* nor a consistently need-reduction one, but borrows freely from both traditions. It is based largely on the concepts advanced by Berlyne (ref. 5) and thus focuses on uncertainty as an important intervening variable. For all organisms there is some optimum level of uncertainty greater than zero, and displacements from this level elicit responses which have been instrumental in the past in reducing incertainty. Information acquisition and symbolic information processing activities are two responses assumed to have high habit strength under conditions of uncertainty and thus have a high probability of being elicited. Since such responses tend to restore the organism to its preferred uncertainty level they are reinforced.

To account for the obvious importance of "cost" variables the concept of an expectation for uncertainty reduction is introduced. Information search is reinforced only if the uncertainty reduction which results from search exceeds the expected rate of uncertainty reduction. The latter variable is assumed to be a positive function of information cost, importance of the problem, and initial level of uncertainty.

The major appeal of the "uncertainty" theory is its potential for handling affective states which have been sorely neglected by investigators in the study of decision making. Its major deficiency is the relative neglect of "rational" factors in choice behavior. Organisms are neither computers that impartially weigh and evaluate costs, payoffs, and probabilities, nor irrational, impetuous, seekers of equilibrium in drive states, whatever the costs. Presumably, they exhibit some facets of both idealization in their decision making behavior; the task is to determine the nature of the compromise.

REFERENCES

1. Abelson, R. P., "Modes of Resolution of Belief Dilemma," Journal of Conflict Resolution, Vol 3, pp 343-353, 1959.

2. Abelson, R. P., and M. J. Rosenberg, "Symbolic Psychologic: a Model of Attitudihal Cognition," Behavioral Science, Vol 3, pp 1-13, 1958.

3. Alluisi, E. A., P. F. Muller, and P. M. Fitts, Rate of Handling Information Presentation, WADC Technical Note 55-745, Wright Air Development Center, Wright-Patterson Air Force Base, Ohio, 1955.

4. Attneave, F., Applications of Information Theory to Psychology, Holt, New York, 1959.

5. Berlyne, D. E., Conflict, Arousal, and Curiosity, McGraw-Hill, New York, 1960.

6. Berlyne, D. E., "Conflict and the Orientation Reaction," Journal of Experimental Psychology, Vol 42, pp 476-483, 1961.

7. Broadbent, D. E., Perception and Communication, Pergamon, New York, 1958.

8. Bruner, J. S., J. J. Goodnow, and G. A. Austin, A Study of Thinking, Wiley, New York, 1956.

9. Cartwright, D., "Relation of Decision-Time to Categories of Response," American Journal of Psychology, Vol 54, pp 174-196, 1941.

10. Cartwright, D., "Decision-Time in Relation to the Differentiation of the Phenominal Field," Psychological Review, Vol 48, pp 425-442, 1941.

11. Davidson, D., and P. Suppes in collaboration with S. Siegel, Decision-Making, an Experimental Approach, Stanford University Press, Stanford, California, 1957.

12. Driscoll, J. M., Subjective Uncertainty in Decision Making, Unpublished Doctoral Dissertation, University of Delaware, 1964.

13. Driscoll, J. M., and J. T. Lanzetta, "The Effects of Problem Uncertainty and Prior Arousal on Predecisional Information Processes, Psychological Reports, (in press), 1964.

14. Edwards, W., "Reward Probability, Amount and Information as Determiners of Sequential Two-Alternative Decisions," Journal of Experimental Psychology, Vol 52, pp 177-187, 1956.

15. Festinger, L., "Studies in Decision: I. Decision-Time, Relative Frequency of Judgment and Subjective Confidence as Related to Physical Stimulus Difference," Journal of Experimental Psychology, Vol 32, pp 291-306, 1943.

16. Festinger, L., Theory of Cognitive Dissonance, Row, Peterson, Evanston, Illinois, 1957.

17. Glanzer, M., "Curiosity, Exploratory Drive and Stimulus Satiation," Psychological Bulletin, Vol 55, pp 302-315, 1958.

18. Harvey, O. J., D. Hunt, and H. Schroder, Conceptual Systems and Personality Organization, Wiley, New York, 1961.

19. Hoge, R., The Effects of Attitudes on Predecisional Behavior, Unpublished Master's Thesis, University of Delaware, 1963.

20. Hyman, R., "Stimulus Information as a Determinant of Reaction Time," Journal of Experimental Psychology, Vol 45, pp 188-196, 1953.

21. Irwin, F. W., and W. A. S. Smith, "Value, Cost, and Information as Determiners of Decision," Journal of Experimental Psychology, Vol 54, pp 229-232, 1957.

22. Kochin, M., and E. Galanter, "The Acquisition and Utilization of Information in Problem Solving and Thinking," Information & Control, Vol 1, pp 267-288, 1958.

23. Lanzetta, J. T., and C. Hawkins, "Information Search in Decision Making," Paper to be read at annual meetings of the American Psychological Association, Los Angeles, California, September 1964.

24. Lanzetta, J. T., and V. T. Kanareff, "Information Cost, Amount of Payoff and Level of Aspiration as Determinants of Information-Seeking in Decision Making," Behavioral Science, Vol 47, pp 208-214, 1962.

25. Leuba, C., "Toward Some Integration of Learning Theories: the Concept of Optimal Stimulation," Psychological Reports, Vol 1, pp 27-33, 1955.

26. March, J. G., and H. Simon in collaboration with H. Guetzkow, Organizations, Wiley, New York, 1958.

27. Marschak, J., "Towards an Economic Theory of Organization and Information," In R. M. Thrall, C. H. Coombs, and R. L. Davis (Editors), Decision Processes, Wiley, New York, 1954.

28. Miller, G. A., "The Magical Number Seven, Plus or Minus Two," Psychological Review, Vol 63, pp 81-97, 1956.

29. Miller, G., E. Galanter, and K. Pribram, Plans and the Structure of Behavior, Holt, New York, 1960.

30. Miller, L., and J. T. Lanzetta, "Choice among Four Alternative Gambles as a Function of Monetary Level of Play," Psychological Reports Monograph Supplement, Vol 6, Number VII, 1962.

31. Neimark, E., "Information Gathering in Diagnostic Problem-Solving," Unpublished manuscript, Department of Psychology, New York University, 1959.

32. Pavlov, I. P., Conditioned Reflexes, Clarendon Press, Oxford, England, 1927.

33. Popov, N. A., "Electroencephalographic Observations of Human Cortical Reaction," Année Psychologique, Vol 53, pp 415-429, 1953.

34. Pruitt, D. G., An Exploratory Study of Individual Differences in Sequential Decision Making, Unpublished Doctoral Dissertation, Yale University, 1957.

35. Savage, L. J., The Foundations of Statistics, Wiley, New York, 1954.

36. Schroder, H., and S. Streufert, The Measurement of Four Systems Varying in Level of Abstractness (Sentence Completion Threshold), ONR Technical Report Number 11, Project NR 171-055, 1962.

37. Scodel, A., P. Ratoosh, and J. S. Minas, "Some Personality Correlates of Decision Making under Conditions of Risk," Behavioral Science, Vol 4, pp 19-28, 1959.

38. Sharpless, S., and H. Jasper, "Habituation of the Arousal Reaction, Brain, Vol 79, pp 655-680, 1956.

39. Sieber, J. E., and J. T. Lanzetta, "Conflict and Conceptual Structures as Determinants of Decision Making Behavior," Journal of Personality (in press), 1964.

40. Simon, H. A., Models of Man, Wiley, New York, 1957.

41. Tanner, W. P., Jr., and J. A. Swets, "A Decision Making Theory of Visual Detection," Psychological Review, Vol 61, pp 401-409, 1954.

BLANK PAGE

DISCUSSION OF LANZETTA'S PAPER

HUTTENLOCHER (Discussant). I would like to comment on the interesting result that people stop gaining information short of a maximum level, even when they were not charged for asking questions. You mentioned time and effort as possible "stoppers" of information acquisition. I wondered whether the subjects were actually sure of the answer at some point short of complete question asking? Have you counted that out, i.e., that the probabilities were actually quite low, so they were not short?

I also wondered whether the subjects really didn't anticipate any further information? Is there, for a particular task, some uncertainty level at which people feel they don't want any more information? There are certainly situations where people try to get as much information as possible, and I believe it might be interesting to discuss the relevance of the task to the number of questions asked.

LANZETTA. The issue of subjects not acquiring all the information, even under zero costs, is one that Ward has raised previously about the first study under the one cost condition. I don't know what it is; we have done some interviewing of subjects, and it seems to be, in part for some subjects, a function of their conception of what a good decision maker is. A good decision maker is not one who needs all the information. Any damn fool can make a decision when all the data is in. So, you get some subjects that presumably terminate search just for, well, a kind of consistency with their self concept, the concept of what decision making ought to be. For others, I think you probably do find they do not anticipate that they will gain very much. That is very plausible in the first study, less so in the concept attainment task where they have a sheet of paper and may cross off the wrong concepts, and the amount of uncertainty reduction is constant per query and is almost as high as half a bit per query. Even here, they terminate prior to reaching that threshold. This still may be reflecting not so much that they don't anticipate getting anything on the next query, but that they anticipate, and know, from prior problem experience in our situation, that even if they take all the queries they will not reduce the uncertainty to zero. So, if you are going to have to take a risk, maybe you ought to stop at this risk level. There may be operating, in essence, a preference for a particular risk level which gets reflected in a preference for an uncertainty level under conditions where uncertainty cannot be reduced to zero. Now, since I think most of the decisions that a lot of us are interested in are probably of that sort, that is, situations in which, in fact, you cannot reduce your uncertainty to zero, this may be quite relevant.

HUTTENLOCHER. It would be interesting if this being able to make a guess early has the higher utility.

ROBY. In a task we have used, a kind of mechanized "20 Questions" in which uncertainty can be reduced to zero, subjects will stop at about one bit. It is, I think, a kind of sporting feeling, that it's more fun.

LANZETTA. If you increased the importance of the decision you would probably be able to force search further.

