Home

Decision Making and Problem Solving

by Herbert A. Simon and Associates

Associates: George B. Dantzig, Robin Hogarth, Charles R. Piott, Howard Raiffa, Thomas C. Schelling, Kennth A. Shepsle, Richard Thaier, Amos Tversky, and Sidney Winter.

Simon was educated in political science at the University of Chicago (B.A., 1936, Ph.D., 1943). He has held research and faculty positions at the University of California (Berkeley), Illinois Institute of Technology and since 1949, Carnegie Mellon University, where he is the Richard King Mellon University Professor of Computer Science and Psychology. In 1978, he received the Alfred Nobel Memorial Prize in Economic Sciences and in 1986 the National Medal of Science.

Introduction

The work of managers, of scientists, of engineers, of lawyers--the work that steers the course of society and its economic and governmental organizations--is largely work of making decisions and solving problems. It is work of choosing issues that require attention, setting goals, finding or designing suitable courses of action, and evaluating and choosing among alternative actions. The first three of these activities--fixing agendas, setting goals, and designing actions--are usually called *problem solving;* the last, evaluating and choosing, is usually called *decision making.* Nothing is more important for the well-being of society than that this work be performed effectively, that we address successfully the many problems requiring attention at the national level (the budget and trade deficits, AIDS, national security, the mitigation of earthquake damage), at the level of business organizations (product improvement, efficiency of production, choice of investments), and at the level of our individual lives (choosing a career or a school, buying a house).

The abilities and skills that determine the quality of our decisions and problem solutions are stored not only in more than 200 million human heads, but also in tools and machines, and especially today in those machines we call computers. This fund of brains and its attendant machines form the basis of our American ingenuity, an ingenuity that has permitted U.S. society to reach remarkable levels of economic productivity.

There are no more promising or important targets for basic scientific research than understanding how human minds, with and without the help of computers, solve problems and make decisions effectively, and improving our problem-solving and decision-making capabilities. In psychology, economics, mathematical statistics, operations research, political science, artificial intelligence, and cognitive science, major research gains have been made during the past half century in understanding problem solving and decision making. The progress already achieved holds forth the promise of exciting new advances that will contribute substantially to our nation's capacity for dealing intelligently with the range of issues, large and small, that confront us.

Much of our existing knowledge about decision making and problem solving, derived from this research, has already been put to use in a wide variety of applications, including procedures used to assess drug safety, inventory control methods for industry, the new expert systems that embody artificial intelligence techniques, procedures for modeling energy and environmental systems, and analyses of the stabilizing or destabilizing effects of alternative defense strategies. (Application of the new inventory control techniques, for example, has enabled American corporations to reduce their inventories by hundreds of millions of dollars since World War II without increasing the incidence of stockouts.) Some of the knowledge gained through the research describes the ways in which people actually go about making decisions and solving problems; some of it prescribes better methods, offering advice for the improvement of the process.

Central to the body of prescriptive knowledge about decision making has been the theory of subjective expected utility (SEU), a sophisticated mathematical model of choice that lies at the foundation of most contemporary economics, theoretical statistics, and operations research. SEU theory defines the conditions of perfect utility-maximizing rationality in a world of certainty or in a world in which the probability distributions of all relevant variables can be provided by the decision makers. (In spirit, it might be compared with a theory of ideal gases or of frictionless bodies sliding down inclined planes in a vacuum.) SEU theory deals only with decision making; it has nothing to say about how to frame problems, set goals, or develop new alternatives.

Prescriptive theories of choice such as SEU are complemented by empirical research that shows how people actually make decisions (purchasing insurance, voting for political candidates, or investing in securities), and research on the processes people use to solve problems (designing switchgear or finding chemical reaction pathways). This research demonstrates that people solve problems by selective, heuristic search through large problem spaces and large data bases, using means-ends analysis as a principal technique for guiding the search. The expert systems that are now being produced by research on artificial intelligence and applied to such tasks as interpreting oil-well drilling logs or making medical diagnoses are outgrowths of these research findings on human problem solving.

What chiefly distinguishes the empirical research on decision making and problem solving from the prescriptive approaches derived from SEU theory is the attention that the former gives to the limits on human rationality. These limits are imposed by the complexity of the world in which we live, the incompleteness and inadequacy of human knowledge, the inconsistencies of individual preference and belief, the conflicts of value among people and groups of people, and the inadequacy of the computations we can carry out, even with the aid of the most powerful computers. The real world of human decisions is not a world of ideal gases, frictionless planes, or vacuums. To bring it within the scope of human thinking powers, we must simplify our problem formulations drastically, even leaving out much or most of what is potentially relevant.

The descriptive theory of problem solving and decision making is centrally concerned with how people cut problems down to size: how they apply approximate, heuristic techniques to handle complexity that cannot be handled exactly. Out of this descriptive theory is emerging an augmented and amended prescriptive theory, one that takes account of the gaps and elements of unrealism in SEU theory by encompassing problem solving as well as choice and demanding only the kinds of knowledge, consistency, and computational power that are attainable in the real world.

