Toward Understanding and Improving Decisions[1]

Paul Slovic

Decision Research, A Branch of Perceptronics
Eugene, Oregon

"The capacity of the human mind for formulating and solving complex problems is very small compared with the size of the problems whose solution is required for objectively rational behavior in the real world--or even for a reasonable approximation to such objective rationality."

Herbert Simon

The development of modern technology has changed radically the hierarchy of needed human skills. Strength and motor performance have become less important. So have perceptual skills although these will never be unimportant. Intellectual skills, especially those of judgment and decision making have become the crucial human elements.

The difficulties of decision making are usually blamed on the inadequacy of the available information; therefore, much technological sophistication has been mobilized to remedy this problem. Computers and other electronic devices supply the decision maker with an abundance of data. However, even the best attainable information often leaves a mass of uncertainties and doubts. It has become evident that a key element in decision making is the ability to interpret and integrate information items, the reliability and validity of which are imperfect. Typically, decision makers are left to their own devices. More likely than not they will proceed in much the same manner that has been relied upon since antiquity--by following their intuition.

In E. A. Fleishman (Ed.), Human Performance and Productivity, in press.

But things have begun to change. Specialists from many disciplines have begun to study information processing and decision making. Their efforts, and mine in this paper, center around two broad questions: "What are decision makers doing?" and "What should they be doing?" The first is a psychological problem, one of understanding how people make decisions and relating this knowledge to the mainstream of cognitive psychology. The second problem is a practical one and involves the attempt to make decision making more effective and efficient.

AIMS AND ORGANIZATION OF THE PAPER

Decision makers of the future will be supplied with many techniques, simple and complex, to help them. The purpose of this paper is to preview these decision-aiding technologies and to outline some of the behavioral considerations underlying their development and their potential for successful application.

The paper begins with an overview of research that describes the shortcomings of unaided decisions. This work has led to the sobering conclusion that, in the face of uncertainty, man may be an intellectual cripple, whose intuitive judgments and decisions violate many of the fundamental principles of optimal behavior. These intellectual deficiencies underscore the need for decision-aiding techniques; the prospects for such techniques are outlined in the second half of the paper.

A NEW IMAGE OF HUMAN CAPABILITIES

The traditional view of human beings' higher mental processes assumes that we are intellectually gifted creatures. Shakespeare referred to man as ". . . noble in reason, infinite in faculties . . . the beauty of the world, the paragon of animals." A more recent expression of this esteem was provided by economist Frank Knight: "We are so built that what seems reasonable to us is likely to be confirmed by experience or we could not live in the world at all." (Knight, 1921, p. 227). Given appropriate information on which to take action, why should such a creature need decision aids?

The answer lies with a rather different picture of human capabilities that has emerged during the computer era from the concern with information processing by people and machines. Miller (1956) in his famous study of classification and coding, showed that there are severe limitations on people's ability to attend to and process sensory signals. About the same time, close observation of performance in concept formation tasks led Bruner, Goodnow and Austin (1956) to conclude that their subjects were experiencing a condition of "cognitive strain" and were trying to reduce it by means of simplification strategies. The processing of conceptual information is currently viewed as a serial process that is constrained by limited short-term memory and a slow storage in long-term memory (Newell & Simon, 1972).

In the study of decision making, too, the classic view of behavioral adequacy, or rationality, has been challenged on psychological grounds. For example, Simon's theory of "bounded rationality" asserts that cognitive limitations force decision makers to construct simplified models in order to cope with their problems. Simon argued that the decision maker

> . . . behaves rationally with respect to this [simplified] model, and such behavior is not even approximately optimal with respect to the real world. To predict his behavior, we must understand the way in which this simplified model is

constructed, and its construction will certainly be related to his psychological properties as a perceiving, thinking, and learning animal (Simon, 1957, p. 198).

Recent laboratory experiments have provided dramatic support for the concept of bounded rationality and have demonstrated its impact in a variety of judgmental and decision making situations. This research, to be reviewed below, is organized around several basic problems of concern to decision makers. First, they need to know what will happen or how likely it is to happen, and their use of information to answer these questions involves them in inference, prediction, probability estimation and diagnosis. They must also evaluate the worth of objects, and this often requires them to combine information from several component attributes of the object into an overall judgment. Finally, they are called upon to integrate their opinions about probabilities and values into the selection of some course of action. What is referred to as "weighing risks against benefits" is an example of the latter combinatorial process.

Studies of Probabilistic Information Processing

Because of the importance of probabilistic reasoning to decision making, a great deal of recent experimental effort has been devoted to understanding how people perceive and use the probabilities of uncertain events. By and large, this research indicates that people systematically violate the principles of rational decision making when judging probabilities, making predictions, or otherwise attempting to cope with probabilistic tasks. Frequently these violations can be traced to the use of judgmental heuristics or simplification strategies (Tversky &

Kahneman, 1974). These heuristics may be valid in some circumstances but in others they lead to biases that are large, persistent, and serious in their implications for decision making.

Misjudging sample implications. One example of the errors people make when dealing intuitively with probabilistic phenomena comes from a study by Tversky and Kahneman (1971) who analyzed the kinds of decisions psychologists make when planning scientific experiments and interpreting their results. Despite extensive formal training in statistics, psychologists usually rely on their educated intuitions when they make decisions about how large a sample of data to collect or whether they should repeat an experiment to make sure their results are reliable. After questioning a large number of psychologists about their research practices and studying the designs of experiments reported in psychological journals, Tversky and Kahneman concluded that these scientists seriously underestimated the error and unreliability inherent in small samples of data. As a result, they (1) had unreasonably high expectations about the replicability of results from a single sample, (2) had undue confidence in early results from a few subjects, (3) gambled their research hypotheses on small samples without realizing the extremely high odds against detecting the effects being studied, and (4) rarely attributed any unexpected results to sampling variability because they found a causal explanation for every observed effect.

Tversky and Kahneman summarized these results by asserting that people's intuitions seemed to satisfy a "law of small numbers," which means that the "law of large numbers" applies to small samples as well as to large ones. The "law of large numbers" says that very large samples will be highly representative of the population from which they are drawn. For the scientists in this study, small samples were also expected to be highly representative of the population. Since knowledge of logic or probability theory did not make the scientist any less susceptible to these cognitive biases, Tversky and Kahneman concluded that the only effective precaution is the use of formal statistical procedures, rather than intuition, to design experiments and evaluate data.

In a related study using Stanford undergraduates as subjects, Kahneman and Tversky (1972) found that many of these individuals did not understand the fundamental principle of sampling--that the variance of a sample decreases as the sample size gets larger. They concluded that "For anyone who would wish to view man as a reasonable intuitive statistician, such results are discouraging."

Errors of prediction. Kahneman and Tversky (1973) contrasted the rules that determined people's intuitive predictions with the normative principles of statistical prediction. Normatively, the prior probabilities or base rates, which summarize what we knew before receiving evidence specific to the case at hand, are relevant even after specific evidence is obtained. In fact, however, people seem to rely almost exclusively on specific information and neglect prior probabilities.