BERLYNE. This may relate to the point I was making this morning, this great overlap of extrinsic and intrinsic exploratory functions, even though it is a good thing to separate them experimentally. When you're gambling, part of the reward comes from the money or whatever else it is. You also get some reward from reduction of curiosity, whether you win or not. If you still have residual uncertainty you don't deprive yourself of this extra reward. It's a pity that the subject could not be rewarded without knowing whether he has been or not. I think your last suggestion would be interesting to try out. If you made it so very, very important, in the way of a gain or loss or both, dependent on the outcome, then I suspect subjects would reduce their uncertainties as much as possible. It is not just a matter of uncertainty. The matter of conflict goes up as uncertainty goes up, all things being equal. Also uncertainty can be held constant if what is at stake goes up. The idea of tolerating uncertainty for the fun of getting the answer, probably only works if you have moderate uncertainty. If you have to wonder whether or not you're going to be shot in the next five minutes, the fun of finding out might be a minimum consideration.

EDWARDS. You can reward people without having them know it, and I'm engaged in doing such an experiment. I can't tell you yet how it came out, since it hasn't been completed.

BERLYNE. Can you tell us how it's done?

EDWARDS. This is a one-armed bandit situation, and the subject on any trial has a choice of observing or predicting. If he chooses to observe, he sees which event of the two he might have predicted was in fact the one that would have occurred. If he chooses to predict, he gets no formal return whatsoever, but a counter off in the next room counts whether he's right or wrong and if he's right he gets paid later on. If he guesses wrong, he loses ten cents, later on.

BERLYNE. He knows that it's going to matter and that he better darn well be right, because it's going to matter to him and he knows what the payoff function is; he just doesn't know anything until he's all done.

I'd like to say a word about John's more general feeling about the degree to which the subjects were rational, particularly in the first experiment. This is one of these experiments in which there is the inside view and the outside view, and the inside view is known only to the experimenter and God. The question being asked is irrelevant and the information being acquired is irrelevant, and nothing, in fact, of all of the elaborate cognitive structure being presented to the subject means a single solitary thing except the number of times that he pushes that button and pays a chip for one more item of information. But of course the subject doesn't know that, and since he has only 25 trials, he doesn't have much of an opportunity of finding out. Instead, he has an elaborate, complicated story that gets presented to him on a ground glass screen, and every time he puts a chip in, he gets an increment of two or three further sentences contributing to this story. It seems to me that when the subject's conception of what the world is all about is so grossly different from what the experimenter's conception of what the world is all about, it is not markedly surprising that rationality defined in the experimenter's world has relatively little relation to what our subject rationally or irrationally chooses to do.

LANZETTA. That would have plausibility except that I presented some data on subjective probabilities which showed behavior was not "rational" using subjects' own estimates of the probabilities of being correct, given the information items. Obviously, I don't consider the objective probabilities a very good definition. They were not employed to test the rationality assumption. Frankly, they were employed to minimize individual variability in the utilization of information.

Certainly, in the third study, what you are saying would not hold at all. The experimenter's view of the situation and the subject's were exactly identical and we got essentially the same phenomena.

EDWARDS. Furthermore, in the subject's conception of what his world is like, subjective probability must seem like rather a strange animal because it's as true that some number representing how likely he is to be right as a function of the number of queries he makes is irrelevant, if he is really thinking about the complex story. It's just as true for the subject that the number is irrelevant as the number inside your head is irrelevant. It isn't a number type of problem for the subject.

SHUFORD. I'd like to suggest another reason for why subjects did not look at everything. You could have designed the experiment so that the optimal amount of information to be gathered varied over the whole range from zero to 6. Maybe the subject has a possibility of knowing or learning that this varies. Instead, in these experiments, if each time the optimal amount of information that he gathered is all of it, it is pretty obvious to the subject. The subject wants to please or he wants to ruin you, one or the other. Say he wants to please. One of the things you are measuring is how much information he looks at, and I would suggest that the subject is going to vary under these conditions where each time the optimal thing is one fixed extreme value, and that the subject is going to vary this from time to time just in order to please you or get some kind of differences here. Is there any further experimentation concerning that point?

LANZETTA. We have plotted data showing these trial by trial effects. These were attributable to the nature of the problem and the perceived uncertainty reflecting the subject's presumed response tendency.

SHUFORD. It just seems as though you have it loaded against him looking at all of it all the time.

LANZETTA. Frankly, that wasn't the thing of most interest, whatever we postulate as the variable for accounting for why he did not acquire all of the information. The real issue is how to account for the fact that he took less and less as the cost-payoff schedule went up.

EDWARDS. The information cost more and more.

LANZETTA. Yes, but the payoff is simultaneously going up more and more. You can see the important thing is not only the fact of failure to acquire all of the information when it was of zero cost. I think that is very interesting, and can be accounted for in a variety of ways, but I think more interesting is the finding that their behavior varies as a function of perceived uncertainty, cost-payoff, time pressure, level of aspirations, etc.

EDWARDS. If these subjective probabilities you got from your subjects were, in fact, relevant to their behavior, then you could, of course, predict what you found, because they would simply say that the probability of getting the prize is not as great as the official value that you use, and consequently, the incurring of a cost in order to increase that probability isn't going to be sufficiently worthwhile. Subjectively expected utility will decrease, as your arithmetic shows, by virtue of that.

LANZETTA. That is what I pointed out, that in a limited sense they do seem to be consistent with some notion of at least being aware of the expected value of the information, though certainly, not in any detailed sense consistent with the prediction.

HUTTENLOCHER. The finding that less queries were made with a higher level of aspiration struck me as somewhat unexpected.

LANZETTA. Those differences were not significant. The variability thing is really, in some sense, more interesting and you seem to get greater variability under 50% level of aspiration. I'd like to ask Jack if this seems to be something you would predict from achievement motivation. Would you predict much greater variability?

ATKINSON. May I ask a question about the experiment? You told the subjects what percentage of college students could get the solution?

LANZETTA. That was our level of aspiration in the first study.

ATKINSON. And were the subjects random samples of students or volunteers?

LANZETTA. No, they were not. In this experiment they were obtained out of ROTC and gym classes.

ATKINSON. Let me say a word about this in terms of the theory of achievement motivation. The strength of motivation to succeed at the task should have been highest in your 50% level of aspiration group, if the students believed that they were typical of the student population you were talking about in stating the norm. We tell students that 50% of college students can get the solution to this and later on if you ask them about their own abilities, they all think they are in the upper quarter. There's a little slippage between using a college reference norm and inducing a subjective probability of success in a particular group of subjects. My hunch was, as you presented data, if anything, your students thought of themselves as better than average students, which might mean, subjectively, your 25% condition might have been the more motivating condition of the experiment. But if so, then following what Dr. Berlyne said this morning, the competition involved in these conflicting decisions should have been greatest there, and produced more search for information.

LANZETTA. Well, actually it did.

ATKINSON. That sounded fine, until you showed the greater variability in the 50% condition, and the unique feature of that left me up in the air as to what I should conclude.

BERLYNE. Somebody drew attention to the magic number of one bit, that is, one bit residual uncertainty. The subjects in your situation left it at one bit as the favorite uncertainty. One bit of uncertainty came up in Ward's experiment on gambling preferences.

EDWARDS. I know what you're referring to, but I'm reluctant to say that is an uncertainty type of effect.

BERLYNE. It's still one bit again.

LANZETTA. Somebody tested this in a psychophysical judgment type test and got something like .85 bits.

BERLYNE. I have come out with a magic number.

FLOOR. I'm amused by the possibility of designing questions or problems such that all 6 of the possible solutions are acceptable as being correct, if the person has been curious enough to ask five queries and get five answers. It seems to me that either the stories must be so incredibly complex that the subject never really knows what is going on at all, or some feeling of implausibility might arise in the subject's mind. I mean, for instance, if a subject chose what he thought to be the least likely of the alternatives, something that just couldn't be, and was told, "Sure enough, you're right," you'd think a doubt might arise.

LANZETTA. It might if he were willing to risk that. I ought to clear up the fact that because we consider every alternative equally likely, that is, equally likely to be reinforced, that doesn't mean the subjects did. As a matter of fact, this is what our uncertainty measure measures. From the subject's point of view, these were not all equally plausible answers. On some of the problems there was a very high prominence accorded to some responses, that is, a large majority of subjects saw that as the most likely correct decision to make under the circumstances presented. I'm sorry I even raised the issue except that it is important in the experiment. All this is, in essence, is an attempt to say, "Well, look, I know subjects are going to differ in intelligence, etc., but I'd like to be able to keep that kind of confounding variable out of our information acquisition studies;" that is, we'd like to study their behavior unconfounded by their ability to process the information, so we simply reinforce them. In essence, we say that we will assume they did process the information correctly and come to some correct decision. All we can do is build evidence through observation of the apparent conflict when the subject is faced with a choice. You could anticipate the subject would pay no attention to the response box, just reach over and slap a key, while in fact, this is not the way it seems to go. The subject sits there and has a chip in his hand and starts putting it in and then brings it back; then goes over here and starts to make a response; then changes his mind and goes back and puts the chip in the slot and looks at the new information, obviously debating over the choice of an alternative.

FLOOR. I have done some work which suggests that there are important personality correlations with the kind of variability you found, and I note that you did find variability to be at least in part a function of aspiration. You attempted to induce it. Do you have any other observations as such, insights from your experiment that might be interesting?

LANZETTA. No, but you are suggesting something we haven't tried, that is, to try the correlations between these various time and personality measures against variability rather than against the search costs themselves. We have not tried that.

FLOOR. Being variable sometimes seems to be an important way of reacting to an experiment. The subject is testing the experimenter, while the experimenter is testing him.

