

1957

Background of Decision Making

Herbert A. Simon

Follow this and additional works at: <https://digital-commons.usnwc.edu/nwc-review>

Recommended Citation

Simon, Herbert A. (1957) "Background of Decision Making," *Naval War College Review*: Vol. 10 : No. 9 , Article 2.
Available at: <https://digital-commons.usnwc.edu/nwc-review/vol10/iss9/2>

This Article is brought to you for free and open access by the Journals at U.S. Naval War College Digital Commons. It has been accepted for inclusion in Naval War College Review by an authorized editor of U.S. Naval War College Digital Commons. For more information, please contact repository.inquiries@usnwc.edu.

SPECIAL ATTENTION TO THE READER

The material contained herein is furnished to the individual addressee for his private information and education only. The frank remarks and personal opinions of many Naval War College guest lecturers are presented with the understanding that they will not be quoted; you are enjoined to respect their privacy. Under no circumstances will this material be republished or quoted publicly, as a whole or in part, without specific clearance in each instance with both the author and the Naval War College.

NAVAL WAR COLLEGE REVIEW was established in 1948 by the Chief of Naval Personnel in order that officers of the service might receive some of the educational benefits of the resident students at the Naval War College. Distribution is in accordance with BUPERS Instruction 1552.5 of 23 June 1954. It must be kept in the possession of the subscriber, or other officers eligible for subscription, and should be destroyed by burning when no longer required.

The thoughts and opinions expressed in this publication are those of the author, and are not necessarily those of the Navy Department or of the Naval War College.

NAVAL WAR COLLEGE REVIEW

**Issued Monthly
U. S. Naval War College
Newport, R. I.**

BACKGROUND OF DECISION MAKING

A lecture delivered
at the Naval War College
on 26 August 1967 by
Professor Herbert A. Simon

Rather than pose as an expert (which, obviously, I am not) on naval decision making, or any other kind of decision making for that matter, I thought that I would explore with you this morning how one might go about training the decision processes — our own processes, or the processes of the other men in the organizations for whom we are responsible. Then, I will follow that up by some observations on current research on the decision making process which I think carries some very important implications over the next few years for decision making in military as well as civilian organizations. As a matter of fact, an important segment of the fundamental research now going on in this area — since the war, particularly — has been carried on under military sponsorship. At Carnegie Tech, for example, we have had for some years an Office of Naval Research contract for the study of logistic decisions and now are carrying on work with The RAND Corporation under Air Force sponsorship.

The Decision-Making Process

You recall the hero of one of Moliere's plays, who was surprised to be told sometime in middle age that he had been talking prose all his life. I think that talking about decision making is a little bit like telling somebody how to talk prose. We have all been making decisions all of our lives; hence, I do not think it is likely that someone can come along and tell us a lot of things about decision making that we do not already know and that are not familiar facts about our everyday decision-making processes. Sometimes, however, it is useful to review even familiar facts. Sometimes it is useful not only to talk prose but to be aware that one is talking prose, and to look a little bit at one's own processes when he makes decisions.

One can find analyses of the decision-making process that will seem reasonable and common sense by going back in both the literature of psychology and the military literature. In psychology, people who are interested in decision making often go back to a famous passage in one of John Dewey's books, in which he laid out the elements involved in orderly problem solving or decision making. Likewise, in the military there is a process that has become formalized over the years as "the estimate of the situation." This is, basically, an outline of the processes and stages that one goes through in solving a problem in an orderly manner. I understand that considerable attention will be devoted here in your courses of training to the use of the estimate of the situation as a tool in military decision making and military planning.

There is nothing very abstruse or peculiar about either Dewey's problem-solving steps or the theory underlying the estimate of the situation. In order to make a reasonable decision, there has to be an objective or a goal — something we are trying to accomplish. So, somewhere along the process — and usually early in the process — it is a good idea to get clear as to what that objective or goal is. I am sure that many of you know, and maybe practice, the familiar ploy by which you can set almost any committee meeting into an uproar; that is, after there has been an hour and a half or two hours of fruitless discussion about this or that, to turn to the other people at the table and say, "Now, gentlemen, what is the problem?" In naval conferences perhaps that question gets asked earlier than the end of the first hour. Believe me, I have sat in plenty of committee meetings where the question upset the whole proceedings!

Let's suppose that we are solving problems in an orderly fashion. First, we have stated the objectives. Second, according to most of these schemes of problem solving, we are supposed to find the alternative courses of action: what are the things that we might do? Third, we are to evaluate each course of action: what will happen? what are the consequences? will I achieve my

objectives? what side effects will it have for good or for bad? Fourth, having made evaluations and perhaps tentatively agreed on a course of action, we have to do some final checking against particular things that we want to be sure to take care of. Finally, but not always the easiest part of the decision-making process, we act on our decision.