The growing realization that coping with complexity is central to human decision making strongly influences the directions of research in this domain. Operations research and artificial intelligence are forging powerful new computational tools; at the same time, a new body of mathematical theory is evolving around the topic of computational complexity. Economics, which has traditionally derived both its descriptive and prescriptive approaches from SEU theory, is now paying a great deal of attention to uncertainty and incomplete information; to so-called "agency theory," which takes account of the institutional framework within which decisions are made; and to game theory, which seeks to deal with interindividual and intergroup processes in which there is partial conflict of interest. Economists and political scientists are also increasingly buttressing the empirical foundations of their field by studying individual choice behavior directly and by studying behavior in experimentally constructed markets and simulated political structures.

The following pages contain a fuller outline of current knowledge about decision making and problem solving and a brief review of current research directions in these fields as well as some of the principal research opportunities.

Decision Making

SEU THEORY

The development of SEU theory was a major intellectual achievement of the first half of this century. It gave for the first time a formally axiomatized statement of what it would mean for an agent to behave in a consistent, rational matter. It assumed that a decision maker possessed a utility function (an ordering by preference among all the possible outcomes of choice), that all the alternatives among which choice could be made were known, and that the consequences of choosing each alternative could be ascertained (or, in the version of the theory that treats of choice under uncertainty, it assumed that a subjective or objective probability distribution of consequences was associated with each alternative). By admitting subjectively assigned probabilities, SEU theory opened the way to fusing subjective opinions with objective data, an approach that can also be used in man-machine decision-making systems. In the probabilistic version of the theory, Bayes's rule prescribes how people should take account of new information and how they should respond to incomplete information.

The assumptions of SEU theory are very strong, permitting correspondingly strong inferences to be made from them. Although the assumptions cannot be satisfied even remotely for most complex situations in the real world, they may be satisfied approximately in some microcosms--problem situations that can be isolated from the world's complexity and dealt with independently. For example, the manager of a commercial cattle-feeding operation might isolate the problem of finding the least expensive mix of feeds available in the market that would meet all the nutritional requirements of his cattle. The computational tool of linear programming, which is a powerful method for maximizing goal achievement or minimizing costs while satisfying all kinds of side conditions (in this case, the nutritional requirements), can provide the manager with an optimal feed mix--optimal within the limits of approximation of his model to real world conditions. Linear programming and related operations research techniques are now used widely to make decisions whenever a situation that reasonably fits their assumptions can be carved out of its complex surround. These techniques have been especially valuable aids to middle management in dealing with relatively well-structured decision problems.

Most of the tools of modern operations research--not only linear programming, but also integer programming, queuing theory, decision trees, and other widely used techniques--use the assumptions of SEU theory. They assume that what is desired is to maximize the achievement of some goal, under specified constraints and assuming that all alternatives and consequences (or their probability distributions) are known. These tools have proven their usefulness in a wide variety of applications.

THE LIMITS OF RATIONALITY

Operations research tools have also underscored dramatically the limits of SEU theory in dealing with complexity. For example, present and prospective computers are not even powerful enough to provide exact solutions for the problems of optimal scheduling and routing of jobs through a typical factory that manufactures a variety of products using many different tools and machines. And the mere thought of using these computational techniques to determine an optimal national policy for energy production or an optimal economic policy reveals their limits.

Computational complexity is not the only factor that limits the literal application of SEU theory. The theory also makes enormous demands on information. For the utility function, the range of available alternatives and the consequences following from each alternative must all be known. Increasingly, research is being directed at decision making that takes realistic account of the compromises and approximations that must be made in order to fit real-world problems to the informational and computational limits of people and computers, as well as to the inconsistencies in their values and perceptions. The study of actual decision processes (for example, the strategies used by corporations to make their investments) reveals massive and unavoidable departures from the framework of SEU theory. The sections that follow describe some of the things that have been learned about choice under various conditions of incomplete information, limited computing power, inconsistency, and institutional constraints on alternatives. Game theory, agency theory, choice under uncertainty, and the theory of markets are a few of the directions of this research, with the aims both of constructing prescriptive theories of broader application and of providing more realistic descriptions and explanations of actual decision making within U.S. economic and political institutions.

LIMITED RATIONALITY IN ECONOMIC THEORY

Although the limits of human rationality were stressed by some researchers in the 1950s, only recently has there been extensive activity in the field of economics aimed at developing theories that assume less than fully rational choice on the part of business firm managers and other economic agents. The newer theoretical research undertakes to answer such questions as the following:

- Are market equilibria altered by the departures of actual choice behavior from the behavior of fully rational agents predicted by SEU theory?
- Under what circumstances do the processes of competition "police" markets in such a way as to cancel out the effects of the departures from full rationality?
- In what ways are the choices made by boundedly rational agents different from those made by fully rational agents?

Theories of the firm that assume managers are aiming at "satisfactory" profits or that their concern is to maintain the firm's share of market in the industry make quite different predictions about economic equilibrium than those derived from the assumption of profit maximization. Moreover, the classical theory of the firm cannot explain why economic activity is sometimes organized around large business firms and sometimes around contractual networks of individuals or smaller organizations. New theories that take account of differential access of economic agents to information, combined with differences in self-interest, are able to account for these important phenomena, as well as provide explanations for the many forms of contracts that are used in business. Incompleteness and asymmetry of information have been shown to be essential for explaining how individuals and business firms decide when to face uncertainty by insuring, when by hedging, and when by assuming the risk.

Most current work in this domain still assumes that economic agents seek to maximize utility, but within limits posed by the incompleteness and uncertainty of the information available to them. An important potential area of research is to discover how choices will be changed if there are other departures from the axioms of rational choice--for example, substituting goals of reaching specified aspiration levels (satisficing) for goals of maximizing.

