NOTICE WARNING CONCERNING COPYRIGHT RESTRICTIONS:
The copyright law of the United States (title 17, U.S. Code) governs the making of photocopies or other reproductions of copyrighted material. Any copying of this document without permission of its author may be prohibited by law.

PROBLEM FORMULATION AND ALTERNATIVE GENERATION IN THE DECISION MAKING PROCESS

Technical Report AIP - 43

Herbert A. Simon

Deparment of Psychology
Carnegie Mellon University
Pittsburgh, Pa. 15260

30 June 1988

This research was supported by the Computer Sciences Division, Office of Naval Research and DARPA under Contract Number N00014-86-K-0678; and the Defense Advanced Research Projects Agency, Department of Defense, ARPA Order 3597, monitored by the Air Force Avionics Laboratory under contract F-33615-81-K-1539. Reproduction in whole or in part is permitted for purposes of the United States Government. Approved for public release; distribution unlimited.

006.3

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION: Unclassified

1b. RESTRICTIVE MARKINGS:

2a. SECURITY CLASSIFICATION AUTHORITY:

2b. DECLASSIFICATION / DOWNGRADING SCHEDULE:

3. DISTRIBUTION/AVAILABILITY OF REPORT: Approved for public release; Distribution unlimited

4. PERFORMING ORGANIZATION REPORT NUMBER(S): AIP - 43

5. MONITORING ORGANIZATION REPORT NUMBER(S):

6a. NAME OF PERFORMING ORGANIZATION: Carnegie-Mellon University

6b. OFFICE SYMBOL (If applicable):

7a. NAME OF MONITORING ORGANIZATION: Computer Sciences Division, Office of Naval Research

6c. ADDRESS (City, State, and ZIP Code): Department of Psychology, Pittsburgh, Pennsylvania 15213

7b. ADDRESS (City, State, and ZIP Code): 800 N. Quincy Street, Arlington, Virginia 22217-5000

8a. NAME OF FUNDING/SPONSORING ORGANIZATION: Same as Monitoring Organization

8b. OFFICE SYMBOL (If applicable):

9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER: N00014-86-K-0678

8c. ADDRESS (City, State, and ZIP Code):

10. SOURCE OF FUNDING NUMBERS: p4000ub201/7-4-86

PROGRAM ELEMENT NO.	PROJECT NO.	TASK NO.	WORK UNIT ACCESSION NO
N/A	N/A	N/A	N/A

11. TITLE (Include Security Classification): PROBLEM FORMULATION AND ALTERNATIVE GENERATION IN THE DECISION MAKING PROCESS

12. PERSONAL AUTHOR(S): Simon, Herbert A.

13a. TYPE OF REPORT: Technical

13b. TIME COVERED: FROM 86Sept15 TO 91Sept14

14. DATE OF REPORT (Year, Month, Day): 1988 June 30

15. PAGE COUNT: 10

16. SUPPLEMENTARY NOTATION:

17. COSATI CODES

FIELD	GROUP	SUB-GROUP

18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number): cognitive psychology, economics, decision making, subjective expected utility

19. ABSTRACT (Continue on reverse if necessary and identify by block number)

Classical and neoclassical economic theory, as well as statistical decision theory, through their neglect of human bounded rationality -- the vast disparity between human computing capabilities and the complexity of our world -- both give a seriously distorted picture of human decision making and omit at least three components of the decision making process that are of central importance. In this paper, I will outline what is known, today, about these neglected aspects of human decision-making. A great deal is known, mainly as a result of the progress of cognitive science in the last generation. Economics can make rapid progress by drawing upon this storehouse of new knowledge to reconstruct and expand its foundations.

