A BEHAVIORAL MODEL OF RATIONAL CHOICE

Herbert A. Simon

P-365

20 January 1953

The RAND Corporation

SANTA MONICA · CALIFORNIA

A BEHAVIORAL MODEL OF RATIONAL CHOICE

Herbert A. Simon

~~Summary~~: A model is proposed for the description of rational choice by organisms of limited computational ability.( )

The "flavor" of various models of rational choice is primarily due to the specific kinds of assumptions that are introduced as to the "givens" or constraints within which rational adaptation must take place. Among the common constraints--which are not themselves the objects of rational calculation--are the set of alternatives open to choice, the relationships that determine the payoffs as a function of the alternative that is chosen, and the preference-orderings among payoffs. The selection of particular constraints and the rejection of others for incorporation in the model of rational behavior involves implicit assumptions as to what variables the rational organism "controls"--and hence can optimize as a means to rational adaptation--and what variables it must take as fixed. It also involves assumptions as to the character of the variables that are fixed. For example, by making different assumptions about the amount of information the organism has with respect to the relation between alternatives and payoffs, optimization might involve selection of a certain maximum, the maximum of an expected value, or a minimax.

Another way of characterizing the givens and the behavior variables is to say that the latter refer to the organism itself, the former to its

---

*The ideas embodied in this memorandum were initially developed in a series of discussions with Herbert Bohnert, Norman Dalkey, Gerald Thompson, and Robert Wolfson during the summer of 1952. These collaborators deserve a large share of the credit for whatever merit this approach to rational choice may possess.

environment. But if we adopt this viewpoint, we must be prepared to accept the possibility that what we call "the environment" may lie, in part, within the skin of the biological organism. That is, some of the constraints that must be taken as givens in an optimization problem may be physiological and psychological limitations of the organism (biologically defined) itself. For example, the maximum speed at which an organism can move establishes a boundary on the set of its available behavior alternatives. Similarly, limits on computational capacity may be important constraints entering into the definition of rational choice under particular circumstances. It is the purpose of this memorandum to explore possible ways of formulating the process of rational choice in situations where we wish to take explicit account of the "internal" as well as the "external" constraints that define the problem of optimization for the organism.

Whether our interests lie in the normative or in the descriptive aspects of rational choice, the construction of models of this kind should prove instructive. Because of the psychological limits of the organism (particularly with respect to computational and predictive ability), actual human rationality-striving can at best be an extremely crude and simplified approximation to the kind of global rationality that is implied, for example, by game-theoretical models. While the approximations that organisms employ may not be the best—even at the levels of computational complexity they are able to handle—it is probable that a great deal can be learned about possible mechanisms from an examination of the schemes of approximation that are actually employed by human and other organisms.

In describing the proposed model, we shall begin with elements it has in common with the more global models, and then proceed to introduce simplifying assumptions and (what is the same thing) approximating procedures.

Primitive Terms and Definitions

1. A point set, A (the set of behavior alternatives).

2. A point set, $\breve{A}$. $\breve{A} \subseteq A$. (the set of behavior alternatives that the organism "considers").

3. A point set, S. (the set of possible future states, or "outcomes" of choice).

4. A real function V (s) on the elements $s \in S$. (the pay-off. For many purposes, we need only an ordering relation on pairs of elements of S, but for the moment we postulate a cardinal pay-off.)

5. For each element $a \in A$, a mapping of a on to S_a, $S_a \subseteq S$. (S_a is the set of possible outcomes of a).

6. For each $a \in A$ and $s \in S$ a real function, $P_a(s)$, with $P_a(s) \geq 0$; $\sum_{s \text{ in } S_a} P_a(s) = 1$. (the probability that s will occur if a is chosen)

Attention is directed to the three-fold distinction drawn by the definitions among A, S, and V. In the representation of a game, in reduced form, by its payoff matrix, the set S consists of the cells of the matrix, A the strategies of the first player, and V the values in the cells. The set S_a is then the set of cells in the a^{th} row. By keeping in mind this interpretation, the reader may compare the present formulation with "classical" game theory.