Applying the new assumptions about choice to economics leads to new empirically supported theories about decision making over time. The classical theory of perfect rationality leaves no room for regrets, second thoughts, or "weakness of will." It cannot explain why many individuals enroll in Christmas savings plans, which earn interest well below the market rate. More generally, it does not lead to correct conclusions about the important social issues of saving and conservation. The effect of pensions and social security on personal saving has been a controversial issue in economics. The standard economic model predicts that an increase in required pension saving will reduce other saving dollar for dollar; behavioral theories, on the other hand, predict a much smaller offset. The empirical evidence indicates that the offset is indeed very small. Another empirical finding is that the method of payment of wages and salaries affects the saving rate. For example, annual bonuses produce a higher saving rate than the same amount of income paid in monthly salaries. This finding implies that saving rates can be influenced by the way compensation is framed.

If individuals fail to discount properly for the passage of time, their decisions will not be optimal. For example, air conditioners vary greatly in their energy efficiency; the more efficient models cost more initially but save money over the long run through lower energy consumption. It has been found that consumers, on average, choose air conditioners that imply a discount rate of 25 percent or more per year, much higher than the rates of interest that prevailed at the time of the study.

As recently as five years ago, the evidence was thought to be unassailable that markets like the New York Stock Exchange work efficiently--that prices reflect all available information at any given moment in time, so that stock price movements resemble a random walk and contain no systematic information that could be exploited for profit. Recently, however, substantial departures from the behavior predicted by the efficient-market hypothesis have been detected. For example, small firms appear to earn inexplicably high returns on the market prices of their stock, while firms that have very low price-earnings ratios and firms that have lost much of their market value in the recent past also earn abnormally high returns. All of these results are consistent with the empirical finding that decision makers often overreact to new information, in violation of Bayes's rule. In the same way, it has been found that stock prices are excessively volatile--that they fluctuate up and down more rapidly and violently than they would if the market were efficient.

There has also been a long-standing puzzle as to why firms pay dividends. Considering that dividends are taxed at a higher rate than capital gains, taxpaying investors should prefer, under the assumptions of perfect rationality, that their firms reinvest earnings or repurchase shares instead of paying dividends. (The investors could simply sell some of their appreciated shares to obtain the income they require.) The solution to this puzzle also requires models of investors that take account of limits on rationality.

THE THEORY OF GAMES

In economic, political, and other social situations in which there is actual or potential conflict of interest, especially if it is combined with incomplete information, SEU theory faces special difficulties. In markets in which there are many competitors (e.g., the wheat market), each buyer or seller can accept the market price as a "given" that will not be affected materially by the actions of any single individual. Under these conditions, SEU theory makes unambiguous predictions of behavior. However, when a market has only a few suppliers --say, for example, two--matters are quite different. In this case, what it is rational to do depends on what one's competitor is going to do, and vice versa. Each supplier may try to outwit the other. What then is the rational decision?

The most ambitious attempt to answer questions of this kind was the theory of games, developed by von Neumann and Morgenstern and published in its full form in 1944. But the answers provided by the theory of games are sometimes very puzzling and ambiguous. In many situations, no single course of action dominates all the others; instead, a whole set of possible solutions are all equally consistent with the postulates of rationality.

One game that has been studied extensively, both theoretically and empirically, is the Prisoner's Dilemma. In this game between two players, each has a choice between two actions, one trustful of the other player, the other mistrustful or exploitative. If both players choose the trustful alternative, both receive small rewards. If both choose the exploitative alternative, both are punished. If one chooses the trustful alternative and the other the exploitative alternative, the former is punished much more severely than in the previous case, while the latter receives a substantial reward. If the other player's choice is fixed but unknown, it is advantageous for a player to choose the exploitative alternative, for this will give him the best outcome in either case. But if both adopt this reasoning, they will both be punished, whereas they could both receive rewards if they agreed upon the trustful choice (and did not welch on the agreement).

The terms of the game have an unsettling resemblance to certain situations in the relations between nations or between a company and the employees' union. The resemblance becomes stronger if one imagines the game as being played repeatedly. Analyses of "rational" behavior under assumptions of intended utility maximization support the conclusion that the players will (ought to?) always make the mistrustful choice. Nevertheless, in laboratory experiments with the game, it is often found that players (even those who are expert in game theory) adopt a "tit-for-tat" strategy. That is, each plays the trustful, cooperative strategy as long as his or her partner does the same. If the partner exploits the player on a particular trial, the player then plays the exploitative strategy on the next trial and continues to do so until the partner switches back to the trustful strategy. Under these conditions, the game frequently stabilizes with the players pursuing the mutually trustful strategy and receiving the rewards.

With these empirical findings in hand, theorists have recently sought and found some of the conditions for attaining this kind of benign stability. It occurs, for example, if the players set aspirations for a satisfactory reward rather than seeking the maximum reward. This result is consistent with the finding that in many situations, as in the Prisoner's Dilemma game, people appear to satisfice rather than attempting to optimize.