20. DISTRIBUTION/AVAILABILITY OF ABSTRACT: ☐ UNCLASSIFIED/UNLIMITED ☒ SAME AS RPT ☐ DTIC USERS

21. ABSTRACT SECURITY CLASSIFICATION:

22a. NAME OF RESPONSIBLE INDIVIDUAL: Dr. Alan L. Meyrowitz

22b. TELEPHONE (Include Area Code): (202) 696-4302

22c. OFFICE SYMBOL: N00014

DD FORM 1473, 84 MAR

83 APR edition may be used until exhausted.
All other editions are obsolete.

SECURITY CLASSIFICATION OF THIS PAGE

Unclassified

UNIVERSITY LIBRARIES
CARNEGIE MELLON UNIVERSITY
PITTSBURGH, PA 15213-3890

Problem Formulation and Alternative Generation in the Decision Making Process

Herbert A. Simon
Carnegie-Mellon University

The standard theory of decision making, widely employed in economics and statistics, is concerned with the rational choice of an action when the set of alternatives for choice and the evaluation function (utility function) are assumed to be given. In the *certainty* condition, the consequences, in terms of the evaluation function, of choosing any alternative are assumed to be known. In the *uncertainty* condition, the probability distribution of the outcomes of choosing any alternative is assumed to be known. In the standard theory, rationality is assumed to consist, under the circumstances of uncertainty, of choosing the alternative that will produce the highest expected value of the outcome. This procedure is usually called *subjective expected utility* (SEU) maximization.

The SEU theory may be adopted as a description of how people actually make decisions (the positive theory), or a prescription of how they ought (rationally) to make decisions (the normative theory). The evidence is overwhelming that it is not a correct descriptive theory, either as applied to people making the ordinary decisions of everyday economic life, or as applied to managers conducting the affairs of firms.

Limitations of the SEU Theory

Its defects are numerous and serious, along several dimensions. First, behaving according to the SEU rule would impose computational demands upon economics actors that they simply cannot sustain. In most real-world situations, the consequences of actions can be calculated in only the roughest and most approximate ways, and it is ludicrous to suppose that we very often could compute joint probability distributions of our actions for all sets of possible future contingencies. The empirical evidence shows that people substitute for these wholly infeasible computations simple rules of thumb or heuristics, like the heuristic of satisficing -- settling for "satisfactory" courses of action.

Second, the empirical evidence shows that people do not have comprehensive, consistent utility functions that would enable them to order any set of alternatives that is set before them. Preferences are, in fact, highly contingent upon environmental context and upon the focus of attention of the actors to particular features of the environment. (These characteristics of choice can also be explained in terms of computational limits, in this case the inability of actors to take account simultaneously of all features of present and future environments.)

Third, the SEU theory does not provide an explanation of three critical dimensions of choice. It does not explain how the utility function (if there were one) could obtain its particular content and shape -- it includes no theory of the origin of values. It does not explain how the occasions for decision arise; how items are placed on the agenda for decision making. And it also does not explain where the alternatives of choice come from; it simply postulates them as givens of the situation. In previous papers, I have criticized all of these deficiencies of SEU theory, and have made some proposals for an alternative theory of rational decision theory -- a theory of *bounded* rationality -- that takes into account human computational limits as well as abilities.

In the present paper, I would like to focus on just two aspects of the decision making process -- the ways in which problems are formulated and alternatives generated. In particular, I would like to sketch out a positive theory of how problems get on the agenda and are defined, and how the alternatives of choice are generated. I must emphasize the word "sketch," since we are far from having the body of empirical evidence that would permit a definitive theory to be defined.

All of the criticisms of the descriptive version of SEU theory that I have just set out apply also to the normative theory. Human "irrationality" -- i.e., departure from the norms of the SEU theory -- is not just some form of human stubbornness, or refusal to be rational. Rather, it is a consequence of inability: inability to carry out the computations called for by the theory. That inability prevents us from following the dictates of SEU theory even if we accept it as a norm. Nor does the availability of computers remove the difficulty. Whatever their computational powers, present or future computers are no match for the complexity of the real world. They (and we) are forever condemned to carrying out our reasoning with highly simplified models of tiny parts of the entire reality that confronts us. Hence, my discussion of problem formulation and alternative generation will have relevance for the normative as well as for the descriptive theory of rational decision making.