BLANK PAGE

SOME NEGLECTED VARIABLES IN CONTEMPORARY CONCEPTIONS OF DECISION AND PERFORMANCE*

John W. Atkinson and Dorwin Cartwright**
University of Michigan

A few decades ago, psychologists frequently got into heated debates over the desirability of attempting to construct a theory of an active organism. The debate has subsided, and as we review the various conceptions of determinants of performance and decision that have been proposed, it is quite clear that we have so far settled for a theory of a reactive organism. Our attempts to develop formal statements of how the impetus to act in a particular situation is determined invariably presume that initiation of a tendency to respond resides in a stimulus. We tend to begin our conceptual analysis with a stimulus. Just as Aristotle found it impossible to conceive the cause of terrestrial motion without positing the action of external force, we find it extremely difficult to conceive the possibility of an active tendency to perform a response without an immediately present stimulus, somewhere inside or outside of the subject we study, as the cue or trigger or instigation. Yet neurophysiologists have thrown away the stagnant picture of the brain as a machine designed to turn off excitations from outside. They are beginning to yield a new picture, one in which an active brain selects what shall and what shall not be stimuli for it (ref. 17). Does psychology lag behind in its conceptual and experimental analysis of the determinants of preference and performance? We think so.

Our intention in this paper is to propose some ways of shaking off the conceptual shackles of the traditional stimulus-bound view of the determinants of decision and performance. Our argument points towards an adequate theory of an active organism.

---

*This is a revised version of the paper given at the symposium by John W. Atkinson.

**The basic argument of this paper was the focus of a series of discussions between the authors during 1960-61 which explored the relatedness of the conception of motivation which has evolved in studies of achievement motivation and Lewinian Field Theory. One of us (JWA) was then supported in a historical survey of theories of motivation by a Guggenheim Fellowship while the other (DC) was a Fellow at the Center for Advanced Study in the Behavioral Sciences, the locus of our joint activity. We are indeed grateful to both sponsors for providing this opportunity for collaboration. This paper is a particular phrasing of the general position developed in these discussions. It was prepared for oral presentation by the first author and was designed both to suit the requirements of a symposium on predecisional processes and to make explicit the relationship between contemporary research on personality dispositions and some of the conceptual problems of decision theory. We express our appreciation to our colleague David Birch whose criticism has helped to clarify the argument.

The ideas we will present have evolved in a program of research on achievement motivation and performance (ref. 2, chapters 20, 22, 42; refs. 8, 9,). The aim of this work has been twofold: (a) to improve technology for assessment of individual differences in personality; and (b) to develop a conceptual scheme which explains how dispositions of personality, assumed to be relatively general and stable characteristics of a person, combine with more specific and transient situational influences to determine the direction, vigor, and persistence of action. We believe, in the words of a popular song, "you can't have one without the other". Certainly there will be no useful technology for assessment of individual differences in personality until there is some principle relating what has been assessed to behavior. Nor, we believe, are we ever likely to have very useful explanatory schemes for decision making and performance until assessment of relatively stable properties of the persons we study becomes as integral a part of experimental analysis of decision making and performance as the assessment of the mass of an object is to the study of motion.

There has been a considerable evolution in thought about what causes the impetus to perform a particular response. In the early days of behaviorism [e.g., Thorndike (ref. 21), Watson (ref. 23), and Pavlov (ref. 16)] the stimulus alone was thought to provide the impetus or motivation of the response. The important issue seemed merely to account for the strength of the connection between S and R. This, it was believed, would yield a useful explanation of the probability of occurrence and latency of the response. The so-called "purposivists" of these early days [e.g., McDougall (ref. 15); see also Woodworth (ref. 24)] argued vainly for quite a long period that a particular S-R incident had to be viewed within the broader context of goal-seeking activity which they held to be generally characteristic of living organisms. The "purposivists" emphasized the importance of the anticipated consequence of a response (the goal) as a determinant of the impetus to perform the response. The programmatic principles of performance presented by both Tolman and Lewin, beginning in the 1920's, focused attention upon (a) the degree to which a given consequence was anticipated, and (b) degree of preference for that type of consequence, as immediate determinants of performance. Tolman (ref. 22) introduced the concepts of expectancy of the goal and demand for the goal; Lewin (ref. 13) referred to the potency, psychological distance, and valence of a goal as variables which intervene between observed stimulus situation and response. The first formal conception of the determinants of the tendency to respond within the S-R orientation, by Hull (ref. 11), took the form ${}_SE_R = D \times {}_SH_R$. The tendency to respond (${}_SE_R$) was attributed to general excitability (D) and the strength of association between specific S and R (${}_SH_R$), with no explicit reference to anticipated consequences of the R. But the difficulties with this conception became immediately apparent, and in more recent formal statements of S-R behavior theory [e.g., Hull (ref. 12) and Spence (ref. 20)], the anticipated consequence or goal has been introduced as a determinant of the impetus to respond. Habit strength has begun to recede as a variable of predominant interest and significance in research. Spence, for example, now proposes that the strength of the tendency to respond (E) should be conceived as a function of $(D+K) \times H$, where D (drive) represents the effect of deprivation, K (incentive) represents the combined effect of frequency of reward, delay of reward, and amount of reward, and H (habit) represents the number of times S and R have occurred contiguously. Formerly it had been assumed that frequency, amount, and delay of reward merely influenced the growth of habit. [See Thorndike (ref. 21), and Hull (ref. 11)]. It has already been proposed by Birch, Burnstein, and Clark (ref. 4) that all of the motivating effects implied by $(D+K)$ may be attributed to the frequency, delay, and amount of reward in previous training as these several variables influence the strength of anticipatory goal reaction, the proposed physical mechanism of K.

It is now more apparent than ever before that there is considerable agreement that the anticipated goal or consequence of a response is an important determinant of the tendency to perform the response. There remains, if one examines closely the implications of expectancy times value formulations of the sort proposed by Lewin, by Tolman, and by Rotter (ref. 18), and contemporary decision theory (ref.7) in relation to a conception like that of Spence, the important issue of whether the anticipated goal influences performance selectively, as the expectancy-value theories propose, or non-selectively, as proposed in the equation $E=(D+K)\times {}_{S}H_{R}$. The implications of a non specific drive are fully developed by Brown (ref. 5).

The mode of thought involved in the development of these conceptions and their use as guides in design of research is essentially the same. Given a stimulus situation, there is then and there aroused in the subject of study a tendency to perform a response, T_R. The strength of T_R is accounted for in terms of a combination of certain variables A, B, C, etc., which represent the present effects of past training and the organic condition of the subject at the time. One appeals to the immediate stimulus situation to account for the activation of an otherwise latent associative disposition, whether it be conceived as the link between S and R (habit) or between R and its previous consequence or goal (expectancy). If the stimulus situation is such as to arouse simultaneously two or more tendencies, T_{R1} and T_{R2}, there is stated (or implied) a model linking the relative strengths of tendencies to actual measurements of probability or latency of response. It is generally agreed that the stronger or dominant tendency is most likely to occur as a response and that latency of the response will depend upon the relative strengths of the competing tendencies. [See, for example, Cartwright and Festinger (ref. 6).]

The general conception evolved in empirical studies of achievement motivation and performance is of the same form. In the simplest case, it is assumed that the tendency to perform a particular response (viewed explicitly as a goal-directed tendency) is a multiplicative function of *Motive* (M_G), a relatively stable disposition of personality which refers to a class of goal objects or activities, *Expectancy* ($E_{R,g}$) that a given response will lead on to a goal of that class, and *Incentive Value* of the particular goal that is expected (I_g). The product of these three variables is referred to as the strength of motivation to perform the response, or more simply, as the *tendency to perform the response*, $T_{R,G}$.

$$T_{R,G} = M_G \times E_{R,g} \times I_g \quad (1)$$

$$T_{R,G} = E_{R,g} \times (M_G \times I_g) \quad (2)$$

The use of two subscripts to describe the tendency to respond, one referring to the act in question, the other to the expected consequence, or goal of the act, is an attempt to convey the idea that performance of a particular response is to be understood as an expression of the goal-seeking trend of behavior.

The second statement (above) shows more clearly how this conception corresponds with other versions of the expectancy-value type of theory. We, in fact, have explicitly stated what might be considered a theory of utility or valence in the proposition that utility or valence of an object or activity (called a goal when it is the expected consequence of some other response) is a multiplicative function of a disposition M_G, which refers to a whole set of objects and activities, and I_g, the relative value or preference for some element within this set. In our thinking, the **Motive** construct (M_G) refers to preferences between classes of objects and activities, and **Incentive Value** (I_g) refers to preference within a homogeneous set (ref. 2, chapters 20, 22).

We say this much about the general form of various conceptions of the immediate determinants of the tendency to respond and the particular conception evolved in work dealing with effects of individual differences in general dispositions called motives, so that you can better appreciate the possible relevance of suggestions arising in this work for the conceptualization of problems of more immediate interest to you.

Now let us proceed to use more general terms: T_R to represent the final strength of a tendency to perform a response (the impetus to act); and A, B, C to represent any three determinants of this tendency which, we shall assume, combine multiplicatively. Thus:

$$T_R = A \times B \times C \tag{3}$$

Each of us may now substitute his own favorite constructs for A, B, C, and even for T.

Studies of achievement motivation have yielded encouraging results when we apply this kind of conception to prediction of differences in level of performance (speed, response output) and selectivity when several tendencies are simultaneously aroused and preference is observed. However, we have learned something new from recent studies of persistence in problem-solving activity, that is, in studies of the duration of the attempt to succeed at a certain task when there is no knowledge of results or when success has not, in fact, been attained. French and Thomas (ref. 10), for example, found that Air Force personnel who were strong in assessed strength of achievement motive persisted longer in the attempt to solve a complex problem than men who were weak in assessed strength of achievement motive. Atkinson and Litwin (ref. 3) found essentially the same thing when time spent working on a final examination before turning it in was measured. But Feather (refs. 8 and 9), in a recent test of some of the more subtle implications of the theory of achievement motivation [see Atkinson (ref. 1)], discovered that it was really impossible to derive a hypothesis concerning differential persistence of activity on an insoluble puzzle (R_1) that could be attributed to differential strength of T_{R1}, as we conceive it here, without a very explicit assumption about the strength of T_{R2}, the tendency to perform an alternative task to which the subject might turn whenever he decided to give up in the attempt to solve the initial puzzle. The general implications of this now fairly obvious point are worth very careful consideration. One expects R_1 an activity in progress, to persist as long as $T_{R1} > T_{Rx}$, where T_{Rx} is the tendency to perform any other response as an alternative to continued performance of the initial activity. When $T_{Rx} > T_{R1}$, the initial activity will cease and R_x, whatever it is, will replace it.