With these elements, we can define procedures of rational choice corresponding to the ordinary game-theoretical and probabilistic models.

A. <u>Game-theoretical Choice Process</u>. Select an $\hat{a}$, $\hat{a} \in A$, such that

$$\hat{V}(\hat{a}) = \min_{s \in S_{\hat{a}}} V(s) = \max_{a \in A} \min_{s \in S_a} V(s)$$

The terms in (2) and (6) do not play any role here.

B. <u>Probabilistic Choice Process</u>. Select an $\hat{a}$, $\hat{a} \in A$, such that:

$$\hat{V}[\hat{a}] = \sum_{s \in S_{\hat{a}}} V(s) P_{\hat{a}}(s) = \max_{a \in A} \sum_{s \in S_a} V(s) P_a(s)$$

C. Choice Under Certainty. Suppose each a maps on to a single $s_a \in S$. then B reduces to:

Select an $\hat{a}$, $\hat{a} \in A$ such that:

$$\hat{V}[\hat{a}] = V(s_{\hat{a}}) = \underset{a \text{ in } A}{\mathrm{Max}}\ V(s_a)$$

The Essential Simplifications

We now introduce some modifications that appear to correspond to observed behavior processes in humans, and that lead to substantial computational simplifications in the making of a choice. There is no implication that human beings use all of these modifications and simplifications all the time. The point is rather that these are procedures which, it is believed, are sometimes employed by human beings in complex choice situations to find an approximate model of manageable proportions.

I. "Flat" Pay-off Functions.

One route to simplification is to assume that V (s) necessarily assumes one of three values, $\left\{ \begin{array}{c} +1 \\ 0 \\ -1 \end{array} \right\}$, for all $s \in S$. Depending on the circumstances, we might want to interpret these three values, as (a) $\left\{ \begin{array}{l} \text{win} \\ \text{draw} \\ \text{lose} \end{array} \right\}$ or (b) $\left\{ \begin{array}{l} \text{very satisfactory} \\ \text{acceptable} \\ \text{unsatisfactory} \end{array} \right\}$.

As an example of (a), let S represent the possible positions in a chess game at White's 20th move. Then a (+ 1) position is one in which White possesses a strategy leading to a win whatever Black does. A (0) position is one in which White can enforce a draw, but not a win. A (-1) position is one in which Black can force a win.

As an example of (b) let S represent possible prices for a house an individual is selling. He may regard \$15,000 as an "acceptable" price, anything over this amount as "very satisfactory", anything less as "unsatisfactory". In psychological theory we would fix the zero point at the "aspiration level", in economic theory we would fix the zero point at the price which evokes indifference between selling and not selling (an opportunity cost concept).

The objection may be raised that, although $16,000 and $25,000 are both "very satisfactory" prices for the house, a rational individual would prefer to sell at the higher price, and hence, that the simplified pay-off function is an inadequate representation of the situation. The objection may be answered in several different ways, each answer corresponding to a class of situations in which the "flat" function might be appropriate.

First, the individual may not be confronted simultaneously with a number of buyers offering to purchase the house at different prices, but may receive a sequence of offers, and may have to decide to accept or reject each one before he receives the next. (Or, more generally, he may receive a sequence of pairs or triplets or n-tuples of offers, and may have to decide whether to accept the highest of an n-tuple before the next n-tuple is received.) Then, if the elements s correspond to n-tuples of offers, V (s) would be 1 whenever the highest offer in the n-tuple exceeded the "acceptance price" the seller had determined upon at that time. We could then raise the further question of what would be a rational process for determining the acceptance price.[1]

---

[1] I propose to deal with the problem of a rational process for determining aspiration levels in a subsequent paper. See also the discussion below of the existence and uniqueness of solution.