The Agenda

We can form a first approximation to a theory of how items get on the agenda for decision by observing the behavior of birds and mammals. Living creatures are endowed with needs and wants (hunger, thirst, sex drives, fright) that express themselves through messages sent from sense organs and internal organs to the brain. Human beings have these same needs and wants, and others as well, but let us keep for the moment to a simple picture.

Agenda-setting in Simpler Organisms

When a creature feels hunger pangs, and especially when messages relating to other drives are absent, it turns its attention to satisfying its hunger. But in a state of nature, it is usually not confronted with a supermarket full of foods among which it can choose. Instead, it must search for food -- that is, it must create the alternatives that will have utility for it. Its search is directed by its knowledge of the environment, of where food is likely to be found. And it may have food preferences, looking first at places where preferred foods are most often encountered.

If two or more needs express themselves at the same time, then the organism must decide which to put first on the agenda, since its repertory of search activities can usually allow it to deal with only one of them at a time. These priorities are usually settled by simple rules -- attend first to the need whose inventory of satisfiers is more nearly exhausted. Hence, because of typical storage capacities, thirst generally takes precedence over hunger, and the need for oxygen over both of the others.

Birds and mammals set their agendas in very much the manner of the familiar two-bin inventory system of industrial practice. For each need or want there is an "order point" and an "order quantity." At some level of deprivation, signals are sent to the central nervous system that secure attention to the want -- unless more urgent signals are present. If the want is not attended to immediately, the signals become gradually more insistent until the want secures first priority. Then a burst of activity is triggered, which ordinarily persists until the need is met and the inventory restored to a satisfactory level.

Notice that this system for fixing the agenda requires nothing like a comprehensive utility function. To be sure, the urgency of needs is compared, but only for the purpose of setting search priorities, and a simple mechanism that will signal deprivation and gradually increase the intensity of its signal is all that is required. Nothing needs to be maximized. This crude procedure will work satisfactorily (not optimally) as long as the organism has ample time to carry out its searches in order to meet all of its wants before inventories are completely exhausted. With some slack available, searches can be interrupted in the face of more urgent demands (e.g., the need to escape from danger).

With this kind of agenda setting, there are some possibilities for tuning the system through evolution. Camels can acquire, through natural selection, an enhanced capacity for inventorying water. The urgency of signals to attend to wants can be adjusted to the real-time demands for satisfying these wants. Danger signals can, as I have just suggested, acquire a high priority. But there is no need to solve any sophisticated optimization problem in order to set the agenda for an organism that can survive.

Agenda-Setting in Man

It will be objected that we human beings are much more complex than the creatures I have been describing, along at least two dimensions. We have many needs and wants that are not instinctual, but learned. And we can project ourselves into the future to do some kind of planning.

A mere increase in number of wants does not appear to complicate the agenda-setting problem, provided that satisfaction of all of them is not essential for survival. Those that do not clamor loudly enough, though unsatisfied, simply never get on the agenda. For example, I never seem to find on my agenda the task of searching for a pipe organ (although I think I would enjoy one if I had it). We can classify most potential agenda items as either problems or opportunities. Problems are items that, if not attended to, will cause trouble. Opportunities are items that, if attended to, may increase our satisfaction or profit or probability of surviving.

Nor must it be supposed that there is a definite list of opportunities, or even problems, among which the priorities are worked out. Neither problems nor opportunities can be considered for the agenda unless they are noticed, and except for those that attract attention by means of an internal signaling system, they must be picked out from a complex external sensory environment. Until they are noticed, opportunities are not opportunities.

The most realistic assumption about the world in which we actually live is that in any given period of time we notice only a tiny fraction of the opportunities that are objectively present, and even only a small part of the problems. A major initial step -- and by no means an assured one -- in the process of technological or social invention is to extract opportunities from the confusion of the environment -- to attend to the right cues.