Having appreciated this much, it was but a small step to recognition of something considerably more important. Whenever we attempt to measure the strength of a tendency to respond using the latency of response as our performance criterion, we present a stimulus and start a clock. We stop the clock and take a reading when the response of interest occurs. We then tend to attribute differences in latency of the response, let us say R_3, to differences in strength of T_{R3} produced by experimental influences on A_3, B_3, C_3, any variables which refer uniquely to strength of T_{R3}. When we do this, we ignore completely the possibility that the subject of study is already doing something when the stimulus S_3 for the response of interest, R_3, is presented; *for the same time measurement is called persistence when duration of an activity in progress is the matter of central interest and latency when the initiation of some new response is the matter of central interest.*

We can no more attribute the latency of response in a single stimulus presentation solely to variables which influence the tendency to perform that response than we can, as Feather discovered, attribute persistence of an activity in progress solely to variables affecting the strength of the tendency to perform that response. The question arises: have we conventionally treated the subject as at rest (inactive) in studies employing latency (or reciprocal latency = speed) to measure the strength of the tendency to respond? The answer would appear to be yes in that we have not formulated the principles of performance in a way that takes systematic account of the initial activity, that is, the activity already in progress at the time the stimulus of interest to the experimenter is first presented. We have recognized the problem to the extent that it is common experimental practice to establish a cooperative attitude in the human subject or to minimize extraneous influences in the animal experiment whenever latency of response or decision time is measured. But this is a defeatist strategy, evidence that we have conceptual schemes that always want to begin with the assumption of an organism at rest. Our first proposal is to consider activity in progress, the initial activity, systematically as a determinant of subsequent performance in our conceptual scheme and in experimentation.

A commonplace example will make the point of criticism we intend in reference to conventional attempts at systematic principles of performance and decision-making. Consider the case of children playing in the living room with their toys when their mother calls from the kitchen, "Dinner is ready." Now imagine that there is a deaf psychologist in the living room studying the persistence of play and a blind psychologist in the kitchen studying latency of response to mother's call. The deaf psychologist in the living room will try to account for observed differences between children in their persistence at play in terms of variables which have to do with the tendency to play. The blind psychologist in the kitchen will attribute variations in latency of response to mother's call to variables having to do with the tendency to eat. Why? Because the context of events embraced in conventional conceptions of decision-making and performance, or at least our use of these schemes, tends to be restricted to the immediate effects of the environmental stimulus in which we, the experimenters, are interested and which we control.

We need to accomplish a further transition in our thought about the factors to be considered systematically as immediate determinants of a particular response. The early conception that observed stimulus was the cause of the response was inadequate. It became necessary to broaden the context of factors to be given systematic consideration to embrace the idea that response to a stimulus is an incident in a larger unit, the goal-directed trend of the behavior at the time. This enlargement of the conception is evident when we now assert that expectancy of a consequence and value of the consequence are determinants of the tendency to respond; or when we assert that the magnitude of the anticipatory goal reaction must be represented in the equation for effective reaction potential. But this step did not require us to consider seriously the implication of what we now propose as a basic premise: a living organism is constantly active, doing something that we have to conceptualize as the overt expression of a tendency to respond, even, we would assert, when it is asleep. Hence the activity already in progress (specifically the tendency sustaining that activity) must appear in the principle of performance as a potent determinant of the performance of any subsequent response. Given this basic premise, the fundamental decision problem which deserves careful experimental scrutiny is not that of a subject implicitly assumed to be at rest when stimuli for option A versus option B are presented. The fundamental decision problem is that presented by the option of continuing an activity already in progress or performing, instead, some other activity when the appropriate stimulus for it has been presented.

Having once made explicit the premise of a constantly active subject, it comes as no surprise that there should be so little correlation between particular environmental stimuli and particular reactions that Skinner (ref. 19) should have found it useful to invent the category, operant behavior, to embrace responses which appear to be spontaneously emitted rather than elicited, as is a reflex, by an environmental stimulus. When T_{R1}, the tendency sustaining an initial activity is very strong, the presentation of a stimulus which excites T_{R2} may never yield R_2, the appropriate response to that stimulus. On the other hand, when T_{R1}, the tendency sustaining activity in progress is very weak, and S_2 is presented, the excitation of T_{R2} should be followed almost immediately by the occurrence of R_2. The extreme variations in latency of a particular response to a stimulus should be amenable to systematic explanation if we begin our attempt at explanation with the premise that a living organism is already actively doing something when the stimulus of interest is presented.

Let us consider the simplest case, where the subject is engaged in an activity designated R_1, which is sustained by T_{R1}, when the appropriate stimulus S_2 for R_2 is presented. Our conception begins with the assertion:

$$R_2 \text{ occurs when } T_{R2} > T_{R1} \tag{4}$$

(We ask you to assume the existence of a more inclusive model which coordinates strength of tendencies with measurements of probability and latency of response according to which the probability and latency of R_2 depend upon the relative strengths of T_{R2} and T_{R1}. Our present intention is not to present such a finished model of action, but to present certain guiding ideas in reference to variables that have been neglected in the traditional conception of decision-making and performance.)

We may substitute for T_{R2} and T_{R1} a general conception of the determinants of each of these tendencies:

$$R_2 \text{ occurs when } (A_2 \times B_2 \times C_2) > (A_1 \times B_1 \times C_1) \tag{5}$$

We may now state two sets of hypotheses about the latency or probability of R_2. The first set of hypotheses refers to the effect on performance of R_2 of variables which are ordinarily considered determinants of T_{R2}, viz. A_2, B_2, C_2. The second set of hypotheses is novel. It refers to the effect on performance of R_2 of variables which determine the strength of T_{R1}, the activity already in progress, viz. A_1, B_1, C_1. These two sets of hypotheses are, at the same time, hypotheses about the determinants of persistence of R_1, the activity in progress when the stimulus for R_2 is presented.

In research which explores the behavioral consequences of individual differences in the strength of basic social motives, we might attempt to create a controllable experimental analog of the following type of life situation.* Consider

---

*The writers, in collaboration with David Birch and Bernard Weiner, are now conducting such an experiment.

the case of a college professor at work at his desk on a difficult paper. Let us suppose he is motivated to achieve. Suddenly there is a call from a colleague and friend announcing that coffee is ready in the coffee room. Let us assume his past experiences in the coffee room have amounted to good fellowship in affiliation with friends. Let R_2 be the act of walking to the coffee room. (We certainly mean to take at least as many liberties with the concept of response as current practice appears to allow.)

$$R_2 \text{ occurs when } T_{R2,aff} > T_{R1,succ} =$$

$$(M_{aff} \times E_{R2,aff} \times I_{aff}) > (M_{succ} \times E_{R1,succ} \times I_{succ}) \qquad (6)$$

In this example, the functional significance of the call from the friend, the environmental stimulus, is mapped into the cognitive expectation that walking to the coffee room will have, as a consequence, affiliation with friends ($E_{R2,\ aff}$), and the incentive value of this goal (I_{aff}). The task at which the professor works is the stimulus situation sustaining an expectancy that work at the task will be followed by success ($E_{R1,\ succ} \times I_{succ}$). The expression "$R_2$ *occurs when* $T_{R2,aff} > T_{R1,succ}$" is meant to imply that the professor will go to the coffee room as soon as the tendency to affiliate ($T_{R2,aff}$) dominates the tendency to achieve ($T_{R1,succ}$) whether this comes about as a consequence of random fluctuations in the strengths of tendencies about a mean value or, possibly, as a consequence of growth in strength of the tendency to affiliate over time.

We may bring to the right side of the expression all those variables which are descriptive of the initial condition of the subject *before* the stimulus of interest (the call to coffee hour) is presented:

$$R_2 \text{ occurs when } (E_{R2,aff} \times I_{aff}) > \frac{(M_{succ} \times E_{R1,succ} \times I_{succ})}{M_{aff}} \qquad (7)$$

When we do this, we begin to see clearly the kinds of systematic assertions we might make as to the kind of personality which will be slow to leave work and the kind of personality which will dash to the coffee room, for both the motive to achieve success (M_{succ}) and the motive to affiliate (M_{aff}) are assumed to be relatively general and stable characteristics of the person. In this situation, latency of R_2 should be proportionate to the strength of achievement motive (M_{succ}) and inversely proportionate to the strength of affiliative motive (M_{aff}). Just the reverse would be expected should we consider the matter of the professor coming back to work from the coffee room.

If we were to adopt the Lewinian term *force* in reference to the directional influence attributed to the environmental stimulus when its effect is mapped into Expectancy x Incentive Value, we might ask: How much environmental force, or inducement, is required to move this person into the coffee room?*

---

*It may be useful here to indicate briefly how the conceptions of the present paper relate to the treatment of similar problems of Lewin. First, we believe that statement (1) given above, although different in some details, is similar in approach to Lewin's basic assumptions. Thus, he says, "the force for a locomotion depends on the need or tension of the person; on those nonpsychological factors which affect the existence of the valence; and on the relative position of the person and the valence" (ref. 13, p. 107). These three determinants of force are analogous to the determinants of motivation: motive, incentive, and expectancy.

The force required is proportionate to the strength of motive sustaining the initial activity and inversely proportionate to the strength of motive which, so to speak, answers the call of the incentive signified by the environmental stimulus.

We call your attention to this last assertion because it illustrates how our current thinking about this problem has been influenced by the striking analogy provided in the Newtonian principles of motion. In asserting that the environmental force, or inducement, required to provoke R_2 is proportionate to the strength of motive sustaining activity in progress, we have said, in effect, that the motive sustaining initial activity functions just like inertial mass, as resistance to change, in the second law of motion. When we study persistence of an activity, we are studying something comparable to the resistance of a piano or a chair to change. Both objects are caught up in the earth's gravitational attraction, but differentially because they differ in mass.

In the assertion that the environmental force, or inducement, required to provoke R_2 is inversely proportionate to the strength of motive which answers the call of the incentive signified by the environmental stimulus, the analogy with the physical concept of gravitational mass is striking. In our conception of the determinants of the newly aroused tendency to walk to the coffee room, we assume that strength of motive, conceived like the mass of an object as a stable property of the person, combines multiplicatively with the incentive value of the expected goal. Gravitational force, you may recall, is attributed to the product of the masses of two objects divided by the square of the distance between them times a constant.