Second, even if there were a more general pay-off function, W (s), capable of assuming more than three different values, the simplified V (s) might be a satisfactory approximation to W (s). Suppose, for example, that there were some way of introducing a cardinal utility function, defined over S, say U (s). Suppose further that U (W) is a monotonic increasing function with a strongly negative second derivative. Then $V(s) = V\{W(s)\}$ might be the following approximation:

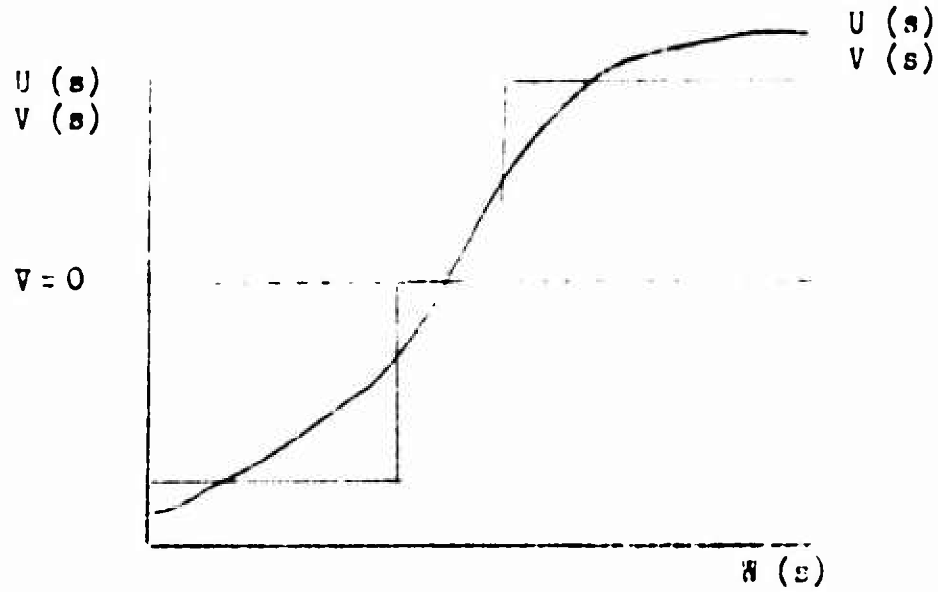

When a simplified V (s), assuming only the values $\left\{\begin{matrix} +1 \\ 0 \\ -1 \end{matrix}\right\}$, is admissible, under the circumstances just discussed or under other circumstances, then a rational decision-process could be defined as follows:

D. α. Look for a subset $S' \subseteq S$ such that $V(s) = 1$ for all $s \in S'$.

β. Look for an $\underline{a} \in \breve{\underline{A}}$ that maps on an $S_a \subseteq S'$.

The procedure does not, of course, guarantee the existence or uniqueness of an $\underline{a}$ with the desired properties.

Parenthetically, it may be noted that if we start with a more general pay-off function, W (s), it is not necessary to introduce V (s) explicitly.

Working directly with W (s), we might introduce comparable rules of the following general forms:

E. -α-. Look for a $S' \subset S$ such that $V(s) \geq k$ for all $s \in S'$, where k is some constant; or

F. -α-. Look for a $S_a \subset S$ such that

$$\sum_{s \in S_a} V(s)\, P_a(s) \geq k \quad ; \text{ or}$$

G. -α-. Look for a $S_a \subset S$ such that

$$\sum_{s \in S_a} P_a(s) < \mathcal{E}.$$

$$V(s) < k$$

II. Information Gathering.

One element of realism we may wish to introduce is that, while V (s) may be known in advance, the mapping of A on subsets of S may not. In the extreme case, at the outset each element, a, may be mapped on the whole set, S. We may then introduce into the decision-making process information-gathering steps that produce a more precise mapping of the various elements of A on non-identical subsets of S. If the information-gathering process is not costless, then one element in the decision will be the determination of how far the mapping is to be refined.