The Prisoner's Dilemma game illustrates an important point that is beginning to be appreciated by those who do research on decision making. There are so many ways in which actual human behavior can depart from the SEU assumptions that theorists seeking to account for behavior are confronted with an embarrassment of riches. To choose among the many alternative models that could account for the anomalies of choice, extensive empirical research is called for--to see how people do make their choices, what beliefs guide them, what information they have available, and what part of that information they take into account and what part they ignore. In a world of limited rationality, economics and the other decision sciences must closely examine the actual limits on rationality in order to make accurate predictions and to provide sound advice on public policy.

EMPIRICAL STUDIES OF CHOICE UNDER UNCERTAINTY

During the past ten years, empirical studies of human choices in which uncertainty, inconsistency, and incomplete information are present have produced a rich collection of findings which only now are beginning to be organized under broad generalizations. Here are a few examples. When people are given information about the probabilities of certain events (e.g., how many lawyers and how many engineers are in a population that is being sampled), and then are given some additional information as to which of the events has occurred (which person has been sampled from the population), they tend to ignore the prior probabilities in favor of incomplete or even quite irrelevant information about the individual event. Thus, if they are told that 70 percent of the population are lawyers, and if they are then given a noncommittal description of a person (one that could equally well fit a lawyer or an engineer), half the time they will predict that the person is a lawyer and half the time that he is an engineer--even though the laws of probability dictate that the best forecast is always to predict that the person is a lawyer.

People commonly misjudge probabilities in many other ways. Asked to estimate the probability that 60 percent or more of the babies born in a hospital during a given week are male, they ignore information about the total number of births, although it is evident that the probability of a departure of this magnitude from the expected value of 50 percent is smaller if the total number of births is larger (the standard error of a percentage varies inversely with the square root of the population size).

There are situations in which people assess the frequency of a class by the ease with which instances can be brought to mind. In one experiment, subjects heard a list of names of persons of both sexes and were later asked to judge whether there were more names of men or women on the list. In lists presented to some subjects, the men were more famous than the women; in other lists, the women were more famous than the men. For all lists, subjects judged that the sex that had the more famous personalities was the more numerous.

The way in which an uncertain possibility is presented may have a substantial effect on how people respond to it. When asked whether they would choose surgery in a hypothetical medical emergency, many more people said that they would when the chance of survival was given as 80 percent than when the chance of death was given as 20 percent.

On the basis of these studies, some of the general heuristics, or rules of thumb, that people use in making judgments have been compiled---heuristics that produce biases toward classifying situations according to their representativeness, or toward judging frequencies according to the availability of examples in memory, or toward interpretations warped by the way in which a problem has been framed.

These findings have important implications for public policy. A recent example is the lobbying effort of the credit card industry to have differentials between cash and credit prices labeled "cash discounts" rather than "credit surcharges." The research findings raise questions about how to phrase cigarette warning labels or frame truth-in-lending laws and informed consent laws.

METHODS OF EMPIRICAL RESEARCH

Finding the underlying bases of human choice behavior is difficult. People cannot always, or perhaps even usually, provide veridical accounts of how they make up their minds, especially when there is uncertainty. In many cases, they can predict how they will behave (pre-election polls of voting intentions have been reasonably accurate when carefully taken), but the reasons people give for their choices can often be shown to be rationalizations and not closely related to their real motives.

Students of choice behavior have steadily improved their research methods. They question respondents about specific situations, rather than asking for generalizations. They are sensitive to the dependence of answers on the exact forms of the questions. They are aware that behavior in an experimental situation may be different from behavior in real life, and they attempt to provide experimental settings and motivations that are as realistic as possible. Using thinking-aloud protocols and other approaches, they try to track the choice behavior step by step, instead of relying just on information about outcomes or querying respondents retrospectively about their choice processes.

Perhaps the most common method of empirical research in this field is still to ask people to respond to a series of questions. But data obtained by this method are being supplemented by data obtained from carefully designed laboratory experiments and from observations of actual choice behavior (for example, the behavior of customers in supermarkets). In an experimental study of choice, subjects may trade in an actual market with real (if modest) monetary rewards and penalties. Research experience has also demonstrated the feasibility of making direct observations, over substantial periods of time, of the decision-making processes in business and governmental organizations--for example, observations of the procedures that corporations use in making new investments in plant and equipment. Confidence in the empirical findings that have been accumulating over the past several decades is enhanced by the general consistency that is observed among the data obtained from quite different settings using different research methods.

There still remains the enormous and challenging task of putting together these findings into an empirically founded theory of decision making. With the growing availability of data, the theory-building enterprise is receiving much better guidance from the facts than it did in the past. As a result, we can expect it to become correspondingly more effective in arriving at realistic models of behavior.

Problem Solving

The theory of choice has its roots mainly in economics, statistics, and operations research and only recently has received much attention from psychologists; the theory of problem solving has a very different history. Problem solving was initially studied principally by psychologists, and more recently by researchers in artificial intelligence. It has received rather scant attention from economists.

CONTEMPORARY PROBLEM-SOLVING THEORY

Human problem solving is usually studied in laboratory settings, using problems that can be solved in relatively short periods of time (seldom more than an hour), and often seeking a maximum density of data about the solution process by asking subjects to think aloud while they work. The thinking-aloud technique, at first viewed with suspicion by behaviorists as subjective and "introspective," has received such careful methodological attention in recent years that it can now be used dependably to obtain data about subjects' behaviors in a wide range of settings.