In Newtonian physics, there is only one kind of mass; so the ratio of inertial mass to gravitational mass always equals 1. An equivalent situation exists in psychology (given, of course, the assumptions of the present scheme) when both the initial activity R_1, and T_{R2}, the tendency elicited by a stimulus while R_1 is in progress, are both instrumental to the same kind of goal. In our example of the college professor, if the second tendency, T_{R2}, were also an achievement-related activity, the achievement motive (M_{succ}) would appear as both the "inertial motive" in the numerator and as "gravitational motive" in the demoninator. The occurrence of R_2 would then depend only on the relative strengths of the environmental influences, the product of Expectancy x Incentive, for each of the two tendencies. That is:

$$R_2 \text{ occurs when } (E_{R2,succ} \times I_{succ}) > (E_{R1,succ} \times I_{succ}) \qquad (8)$$

This is the special case that was considered by Norman Feather (refs. 8, 9) in the study that spurred the treatment of more general problems which are considered here.

If we consider intuitively, for a moment, the case of the professor strongly motivated to achieve at his desk, interrupted by a friend's call to the coffee hour, we can appreciate the need to take account of one more influence on the tendency to respond to a stimulus. If this professor is at all like most of us, he will probably continue to think about the problem he was working on even though he has left his desk to go for coffee. What is more, he will probably hasten back to his desk a lot sooner than he might have if he had not been interrupted at an important task. Specifically, we must ask: What do our contemporary conceptions say about the tendency to respond when the environmental stimulus responsible for elicitation of a habit or an expectancy is withdrawn? Does the tendency immediately drop to zero strength? Our conventional Aristotelian principles of performance, which demand the presence of a stimulus to elicit or cause the response tendency,

BLANK PAGE

implicitly say yes. But both Freud, in his assertion that a wish persists until it is satisfied, and Lewin, in his treatment of the effect of interruption prior to goal-attainment on spontaneous recall and resumption of activities, have answered no, the tendency persists until it is satisfied or gradually dissipated.

Here again we have found it very helpful to consider the analogous problem in physics. Aristotle held that an external force was needed to cause and sustain motion. But Galileo and later Newton appreciated that under idealized conditions the motion imparted to an object by an external influence would persist indefinitely. The principle of inertia became the first principle of motion. It was acknowledged that external force did not cause motion but change of motion in objects already having momentum attributable to earlier, but not immediately present, external influences. Of course, under natural conditions physical objects do slow down and stop eventually after the push that got them started is withdrawn because there are frictions, forces opposing the momentum that has been established. There are undoubtedly similar frictions in the domain of behavior. But, if we were to assume that a goal-directed tendency, once aroused, persists until satisfied or dissipated, as we now propose, it would have several important implications. It would begin to yield a picture of an active organism for whom the immediate stimulus situation does not function to *elicit* certain response tendencies from a state of rest but rather selectively strengthens or enhances already active tendencies to respond in certain ways. It would also direct experimental interest to a number of problems that have been neglected, e.g., what factors account for the dissipation of a tendency; how do substitute activities get initiated, etc.

In an effort to incorporate this idea into the principle of performance in studies of achievement motivation, we are guided by the conception of force as the cause of change in momentum in classical mechanics. That is:

$$Force = mv_f - mv_i \quad (9)$$

Where mv_f is final momentum and mv_i is initial momentum (the inertial tendency). In psychology, we focus interest on the final strength of a tendency, T_R. In physics, the final momentum is equal to the initial momentum plus the force which produces acceleration to the final momentum.

$$mv_f = Force + mv_i \quad (10)$$

If the initial momentum, mv_i, is zero, that is, if the object is at rest, then the final momentum is proportionate to the force applied. This was essentially Aristotle's position and is, we think, essentially the conventional view of the effect of the immediate stimulus in the determination of a tendency to respond. Certainly it is implied when we assert (as on an early page) that the tendency to perform a response is a multiplicative function of ***Motive***, an ***Expectancy aroused by immediate situational cues***, and the ***Incentive Value*** of the expected goal. Now, however, given the assumption that a previously aroused and unsatisfied tendency to attain some goal persists, we would restate the conception of the immediate determinants of strength of motivation to perform a response this way:

$$T_{R,G} = (Motive_G \times Expectancy_{R,g} \times Incentive\ Value_g) + {}^{*}T_G \quad (11)$$

In this assertion, ${}^{*}T_G$ refers to the persistent unsatisfied motivation, or what might be called, the inertial tendency. In using the capital G as subscript we mean to imply this: when a person, like the professor in our example has had motivation to achieve aroused in the performance of a particular task and then is interrupted before he has successfully completed the task, what persists is his general tendency to achieve, $^{*}T_G$. We mean to imply that this persistent tendency will be added to any subsequently aroused tendency to perform any response which is expected to have, as a consequence, success. Thus, to take another example, sexual motivation aroused in reference to one particular activity and one particular goal object, if unsatisfied, will persist and influence any other instance of sexual activity regardless of the activity or the goal object. It is, in other words, the persistence of a particular kind of motivation that we have sought to capture. How long a tendency once aroused by a stimulus will persist given the frictions in the nervous system, whatever they are, is an empirical question. Our intention is to call your attention to the fact that in the past we have implicitly assumed there is no inertial tendency.

The experimental facts in research on achievement motivation which have seemed to require this kind of assumption about the persistence of motivation are those which show that experimental arousal of achievement motivation by achievement-oriented instructions and performance of tasks immediately prior to the writing of thematic apperceptive stories tends to increase *by a constant amount* the average frequency of achievement-related imaginative responses to particular pictures obtained in a control condition [see McClelland, Atkinson, Clark, and Lowell (ref. 14, p. 204)]. We interpret this to mean that the tendency to achieve is aroused and specifically expressed in performance of preliminary tasks employed to induce motivation, but because there is no knowledge of results, no experience of success, the general tendency persists and immediately afterwards influences the strength of achievement-related response to each of several pictures used as eliciting stimuli, as stated in expression (11) above. Consistent with this view is the fact that when achievement motivation is aroused and subjects are given false norms to induce feelings of success in reference to tasks they have performed immediately before a thematic apperception test is administered to assess the strength of achievement motivation, the level of achievement-related response to the pictures is comparable to the control condition (ref. 14, p. 184). We mention these results merely to indicate that the suggested inclusion of inertial tendencies in formal conceptions of the determinants of decision and performance is not without some empirical foundation in recent research, as well as in the classic studies of the Zeigarnik effect.*

---

*In pursuing the analogy between motivational tendencies and Newtonian conceptions of physical force, we follow a course that Lewin explicitly considered but rejected as unpromising. He wrote: "In physics, force is related to acceleration rather than to velocity...To correlate acceleration to force would imply that in psychology as well as in physics inertia exists; that means that locomotion would go on with the same velocity after the forces which initiated the locomotion would cease to exist...In my mind, such an assumption cannot be made in psychology. However, the proof of this negative statement is not as easy as it might seem.... It seems to me within the reach of present-day psychology to decide these questions" (ref. 13, pp. 148-150). It should be noted, of course, that Lewin was greatly interested in persistence of motivation. In his scheme, the concept of tension accounts for persistence. In ours, the postulate of inertia in reference to motivation accomplishes the same objective. Only further conceptual analysis, together with empirical research, can decide between these two approaches. Of critical importance will be the treatment of the phenomena of "consumption," which Lewin coordinated to the release of tension. Unfortunately, the analysis of this problem in terms of the present conception, would take us too far afield for present purposes.

Let us return to the college professor at work at his desk when the friend calls him to coffee. We now introduce the concept of inertial motivation, or inertial tendency, ${}^{*}T_G$, in this example:

$$R_2 \text{ occurs when } T_{R2,aff} > T_{R1,succ} =$$

$$(M_{aff} \times E_{R2,aff} \times I_{aff}) + {}^{*}T_{aff} > T_{R1,succ} =$$

$$(E_{R2} \times I_{aff}) > \frac{T_{R1,succ} - {}^{*}T_{aff}}{M_{aff}} \qquad (12)$$

What we have added to our earlier statement is the implication that, all other things equal, the external force, or inducement, required to get the professor away from his desk and to the coffee room will be less the greater the magnitude of ${}^{*}T_{aff}$, his persistent unsatisfied tendency to affiliate.

Perhaps a more general statement, to end with, will call attention to what we consider the neglected variables in contemporary formal conceptions of decision and performance. Let R_1 and R_2 refer, respectively, to the initial activity (activity in progress) and the response to a stimulus presented while R_1 is occurring. Let M_A refer to one class of goal objects or activities; I_a to the incentive value of a particular goal of that class; M_B to another class of goal object or activities; and I_b to the incentive value of a particular goal of that class.

$$R_2 \text{ occurs when } T_{R2,B} > T_{R1,A} =$$

$$(M_B \times E_{R2,b} \times I_b) + {}^{*}T_B > (M_A \times E_{R1,a} \times I_a) + {}^{*}T_A =$$

$$(E_{R2,b} \times I_b) > \frac{[(M_A \times E_{R1,a} \times I_a) + {}^{*}T_A] - {}^{*}T_B}{M_B} \qquad (13)$$

This, to our way of thinking, represents a relatively simple instance of the fundamental decision problem--to continue in the initial activity or to initiate an alternative, where there is only one kind of rewarding consequence for each activity and hence only one kind of motive involved in the determination of each tendency. Traditionally, experimental psychology has focused primary interest upon the effect of the controlled environmental stimulus. This is what remains on the left when we bring to the right side of the expression all of the variables which constitute a description of the state of the subject before the stimulus of interest is presented. On the right side are represented the classes of variables that need to be given more articulate representation in our conception of decision and performance. The quantity $[(M_A \times E_{R1,a} \times I_a) + {}^{*}T_A]$ is one way of representing the determinants of the strength of the tendency sustaining activity in progress, it being assumed that the subject of study is always doing something. The possibility that the tendency we expect to arouse by the environmental stimulus may already be active is represented by ${}^{*}T_B$, what we have called the inertial tendency. And in the denominator, M_B represents the latent attribute of personality which interacts with the subsequent environmental stimulus in the determination of T_{R2}.