For example, Kahneman and Tversky asked subjects to judge the likelihood that an individual, Tom W., is a graduate student in a particular field of specialization. The judges in this study were all graduate students in psychology. The only information they had available to them was the following brief description written several years earlier by a psychologist on the basis of some projective tests:

> Tom W. is of high intelligence, although lacking in true creativity. He has a need for order and clarity, and for neat and tidy systems in which every detail finds its appropriate place. His writing is rather dull and mechanical, occasionally enlivened by somewhat corny puns and by flashes of imagination of the sci-fi type. He has a strong drive for competence. He seems to have little feel and little sympathy for other people, and does not enjoy interacting with others. Self-centered, he nonetheless has a deep moral sense.
>
> Tom W. is currently a graduate student. Please rank the following nine fields of graduate specialization in order of the likelihood that Tom W. is now a student in that field. Let rank 1 be the most probable choice.
>
> _____ Business Administration
>
> _____ Computer Sciences
>
> _____ Engineering
>
> _____ Humanities and Education
>
> _____ Law
>
> _____ Library Sciences
>
> _____ Medicine
>
> _____ Physical and Life Sciences
>
> _____ Social Science and Social Work

In this study, people ranked the graduate programs on the basis of the similarity between the brief description and typical student in each program. What was remarkable was that the prior probabilities, as determined by the base rates for these graduate programs, had no influence whatsoever upon the judgments. Computer Sciences and Engineering were judged to be the most probable fields for Tom W., even though these fields have relatively few students in them. This is especially surprising considering the fact that the judges recognized the thumbnail personality sketch as having little or no validity. In addition, all of these judges had been exposed to the notion of base-rate prediction in their statistical training, and they used the base rate in a condition where no other information was provided. The important result here is the apparent inability of the judges to integrate the similarity ordering with the base-rate information in a situation where base rate should have been predominant. In other words, the judges knew the description was of low validity and they knew that base rates differed, yet they were unable to put this knowledge into practice. As a result, their judgments did not properly reflect their underlying beliefs.

Another normative principle is that the variance of one's predictions should be sensitive to the validity of the information on which the predictions are based. If validity is not perfect, predictions should be regressed toward some central value. Furthermore, the lower the validity of the information on which predictions are based, the greater the regression should be. Kahneman and Tversky (1973) observed that otherwise intelligent people have little or no intuitive understanding of the concept of regression. They fail to expect regression in many situations when it is bound to occur and, when they observe it, they typically invent complex but spurious explanations. People fail to regress their predictions towards a central value even when they are using information that they themselves consider of low validity.

A third principle of prediction asserts that, given input variables of stated validity, accuracy of prediction decreases as redundancy increases. Kahneman and Tversky (1973) found, however, that people have greater confidence in predictions based on highly redundant or correlated predictor variables, since these tend to agree with one another in their implications. Thus, the effect of redundancy on confidence is opposite what it should be.

Availability bias. Another form of judgmental bias can be traced to the use of the "availability heuristic" (Tversky & Kahneman, 1973) whereby an event is judged likely or frequent if it is easy to imagine or recall relevant instances. Generally, instances of frequent events are typically easier to recall than instances of less frequent events, and likely occurrences are usually easier to imagine than unlikely ones. Thus, availability is often an appropriate cue for judging frequency and probability. However, since availability is also affected by subtle factors unrelated to likelihood, reliance on it may result in systematic overestimation of probabilities for familiar, recent, emotionally salient, or otherwise memorable or imaginable events.

Availability helps explain many distortions in our perceptions of risk. Consider fears about grizzly bear attacks in our national parks. Although many people are concerned about the dangerousness of grizzlies, the rate of injury is only 1 per 2 million visitors and the rate of death is even lower (Herrero, 1970). Sensational media reports contribute to the imaginability of death at the claws of an enraged grizzly but the media ignore the multitude of favorable public experiences. The motion picture, "Jaws," has done a similar service for the availability (and perceived likelihood) of shark attacks. Some nuclear power proponents feel that the risks of that technology are exaggerated in the public's eye because of excessive media coverage and association with the vivid, imaginable, memorable dangers of nuclear war. As Zebroski (1976) notes, "fear sells;" the media dwell on potential catastrophes, not on the successful day-to-day operations of a power plant.

Availability bias is further illustrated in a study by Lichtenstein, Slovic, Fischhoff, Layman & Combs (1978) which found that (1) the probabilities of dramatic, well-publicized events such as botulism, tornadoes, motor vehicle accidents, homicides, and cancer were overestimated and (2) unremarkable or less dramatic events such as asthma, diabetes, and emphysema were underestimated. In addition to demonstrating availability bias, this study shows that intelligent individuals do not have valid perceptions about the frequency of hazardous events to which they are exposed.

Anchoring bias. Bias also occurs when a judge attempts to ease the strain of processing information by following the heuristic device of "anchoring and adjustment." In this process, a natural starting point or anchor is used as a first approximation to the judgment. This anchor is then adjusted to accommodate the implications of additional information. Typically, the adjustment is crude and imprecise and fails to do justice to the importance of additional information. Recent work by Tversky and Kahneman (1974) demonstrates the tendency for adjustments to be insufficient. They asked subjects questions such as "What is the percentage of people in the U.S. today who are age 55 or older?" Subjects were given starting percentages that were randomly chosen and were asked to adjust these percentages until they reached their best estimate. Because of insufficient adjustment, subjects whose starting points were high ended up with higher estimates than those who started with low values.

Application of the anchoring and adjustment heuristic is hypothesized to produce a bias that occurs when people attempt to calibrate the degree to which they are uncertain about an estimate or prediction. Specifically, in a number of studies subjects were given almanac questions such as the following:

How many foreign cars were imported into the United States in 1968?

(a) Make a high estimate such that you feel there is only a 1% probability the true answer would exceed your estimate.

(b) Make a low estimate such that you feel there is only a 1% probability the true answer would be below this estimate.

In essence, the person is being asked to estimate an interval such that there is a 98% chance that the true answer will fall within the interval. The spacing between the high and low estimates is an expression of the person's uncertainty about the quantity in question. We cannot say that this single pair of estimates is right or wrong. However, if the person were to make many such estimates or if a large number of persons were to answer this question, we should expect the range between upper and lower estimates to include the truth about 98% of the time--if the subjective probabilities were unbiased. What is typically found, however, is that the 98% confidence range fails to include the true value from 25 to 40% of the time, across many subjects answering many kinds of almanac questions (Lichtenstein, Fischhoff & Phillips, 1977). In other words, subjects' confidence bands are much too narrow, given their state of knowledge. This bias persists even when subjects are given feedback about their overly narrow confidence bands and are urged to widen the bands on a new set of estimation problems.

These studies indicate that people believe they have a much better picture of the truth than they really do. Why this happens is not entirely clear. Slovic (1972) hypothesized that people approach these problems by searching for a calculational scheme or algorithm by which to make a best estimate. They may then adjust this estimate up and down to get a 98% confidence range. For example, in answering the above question, one might proceed as follows:

> I think there were about 180 million people in the U.S. in 1968; there is about one car for every three people thus there would have been about 60 million cars; the lifetime of a car is about 10 years; this suggests that there should be about 6 million new cars in a year but since the population and the number of cars is increasing, let's make that 9 million for 1968; foreign cars make up about 10% of the U.S. market, thus there were probably about 900,000 foreign imports; to set my 98% confidence band, I'll add and subtract a few hundred thousand cars from my estimate of 900,000.

People's estimates seem to assume that their computational algorithms are 100% correct. However, there are two sources of uncertainty that plague these algorithms. First, there is uncertainty associated with every judgment in the algorithm and there is uncertainty about the algorithm itself. That is, the whole calculational scheme may be incorrect. It is apparently quite difficult to carry along these several sources of uncertainty and translate them intuitively into a confidence band. Once the "best guess" is arrived at as an anchor (e.g. the 900,000 figure above), the adjustments are insufficient in magnitude, failing to do justice to the many ways in which the estimate can be in error.

The research just described implies that our estimates may be grossly in error--even when we attempt to acknowledge our uncertainty. This may have profound implications for many important judgments.