Now in the case of the simplified pay-off functions, $\begin{Bmatrix} +1 \\ 0 \\ -1 \end{Bmatrix}$, the information-gathering process can be streamlined in an important respect. First, we suppose that the individual has initially a very coarse mapping of A on S. Second, he looks for an $S' \subset S$ such that $V(s) = 1$ for $s \in S'$. Third, he gathers information to refine that part of the mapping of A on S in which elements of S' are involved. Fourth, having refined the mapping, he looks for an a that maps on to a subset of S'.

Under favorable circumstances, this procedure may require the individual to gather only a small amount of information--an insignificant part of the whole mapping of elements of A on individual elements of S. If the search for an a having the desirable properties is successful, he is certain that he cannot better his choice by securing additional information.

It appears that the decision process just described is one of the important means employed by chess players to select a move in the middle and end game. Let A be the set of moves available to White on his 20th move. Let S be a set of positions that might be reached, say, by the 30th move. Let S' be some subset of S that consists of clearly "won" positions. From a very rough knowledge of the mapping of A on S, White tentatively selects a move, a, that (if Black plays in a certain way) maps on S'. By then considering alternative replies for Black, White "explores" the whole mapping of a. His exploration may lead to points, s, that are not in S', but which are now recognized also as winning positions. These can be adjoined to S'. On the other hand, a sequence may be discovered that permits Black to bring about a position that is clearly not "won" for White. Then White may reject the original point, a, and try another.

Whether this procedure leads to any essential simplification of the computation depends on certain empirical facts about the game. Clearly all positions can be categorized as "won", "lost" or "drawn" in an objective sense. But from the standpoint of the player, positions may be categorized as "clearly won", "clearly lost", "clearly drawn", "won or drawn", "drawn or lost", and so forth--depending on the adequacy of his mapping. If the "clearly won" positions represent a significant subset of the objectively "won" positions, then the combinatorics involved in seeing whether a position can be transformed into a clearly won position, for all possible replies by Black, may not be unmanageable.[2] The advantage of this procedure over the more common notion (which may, however, be applicable in the opening) of a general valuation function for positions, taking on values from -1 to 1, is that it implies much less complex and subtle evaluation criteria. All that is required is that the valuation function be reasonably sensitive in detecting when a position in one of the three states--won, lost, or drawn--has been transformed into a position in another state. The player, instead of seeking for a "best" move, needs only to look for a "good" move.

Within the usual game-theoretical framework, it is difficult to define such terms as "attack", "plan of attack", "initiative", etc. In the present framework, a plan of attack is an alternative, a that maps the present position on some set of future positions which is regarded as satisfactory. A player has the initiative when he has a plan of attack that he thinks leads only to satisfactory positions. Since his calculations have been incomplete, and since he may have misvalued certain future positions (has regarded as "clearly won" positions that are defensible by his opponent), the opponent operates on the defensive by trying to find those paths that the attacker has inaccurately analyzed.

The term "counter-attack" suggests that there is yet another approximating mechanism that is involved in the computations. The attacker may

---

[2] This possibility has been realized in two chess combinations of about eight moves in length with which I have experimented. Of 5^{16} or more legal sequences of plays available, only about 100 needed to be explored in each case for a complete analysis of the position.

consider only one part of the board in determining what positions are clearly won, and may ignore other parts of the position. His opponent may expose the fallacy of such a ceteris paribus assumption by developing an attack in this other area. A chessboard is just large enough so that there is a somewhat loose "coupling" between the two wings. However, these considerations go beyond the present model. They are mentioned simply to indicate that many other approximating mechanisms may be involved in choice besides those explicitly introduced here.

III. Partial Ordering of Payoffs.

Instead of a scalar pay-off function, V (s), we might have a vector function, V (s) ; where V has the components V_1, V_2, ... A vector pay-off function may be introduced to handle a number of situations:

(1) In the case of a decision to be made by a group of persons, the components may represent the pay-off functions of the individual members of the group;

(2) In the case of an individual, he may be trying to implement a number of values that do not have a common denominator--e.g., he compares two jobs in terms of salary, climate, pleasantness of work, prestige, etc.;

(3) Where each behavior alternative, a, maps on a set of n possible consequences, S_a; we may replace the model by one in which each alternative maps on a single consequence, but each consequence has as its payoff the n-dimensional vector whose components are the payoffs of the elements of S_a.