We call attention to two neglected variables in contemporary conceptions of the determinants of decision and performance: first, *activity already in progress* when the stimulus of interest is presented; second, the possibility that the particular tendency to be aroused by a stimulus may already be active, which we have called *the inertial tendency*. Both of these suggestions are critical of the implicit assumption in most research that the subject is at rest when we begin to study him. We all agree, do we not, that the subject we study is at least as active as a rolling ball?

REFERENCES

1. Atkinson, J. W., "Motivational Determinants of Risk-Taking Behavior," Psychological Review, Vol 64, pp. 359-372, 1957.

2. Atkinson, J. W. (Editor), Motives in Fantasy, Action, and Society; A Method of Assessment and Study, D. Van Nostrand, Princeton, N. J., 1958.

3. Atkinson, J. W.,and G. H. Litwin, "Achievement Motive and Test Anxiety Conceived as Motive to Approach Success and Motive to Avoid Failure," Journal of Abnormal & Social Psychology, Vol 60 , pp 52-63, 1960.

4. Birch, D., E. Burnstein, and R. A. Clark, "Response Strength as a Function of Hours of Food Deprivation Under a Controlled Maintenance Schedule," Journal of Comparative & Physiological Psychology, Vol 51, pp. 350-354, 1958.

5. Brown, J. S., The Motivation of Behavior, McGraw-Hill, New York, N. Y., 1961.

6. Cartwright, D.,and L. Festinger, "A Quantitative Theory of Decision," Psychological Review, Vol 50, pp. 595-621, 1943.

7. Edwards, W., "The Theory of Decision Making," Psychological Bulletin, Vol 51, pp. 380-417, 1954.

8. Feather, N. T., "The Relationship of Persistence at a Task to Expectation of Success and Achievement-Related Motives," Journal of Abnormal & Social Psychology, Vol 63, pp. 552-561, 1961.

9. Feather, N. T., "The Study of Persistence," Psychological Bulletin, Vol 59, pp. 94-115, 1962.

10. French, E. G., and F. H. Thomas, "The Relation of Achievement Motivation to Problem-Solving Effectiveness," Journal of Abnormal & Social Psychology, Vol 56, pp. 45-48, 1958.

11. Hull, C. L., Principles of Behavior, Appleton-Century, New York, N. Y., 1943.

12. Hull, C. L., A Behavior System, Yale University Press, New Haven, Connecticut, 1952.

13. Lewin, K., "The Conceptual Representation and Measurement of Psychological Forces," Contributions to Psychological Theory, Vol I, No. 4, Duke University Press, Durham, N. C., 1938.

14. McClelland, D. C., J. W. Atkinson, R. A. Clark, and E. L. Lowell, The Achievement Motive, Appleton-Century-Crofts, New York, N. Y., 1953.

15. McDougall, W., Outline of Psychology, Scribners, New York, N. Y., 1923.

16. Pavlov, I. P., Conditioned Reflexes, Oxford University Press, London, 1927.

17. Pribram, K. H., "A Review of Theory in Physiological Psychology," in P. R. Farnsworth, (Editor), Annual Review of Psychology, Vol 11, pp. 1-41, Annual Reviews, Inc., Palo Alto, California, 1960.

18. Rotter, J. B., Social Learning and Clinical Psychology, Prentice-Hall, New York, N. Y., 1954.

19. Skinner, B. F., The Behavior of Organisms: An Experimental Analysis, Appleton-Century-Crofts, New York, N. Y., 1938.

20. Spence, K. W., Behavior Theory and Conditioning, Yale University Press, New Haven, Connecticut , 1956.

21. Thorndike, E. L., Animal Intelligence, MacMillan, New York, N. Y., 1911.

22. Tolman, E. C., Purposive Behavior in Animals and Men, Appleton-Century, New York, N. Y., 1932.

23. Watson, J. B., Psychology from the Standpoint of a Behaviorist, Lippincott, Philadelphia and London, 1919.

24. Woodworth, R. S., Dynamic Psychology, Columbia University Press, New York, N. Y., 1919.

DISCUSSION OF ATKINSON'S PAPER

BERLYNE (Discussant). Those of us who have participated in today's program are people who have not been working in the mainstream of current decision theory, and I suppose that we were invited to participate in this symposium in order that we could point out what is wrong with current decision theory, and suggest ways of improvement. Appreciating very much the privilege of attending the symposium, we are, of course, aware of our responsibilities as guests. We feel an obligation to oppose, if only out of politeness, and we must be very careful not to express agreement too bluntly. Now from this point of view, I might say that Dr. Atkinson has been rather less tactful than I, since he has put forth a model which has a good deal in common with a prevalent kind of decision theoretic model, and which is reconcilable with it without too much difficulty.

This kind of model, which depends on the concepts of expectation or subjective probability of outcome, on the one hand, and utility of outcome, on the other (perhaps more generally one could say motivational variables), has a long and respectable history. Its immediate forebearers are Tolman and Lewin, but it goes back to the philosophical psychology of previous centuries. It has various obvious virtues. First of all, it does make contact with normative models, because the variables to which it ascribes actual behavior are the very variables that people ought to take into account if they are to behave rationally. It also makes contact with several disciplines that are concerned with descriptive models, particularly economics and decision theory; and yet in these very virtues there may lie danger, because these other disciplines I mentioned are ones that concentrate on extremely complex forms of behavior in human beings. Now, in studying the most advanced forms of behavior of which the human organism is capable, we ought, as I suggested this morning, to bear in mind the relations between these of behavior and the simplest forms of behavior found in the human being and sub-human animals. But one mistake we can make, and some people do make, is to develop models that slur over these distinctions. This can occur in either of two ways. We can either base a model on obvious facts about complex behavior, and apply it to simple behavior, or do the opposite as some of the earlier and more naive forms of behaviorism and reflexology did.

I am not suggesting that Dr. Atkinson is guilty of these mistakes. I am suggesting that having put forth a model like this, and having demonstrated its fruitfulness in research, he is sooner or later going to come up against these dangers, and he is therefore going to face some problems which can be avoided if we are aware of the danger.

Take this concept of expectancy. It is possible to define expectancy in various ways. At one extreme we can define expectancy in such a broad way that virtually any behavior is an illustration of it. This seems to be the line that Tolman took. At the other extreme, we can define expectancy so narrowly, that we would only recognize such a thing in a human being, who is aware of what is expected, that is to say, who is capable of stating what he is expecting, and how strongly perhaps, or in verbal or some other symbolic form can answer questions about expectancy. In between these are various compromises. We can recognize expectancy in lower animals, at least in animals like cats, dogs and rats, without stating that all of the behavior of which these animals are capable depends on expectancy. We might have to apply more stringent criteria, such as the possibility of being surprised when something turns out different from what is expected, or the presence of preparatory and avoidant responses before what is expected is due. On this thought, I would like to suggest this compromise solution: that we reserve these intellectualistic sounding terms like "expectancy" and "probability" for forms of behavior, which, even though they may not represent the peak of intellectual achievement, are a bit more complex than a great deal of animal behavior, or, for that matter, a great deal of human behavior.

Another question that Dr. Atkinson's paper raises, is the question of the role of variables that are measured principally through verbal operations. Intervening variables also always raise logical problems, which in human behavior we can short-cut very often; we can measure a thing like expectancy, for example, by asking the subject questions and recording his verbal answers. Certainly this has been done with utility, as we know from the vast literature on the problems of scaling utility; several writers have had a great deal of success, but they have achieved success through asking people verbally to indicate verbally which object they prefer of a certain set. The question still remains of the relation between utility, as expressed verbally, and reinforcement or reward value as it applies to instrumental responses. I am aware that there are experimental results showing that what people say they like has something to do with how strongly they strive to obtain it, but precisely how close the relation is between reinforcement value and expected utility, is a question that is by no means resolved definitively. We have lately been obtaining data on a rather different but, I think, analogous question. If you show people different figures and ask them which they prefer, which of course is an age-old technique for experimental esthetics, is the figure they prefer the figure they spend more time looking at spontaneously? Our data suggests that it is not; in fact, on the contrary, there seems to be a tendency for them to say that they prefer those figures they spend less time looking at. This, of course, is not quite the same as the problem of the relation between verbally expressed utility and reinforcement values, but it might, I suggest, strike a cautionary note.

Dr. Atkinson has suggested some of the aspects that current decision theory neglects, and they overlap considerably with some of the points I was trying to make myself. Perhaps I could re-echo some of these points as they appear from his contribution. He has pointed out for example that decisions do not occur in a vacuum. A person isn't sitting still until he makes a decision: decision processes have antecedents. In most experimental work, or, for that matter, most theoretical work on decision making, either the subject is observed and considered from the point at which he already has a set of decisions and some information about outcomes, or, he is artificially given a set of decisions and information about outcomes by the experimenter. Dr. Atkinson's model enables us to study the question of where people, in real life situations, get these sets of decisions and expectations regarding outcomes from. Then, he has pointed out, as I tried to point out myself, that when people are making decisions, they are not dispassionately contemplating a set of alternative possibilities; they are holding in readiness responses corresponding to all the possibilities that they have to recognize. These responses are in competition with each other, and this leads to stress and strain, and affectivity. Finally, if there is response competition, or competition between tendencies corresponding to responses, there is very often a bias in favor of a particular response which must be overcome if that response is not the correct or optimal one, and if the subject is to arrive at the optimal response. There may be a bias in favor of what the subject has just been doing, or there may be biases arising from some other source. Whatever it is, they complicate the picture of decision making considerably.

ATKINSON. I might take just a minute to remark on what Dr. Berlyne said. I would agree with him, I think in his distrust of verbal reports as adequate means of measuring the strength of expectancies, or subjective probabilities, or of anticipated satisfaction in some consequence of action. I think the terms should be used in reference to observations that could be made of actions under much more controlled circumstances, given our general awareness of how complexly determined any statement about one's self is. I have in mind such things as inferring the strength of motive, or expectancy, or incentive from knowledge of antecedents or from actions, including imaginative actions, under fairly controlled circumstances.