We have only the beginnings of a theory of how opportunities (or problems) are noticed.

The greatest progress has been made in the domain of scientific discovery. It is well known that one of the mechanisms that focuses human attention on important problems is surprise. Fleming, in his laboratory, noticed a Petri dish in which the bacteria were disintegrating by lysis. He was surprised -- there was no obvious reason why the bacteria should be dying. On the edge of the dish, near where the lysis was occurring, was a mold, of the genus Penicillium.

What are the conditions for surprise? We are surprised when we are knowledgeable about a situation and something unusual (contrary to our knowledge) is occurring in that situation. Fleming was knowledgeable about bacteria and molds, and nothing in his knowledge led him to expect that bacteria would die in the presence of a mold. "Accident," as Pasteur said, "happen to the prepared mind." Surprise put the problem (or opportunity) of explaining why the bacteria were dying on Fleming's research agenda; it would not have been noticed by anyone without his knowledge. And his story is not an isolated one in the history of science. A great many discoveries, including many of the first order of magnitude, secured their place on the agenda through surprise.

Attention in an Information-Rich World

What we need, and lack, is a generalization from the surprise mechanism to a more comprehensive theory of what it is that focuses human attention on specific parts of the environment. It is especially characteristic of the contemporary environment of industrialized nations that all of us are surrounded by -- even drowned in -- a sea of information, only an infinitesmal part of which can be attended to. However we may wish to have certain kinds of information that are not available (e.g., reliable horoscopes), on the whole, the scarce factor in our decision making is not information but attention. What we attend to, by plan or by chance, is a major determinant of our decisions.

The general scarcity of attention suggests that people and organizations can enhance the quality of their decision making by searching systematically, but selectively, among potential information sources to find those that deserve most careful attention, and that might provide items for the agenda. This is a major function of so-called "intelligence" units in organizations, and also of research and development units, and even planning units.

For example, it is seldom that a company laboratory is the major source of basic discoveries from which new products can be developed. More often, the laboratory serves as an intelligence link to the world of academic and other science from which ideas may be drawn. Its task is to observe and communicate with that world, and to notice opportunities that are presented by it. Of course, the laboratory itself is also a window on the world of opportunity, but a rather narrow window unless supplemented by close interaction with the scientific community.

A common responsibility of planning units, not always explicitly recognized in the definition of their function, is early recognition of problems. One mechanism for problem recognition is to build computational models of the system of interest and to use them to make predictions. But again, this is only one mechanism of several. Selective surveillance of information available in the environment may provide a more reliable early warning system than prediction.

Perhaps I have said enough to demonstrate that a theory of agenda formation -- which is, in turn, a theory of attention focusing -- is an essential component in a theory of rational decision. It is a component that is missing from the SEU framework, and that consequently has not been much developed in economics. If we wish to borrow ideas that may be useful for economics and decision theory, we must look for them in artificial intelligence and cognitive science -- for example, in recent research on the processes of scientific discovery.

Problem Formulation

Perhaps even less is known today about the mechanisms of problem formulation than about agenda-setting processes. Of course, if the item placed on the agenda by the attention-directing mechanisms is of a familiar kind, standard procedures may be available for casting it in the form of a solvable problem. We all learned in secondary school an algorithm for solving linear algebra equations. If we can formulate a problem as an equation, then we know how to solve it.

Or, to return to items placed on the agenda by surprise, scientists have a rather standard procedure for exploiting surprises. If a surprising phenomenon is encoountered, first try to characterize the scope of the phenomenon. If bacteria are dying in the presence of a mold, what kinds of bacteria are affected? (Fleming found that many kinds were.) What kinds of mold? (Evidently only the mold Penicillium.) And when the scope of the phenomenon has been defined, try to find its mechanism. (Can we extract from Penicillium, by crushing, treating with alcohol, heating, etc., a substance that retains, or even enhances, its effect upon bacteria? If we find such a substance, can we purify it and characterize it chemically? A whole sequence of experiments, first by Fleming, then by Florey and Chaim, achieved just this.)