Hindsight bias. A series of experiments by Fischhoff (1975a, b; Fischhoff & Beyth, 1975) has examined the phenomenon of hindsight. Fischhoff found that being told some event has happened increases our feeling that it was inevitable. We are unaware of this effect, however, and tend to believe that this inevitability was apparent in foresight, before we knew what happened. In retrospect, we tend to believe that we (and others) had a much better idea of what was going to happen than we actually did have. Fischhoff (1975b) shows how such misplaced belief that we "knew it all along" can seriously prejudice the evaluations of decisions made in the past and limit our ability to learn by experience. Hindsight bias may also lead us to underestimate the informativeness of facts gleaned from intelligence operations (Fischhoff, 1977) and research studies (Slovic & Fischhoff, 1977).

Overconfidence. An important criterion for evaluating judgments of probability is their degree of calibration. A probability assessor is well calibrated if, for all statements assigned a given probability, the proportion that is true is equal to the probability assigned. For example, if you are well calibrated, then across the many statements to which you assign a probability of .80, 80% of them should turn out to be true. In the past few years, numerous laboratory and real-world experiments have studied calibration (Lichtenstein, Fischhoff & Phillips, 1977; Lichtenstein & Fischhoff, 1977). Across a wide variety of tasks and subjects, one result has consistently occurred. People are overconfident; they tend to estimate much higher probabilities than are warranted. Fischhoff, Slovic and Lichtenstein (1977) studied cases of extreme overconfidence in a task in which people judged the odds that their answers to general knowledge questions were correct. Subjects were wrong frequently on answers they judged almost certain (odds of 50 : 1 or greater) to be correct. Feelings of certainty were so strong that subjects were willing to bet on the correctness of their knowledge. Because of their great overconfidence, the bets they accepted were disadvantageous to them and they lost considerable money. The psychological basis for unwarranted certainty seems to derive from the fact that people reach conclusions about answers by reconstructing their knowledge from fragments of information, much as a paleontologist infers the appearance of a dinosaur from fragments of bone. For example, a person who is "absolutely certain" that the potato is native to Ireland and not Peru may base this judgment on the ready association "Irish potato" and the knowledge that a great potato famine caused mass emigration from Ireland to America. Unfortunately, we appear to be insufficiently critical of the assumptions and reasoning on which our opinions are based--indeed, we typically feel that we have direct access to our knowledge and thus we are unaware that we are making inferences. The potato, by the way, is native to Peru.

Problems of Decision Making

Consider next the integration of information from diverse sources into an overall judgment of value or a decision about a course of action. Here, too, we observe that cognitive limitations lead people to take actions that are inconsistent with their underlying values and opinions.

The failure of one's decisions to reflect personal opinions can be considered one of the most fundamental aspects of nonoptimal decision making. One example of this comes from an experiment by Lichtenstein and Slovic (1973) conducted on the floor of the Four Queens Casino in Las Vegas. Consider the following pair of gambles used in the experiment:

Bet A

11/12 chance to win 12 chips

1/12 chance to ~~win~~ lose 24 chips

Bet B

2/12 chance to win 79 chips

10/12 chance to lose 5 chips

where the value of each chip has been previously fixed at, say, 25¢. Notice that Bet A has a much better chance of winning, but Bet B offers a higher winning payoff. Subjects were shown many such pairs of bets. They were asked to indicate, in two ways, how much they would like to play each bet in a pair. First they made a simple choice A or B. Later, they were asked to assume they owned a ticket to play each bet, and they were to state the lowest price for which they would sell this ticket.

Presumably, these selling prices and choices are both governed by the same underlying quality, the subjective attractiveness of each gamble. Therefore, people should state a higher selling price for the gamble that they prefer in the choice situation. However, the results indicated that subjects often chose one gamble, yet stated a higher selling price for the other gamble. For the particular pair of gambles shown above, Bets A and B were chosen about equally often. However, Bet B received a higher selling price about 88% of the time. Of the subjects who chose Bet A, 87% gave a higher selling price to Bet B, thus exhibiting an inconsistent preference pattern. Grether and Plott (1978), two skeptical economists, recently replicated this study with several variations designed to show that the observed inconsistencies were artifactual. They obtained essentially the same results as were found by Lichtenstein and Slovic.

What accounts for the inconsistent pattern of preferences? Lichtenstein and Slovic concluded that people use different cognitive strategies for setting prices than for making choices. People choose Bet A because of its good odds, but they set a higher price for B because of its large winning payoff. Specifically, it was found that, when making pricing judgments, people who find a gamble basically attractive use the amount to win as a natural starting point. They then adjust the amount to win downward to take into account the less-than-perfect chance of winning and the fact that there is some amount to lose as well. Typically, this adjustment is insufficient and that is why winning payoffs lead people to set prices that are inconsistent with their choices. Because the pricing and choice responses are inconsistent, it is obvious that at least one of these responses does not accurately reflect what the decision maker believes to be the most important attribute in a gamble.

A "compatibility" effect seems to be operating here. Since a selling price is expressed in terms of monetary units, subjects apparently found it easier to use the monetary aspects of the gamble to produce this type of response. Such a bias did not exist with the choices, since each attribute of one gamble could be directly compared with the same attribute of the other gamble. With no reason to use payoffs as a starting point, subjects were free to use any number of strategies to determine their choices. The overdependence on payoff cues when pricing a gamble suggests a general hypothesis to the effect that the compatibility or commensurability between a dimension of information and the required response affects the ease with which that information can be used and, ultimately, its importance in determining the response. This hypothesis received support in an experiment by Slovic and MacPhillamy (1974) who found that dimensions common to each alternative in a choice situation had greater influence on decisions than did dimensions that were unique to a particular alternative. Interrogation of the subjects after the experiment indicated that most did not wish to give more weight to the common dimension and were unaware that they had done so.

The message in these experiments is that the amalgamation of different types of information and different types of values into an overall judgment or decision is a difficult cognitive process and, in our attempts to ease the strain of processing information, we often resort to judgmental strategies that may do an injustice to our underlying values. In other words, even when all the relevant events, probabilities, and outcomes are known and made explicit, as in the gambling situation, subtle aspects of the decision we have to make, acting in combination with our intellectual limitations, may bias the balance we strike among the attributes.

When the decision is not well structured, that is, when all the relevant aspects are not explicitly specified, further difficulties arise. Foremost among these is the neglect of one or more crucial factors whose relevance only becomes apparent, unfortunately, after the decision has been made. An example of this is provided by Birkin and Ford (1973), who examined the after-effects of the "Zero Defects" program. This program, adopted by more than 12,000 industrial firms, attempted to attack the problem of defective workmanship by motivating employees to do the job right the first time. The program was based on the following rationale: "Because of the complexity of today's products and because of the drastic consequences of product failure, management should use all means possible to get a job done right the first time." Once the program was implemented, many firms discovered they could not live with the consequences of making quality a primary goal. As quality rose, productivity declined, production deadlines were missed, and amounts of spoiled and scrapped goods increased. A high percentage of firms dropped the program.

Random error. We're all familiar with the effects of random error in activities that involve motor skills--playing golf is one such activity that comes to mind. Random error is the mysterious lack of control that causes two drives, seemingly executed the same way, to end up in different parts of the fairway. We're less aware that similar lack of control affects our decision making behaviors as well as our golf games. In fact, it's only quite recently that decisions have been studied in a way that illustrates this problem.

Goldberg (1970) described the problem of error and unreliability by noting that:

> He [the judge] 'has his days': Boredom, fatigue, illness, situational and interpersonal distractions all plague him, with the result that his repeated judgments of the exact same stimulus configuration are not identical. He is subject to all those human

frailties which lower the reliability of his judgments below unity. (p. 423).