The laboratory study of problem solving has been supplemented by field studies of professionals solving real-world problems--for example, physicians making diagnoses and chess grandmasters analyzing game positions, and, as noted earlier, even business corporations making investment decisions. Currently, historical records, including laboratory notebooks of scientists, are also being used to study problem-solving processes in scientific discovery. Although such records are far less "dense" than laboratory protocols, they sometimes permit the course of discovery to be traced in considerable detail. Laboratory notebooks of scientists as distinguished as Charles Darwin, Michael Faraday, Antoine-Laurent Lavoisier, and Hans Krebs have been used successfully in such research.

From empirical studies, a description can now be given of the problem-solving process that holds for a rather wide range of activities. First, problem solving generally proceeds by selective search through large sets of possibilities, using rules of thumb (heuristics) to guide the search. Because the possibilities in realistic problem situations are generally multitudinous, trial-and-error search would simply not work; the search must be highly selective. Chess grandmasters seldom examine more than a hundred of the vast number of possible scenarios that confront them, and similar small numbers of searches are observed in other kinds of problem-solving search.

One of the procedures often used to guide search is "hill climbing," using some measure of approach to the goal to determine where it is most profitable to look next. Another, and more powerful, common procedure is means-ends analysis. In means-ends analysis, the problem solver compares the present situation with the goal, detects a difference between them, and then searches memory for actions that are likely to reduce the difference. Thus, if the difference is a fifty-mile distance from the goal, the problem solver will retrieve from memory knowledge about autos, carts, bicycles, and other means of transport; walking and flying will probably be discarded as inappropriate for that distance.

The third thing that has been learned about problem solving--especially when the solver is an expert--is that it relies on large amounts of information that are stored in memory and that are retrievable whenever the solver recognizes cues signaling its relevance. Thus, the expert knowledge of a diagnostician is evoked by the symptoms presented by the patient; this knowledge leads to the recollection of what additional information is needed to discriminate among alternative diseases and, finally, to the diagnosis.

In a few cases, it has been possible to estimate how many patterns an expert must be able to recognize in order to gain access to the relevant knowledge stored in memory. A chess master must be able to recognize about 50,000 different configurations of chess pieces that occur frequently in the course of chess games. A medical diagnostician must be able to recognize tens of thousands of configurations of symptoms; a botanist or zoologist specializing in taxonomy, tens or hundreds of thousands of features of specimens that define their species. For comparison, college graduates typically have vocabularies in their native languages of 50,000 to 200,000 words. (However, these numbers are very small in comparison with the real-world situations the expert faces: there are perhaps 10^{120} branches in the game tree of chess, a game played with only six kinds of pieces on an 8 x 8 board.)

One of the accomplishments of the contemporary theory of problem solving has been to provide an explanation for the phenomena of intuition and judgment frequently seen in experts' behavior. The store of expert knowledge, "indexed" by the recognition cues that make it accessible and combined with some basic inferential capabilities (perhaps in the form of means-ends analysis), accounts for the ability of experts to find satisfactory solutions for difficult problems, and sometimes to find them almost instantaneously. The expert's "intuition" and "judgment" derive from this capability for rapid recognition linked to a large store of knowledge. When immediate intuition fails to yield a problem solution or when a prospective solution needs to be evaluated, the expert falls back on the slower processes of analysis and inference.

EXPERT SYSTEMS IN ARTIFICIAL INTELLIGENCE

Over the past thirty years, there has been close teamwork between research in psychology and research in computer science aimed at developing intelligent programs. Artificial intelligence (AI) research has both borrowed from and contributed to research on human problem solving. Today, artificial intelligence is beginning to produce systems, applied to a variety of tasks, that can solve difficult problems at the level of professionally trained humans. These AI programs are usually called expert systems. A description of a typical expert system would resemble closely the description given above of typical human problem solving; the differences between the two would be differences in degree, not in kind. An AI expert system, relying on the speed of computers and their ability to retain large bodies of transient information in memory, will generally use "brute force"--sheer computational speed and power--more freely than a human expert can. A human expert, in compensation, will generally have a richer set of heuristics to guide search and a larger vocabulary of recognizable patterns. To the observer, the computer's process will appear the more systematic and even compulsive, the human's the more intuitive. But these are quantitative, not qualitative, differences.

The number of tasks for which expert systems have been built is increasing rapidly. One is medical diagnosis (two examples are the CADUCEUS and MYCIN programs). Others are automatic design of electric motors, generators, and transformers (which predates by a decade the invention of the term *expert systems),* the configuration of computer systems from customer specifications, and the automatic generation of reaction paths for the synthesis of organic molecules. All of these (and others) are either being used currently in professional or industrial practice or at least have reached a level at which they can produce a professionally acceptable product.

Expert systems are generally constructed in close consultation with the people who are experts in the task domain. Using standard techniques of observation and interrogation, the heuristics that the human expert uses, implicitly and often unconsciously, to perform the task are gradually educed, made explicit, and incorporated in program structures. Although a great deal has been learned about how to do this, improving techniques for designing expert systems is an important current direction of research. It is especially important because expert systems, once built, cannot remain static but must be modifiable to incorporate new knowledge as it becomes available.

DEALING WITH ILL-STRUCTURED PROBLEMS

In the 1950s and 1960s, research on problem solving focused on clearly structured puzzle-like problems that were easily brought into the psychological laboratory and that were within the range of computer programming sophistication at that time. Computer programs were written to discover proofs for theorems in Euclidean geometry or to solve the puzzle of transporting missionaries and cannibals across a river. Choosing chess moves was perhaps the most complex task that received attention in the early years of cognitive science and AI.