Some problems are very hard as the world presents them, but very easy when they are reformulated properly. The Mutilated Checkboard problem is a celebrated example. Consider a checkerboard (8x8) and 32 dominoes, each of which covers exactly two squares of the board. Clearly, we can entirely cover the checkerboard with the dominoes. Now suppose that two squares are cut out of the checkerboard -- the upper left corner and the lower right corner. Can we cover the remaining 62 squares with 31 dominoes?

We cannot, but the answer is not obvious. None of us would have the patience to demonstrate the impossibility by trying all possible coverings. We must find some other way. Let us abstract the problem, and consider just the number of dominoes, the number of black squares, and the number of red squares. Each domino will cover exactly one black and one red square. But the two squares we remove will be of the same color (they are at opposite ends of a diagonal). Hence there will now be two fewer squares of one color than of the other (let's say 30 black and 32 red). There is no way in which 31 dominoes can cover any but 31 black and 31 red squares, hence a covering is impossible.

Problem representations, like the problems themselves, are not presented to us automatically. They are either retrieved from memory, when we recognize a situation as being of a familiar kind, or they must be discovered through selective search. Problem formulating is itself a problem solving task.

American firms, and European ones, have been fully aware, for some years now, of the challenge presented by Japanese and other Far Eastern competition. The problem is on the agenda, but finding an appropriate problem representation is a more difficult task, and one that has not yet been fully completed. Is the problem one of quality control, of manufacturing efficiency, of managerial style, of worker motivation, of wage levels, of exchange rates, of foreign trade regulations, of investment incentives? The list is endless; and the search for solutions will depend on the diagnosis.

Almost from the beginnings of economics as a field of inquiry, economists have given attention to the problem of unemployment. Satisfactory solutions to that problem in market economies have surely not been found, and a major reason is that we do not even have a good formulation of the problem. Within classical and neoclassical models, unemployment can be introduced only by way of highly arbitrary, and empirically unsubstantiated, assumptions. We can have no confidence (and few of us, today, have such confidence) that such models provide a problem formulation within which solutions to the unemployment problem can be found.

In schools of business and of management, the SEU framework is taught as a general formulation for problems of management decision making. If it is as inadequate a formulation as I have argued that it is, then it seems unlikely to provide the right framework for managers' representations of the problems they must solve.

From all of these examples, and others that could be proposed, we can easily conclude that developing a veridical theory of problem representation must stand high on the agenda of decision-making research.

Generating Alternatives for Search

Having discussed the problems of agenda setting and of problem formulation, we have still not reached the threshold of the classical SEU model of decision making. Since in the real world, alternatives are seldom given to the decision maker, we must consider how they are invented or discovered.

Of course, there are situations in the real world where the alternatives are, for all practical purposes, given. The shelves of the market define the range of foodstuffs among which we can choose (although even here, it may be an onerous task to scan all of them, and different markets may carry different brands). There is only a finite number, almost manageable, of different makes of automobiles to choose among, and we can find out what they are. (But did we consider the whole list when we bought our last car? And if we know what kinds of cars are available, we may not know, without extensive search, at what prices we can buy them.)

House hunting and job hunting are market activities that normally require extensive search among an ill-bounded set of alternatives. A graduating student, searching for his or her first job must not only have procedures for discovering prospective employers, but stop rules for determining when the search should end, and procedures for obtaining relevant information about each employment opportunity.

In short, most economic activity is problem solving activity in which alternatives are not given but are generated through selective search. The problem-solving character of alternative generation is most apparent when it takes the form of design. Many economic products are not manufactured for the open market or to be sold from the shelf, but are designed specially on contract with a particular customer. And even shelf goods have to be conceived and designed, a task that becomes central in industries, like clothing or pharmaceuticals, where new products are constantly coming into the market.