There are a number of studies demonstrating the presence of random error in the judgments of experts. One of the most significant of these studies was done by Garland (1960) who measured the reliability of radiologists as they attempted to detect the presence of lung disease on X-ray films. Garland found that radiologists changed their minds in about 20% of the cases when reading the same film on two separate occasions.

Another example of inconsistency comes from an unpublished study of expert horserace handicappers, which Bernard Corrigan and I conducted at the Oregon Research Institute. We were interested, not in horserace predictions but in the stresses caused by information overload. Horseracing provided an appropriate context in which to study this. We expect that the results will generalize to any domain in which the integration of large masses of quantitative information is performed by means of skilled human judgment.

Our judges in this study were eight individuals, carefully selected for their expertise as handicappers. Each judge was presented with a list of 88 variables taken from the horses' past-performance charts. The judges were asked to indicate which five variables out of the 88 they would wish to use when handicapping a race, if they were limited to just five variables. They were then asked to indicate which 10, which 20, and which 40 they would use if 10, 20, or 40 items of information were available.

All the handicappers judged each of 45 races under all four information conditions. First they saw five variables and ranked the top five horses in the race in the order they thought the horses would finish. They then received their preselected 10-variable set and reranked the horses. Next they ranked them again using 20 and finally 40 variables. All handicappers had their own personalized set of 5, 10, 20 and 40 variables. Five of the races were repeated at the end of the experiment. By examining a handicapper's two rankings for the same race, we were able to assess the degree of inconsistency in that person's judgment policy.

The results indicated that, on the average, accuracy of prediction was as good with five variables as it was with 10, 20 or 40. However, every handicapper became more confident in the accuracy of the judgments as amount of information increased. Examination of judgments for the repeated races showed that inconsistency increased sharply as the amount of available information increased. With 5 predictors, 22% of the first-place choices were changed on the second replication; with 40 predictors, 39% of the judgments changed. These results should give pause to those who believe they are better off getting as much information as possible prior to making a decision.

Are Important Decisions Biased?

The results described above contradict our traditional image of the human intellect. It is, therefore, important to determine whether these inadequacies in decision making exist outside the laboratory in situations where experts use familiar sources of information to make decisions that are important to themselves and others.

Much evidence suggests that the laboratory results will generalize. Cognitive biases appear to pervade a wide variety of socially important judgments in which intelligent individuals serve as decision makers, often under conditions that maximize motivation and involvement. For example, the subjects studied by Tversky and Kahneman (1974) were scientists, highly trained in statistics, evaluating problems similar to those they faced in their own research. The overdependence on specific evidence and neglect of base rates observed in laboratory studies have also been found among psychometricians responsible for the development and use of psychological tests (Meehl & Rosen, 1955) and among intelligence officers evaluating military information reports (Samet, 1975). The latter based their evaluations primarily on a report's content, neglecting the base-rate reliability of the report's source. Flood-plain residents misjudge the probability of floods in ways readily explained in terms of availability bias (Kates, 1962; Slovic, Kunreuther & White, 1974). A study by Wohlstetter (1962) of American unpreparedness at Pearl Harbor found the U.S. Congress and military investigators guilty of hindsight bias in their judgment of the Pearl Harbor command staff's negligence. A classic case of the "law of small numbers" is the discovery by Berkson, Magath, and Hurn (1940) that aspiring lab technicians were expected by their instructors to show greater accuracy in performing blood cell counts than was possible given sampling variation. These instructors marveled that the best students (those who would not cheat) had the greatest difficulty in producing acceptable counts.

The anchoring and insufficient adjustment that Tversky and Kahneman observed with their almanac questions could well contribute to errors that plague projected cost estimates. For example, one congressional study noted that the cost of major weapon systems was running nearly 50% ahead of original estimates. In one case where the original estimate for six submarine rescue vehicles was $18 million, the actual cost was close to $460 million--a value that most certainly would have been viewed as impossible when the original estimates were made. This gigantic overrun, like many others, was blamed on a failure to foresee development problems. The moral seems to be that there are many ways our estimates can go wrong, and it is difficult to incorporate our uncertainty about these possible sources of error into our judgments.

In case studies of policy analyses, Albert Wohlstetter (1974) found that American intelligence analysts consistently underestimated Soviet missile strength, a bias possibly due to anchoring.

Finally, I'd like to point out a particularly painful example of anchoring and insufficient adjustment from my own experience. A few years ago a colleague and I agreed to write a chapter for a book. After the project was completed, we were rummaging through our correspondence with the book's editor and were rather dismayed to note the string of optimistic projections and broken promises that is illustrated as follows:

History of the Chapter

On this date	We promised it for this date
Sept. 16, 1968	June, 1969
May, 1969	End of July, 1969
Dec., 1969	End of Jan., 1970
Jan., 1970	Apr., 1970
Apr., 1970	End of June, 1970
But we finally sent the first draft	July 24, 1970.

Many of you may have had the same experience, and we can take some small comfort in a study by Kidd (1970) showing that a similar thing happens when the Central Electricity Generating Board in England and Wales attempts to estimate how long it will take to overhaul its equipment.

Comment

One additional implication of the research on people's limited ability to process probabilistic information deserves comment. Most of the discussions of "cognitive strain" and "limited capacity" that are derived from the study of problem solving and concept formation depict a person as a computer that has the right programs but cannot execute them properly because its central processor is too small. The biases due to availability and anchoring certainly are congruent with this analogy. But the misjudgment of sampling variability and the errors of prediction illustrate more serious deficiencies. Here we see that people's judgments of important probabilistic phenomena are not merely biased but are in violation of fundamental normative rules. Returning to the computer analogy, it appears that people lack the correct programs for many important judgmental tasks.

How could it be that we lack adequate programs for probabilistic thinking? Sinsheimer (1971) argues that the human brain has evolved to cope with certain very real problems in the immediate, external world and thus lacks the proper framework with which to encompass many of the conceptual phenomena. Following Sinsheimer's reasoning, it might be argued that we have not had the opportunity to evolve an intellect capable of dealing conceptually with uncertainty. We are essentially trial-and-error learners, who ignore uncertainty and rely predominantly on habit or simple deterministic rules. When we can afford to learn from our mistakes, this may be a satisfactory way to behave. When we cannot, we must look toward decision aids to help minimize errors of judgment.

DECISION AIDS

Research in both laboratory and field settings strongly supports the view of decision processes as boundedly rational. Given this awareness of our cognitive limitations, what sort of techniques will enhance our capacity for making intelligent decisions?

I have found it useful to consider the repeatability of the task when characterizing decision aids. Near one end of what is really a continuum of repeatability are tasks such as selection or rejection of applicants for jobs. The essential structure of each application (e.g., the types of information available) remains nearly the same from case to case, although the specific details of each application will, of course, change. Toward the other end of the continuum are more unique decisions. The decision to build a supersonic commercial airliner exemplifies this type of problem.

Figure 1 depicts my conception of the relationship between decision repeatability and decision-aiding techniques. When decisions are repeatable, they can be handled quite effectively by precise rules or standard operating procedures (SOPs). Although SOPs, such as rules for reordering supplies in an office, have been around for a long time, there are new and powerful variants, bootstrapping and multiattribute utility analysis, that merit discussion here. When predesignated rules are insufficient, computerized information management systems and realistic experience in a simulated decision environment serve as aids. If the decision task is unique, I believe it is important to consider the time available for deliberations prior to action. If the leadtime is long and the decision is important enough, then decision analysis is the relevant aiding technology. If the leadtime is short, I see no recourse other than to rely on educated intuition. These various types of aids will be discussed below.