The term "expectancy" has that terrible history of being associated with the mentalism that psychology has grown out of. I would prefer to use the term expectancy in reference to habit, linking a particular response and its consequence. In the early papers of Hull in 1930 and 1931, he made a very plausible case for the notion that the stimulus returns from a response might be conditioned to a consequence. He called it a "goal response", such that on a future occasion there might be an anticipatory goal reaction, specifically favoring some particular response. I find it congenial to make this kind of translation, but I do like the term "expectancy" in reference to this associative link because classically now, I think it is fair to say, the term labels the association that Tolman called attention to, in contrast to the association between the environmental stimulus and the compulsion to act which has classically been called habit.

I find no major disagreement with anything else that Dr. Berlyne has said. I do think that the kind of formulation that is evolving in our studies of human motivation is congenial to the mathematical theories of decision that we have been talking about, but there are facts that have suggested that activity in progress and persistence of previously aroused activities need to be introduced somewhere in the principle that we would use to try to explain choice and performance.

LANZETTA. I have three questions that I am curious about. First, I am wondering if your model is really, in many ways an *SEU* model (which you sort of suggested yourself), if you make the assumption that motivation times incentive value is really what people mean by utility? I think if anyone talked about the antecedents of utility they probably would include both something about the desirability of the object, or outcome, and how much of it, that is some quantitative aspect of the outcome. This has been suggested very often: the utility for money being in part a function of the amount of money you have; the utility for a glass of water or for food, certainly being partly of a function of how much food there is, or what the food is, as well as the state of hunger of the organism at the moment. So that is one question: In what sense is this model an *SEU* model, and in what sense is it different?

The second question has to do with the postulation of the notion of inertia for motivation. It seems to me that an alternative possibility (and in a way you almost imply this yourself) is to assume the stimuli are still there, but that they are symbolic. Mediating responses, are in fact, the stimuli which maintain, continue to evoke, or elicit the motivational state. You implied this yourself when you said "On his way to the coffee room he is continuing to think about the goal," or about the responses he was performing. In some sense these stimulus conditions are still represented, and that allows us not to have to postulate the notion of inertia, but more than that, it suggests that to the extent that he does continue to think, that these symbolic stimuli are still present, and to that extent the motivation remains higher.

The third question: I am very unclear about the distinction between $*T_G$ and M_G. I think you said that you would see $*T_G$ as essentially being represented by the tendency to perform the responses which are goal directed, that are relevant to this class of motivations. I guess I would ask, how else would you define M_G except in those terms, as a tendency to perform those responses which are directed toward the goal?

ATKINSON. Let me take the second question first. It certainly is true that the stimulus-response model does better than contemporary decision theory in handling maintained activities after the environmental stimuli that might have initiated them are withdrawn, because stimulus-response theory has, from the outset, been rather articulate about the power of internal stimuli, drive stimuli, proprioceptive returns from responses, as well as environmental stimuli. So one can develop a useful picture of a set of mediating responses, and this is done, I think, within the stimulus-response orientation. When I want to think that way, I have no trouble at all about thinking that the smell of a good hamburger might elicit an anticipatory goal response within me, and once it got going the stimulus return from it, which has been conditioned to it, would sustain it for a while even after the smell of the hamburger is gone. I might, in fact, then respond to another food with a shorter latency than I would have otherwise. No problem. But I direct my intentions to you who want to develop a mathematical conception, mapping and describing the observed behavior, the preferences of an organism, and the time measurement involved in these preferences. If we are to do the job from the outside, then we have to have, as part of our descriptive machinery, references to the initial condition of the subject: concepts like momentum and mass in physics. These concepts refer, on the one hand, to the latent properties of the subject to respond to whatever the call of the environmental stimulus is: that would be motive, analogous to mass, to my way of thinking; and concepts like momentum, to tell us that the person is not at rest with respect to his striving to achieve, or affiliate, or even to eat, but has attained a certain degree of motivation in that direction by virtue of the actions of previous stimuli, and now, looking at him from the outside, he is really moving in a variety of directions. The immediate environ stimulus enhances his tendency to move in one or another of those directions. It adds to this motivational momentum in the same way that an external force adds to momentum in physics. We can do it either way; I think a physiologically oriented behavior theory should continue to try to construct a reasonable network of mediating stimulus-response connections to do the job. Presumably, current information about the activity of the brain can be coordinated with these conceptual attempts to posit the internal stimuli doing the job. Maybe to make my point clear I should say this: If Aristotle had been around when Galileo proved that heavy and light objects are equally accelerated in free fall, if he had to tackle the problems of momentum and inertia, he might have invented some internal push to take the place of the now absent external push. That's essentially what stimulus-response theory does. In our case it seems legitimate because we do know there is a nervous system. But I don't think it is necessary for a mathematical psychology that wants to be, in the Tolman sense, a pure behaviorism.

Now for the first question in reference to utility. I meant that we have been led to the conception that a relatively general latent disposition, called motive, interacts with the relative value of objects belonging to the same class in a multiplicative way to det rmine subjective attractiveness. We have some evidence in our work on achievement motivation that the subjective attractiveness of success of the particular activity, plotted in relation to increasing incentive value, shows a steeper slope for subjects who are strong in achievement motive than the subjects who were weak. There comes from this the general suggestion that the slope of attractiveness, what Lewin would call valence, and I think what the decision theorist would call utility, might be a useful way of measuring the strength of motive, certainly more useful and hopefully more reliable.

The last point, about the relationship of the general inertial tendency $^{*}T_G$ and M_G. For me, M_G is like mass. It is a latent property of the person, and performs two functions of importance. One, it does answer the call, to our way of thinking, of an environmental incentive in about the way the mass of this object answers the call of the earth's mass producing a tendency to approach the earth.

Secondly, it does sustain activity in progress. Our research on persistence leads us to believe it is useful to think of the strength of motive as the general tendency to resist change, in the same sense that inertial mass is a general tendency to resist change. $^{*}T_G$, the inertial tendency, has the same degree of generality as motive. It refers to a class of instances, and my notion is, simply stated, that it should be added, as a constant, to any subsequent elicitations of the tendency to perform a particular response to get to a particular consequence or goal belonging to that class. My intuition here is essentially that the specific tendency be thought of as force, and that general goal direction be thought of as momentum. I think this will become a useful vehicle for explaining why the person who has been aroused to be aggressive, for example, in some situation and has had to inhibit it, might show a disproportionate aggressive reaction to a very remotely related displaced object. This is the kind of intuition: the inertial tendency might attempt to formulate specifically what Freud proposes as the idea that the wish persists until satisfied. I would refer to it as the basis of the content of dreams. I would try to embrace what Lewin tried to handle by introducing something essentially foreign to his system, a tension system. I try to embrace it in the fundamental principle that a tendency toward a goal, once aroused, persists until it is changed, either satisfied by attainment of the goal or further enhanced by the action of the environmental incentive, or, given a reasonable assumption about natural conditions, until it is dissipated by the frictions that are inherent in the behavioral system in the same sense that frictions are inherent in the physical system.

LANZETTA. Jack, I'm afraid I'm still a little bit unclear about $^{*}T_G$. You say it refers to persistent unsatisfied motivation. I guess the question I would like to ask is: as $^{*}T_G$ approaches zero, does M_G change, stay the same or what?

ATKINSON. It would stay the same--my thought being that you can stop the momentum of the rolling ball without changing its mass.

SHUFORD. I would like to comment that it strikes me that the part of decision theory relevant to your formulation here is that which deals with choice under conditions of certainty. In other words, there is no explicit reference to uncertain outcomes of choice.

ATKINSON. Do you mean, does expectancy allow degrees coordinate with subjective probability? I intended that it should.

SHUFORD. Dr. Berlyne referred to the literature on the measurement of utility. I don't see that it is particularly dominated by verbal behavior. I guess I don't understand what the dangers are of verbal responses. Would you explain what this danger is?

ATKINSON. Well, I'll comment quickly and throw the ball across the table. I think that we would all agree that there are better ways of inferring subjective probabilities than just asking the person how certain he feels that the event will occur. In fact, most decision research explicitly assumes this in trying to create conditions under which variations in subjective probability would produce certain preferences rather than asking the person directly.

SHUFORD. Does the difficulty here lie in that we are giving the subject an unspecified test? For example, you might give the subject actual choices between phonograph records, or something like that. We can have him do one of two things: reach out and grab a record, or he can say, "I prefer that record." I don't think you are going to maintain that there is any difference here; so, does it really have to do with verbal behavior, or is it just an incomplete specification of just what we want the subject to do? If that is so, that might apply just as well to "objective" behavior as it does to verbal behavior.

ATKINSON. I don't object to verbal preferences. Maybe Dr. Berlyne has a different attitude on this matter. But I would not trust in my own research, as a measure of expectancy, asking the subject "How many times out of a hundred do you think you can hit this target?" I have done that, that is, obtained verbal reports. I consider it a rather crude way of getting at subjective probability, and it is useful in crude research. I certainly don't mean to include preferences of the sort that have come to be used in many decision studies in my criticism of verbal reports. I mean verbal reports about one's own state. I think this is to be seriously questioned.

BERLYNE. First of all I would like to insist that the methods used by most of the people who have scaled utility are verbal methods in that they involve verbal behavior on the part of the experimenter, and verbal behavior on the part of the subject. It is true that whether the subject points to the object he preferred or names it, doesn't matter: pointing is just as good an example of verbal behavior, according to some definitions such as Skinner's, as is speaking.

The danger of verbal behavior, which certainly is no reason for condemning verbal behavior or methods of measurement depending upon it, arises out of the fact that people don't always do what they say. There is not always a close relationship between what we say and what we do, though it has been established that there is sometimes a close relation. Dr. Atkinson's experiments and others have shown that. All I want to offer for consideration is the point that we have to establish this anew in each case. Even if there is a relation between what people indicate verbally and their non-verbal behavior, the nature of the relation is very often a complex one, and can't be worked out without specific research directed to this problem.