One can find these elements listed in a variety of ways and in different terms, but I think that the list tends to boil down to the basic, very familiar, and very common-sense elements I have mentioned: (1) clarifying, stating, or defining the objective; (2) finding some alternative courses of action that give promise of attaining that objective; (3) evaluating the actions in terms of the objective and in terms of side effects; (4) checking the action out; (5) acting on the decision.

Search Process and Evaluation Processes

Another way of looking at the decision-making process would really boil it down to two sorts of things — two more elementary processes that are involved in all of the steps I have mentioned. First, of all, there are *search processes*. One looks around for plans of action. What are the things that I might do in this situation? One tries to search these out and to elaborate on them. Then there are the search processes involved in discovering the consequences of an action — if I do this, what will happen? Second, there is the *checking or evaluating process*: taking a course of action that has been developed, setting it up against some standards, and then seeing whether it meets the standards. If we really looked in detail at what somebody does when he is going through the stages of the decision-making process, I think we would find that we could further analyze these stages into a whole succession or complex chain of search processes and of checking processes intermixed.

In a sense, what I have said so far perhaps sounds not only like simple and common sense, but perhaps it sounds even simple-

minded. Anyone who has seen an experienced and capable executive go through the processes of making a decision has a feeling that when one mentions these common-sense elements he has left out the important part of the problem.

There are some names for the important part of the problem that is allegedly left out. We say, "Well, yes, you go through these steps; you search and you check. But there is also a big 'X' here: there is a judgment factor. A person has to exercise judgment; he has to have intuition. He cannot just go through these steps in an orderly and logical fashion. In addition, he has to have some intuition if he is to make good decisions." In decisions that involve coming up with a novel course of action or something new that has not been tried before we say, "Well, sure, but there also has to be that spark of creativity."

We have these labels: judgment, intuition, creativity (maybe you can think of some other synonyms for them) to signify that we really do not think that one can solve difficult problems in an intelligent way just by going through the stages that I mentioned earlier.

One of the important and interesting questions about the decision-making process is just this: Is there something more in the form of judgment, intuition, and creativity; or are the stages of the decision-making process I named earlier and the simple processes or search and checking all there is to it? Is that all the human mind is doing in solving problems?

I would like to leave that question open for just a few minutes and talk about the ways in which we might develop, train, and improve at least those parts of the decision-making process we can put our fingers on. Even if there are intangibles that we perhaps do not understand too well, there are some tangibles that we do not understand and that are discernible in human decision making and in good decision making — the stages and processes that I have mentioned. How can an individual or an organization go about improving those processes?

Decisions of Encounter and Set Pieces

That question really ought to be answered in two parts, because the kind of training one does depends on the sort of decisions one is talking about. For this purpose, I find it useful to talk about decisions of *encounter* and decisions that are *set pieces*. The motions behind this distinction — if not the terms — are familiar to anyone in military or civilian life, too. On the one hand, there are decisions that arise because an occasion arises — either is made to occur by the enemy, or by the weather, or by something else — in which action is called for, and a decision has to be made as to what that action is going to be. The set-piece decision occurs in situations in which there is a lot of time to plan out a future course of action. A characteristic example are the decisions that are made in designing a ship. These are quite different from the decisions of encounter, like the tactical decisions that have to be made in a battle.

In distinguishing these two classes of decision I do not mean to imply that very different psychological processes are involved, or that the way the wheels turn in the brain is drastically different in the two cases. I do not think that is so. I think the differences are in the occasions under which the decisions are made: the kinds of information available, and the time element or speed with which the decisions have to be made. It is because of these differences in the way in which the decision occurs, rather than the differences in the basic mental processes, that we are talking about them separately.

The main characteristics of the decision of encounter are that, first, the occasion for making the decision is usually determined by somebody else or by the environment, by nature or the enemy; second, and as a consequence, the decision has to be made by a certain time — usually as soon as possible or day before yesterday — at any rate, it has to be made under extreme time pressure. There is no point in talking about decision-making processes that do not get the decision made in that time.