Insert Figure 1 about here

Aids for Unique Decision Situations

Decision analysis. Decision analysis is a general-purpose technology for making decisions when the stakes are high and both time and resources are ample. The roots of decision analysis can be traced to World War II and the need to solve strategic problems in situations in which experience was either costly or impossible to acquire. The technique developed then was labeled "operations analysis" and later became known as "operations research."

During recent years, a number of closely related offshoots of operations research have been applied to decision problems. These include systems analysis and cost-benefit analysis. Systems analysis is a branch of engineering, whose objective is capturing the interactions and dynamic behavior of complex systems. Cost-benefit analysis attempts to quantify the prospective gains and losses from some proposed action, usually in terms of dollars. If the calculated gain from an act or project is positive, it is said that the benefits outweigh the costs, and its acceptance is recommended (see, for example, the application of cost-benefit analysis to the study of auto safety features by Lave and Weber (1970).

What systems analysis and operations research approaches lacked for many years was an effective normative framework for dealing either with the uncertainty in the world or with the subjectivity of decision makers' values and expectations. The emergence of decision theory provided the general normative rationale missing from these early analytic approaches.

The objective of decision theory is to provide a rationale for making wise decisions under conditions of risk and uncertainty. It is concerned with prescribing the course of action that will conform most fully to the decision maker's own goals, expectations, and values.

Decisions under uncertainty are typically represented by a payoff matrix, in which the rows correspond to alternative acts that the decision maker can select and the columns correspond to possible states of nature. In the cells of the payoff matrix are one set of consequences contingent on the joint occurrence of a decision and a state of nature.

Since it is impossible to make a decision that will turn out best in any eventuality, decision theorists view choice alternatives as gambles and try to choose according to the "best bet." In 1738 Bernoulli defined the notion of a best bet as one that maximizes the "expected utility" of the decision. That is, it maximizes the quantity

$$EU(A) = \sum_{i=1}^{n} P(E_i)U(X_i) \qquad (1)$$

where EU(A) represents the expected utility of a course of action which has consequences $X_1, X_2, \ldots, X_n$ depending on events $E_1, E_2, \ldots, E_n$, $P(E_i)$ represents the probability of the i'th outcome of that action, and $U(X_i)$ represents the subjective value or utility of that outcome.

A major advance in decision theory came when von Neumann and Morgenstern (1953) developed a formal justification for the expected utility criterion. They showed that, if an individual's preferences satisfied certain basic axioms of rational behavior, then that person's decisions could be described as the maximization of expected utility. Savage (1954) later generalized the theory to allow the $P(E_i)$ values to represent subjective or personal probabilities.

Maximization of expected utility commands respect as a guideline for wise behavior because it is deduced from axiomatic principles that presumably would be accepted by any rational person. One such principle, that of transitivity, asserts that, if a decision maker prefers outcome A to outcome B and outcome B to outcome C, it would be irrational for that person to prefer outcome C to outcome A. Persons who are deliberately and systematically intransitive can be used as "money pumps." You can say to them, "I'll give you C. Now, for a penny, I'll take back C and give you B." Since they prefer B to C, they accept. Next you offer to replace B with A for another penny and again they accept. The cycle is completed by offering to replace A by C for another penny; they accept and are 3¢ poorer, back where they started, and ready for another round.

Applied decision theory assumes that the rational decision maker wishes to select an action that is logically consistent with his or her basic preferences for outcomes and feelings about the likelihoods of the events on which those outcomes depend. Given this assumption, the practical problem becomes one of structuring the alternatives and scaling the subjective values of outcomes and their likelihoods so that subjective expected utility can be calculated for each alternative. Another problem in application arises from the fact that the range of possible alternatives is often quite large. Also, each outcome may have multiple facets that must be combined into an overall estimate of worth.

Recently, a methodology called decision analysis has been developed to combine decision theory with the sophisticated modeling of decision problems (i.e., the critical options, events, and consequences) provided by systems analysis (Howard, 1968; 1975). Decision analysis has been applied to such complex and diverse problems as whether or not to modify hurricanes (Howard, Matheson & North, 1972), the selection of experiments for a Mars space mission (Matheson & Roths, 1967), the decision to undergo coronary artery surgery (Pauker, 1976), and the choice of nuclear vs. coal power plants (Barrager, Judd & North, 1976).

A key element of decision analysis is its emphasis on structuring the decision problem and decomposing it into a number of more elementary problems. In this sense, it attempts a simplification process that, unlike the potentially detrimental simplifications the unaided decision maker might employ, maintains all the essential ingredients that are necessary to make the decision and ensures that they are used in a manner logically consistent with the decision maker's basic preferences. Raiffa (1968) expresses this attitude well in the following statement:

> The spirit of decision analysis is divide and conquer: Decompose a complex problem into simpler problems, get your thinking straight in these simpler problems, paste these analyses together with a logical glue, and come out with a program for action for the complex problem. Experts are not asked complicated, fuzzy questions, but crystal clear, unambiguous, elemental hypothetical questions (p. 271).

Decision analysis assumes that all relevant considerations in a decision can be assigned to one or another of four components: initial options, possible consequences, values, and uncertainties. An important tool is the decision tree, which diagrams the stream of uncertain consequences arising from a decision.

Beyond its primary role of serving as a method for the logical solution of complex decision problems, decision analysis has additional advantages as well. The formal structure of decision analysis makes clear all the elements, their relationships and their associated weights that have been considered in a decision problem. Because the model is explicit, it can serve an important role in facilitating communication among those involved in the decision process. With a decision problem structured in a decision analytic framework, one can identify the location, extent, and importance of any areas of disagreement and determine whether such disagreements have any material impact on the indicated decision. In addition, should there be any change in the circumstances bearing on a given decision problem, it is fairly straightforward to reenter the existing problem structure to change values or to add or remove problem dimensions as may be indicated.

It should be emphasized that in no sense does decision analysis replace decision makers with arithmetic or change the role of wise human judgment in decision making. Rather, it provides an orderly and more easily understood structure that helps to aggregate the wisdom of experts on the many topics that may be needed to make a decision, and it aids skilled decision makers by providing them with mathematical techniques to support, supplement, and ensure the internal consistency of their judgments.

It is difficult to convey in a summary such as this the depth of thinking and the logic underlying decision analysis. Any brief description necessarily simplifies the analysis and highlights a chief objection to decision analysis in general--the claim that it oversimplifies the situation and thus misleads. Nevertheless, even those who read a complete analysis may have concerns over its validity. Critics argue that such analyses are inevitably constrained by time, effort, and imagination and must systematically exclude many considerations.

A second major objection to decision analysis is the possibility that it may be used to justify and give a gloss of respectability to decisions made on other and perhaps less rational grounds.

Decision analysts counter these attacks by invoking one of their basic tenets--namely, that any alternative must be considered in the context of other alternatives. What, they ask, are the alternatives to decision analysis, and are they any more immune to the criticisms raised above? The analysts point out that traditional modes of decision making are equally constrained by limits of time, effort, and imagination and are even more likely to induce systematic biases (as illustrated previously). Such biases are much harder to detect and minimize than the deficiencies in the explicit inputs to decision analysis. Furthermore, they argue, if some factors are unknown or poorly understood, can traditional methods deal with them more adequately than decision analysis does? Traditional methods also are susceptible to the "gloss of respectability" criticism noted above. We often resort to expertise to buttress our decisions without really knowing the assumptions and logic underlying the experts' judgments. Decision analysis makes these assumptions explicit. Such explicit data are easy for knowledgeable persons to criticize and the explicitness thus focuses debate on the right issues.