During the past thirty or forty years, research in cognitive science has taught us a great deal about problem-solving processes, including the processes of design. In any problem solving process, we have a goal, or set of goals, formulated as tests to be applied to prospective solutions. We have a generator that produces these prospective solutions, and which can be very simple or highly complex. It can simply produce items, one by one, for test and acceptance or rejection, or it can synthesize prospective solutions, step-by-step, applying tests of progress along the way to direct the search. The more we know about the problem space in which we are searching, the more information we can extract from that space to direct the search, and the more efficient the exploration will be.

The problem solving we understand best concerns well-structured problems. Problems are well structured when the goal tests are clear and easily applied, and when there is a well-defined set of generators for synthesizing solutions. Problems are ill structured to the extent that they lack these characteristics. Many, if not most, of the problems that confront us in the everyday world are ill structured. An architect design a house, an engineer designing a bridge or a power-generating station, a chemist seeking a molecule with desired properties and a way of manufacturing it cheaply -- all of these are solving problems with many ill-defined components.

To the best of our knowledge, the underlying processes involved in ill-defined problem solving are no different from those involved in well-defined problem solving, but the hypothesis that this is so still lacks conclusive demonstration. Sometimes it is argued, to the contrary, that ill-defined problem solving involves processes that are "intuitive," "judgmental," or even "creative," and that these are fundamentally different from the run-of-the-mill, routine, logical or analytic processes employed in well-structured problem solving.

We can refute this argument because we have strong evidence about the nature of intuitive, judgmental, and creative processes that shows how they are achieved. Experts in any domain have stored in their memories a very large number of pieces of knowledge about that domain. For domains in which it has been possible to measure the knowledge, at least crudely, it appears that the expert may have 50,000 or even 200,000 "chunks" of information -- but probably not 5,000,000.

This information is held in memory in a particular way: it is associated with an "index." When the expert is confronted with a situation in his or her domain, various features or cues in the situation will attract attention. A chess player will notice an "open file," "doubled pawns," or a "pinned Knight." Each feature that is noticed will give access to the chunks of information stored in memory that are relevant to that cue. If the accountant sees a low cash balance on the balance sheet, that cue will remind him of whatever he knows about liquidity problems.

The ability, often noticed, of the expert to respond "intuitively," and often very rapidly, to situations -- with a relatively high degree of accuracy and correctness -- is simply the product of this stored knowledge and the problem-solving by recognition that it permits. Intuition, judgment, creativity are basically expressions of these capabilities for recognition and response based upon experience and knowledge. There is nothing more mysterious about them than about our recognizing our friend "instantly" when we meet him on the street.

Nor do we need to postulate two problem-solving styles, the analytic and the intuitive. The power of analysis depends on expert knowledge for its speed and effectiveness. Without knowledge, available by recognition, only tiny, slow, painful steps can be taken in reasoning. We may see relative differences among experts in their reliance on analysis as against recognition, but we may expect to find large components of both, closely intermingled, in virtually all expert behavior.

In sum, economics needs, in addition to a theory of attention direction and a theory of problem formulation, a theory of alternative generation. Fortunately, the fundamentals of such a theory are already available from cognitive science, and can be borrowed for application in economics and management.

Conclusion

Classical and neoclassical economic theory, as well as statistical decision theory, through their neglect of human bounded rationality -- the vast disparity between human computing capabilities and the complexity of our world -- both give a seriously distorted picture of human decision making and omit at least three components of the decision making process that are of central importance. In this paper, I have tried to outline what is known, today, about these neglected aspects of human decision-making paper. A great deal is known, mainly as a result of the progress of cognitive science in the last generation. Economics can make rapdi progress by drawing upon this storehouse of new knowledge to reconstruct and expand its foundations.