As understanding grew of the methods needed to handle these relatively simple tasks, research aspirations rose. The next main target, in the 1960s and 1970s, was to find methods for solving problems that involved large bodies of semantic information. Medical diagnosis and interpreting mass spectrogram data are examples of the kinds of tasks that were investigated during this period and for which a good level of understanding was achieved. They are tasks that, for all of the knowledge they call upon, are still well structured, with clear-cut goals and constraints.

The current research target is to gain an understanding of problem-solving tasks when the goals themselves are complex and sometimes ill defined, and when the very nature of the problem is successively transformed in the course of exploration. To the extent that a problem has these characteristics, it is usually called ill structured. Because ambiguous goals and shifting problem formulations are typical characteristics of problems of design, the work of architects offers a good example of what is involved in solving ill-structured problems. An architect begins with some very general specifications of what is wanted by a client. The initial goals are modified and substantially elaborated as the architect proceeds with the task. Initial design ideas, recorded in drawings and diagrams, themselves suggest new criteria, new possibilities, and new requirements. Throughout the whole process of design, the emerging conception provides continual feedback that reminds the architect of additional considerations that need to be taken into account.

With the current state of the art, it is just beginning to be possible to construct programs that simulate this kind of flexible problem-solving process. What is called for is an expert system whose expertise includes substantial knowledge about design criteria as well as knowledge about the means for satisfying those criteria. Both kinds of knowledge are evoked in the course of the design activity by the usual recognition processes, and the evocation of design criteria and constraints continually modifies and remolds the problem that the design system is addressing. The large data bases that can now be constructed to aid in the management of architectural and construction projects provide a framework into which AI tools, fashioned along these lines, can be incorporated.

Most corporate strategy problems and governmental policy problems are at least as ill structured as problems of architectural or engineering design. The tools now being forged for aiding architectural design will provide a basis for building tools that can aid in formulating, assessing, and monitoring public energy or environmental policies, or in guiding corporate product and investment strategies.

SETTING THE AGENDA AND REPRESENTING A PROBLEM

The very first steps in the problem-solving process are the least understood. What brings (and should bring) problems to the head of the agenda? And when a problem is identified, how can it be represented in a way that facilitates its solution?

The task of setting an agenda is of utmost importance because both individual human beings and human institutions have limited capacities for dealing with many tasks simultaneously. While some problems are receiving full attention, others are neglected. Where new problems come thick and fast, "fire fighting" replaces planning and deliberation. The facts of limited attention span, both for individuals and for institutions like the Congress, are well known. However, relatively little has been accomplished toward analyzing or designing effective agenda-setting systems. A beginning could be made by the study of "alerting" organizations like the Office of Technology Assessment or military and foreign affairs intelligence agencies. Because the research and development function in industry is also in considerable part a task of monitoring current and prospective technological advances, it could also be studied profitably from this standpoint.

The way in which problems are represented has much to do with the quality of the solutions that are found. The task of designing highways or dams takes on an entirely new aspect if human responses to a changed environment are taken into account. (New transportation routes cause people to move their homes, and people show a considerable propensity to move into zones that are subject to flooding when partial protections are erected.) Very different social welfare policies are usually proposed in response to the problem of providing incentives for economic independence than are proposed in response to the problem of taking care of the needy. Early management information systems were designed on the assumption that information was the scarce resource; today, because designers recognize that the scarce resource is managerial attention, a new framework produces quite different designs.

The representation or "framing" of problems is even less well understood than agenda setting. Today's expert systems make use of problem representations that already exist. But major advances in human knowledge frequently derive from new ways of thinking about problems. A large part of the history of physics in nineteenth-century England can be written in terms of the shift from action-at-a-distance representations to the field representations that were developed by the applied mathematicians at Cambridge.

Today, developments in computer-aided design (CAD) present new opportunities to provide human designers with computer-generated representations of their problems. Effective use of these capabilities requires us to understand better how people extract information from diagrams and other displays and how displays can enhance human performance in design tasks. Research on representations is fundamental to the progress of CAD.

COMPUTATION AS PROBLEM SOLVING

Nothing has been said so far about the radical changes that have been brought about in problem solving over most of the domains of science and engineering by the standard uses of computers as computational devices. Although a few examples come to mind in which artificial intelligence has contributed to these developments, they have mainly been brought about by research in the individual sciences themselves, combined with work in numerical analysis.

Whatever their origins, the massive computational applications of computers are changing the conduct of science in numerous ways. There are new specialties emerging such as "computational physics" and "computational chemistry." Computation--that is to say, problem solving--becomes an object of explicit concern to scientists, side by side with the substance of the science itself. Out of this new awareness of the computational component of scientific inquiry is arising an increasing interaction among computational specialists in the various sciences and scientists concerned with cognition and AI. This interaction extends well beyond the traditional area of numerical analysis, or even the newer subject of computational complexity, into the heart of the theory of problem solving.

Physicists seeking to handle the great mass of bubble-chamber data produced by their instruments began, as early as the 1960s, to look to AI for pattern recognition methods as a basis for automating the analysis of their data. The construction of expert systems to interpret mass spectrogram data and of other systems to design synthesis paths for chemical reactions are other examples of problem solving in science, as are programs to aid in matching sequences of nucleic acids in DNA and RNA and amino acid sequences in proteins.