Decision analysts would agree that their craft is no panacea, that incomplete or poorly designed analyses may be worse than no analyses at all, and that analysis may be used to "overwhelm the opposition." It seems clear, however, that the main task for the future is not so much to criticize decision analysis but rather to see how it can be used most appropriately.

Educated Intuition

Decision analysis will require extensive further development before it is ready for use in situations in which unique decisions must be taken with little time for deliberation. Thus, the standard method of decision in these situations will continue to be intuition. Given the pitfalls to which intuitive decisions are susceptible, we have little reason to feel comfortable with this prospect. It would seem desirable to prevent such situations from occurring, whenever possible. Every attempt should be made to foresee contingencies and plan for them in advance. Failing that, conservative decisions, which permit one to take fast corrective action to recover from the inevitable mistakes, would seem advisable.

Since we cannot avoid the necessity of making some important decisions intuitively, we should at least educate decision makers to the pitfalls that await the unwary. For example, one should realize the difficulties of using case-specific information to predict low-base-rate (rare) phenomena and, therefore, should take special precautions to ensure adequate consideration of the base rate. When action is contingent on quantitative estimates that may be susceptible to anchoring bias, the wise decision maker will obtain multiple estimates, based on differing methods, to allow biases to "cancel out." Since feelings of certainty often lead to bold, decisive action, it is important to alert decision makers to the kinds of situations that foster unwarranted confidence. Before taking action in these situations, decision makers should scrutinize the assumptions on which their confidence is based and force themselves to consider scenarios that might make their actions look bad (see, for example, Howard, Merkhofer, Miller and Tani, 1975).

Aids for Repeated Decisions

Bootstrapping. Judgment and decision making have traditionally been viewed as mysterious phenomena, incapable of being described precisely. However, considerable research over the past 15 years has demonstrated that this traditional view is incorrect. The hidden cognitive processes of the judge can be modeled, made explicit, and programmed so that a computer can make judgments that correlate highly with those made by the human. The ability to construct models has important practical consequences. In repeatable decision situations, judges can be replaced by their own models. The benefit from doing this is not merely increased efficiency or freeing the judge for more creative activity. In many cases, the model of the judge makes better predictions than the judge! Dawes (1971) has termed this phenomenon "bootstrapping."

Before discussing bootstrapping in more detail, let's first consider the sorts of models that might be used to simulate the decision maker. These models take two forms, simple and complex. An example of the latter is the simulation by Clarkson (1962) of the portfolio selection process of a bank's trust investment officer. Clarkson followed the officer around for several months and studied his verbalized reflections as he was asked to think aloud while reviewing past and present decisions. Using these verbal descriptions as a guide, the investment process was translated into a sequentially branching computer program. When the validity of the model was tested by comparing its selections with future portfolios selected by the trust officer, the correspondence between actual and simulated portfolios was found to be remarkable good.

Clarkson's work shows that, given patient and intelligent effort, many of the experts's cognitions can be distilled into a form capable of being simulated by a computer. One application of Clarkson-type modeling has been proposed but not yet implemented by researchers at the department of clinical medicine at a leading medical shcool. These researchers are concerned with the difficulty of making decisions with regard to medical tests. In addition to being expensive, the tests are sometimes painful and dangerous. The interpretation of the test results is hindered because they are affected by treatment variables and other aspects of the patient's condition. New tests are continually being developed. As a result of these factors, the average physician often does a poor job of selecting and evaluating tests. It has been proposed that sequential decision trees or flow chart models be developed for the world's leading experts on various sorts of tests--tests for thyroid disorder, liver disease, and so forth. These models could then be programmed into a computer and made accessible to practitioners.

There is yet another approach to modeling--a simpler one that provides less of a sequential analysis and more of a quantified descriptive summary of the way that a decision maker weights and combines information from diverse sources. This approach aims to develop a mathematical model of the decision maker and requires less time and effort on the part of investigator, subject, and computer. It forms a nice compromise between Clarkson's complex, sequentially branching model and the relatively naive approaches of the precomputer era--such as simply asking decision makers how they make their judgments. The rationale behind these mathematical models and techniques for building them are reviewed by Slovic and Lichtenstein (1971).

The basic approach requires the decision maker to make quantitative evaluations of a fairly large number of cases, each of which is defined by a number of quantified cue dimensions or characteristics. A financial analyst, for example, could be asked to predict the long-term price appreciation for each of 50 securities, the securities being defined in terms of cue factors such as their P/E ratios, corporate earnings growth trend, dividend yield, and so forth. The manner in which the analyst weights these various factors can then be described by fitting a linear equation to the judgments.

The resultant equation would be

$$\hat{J}_{pa} = b_1X_1 + b_2X_2 + \ldots b_kX_k \qquad (2)$$

where J_{pa} = predicted judgment of price appreciation; X_1, X_2 . . . X_k are the quantitative values of the defining cue factors (i.e., P/E ratios, earnings, and so forth); and b_1, b_2 . . . b_k are the weights given to the various factors in order to maximize the multiple correlation between the predicted judgments and the actual judgments. These weights are assumed to reflect the relative importance of the factors for the analyst. Equation (2) is known as the linear model.

Psychologists have found linear equations to be remarkably successful in modeling such diverse phenomena as psychiatric and medical diagnoses, and judgments of job performance, graduate school applicants, suicide risk, financial soundness of businesses, price increases of stocks, Air Force cadets, theatrical plays, and trout streams; political scientists have found linear models useful for describing judicial decision processes in workman's compensation and civil liberties court cases (Slovic & Lichtenstein, 1971; Slovic, Fischhoff & Lichtenstein, 1977). U. S. senators have been modeled and their roll-call votes predicted (Wainer, Zill & Gruvaeus, 1973).

More complex, nonlinear, judgmental processes can be modeled by including exponential terms (x^2, x^3, etc.) or cross product terms (e.g., x_1, x_2) into the judge's equation. However, nonlinear processing typically accounts for only a small fraction of the predictable variance in human judgments. Most of the variance is accounted for by linear equations, whose coefficients have provided useful descriptions of the judges cue-weighing policies and have pinpointed the source of inter-judge disagreement and non-optimal cue use (Hammond, Stewart, Brehmer & Steinmann, 1975).

Why do linear models do so well? Dawes and Corrigan (1974) have observed that in most judgment situations (a) the predictor variables are monotonically related to the criterion being judged (or can easily be rescaled to be monotonic) and (b) there is error in the predictors and the judgments. They demonstrated that these conditions practically ensure good fits by linear models.

Now that we've examined the ways that decision makers can be modeled, let's look again at bootstrapping. The rationale behind it is quite simple. As noted earlier in the discussion of random error, human judgment often lacks reliability. Goldberg (1970) observed:

> . . . if the judge's reliability is less than unity, there must be

error in his judgments--error which can serve no other purpose than to attenuate his accuracy. If we could . . . [eliminate] the random error in his judgments, we should thereby increase the validity of the resulting predictions (p. 423).

A model captures the judge's weighting policy and applies it consistently. If there is some validity to this policy to begin with, filtering out the error via the model should increase accuracy. Of course, bootstrapping preserves and reinforces any misconceptions or biases that the judge may have. Implicit in the use of bootstrapping is the assumption that these biases will be less detrimental to performance than the inconsistencies of unaided human judgment.