This representation exhibits a striking similarity among these three important cases where the traditional maximizing model breaks down for lack of a complete ordering of the payoffs. The first case has never been satisfactorily treated--the theory of the n-person game is the most ambitious attempt to deal with it, and the weak welfare principles are attempts to avoid it. The second case is usually handled by superimposing a complete ordering on the points in the vector space ("indifference curves"). The third case has been handled by introducing probabilities as weights for summing the vector components, or by using principles like minimaxing regret.

An extension of the notion of a simplified pay-off function permits us to treat all three cases in much the same fashion. Suppose we regard a payoff as <u>satisfactory</u> provided that $V_i \geq k_i$ for all i. Then a reasonable decision rule is the following:

Look for a subset $S' \subseteq S$ such that $\bar{V}(s)$ is satisfactory for all $s \in S'$ (i.e., $\bar{V}(s)_{s \in S'} \geq k$).

Then look for an $\underline{a} \in A$ such that $S_a \subseteq S'$.

Again existence and uniqueness of solutions are not guaranteed.

In the first of the three cases mentioned above, the satisfactory payoff corresponds to what I have called a <u>viable</u> solution in "A Formal Theory of the Employment Relation" and "A Comparison of Organization Theories".[3] In the second case, the components of <u>a</u> define the <u>aspiration levels</u> with respect to several components of payoff. In the third case (in this case it is most plausible to assume that all the component of k will be equal), k_i may be interpreted as the <u>minimum guaranteed payoff</u>-- also an aspiration level concept.

---

[3] Econometrica, 19:293 - 305 (July, 1951), and Review of Economic Studies, 20:40-49 (1952-53, #1)

Existence and Uniqueness of Solutions

Throughout our discussion we have admitted decision procedures that do not guarantee the existence or uniqueness of solutions. This was done in order to construct a model that parallels as nearly as possible the decision procedures actually used by humans in complex decision-making settings. We now proceed to add supplementary rules to fill this gap.

I. Obtaining a Unique Solution.

In most global models of rational choice, all alternatives are evaluated before a choice is made. In actual human decision-making alternatives are often examined sequentially. We may, or may not, know the mechanism that determines the order of procedure. In any case, we may regard the first satisfactory alternative that is evaluated as such as the one actually selected.

If a chess player finds an alternative that leads to a forced mate for his opponent, he generally adopts this alternative without worrying about whether another alternative also leads to a forced mate. In this case we would find it very hard to predict which alternative would be chosen, for we have no theory that predicts the order in which alternatives will be examined. But in another case discussed above--the sale of a house--the environment presents the seller with alternatives in a definite sequence, and the selection of the first satisfactory alternative has precise meaning.

However, there are certain dynamic considerations, having a good psychological foundation, that we should introduce at this point. Let us consider, instead of a single static choice situation, a sequence of such situations. The aspiration level, which defines a satisfactory alternative may change from point to point in this sequence of trials. A vague principle would be that as the individual, in his exploration of alternatives, finds it easy to discover satisfactory alternatives, his aspiration level rises; as he finds it difficult to discover satisfactory alternatives, his aspiration level falls. Perhaps it would be possible to express the ease or difficulty of exploration in terms of the cost of obtaining better information about the mapping of A on S, or the combinatorial magnitude of the task of refining this mapping. There are a number of ways in which this process could be defined formally.

Such changes in aspiration level would tend to bring about a "near-uniqueness" of the satisfactory solutions and would also tend to guarantee the existence of satisfactory solutions. For the failure to discover a solution would depress the aspiration level and bring satisfactory solutions into existence.