Theories of human problem solving and learning are also beginning to attract new attention within the scientific community as a basis for improving science teaching. Each advance in the understanding of problem solving and learning processes provides new insights about the ways in which a learner must store and index new knowledge and procedures if they are to be useful for solving problems. Research on these topics is also generating new ideas about how effective learning takes place--for example, how students can learn by examining and analyzing worked-out examples.

Extensions of Theory

Opportunities for advancing our understanding of decision making and problem solving are not limited to the topics dealt with above, and in this section, just a few indications of additional promising directions for research are presented.

DECISION MAKING OVER TIME

The time dimension is especially troublesome in decision making. Economics has long used the notion of time discounting and interest rates to compare present with future consequences of decisions, but as noted above, research on actual decision making shows that people frequently are inconsistent in their choices between present and future. Although time discounting is a powerful idea, it requires fixing appropriate discount rates for individual, and especially social, decisions. Additional problems arise because human tastes and priorities change over time. Classical SEU theory assumes a fixed, consistent utility function, which does not easily accommodate changes in taste. At the other extreme, theories postulating a limited attention span do not have ready ways of ensuring consistency of choice over time.

AGGREGATION

In applying our knowledge of decision making and problem solving to society-wide, or even organization-wide, phenomena, the problem of aggregation must be solved; that is, ways must be found to extrapolate from theories of individual decision processes to the net effects on the whole economy, polity, and society. Because of the wide variety of ways in which any given decision task can be approached, it is unrealistic to postulate a "representative firm" or an "economic man," and to simply lump together the behaviors of large numbers of supposedly identical individuals. Solving the aggregation problem becomes more important as more of the empirical research effort is directed toward studying behavior at a detailed, microscopic level.

ORGANIZATIONS

Related to aggregation is the question of how decision making and problem solving change when attention turns from the behavior of isolated individuals to the behavior of these same individuals operating as members of organizations or other groups. When people assume organizational positions, they adapt their goals and values to their responsibilities. Moreover, their decisions are influenced substantially by the patterns of information flow and other communications among the various organization units.

Organizations sometimes display sophisticated capabilities far beyond the understanding of single individuals. They sometimes make enormous blunders or find themselves incapable of acting. Organizational performance is highly sensitive to the quality of the routines or "performance programs" that govern behavior and to the adaptability of these routines in the face of a changing environment. In particular, the "peripheral vision" of a complex organization is limited, so that responses to novelty in the environment may be made in inappropriate and quasi-automatic ways that cause major failure.

Theory development, formal modeling, laboratory experiments, and analysis of historical cases are all going forward in this important area of inquiry. Although the decision-making processes of organizations have been studied in the field on a limited scale, a great many more such intensive studies will be needed before the full range of techniques used by organizations to make their decisions is understood, and before the strengths and weaknesses of these techniques are grasped.

LEARNING

Until quite recently, most research in cognitive science and artificial intelligence had been aimed at understanding how intelligent systems perform their work. Only in the past five years has attention begun to turn to the question of how systems become intelligent--how they learn. A number of promising hypotheses about learning mechanisms are currently being explored. One is the so-called connexionist hypothesis, which postulates networks that learn by changing the strengths of their interconnections in response to feedback. Another learning mechanism that is being investigated is the adaptive production system, a computer program that learns by generating new instructions that are simply annexed to the existing program. Some success has been achieved in constructing adaptive production systems that can learn to solve equations in algebra and to do other tasks at comparable levels of difficulty.

Learning is of particular importance for successful adaptation to an environment that is changing rapidly. Because that is exactly the environment of the 1980s, the trend toward broadening research on decision making to include learning and adaptation is welcome.

This section has by no means exhausted the areas in which exciting and important research can be launched to deepen understanding of decision making and problem solving. But perhaps the examples that have been provided are sufficient to convey the promise and significance of this field of inquiry today.

Current Research Programs

Most of the current research on decision making and problem solving is carried on in universities, frequently with the support of government funding agencies and private foundations. Some research is done by consulting firms in connection with their development and application of the tools of operations research, artificial intelligence, and systems modeling. In some cases, government agencies and corporations have supported the development of planning models to aid them in their policy planning--for example, corporate strategic planning for investments and markets and government planning of environmental and energy policies. There is an increasing number of cases in which research scientists are devoting substantial attention to improving the problem-solving and decision-making tools in their disciplines, as we noted in the examples of automation of the processing of bubble-chamber tracks and of the interpretation of mass spectrogram data.

To use a generous estimate, support for basic research in the areas described in this document is probably at the level of tens of millions of dollars per year, and almost certainly, it is not as much as $100 million. The principal costs are for research personnel and computing equipment, the former being considerably larger.

Because of the interdisciplinary character of the research domain, federal research support comes from a number of different agencies, and it is not easy to assess the total picture. Within the National Science Foundation (NSF), the grants of the decision and management sciences, political science and the economics programs in the Social Sciences Division are to a considerable extent devoted to projects in this domain. Smaller amounts of support come from the memory and cognitive processes program in the Division of Behavioral and Neural Sciences, and perhaps from other programs. The "software" component of the new NSF Directorate of Computer Science and Engineering contains programs that have also provided important support to the study of decision making and problem solving.