Bootstrapping has been explored independently by a number of different investigators (Slovic & Lichtenstein, 1971). One particularly noteworthy demonstration comes from a study of a graduate student admissions committee by Dawes (1971). Dawes built a regression equation to model the average judgment of the four-man committee. The predictors in the equation were overall undergraduate grade point average, quality of the undergraduate school, and a score from from the Graduate Record Examination. To evaluate the validity of the model and the possibility of bootstrapping, Dawes used it to predict the average committee rating for his sample of 384 applicants. He found that it was possible to find a cutting point on the distribution of predicted scores such that no one who scored below that point was invited by the admissions committee. Fifty-five percent of the applicants scored below this point, and thus could have been eliminated by a preliminary screening without doing any injustice to the committee's actual judgments. Furthermore, the weights used to predict the committee's behavior were better than the committee itself in predicting later faculty ratings of the selected students. In a cost-benefit analysis, Dawes estimated that the use of such a linear model to screen applicants to the nation's graduate schools could result in an annual saving of about $18 million worth of professional time.

The potential of judgment modeling for facilitating decision making is unlimited. Wherever expertise exists, there is the possibility of modeling it and using the model for training or for constructing automated decision systems.

Other decision rules. The linear regression model describes the weighting system implicit in the decision maker's behavior. One disadvantage to this approach is that the decision maker, perhaps because of cognitive limitations, may not be weighting information in the desired way. Another disadvantage is that it is not always feasible to obtain the large number of judgments necessary for building the model. These difficulties can be overcome by the use of a multiattribute utility (MAU) model that explicitly states the desired weights for each factor in order to produce some overall judgment. For example, one might wish to define the relative importance of variable X to variable Y as 2 : 1 rather than inferring the values from someone's judgments. MAU procedures are gaining widespread acceptance as rule-based methods for combining component dimensions into an overall evaluation. For a more detailed discussion of this methodology see Fischer (1975), von Winterfeldt and Fischer (1975), Slovic, Fischhoff and Lichtenstein (1977) or Keeney and Raiffa (1977).

Information control systems. Of course, not all repeatable decisions can be handled by rules. When the human element is necessary, performance can be facilitated by computer-based information systems for storing, modifying, retrieving, and displaying data and for performing various sorts of symbolic and arithmetic manipulations. One such system called AESOP (An Evolutionary System for On-line Planning) is described by Doughty and Feehrer (1969).

In one experimental test of AESOP involving allocation of tactical aircraft to various missions, planners were required to make decisions which represented an optimal tradeoff between several criteria, including time over target, minimization of use of recycled aircraft, and minimization of use of recycled aircraft, and minimization of total flying time. Performance of planners assisted by AESOP was superior to that of those who were unassisted. AESOP provided no formal procedures or rules to aid the decision maker. However, its concise displays appeared to help planners comprehend the extent to which their resources would be strained and, therefore, enabled them to develop a better "feel" for their plans.

Simulation. One of the most extensively developed methods for sharpening decision performance is that of simulation. Simulation places the decision maker in situations that are similar in certain important aspects to those they are likely to encounter in the real world. Simulation has the advantage of exposing the decision maker to a rich variety of situations in which the consequences of error are not catastrophic. Performance can be evaluated and immediate feedback provided. On the negative side, simulations must be carefully designed to present the critical aspects of the real decision if proper transfer is to be obtained. For further discussion of simulation approaches see Abt (1970), Driver and Hunsaker (1972) and a review by Nickerson and Feehrer (1975).

Future Work

Decision-aiding technologies are still in an early stage of development. Thus, although decision analysis is undoubtedly the wave of the future, many problems need to be resolved before we can reap its full benefits.

First of all, we need to develop techniques for structuring the decision problem. The logic of decision theory cannot be applied until the alternatives, critical events, and outcomes are specified. We need algorithms for accomplishing this and for simplifying the large, complex decision trees that may result. Crisis situations, where stakes are high, time is short, and the alternatives and information continually changing, pose particularly difficult structuring problems.

Subjective judgments of probability and value are essential inputs to decision analyses. We still do not know the best ways to elicit these judgments. Now that we understand many of the biases to which judgments are susceptible, we need to develop debiasing techniques to minimize their destructive effects. Simply warning a judge about a bias may prove ineffective. Like perceptual illusions, many biases do not disappear upon being identified. It may be necessary to (a) restructure the judgement task in ways that circumvent the bias, (b) use several different methods allowing opposing biases to cancel one another, or (c) correct the judgments externally, based on an estimate of the direction and strength of the bias.

Decision aids must be easy to use. Development of computer graphics techniques is needed to accomplish this goal. Aids also need to be evaluated to determine whether they really are improving quality.

Much progress has been made recently toward understanding judgmental and decision-making processes. We need to continue this pursuit of basic knowledge. Simon (1965), outlining the historical development of writing, the number system, calculus, and other major aids to thought, provided what seems to me a fitting observation with which to conclude this article:

> All of these aids to human thinking, and many others, were devised without understanding the process they aided--the thought process itself. The prospect before us now is that we shall understand that process. We shall be able to diagnose with great accuracy the difficulties of a . . . decision maker . . . and we shall be able to help him modify his problem-solving strategies in specific ways.
>
> We have no experience yet that would allow us to judge what improvement in human decision making we might expect from the application of this new and growing knowledge. . . . Nonetheless, we have reason, I think, to be sanguine at the prospect (p. 92).

Acknowledgement

This research was supported by the Advanced Research Projects Agency of the Department of Defense and was monitored by Office of Naval Research under Contract N00014-79-C-0029 (ARPA Order No. 3668) to Perceptronics, Inc.

References

Abt, C. C. Serious Games, New York: Viking Press, 1970.

Berkson, J., Magath, T. B. & Hurn, M. The error of estimate of the blood cell count as made with the hemocytometer. American Journal of Physiology, 1940, 128, 309-323.

Birkin, S. J. & Ford, J. S. The quantity/quality dilemma: The impact of a zero defects program, in J. L. Cochrane & M. Zeleny, (Eds.), Multiple Criteria Decision Making, Columbia, S. C.: University of S. Carolina Press, 1973, 517-529.

Bruner, J. S., Goodnow, J. J. & Austin, G. A. A Study of Thinking, New York: Wiley, 1956.

Clarkson, G. P. E. Portfolio Selection: A simulation of trust investment, Englewood Cliffs, N. J.: Prentice-Hall, 1962.

Dawes, R. M. A case study of graduate admissions: Applications of three principles of human decision making. American Psychologist, 1971, 26, 180-188.

Dawes, R. M. & Corrigan, B. Linear models in decision making. Psychological Bulletin, 1974, 81, 95-106.

Doughty, J. M. & Feehrer, C. E. The AESOP testbed: Test Series I/2: Summary report, Tech. Report 848, MITRE Corporation, Bedford, Mass., 1969.

Driver, M. J. & Hunsaker, P. L. The Luna I moon colony: A programmed simulation for the analysis of individual and group decision making. Psychological Reports, 1972, 31, 879-888.

Fischer, G. W. Experimental Applications of multiattribute utility models, in D. Wendt & C. A. J. Vlek, (Eds.), Utility, Probability, and Human Decision Making, Dordrecht, The Netherlands: Reidel, 1975.

Fischhoff, B. Hindsight ≠ Foresight: The effect of outcome knowledge on judgment under uncertainty. Journal of Experimental Psychology: Human Perception and Performance, 1975a, 1, 288-299.

Fischhoff, B. Hindsight: Thinking backward?. Psychology Today, 1975b, 8, 70-76.

Fischhoff, B. Perceived informativeness of facts. Journal of Experimental Psychology: Human Perception and Performance, 1977, 3, 349-358.