II. Existence of Solutions: Further Possibilities.

We have already discussed one mechanism by which the existence of solutions, in the long run, is assured. There is another possibility--or at least another way of representing the processes already described. Up to this point no use has been made of the distinction between A, the set of behavior alternatives, and $\breve{A}$, the set of behavior alternatives that the organism considers. Suppose now that the latter is a proper subset of the former. Then, the failure to find a satisfactory alternative in $\breve{A}$ may lead to a search for additional alternatives in A that can be adjoined to $\breve{A}$. This procedure is simply an elaboration of the information-gathering process previously described. (We can regard the elements of A that are not in $\breve{A}$ as elements that are initially mapped on the whole set, S.)

In one organism dynamic adjustment over a sequence of choices may depend primarily upon adjustments of the aspiration level. In another organism, the adjustments may be primarily in the set $\breve{A}$: if satisfactory alternatives are discovered easily, $\breve{A}$ narrows; if it becomes difficult to find satisfactory alternatives, $\breve{A}$ broadens. The more persistent the organism, the greater the role played by the adjustment of $\breve{A}$, relative to the role played by the adjustment of the aspiration level. (It is possible, of course, and even probable, that there is an asymmetry between adjustments upward and downward.)

Further Comments on Dynamics

The models thus far discussed are dynamic only in a very special sense: the aspiration level at time *t* depends upon the previous history of the system (previous aspiration levels and previous levels of attainment). Another kind of dynamic linkage might be very important. The payoffs in a particular trial might depend not only on the alternative chosen in that trial but also on the alternatives chosen in previous trials.

The most direct representation of this situation is to include, as components of a vector pay-off function, the payoffs for the whole sequence of trials. But then optimization would require the selection, at the beginning of the sequence, of a strategy for the whole sequence. Such a procedure would again rapidly complicate the problem beyond the computational capacity of the organism. A possible middle ground is to define for each trial a pay-off function with two components. One would be the "immediate" payoff (consumption), the other, the "position" in which the organism is left for future trials (saving, liquidity).

Let us consider a chess game in which the players are paid off at the end of each ten moves in proportion to arbitrarily assigned values of their pieces left on the board (say, queen, 1; rook, 10; etc.). Then a player could adopt some kind of planning horizon and include in his estimated payoff the "goodness" of his position at the planning horizon. A comparable notion in economics is that of the depreciated value of an asset at the planning horizon. To compute such a value precisely would require the player actually to carry his strategy beyond the horizon. If there is time discounting of payoffs, this has the advantage of reducing the importance of errors in estimating these depreciated values. (Time discounting may sometimes be essential in order to assure convergence of the summed payoffs.)

It is easy to conjure up other dynamic complications, which may be of considerable practical importance. Two more may be mentioned--without attempting to incorporate them formally. The consequences that the organism experiences may change the pay-off function--it doesn't know how well it likes cheese until it has eaten cheese. Likewise, one method for refining the mapping of A on S may be to select a particular alternative and experience its consequences. In these cases, one of the elements of the payoff associated with a particular alternative is the information that is gathered about the mapping or about the pay-off function.

Conclusion

The aim of this paper has been to construct definitions of "rational choice" that are modeled more closely upon the actual decision processes in the behavior of organisms than definitions heretofore proposed. We have outlined a fairly complete model for the static case, and have described one extension of this model into dynamics. As has been indicated in the last section, a great deal remains to be done before we can handle realistically a more completely dynamic system.

In the introduction, it was suggested that definitions of this kind might have normative as well as descriptive value. In particular, they may suggest approaches to rational choice in areas that appear to be far beyond the capacities of existing or prospective computing equipment. It has often been pointed out that a comparison of the I.Q. of a computer with that of a human being is very difficult. If one were to factor the scores made by each on a comprehensive intelligence test, one would undoubtedly find that in those factors on which the one scored as a genius the other would appear a moron--and conversely. A survey of possible definitions of rationality might suggest directions for the design and use of computing equipment with reasonably good scores on some of the factors of intelligence in which present computers are moronic.

HAS:ge