The Office of Naval Research has, over the years, supported a wide range of studies of decision making, including important early support for operations research. The main source of funding for research in AI has been the Defense Advanced Research Projects Agency (DARPA) in the Department of Defense;

important support for research on applications of A1 to medicine has been provided by the National Institutes of Health.

Relevant economics research is also funded by other federal agencies, including the Treasury Department, the Bureau of Labor Statistics, and the Federal Reserve Board. In recent years, basic studies of decision making have received only relatively minor support from these sources, but because of the relevance of the research to their missions, they could become major sponsors.

Although a number of projects have been and are funded by private foundations, there appears to be at present no foundation for which decision making and problem solving are a major focus of interest.

In sum, the pattern of support for research in this field shows a healthy diversity but no agency with a clear lead responsibility, unless it be the rather modestly funded program in decision and management sciences at NSF. Perhaps the largest scale of support has been provided by DARPA, where decision making and problem solving are only components within the larger area of artificial intelligence and certainly not highly visible research targets.

The character of the funding requirements in this domain is much the same as in other fields of research. A rather intensive use of computational facilities is typical of most, but not all, of the research. And because the field is gaining new recognition and growing rapidly, there are special needs for the support of graduate students and postdoctoral training. In the computing-intensive part of the domain, desirable research funding per principal investigator might average $250,000 per year; in empirical research involving field studies and large-scale experiments, a similar amount; and in other areas of theory and laboratory experimentation, somewhat less.

Research Opportunities: Summary

The study of decision making and problem solving has attracted much attention through most of this century. By the end of World War II, a powerful prescriptive theory of rationality, the theory of subjective expected utility (SEU), had taken form; it was followed by the theory of games. The past forty years have seen widespread applications of these theories in economics, operations research, and statistics, and, through these disciplines, to decision making in business and government.

The main limitations of SEU theory and the developments based on it are its relative neglect of the limits of human (and computer) problem-solving capabilities in the face of real-world complexity. Recognition of these limitations has produced an increasing volume of empirical research aimed at discovering how humans cope with complexity and reconcile it with their bounded computational powers. Recognition that human rationality is limited occasions no surprise. What is surprising are some of the forms these limits take and the kinds of departures from the behavior predicted by the SEU model that have been observed. Extending empirical knowledge of actual human cognitive processes and of techniques for dealing with complexity continues to be a research goal of very high priority. Such empirical knowledge is needed both to build valid theories of how the U.S. society and economy operate and to build prescriptive tools for decision making that are compatible with existing computational capabilities.

The complementary fields of cognitive psychology and artificial intelligence have produced in the past thirty years a fairly well-developed theory of problem solving that lends itself well to computer simulation, both for purposes of testing its empirical validity and for augmenting human problem-solving capacities by the construction of expert systems. Problem-solving research today is being extended into the domain of ill-structured problems and applied to the task of formulating problem representations. The processes for setting the problem agenda, which are still very little explored, deserve more research attention.

The growing importance of computational techniques in all of the sciences has attracted new attention to numerical analysis and to the topic of computational complexity. The need to use heuristic as well as rigorous methods for analyzing very complex domains is beginning to bring about a wide interest, in various sciences, in the possible application of problem-solving theories to computation.

Opportunities abound for productive research in decision making and problem solving. A few of the directions of research that look especially promising and significant follow:

- A substantially enlarged program of empirical studies, involving direct observation of behavior at the level of the individual and the organization, and including both laboratory and field experiments, will be essential in sifting the wheat from the chaff in the large body of theory that now exists and in giving direction to the development of new theory.

- Expanded research on expert systems will require extensive empirical study of expert behavior and will provide a setting for basic research on how ill-structured problems are, and can be, solved.

- Decision making in organizational settings, which is much less well understood than individual decision making and problem solving, can be studied with great profit using already established methods of inquiry, especially through intensive long-range studies within individual organizations.

- The resolution of conflicts of values (individual and group) and of inconsistencies in belief will continue to be highly productive directions of inquiry, addressed to issues of great importance to society.

- Setting agendas and framing problems are two related but poorly understood processes that require special research attention and that now seem open to attack.

These five areas are examples of especially promising research opportunities drawn from the much larger set that are described or hinted at in this report.

The tools for decision making developed by previous research have already found extensive application in business and government organizations. A number of such applications have been mentioned in this report, but they so pervade organizations, especially at the middle management and professional levels, that people are often unaware of their origins.

Although the research domain of decision making and problem solving is alive and well today, the resources devoted to that research are modest in scale (of the order of tens of millions rather than hundreds of millions of dollars). They are not commensurate with either the identified research opportunities or the human resources available for exploiting them. The prospect of throwing new light on the ancient problem of mind and the prospect of enhancing the powers of mind with new computational tools are attracting substantial numbers of first-rate young scientists. Research progress is not limited either by lack of excellent research problems or by lack of human talent eager to get on with the job.

Gaining a better understanding of how problems can be solved and decisions made is essential to our national goal of increasing productivity. The first industrial revolution showed us how to do most of the world's heavy work with the energy of machines instead of human muscle. The new industrial revolution is showing us how much of the work of human thinking can be done by and in cooperation with intelligent machines. Human minds with computers to aid them are our principal productive resource. Understanding how that resource operates is the main road open to us for becoming a more productive society and a society able to deal with the many complex problems in the world today.