SHUFORD. Maybe you could explain this for me. Suppose I use some technique of objective measurement. I do this again with the same subject and he behaves differently, and so there is not a very good relation between one objective situation and the other. What is the difference? Secondly, I might want to study verbal behavior: is this a legitimate concern?

BERLYNE. We are really in agreement: there is no difference. That is really the point I am trying to make, that verbal behavior does have special properties, but that the problem of making predictions from verbal behavior to some other kind of behavior is, in principle, the same as making predictions from any kind of behavior to any other kind of behavior. What we must be wary of is thinking that verbal behavior has something magical about it. It has something peculiar about it: it is obviously unusual, and has special properties, but it doesn't give us a magical insight into intervening variables that some other kind of behavior doesn't. The only magical aspect of it, as Eriksen has pointed out very well, is its superior information transmitting capacity.

EDWARDS. Probably it's a matter of cohesiveness. Cohesiveness is really not so much a property of behavior as of intellectual tools, and I guess that means the models that we use to think about it. If several different ways of looking at what I thought were the same thing don't behave cohesively, I would interpret this as meaning that the way I was thinking wasn't right; but I don't think I would want to think of any particular one of these ways as being either more or less valid than any other. I think maybe I'm agreeing with you.

FLOOR (To Atkinson). It seems that the model that you and others have suggested have in common that they use continuous arithmetic processes. We neglect the fact that a decision often has a figure-ground, discontinuous, reversal property. For example, after the person has made a decision, there is a reorganization, suddenly, as if the matrix of water-holes and oil has suddenly been supplanted by a matrix of oil-holes and water. The systems have become somehow disconnected and there has been an abrupt change such as occurs when the hag becomes the young woman in the common figure reversal example. I think one of the things about decison that make it a thing of interest, as Professor Toda has suggested, is that following a decision there is a restriction of openness to information other than the very particular information relevant to that decision. I wonder how the kind of model you suggested would handle this discontinuity of process?

ATKINSON. Let me comment in general first. When a decision is made, certainly the action undertaken is coordinated with the strongest tendency. The issue then is what happens to the other tendencies that are involved in the decison. I'm merely accepting Freud's argument, and all the clinical evidence, that perhaps you would have some dreams for a while until there was dissipation. Furthermore I am, in another way, trying to speak to the problem of post-decisional dissonance. After a tough decision is made, I think perhaps what Festinger has referred to as post-decisional dissonance would be conceptualized here as the inertial tendency of the alternatives that are not satisfied. Something has to be done with them, and many different people now are proposing ways of talking about what is done with these tendencies that do persist for a short time. In reference to Dr. Toda's presentation, after a decision is made we know one thing--that the strongest of the competitors won out; and my thought about this, it's hardly a model, I won't dignify it by calling it that, say, my "refined intuition," is that if there is another environmental stimulus intruding after a decison has been made, it will have to be a relatively strong one, that is, its motivational significance would have to be relatively strong, for this whole process to become conscious. My conception is that the brain is constantly making decisons and weak possibilities are not intruding at all as conscious processes. I rather think that what we tend to focus on, when we talk about decision, is conscious, deliberative, thoughtful decision. The implication in this kind of a conception, is that that would occur when the alternatives are so strong that there is a severe conflict, and this is what dominates consciousness. When activity is in progress, while the competing alternative has weak motivational significance, the brain would handle it without any conscious awareness. This is my intuition--that there is a constant decision going on, and a conscious decision, what we tend to think of as decision, occurs when there must be a delay of some sort before actions may occur. One of the assertions state that probability of R_2 is partly a function of the relative strength of two tendencies; if they are nearly equal in strength, then I would predict long decision time. I think the only difference, and I asked Dr. Toda about this yesterday, is that this kind of fungus-eater wouldn't stop, there would be a decision process going on, and he wouldn't stop until the competing alternative was suffiently strong that it overcame the tendency in progress and then he would stop and there would be conscious deliberation of alternatives.

UNCLASSIFIED
Security Classification

DOCUMENT CONTROL DATA - R&D

(Security classification of title, body of abstract and indexing annotation must be entered when the overall report is classified)

1. ORIGINATING ACTIVITY (Corporate author) Aerospace Medical Research Laboratories, Aerospace Medical Division, Air Force Systems Command, Wright-Patterson Air Force Base, Ohio	2a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED 2b. GROUP N/A

3. REPORT TITLE

PREDECISIONAL PROCESSES IN DECISION MAKING: PROCEEDINGS OF A SYMPOSIUM

4. DESCRIPTIVE NOTES (Type of report and inclusive dates)

Symposium was held at Wright-Patterson Air Force Base in April 1962

5. AUTHOR(S) (Last name, first name, initial)

6. REPORT DATE	7a. TOTAL NO. OF PAGES	7b. NO. OF REFS
December 1964	190	141

8a. CONTRACT OR GRANT NO. b. PROJECT NO. 7183 c. Task No. 718303 d.	9a. ORIGINATOR'S REPORT NUMBER(S) AMRL-TDR-64-77 9b. OTHER REPORT NO(S) (Any other numbers that may be assigned this report)

10. AVAILABILITY/LIMITATION NOTICES

Qualified requesters may obtain copies of this report from DDC.
Available, for sale to the public, from the Clearinghouse for Federal Scientific and Technical Information, CFSTI (formerly OTS), Sills Bldg, Springfield, Virginia 22151.

11. SUPPLEMENTARY NOTES	12. SPONSORING MILITARY ACTIVITY
	Aerospace Medical Research Laboratories, Aerospace Medical Division, Air Force Systems Command, Wright-Patterson AFB, Ohio

13. ABSTRACT

The proceedings of a symposium on "Predecisional Processes in Decision Making" are presented. Predecisional processes may be characterized as the search for, the acquisition of, and the evaluation of information prior to the choice of a course of action. The objectives of the symposium were: an assessment of the adequacy of present decision theories in dealing with human decision making behavior; an assessment of other approaches to decision making situations; and an analysis of predecisional processes. Seven papers are presented that analyze this area from several different theoretical viewpoints. Transcripts of the discussions following each paper are also included.

DD FORM 1 JAN 64 1473

AF-WP-B-AUG 64 400

UNCLASSIFIED
Security Classification

UNCLASSIFIED
Security Classification

14 KEY WORDS	LINK A		LINK B		LINK C	
	ROLE	WT	ROLE	WT	ROLE	WT
Decision making (psychology and psychometrics)						
Psychology						
Reasoning						
Human engineering						
Information processing						

INSTRUCTIONS

1. ORIGINATING ACTIVITY: Enter the name and address of the contractor, subcontractor, grantee, Department of Defense activity or other organization *(corporate author)* issuing the report.

2a. REPORT SECURITY CLASSIFICATION: Enter the overall security classification of the report. Indicate whether "Restricted Data" is included. Marking is to be in accordance with appropriate security regulations.

2b. GROUP: Automatic downgrading is specified in DoD Directive 5200.10 and Armed Forces Industrial Manual. Enter the group number. Also, when applicable, show that optional markings have been used for Group 3 and Group 4 as authorized.

3. REPORT TITLE: Enter the complete report title in all capital letters. Titles in all cases should be unclassified. If a meaningful title cannot be selected without classification, show title classification in all capitals in parenthesis immediately following the title.

4. DESCRIPTIVE NOTES: If appropriate, enter the type of report, e.g., interim, progress, summary, annual, or final. Give the inclusive dates when a specific reporting period is covered.

5. AUTHOR(S): Enter the name(s) of author(s) as shown on or in the report. Enter last name, first name, middle initial. If military, show rank and branch of service. The name of the principal author is an absolute minimum requirement.

6. REPORT DATE: Enter the date of the report as day, month, year; or month, year. If more than one date appears on the report, use date of publication.

7a. TOTAL NUMBER OF PAGES: The total page count should follow normal pagination procedures, i.e., enter the number of pages containing information.

7b. NUMBER OF REFERENCES: Enter the total number of references cited in the report.

8a. CONTRACT OR GRANT NUMBER: If appropriate, enter the applicable number of the contract or grant under which the report was written.

8b, 8c, & 8d. PROJECT NUMBER: Enter the appropriate military department identification, such as project number, subproject number, system numbers, task number, etc.

9a. ORIGINATOR'S REPORT NUMBER(S): Enter the official report number by which the document will be identified and controlled by the originating activity. This number must be unique to this report.

9b. OTHER REPORT NUMBER(S): If the report has been assigned any other report numbers *(either by the originator or by the sponsor)*, also enter this number(s).

10. AVAILABILITY/LIMITATION NOTICES: Enter any limitations on further dissemination of the report, other than those imposed by security classification, using standard statements such as:

(1) "Qualified requesters may obtain copies of this report from DDC."

(2) "Foreign announcement and dissemination of this report by DDC is not authorized."

(3) "U. S. Government agencies may obtain copies of this report directly from DDC. Other qualified DDC users shall request through ________________."

(4) "U. S. military agencies may obtain copies of this report directly from DDC. Other qualified users shall request through ________________."

(5) "All distribution of this report is controlled. Qualified DDC users shall request through ________________."

If the report has been furnished to the Office of Technical Services, Department of Commerce, for sale to the public, indicate this fact and enter the price, if known.

11. SUPPLEMENTARY NOTES: Use for additional explanatory notes.

12. SPONSORING MILITARY ACTIVITY: Enter the name of the departmental project office or laboratory sponsoring *(paying for)* the research and development. Include address.

13. ABSTRACT: Enter an abstract giving a brief and factual summary of the document indicative of the report, even though it may also appear elsewhere in the body of the technical report. If additional space is required, a continuation sheet shall be attached.

It is highly desirable that the abstract of classified reports be unclassified. Each paragraph of the abstract shall end with an indication of the military security classification of the information in the paragraph, represented as *(TS)*, *(S)*, *(C)*, or *(U)*.

There is no limitation on the length of the abstract. However, the suggested length is from 150 to 225 words.

14. KEY WORDS: Key words are technically meaningful terms or short phrases that characterize a report and may be used as index entries for cataloging the report. Key words must be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location, may be used as key words but will be followed by an indication of technical context. The assignment of links, rules, and weights is optional.

UNCLASSIFIED
Security Classification

BLANK PAGE