Fischhoff, B. & Beyth, R. I knew it would happen--Remembered probabilities of once-future things. Organizational Behavior and Human Performance, 1975, 13, 1-16.

Fischhoff, B., Slovic, P. & Lichtenstein, S. Knowing with certainty: The appropriateness of extreme confidence. Journal of Experimental Psychology: Human Perception and Performance, in press.

Garland, L. H. The problem of observer error. Bulletin of the New York Academy of Medicine, 1960, 36, 569-584.

Goldberg, L. R. Man vs. model of man: A rationale, plus some evidence, for a method of improving on clinical inferences. Psychological Bulletin, 1970, 73, 422-432.

Grether, D. M. & Plott, C. R. Economic theory of choice and the preference reversal phenomenon. American Economic Review, in press.

Hammond, K. R., Stewart, T. R. Brehmer, B. & Steinmann, D. O. Social judgment theory, in M. F. Kaplan & S. Schwartz, (Eds.), Human Judgment and Decision Processes, New York: Academic Press, 1975.

Herrero, S. Human injury inflicted by grizzly bears. Science, 1970, 170, 593-597.

Howard, R. A., Matheson, J. E. & North, D. W. The decision to seed hurricanes. Science, 1972, 176, 1191-1202.

Howard, R. A., Merkhofer, M. W. Miller, A. C. & Tani, S. N. A preliminary characterization of a decision structuring process for the task force commander and his staff. Tech. Report MSC-4030, Stanford Research Institute, Palo Alto, Ca., 1975.

Kahneman, D. & Tversky, A. Subjective probability: A judgment of representativeness. Cognitive Psychology, 1972, 3, 430-454.

Kahneman, D. & Tversky, A. On the psychology of prediction. Psychological Review, 1973, 80, 237-251.

Kates, R. W. Hazard and choice perception in flood plain management. Research Paper 78, Department of Geography, University of Chicago, Chicago, Il., 1962

Kidd, J. B. The utilization of subjective probabilities in production planning. Acta Psychologica, 1970, 34, 338-347.

Knight, F. H. Risk, Uncertainty, and Profit. Boston & New York: Houghton-Mifflin, 1921.

Lave, L. B. & Weber, W. E. A benefit-cost analysis of auto safety features. Applied Economics, 1970, 2, 265-275.

Lichtenstein, S. & Fischhoff, B. Do those who know more also know more about how much they know? Organizational Behavior and Human Performance, in press.

Lichtenstein, S., Fischhoff, B. & Phillips, L. D. Calibration of probabilities: The state of the art, in H. Jungermann & G. deZeeuw, (Eds.), Proceedings of the Fifth Research Conference on Subjective Probability, Utility, and Decision Making, Dordrecht, The Netherlands: D. Reidel, 1977.

Lichtenstein, S. & Slovic, P. Response-induced reversals of preferences in gambling: An extended replication in Las Vegas. Journal of Experimental Psychology, 1973, 101, 16-20.

Lichtenstein, S., Slovic, P., Fischhoff, B., Layman, M. & Combs, B. Perceived frequency of low probability, lethal events. Decision Research Report 76-2, 1976.

Matheson, J. E. & Roths, W. J. Decision analysis of space projects: Voyager Mars, in R. Howard, J. E. Matheson, & K. L. Miller, (Eds.), Readings in Decision Analysis, Stanford Research Institute, Palo Alto, Ca., 1976.

Meehl, P. E. & Rosen, A. Antecedent probability and the efficacy of psychometric signs, patterns, or cutting scores. Psychological Bulletin, 1955, 52, 194-216.

Miller, G. A. The magical number seven, plus or minus two: Some limits on our capacity for processing information. Psychological Review, 1956, 63, 81-97.

Newell, A. & Simon, H. A. Human Problem Solving, Englewood Cliffs, N. J.: Prentice-Hall, 1972.

Nickerson, R. S. & Feehrer, C. E. Decision making and training: A review of theoretical and empirical studies of decision making and their implications for the training of decision makers. Tech. Report 73-C-0128-1, Naval Training Equipment Center, Orlando, Fl., 1975.

Raiffa, H. Decision Analysis: Introductory lectures on choice under uncertainty. Reading, Mass.: Addison-Wesley, 1968.

Samet, M. G. Quantitative interpretation of two qualitative scales used to rate military intelligence. Human Factors, 1975, 17, 192-202.

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

Simon, H. A. Models of Man: Social and rational, New York: Wiley, 1957.

Simon, H. A. The Shape of Automation for Men and Management, New York: Harper & Row, 1965.

Sinsheimer, R. F. The brain of Pooh: An essay on the limits of mind. American Scientist, 1971, 59, 20-28.

Slovic, P. From Shakespeare to Simon: Speculations--and some evidence--about man's ability to process information. Oregon Research Institute Research Monograph, 1972, 12, 2.

Slovic, P. & Fischhoff, B. On the psychology of experimental surprises. Journal of Experimental Psychology: Human Perception and Performance, in press.

Slovic, P., Fischhoff, B. & Lichtenstein, S. Behavioral decision theory. Annual Review of Psychology, 1977, 28, 1-29.

Slovic, P., Kunreuther, H. & White, G. F. Decision porcesses, rationality and adjustment to natural hazards. In G. F. White (Ed.), Natural Hazards, Local, National and Global. New York: Oxford Univ. Press, 1974.

Slovic, P. & Lichtenstein, S. Comparison of Bayesian and regression approaches to the study of information processing in judgment. Organizational Behavior and Human Performance, 1971, 6, 649-744.

Slovic, P. & MacPhillamy, D. J. Dimensional commensurability and cue utilization in comparative judgment. Organizational Behavior and Human Performance, 1974, 11, 172-194.

Tversky, A. & Kahneman, D. The belief in the law of small numbers. Psychological Bulletin, 1971, 76, 105-110.

Tversky, A. & Kahneman, D. Availability: A heuristic for judging frequency and probability. Cognitive Psychology, 1973, 5, 207-232.

Tversky, A. & Kahneman, D. Judgment under uncertainty: Heuristics and biases. Science, 1974, 185, 1124-1131.

von Neumann, J. & Morgenstern, O. Theory of Games and Economic Behavior, 3rd Ed., Princeton, N. J.: Princeton University Press, 1953.

von Winterfeldt, D. & Fischer, G. W. Multiattribute utility theory: Models and assessment procedures, in D. Wendt & C. A. J. Vlek, (Eds.), Utility, Probability, and Human Decision Making, Dordrecht, The Netherlands: Reidel, 1975.

Wainer, H., Zill, N. & Gruvaeus, G. Senatorial decision making: II. Prediction. Behavioral Science, 1973, 18, 20-26.

Wohlstetter, A. Legends of the strategic arms race, Part I: The driving machine. Strategic Review, 1974, 67-92.

Wohlstetter, R. Pearl Harbor: Warning and Decision. Stanford, Ca.: Stanford University Press, 1962.

Wohlstetter, R. Cuba and Pearl Harbor: Hindsight and foresight. Memo RM 4328-ISA, RAND Corporation, Santa Monica, Ca., 1965, 36.

Zebroski, E. L. Attainment of balance in risk-benefit perceptions. In D. Okrent (Ed.), Risk-benefit methodology and application: Some papers presented at the Engineering Foundation Workshop, Asilomar, Ca., Report ENG-7598, School of Engineering and Applied Science, UCLA, 1975.

Type of Decision

Unique	Repeated
Long lead time: Decision analysis	Rule-based systems: Bootstrapping Multiattribute utilities
Short lead time: Educated intuition	Computer information systems Simulation

Figure 1. Aids for Major Decisions