The set-piece decision is more often made on occasions determined by the decision maker himself, or by others in his organization. In a large military or business organization, it is often one's commander or one's boss who decides that there is an occasion for a decision — and then the individual makes the decision. In other situations, the individual himself decides that there is something to make a decision about — and he makes it. The time pressure is usually much less than it is in the decision of encounter. There are usually much more elaborate procedures for checking a decision, or for erasing, so to speak, and making corrections in the decision, because you do not take action immediately upon formulating a plan. So, again, in the case of the design of an elaborate structure, or the design of a ship after a set of designs has been prepared, there remains time to review the designs from a variety of standpoints. They can be checked and cross-checked, and can be modified.

Although from this description it would seem that decisions of encounter are the more difficult to make, and that a man would have to be a natural-born decision maker and have good judgment to handle such decisions, in other respects it is more important for people to train themselves — and to train others in their organizations — for decisions of encounter than for the set-piece decisions. This is because of the very fact I have just mentioned: the opportunity for error correction, the interval between decision making and action is so much less in decisions of encounter. Therefore, one had better have a good initial decision-making process if he wants to avoid disasters.

Training for Decisions of Encounter

I think that if a man will look back over his own military training at various stages of his career he will see some obvious things that he can do to become more capable in handling a decision of encounter.

One can train a man so that he has at his disposal a list or repertoire of the possible actions that could be taken under

the circumstances. He knows some of the things that one can do and that appropriately can be done. While that seems very obvious, people are not born with that kind of knowledge. If you compare the behavior of a person who is driving a car for the first month with the behavior of a driver after four or five years of experience, one of the important differences is that when a situation arises in which action needs to be taken the person who is new at the game does not have immediately at his disposal a set of possible actions to consider, but has to construct them on the spot. Constructing possible actions is a time-consuming and difficult mental task. Similarly, the decision maker of experience has at his disposal a check list of things to watch out for before finally accepting a decision.

Both the list of actions and the check list of side effects and other consequences to watch out for tend to be specific to a particular class of decision-making situations. Probably there is not a lot that can be done, on this score at least, to train a man in general for decision making — so that any emergency that arises will be met by instantaneous and appropriate decision. But, people can certainly be trained to respond in a very rapid fashion with possible courses of action and with checks on courses of action in any particular area of human endeavor — as, for example, handling a ship in a tactical situation.

A large part of the difference between the experienced decision maker and the novice in these situations is not any particular intangible like “judgment” or “intuition.” If one could open the lid, so to speak, and see what was in the head of the experienced decision maker, one would find that he had lists of the kinds that I have described; that he had at his disposal repertoires of possible action; that he had at his disposal check lists of things to think about before he acted; and that he had mechanisms in his mind to evoke these, and bring these to his conscious attention when the situations for decision arose. Most of what we do to get people ready to act in situations of encounter consists in drilling these

lists into them sufficiently deeply so that in fact they will be evoked quickly at the time of the decision.

Training for Set-Piece Decisions

When one turns to the more deliberate or set piece kind of decision, training is in some respects less important because of the opportunity for self-correction. I think the kind of training for decision that we want here is less specific to the situation than training for decisions of encounter. One does not have to know all about the technology of the particular situation with which he is dealing. If one is handling a ship and something has to be done about it, he has to know, then and there, the technology of handling a ship. If one is designing something, it is good to know all about the technology of that design, but, if the time pressures are not too severe, one of his techniques can be to sit down and acquire the technical information which he does not have at the moment.

Orderly Problem Solving. So the training one wants to give people to improve their capacity for making set-piece decisions is more general training in orderly problem solving. An initial training goal is developing a habit of looking at a problem in an orderly sequence of steps. One important role of the estimate of the situation is not simply to have a check list in front of a man when he is making a decision and saying, "Do this; then, this; and then, that," but to have it as a training device so that the habit of approaching a decision-making situation can be built into a man, and he responds to the situation automatically in that way — even when he is not using the formal estimate of the situation as his decision-making device.

Alertness to Innovation. On slightly less familiar ground — at least, less familiar in terms of existing training programs for decision making — a great deal can be done (and very little is done) in organizations to develop the skills of actively looking

for occasions for decision and actively looking for novel alternatives of action. If I had to characterize the strengths and weaknesses of typical organizations in decision making, I would say that an organization that engages in a particular kind of activity for any length of time soon develops the check lists and the repertoires of action that enable it to respond promptly — and, usually, more or less efficiently — to situations as they arise. But it is often deficient in its development of prods and of triggers that set it looking for a problem when the problem is not presented by the environment. By and large, organizations do not seem to do as good a job of thinking up problems that have not been presented to them as they do in dealing with the problems that are presented to them with some urgency. Let me give you a concrete example from some research we have been doing on decision-making processes in business.

We have been looking at a number of companies that have recently decided to acquire electronic equipment (digital computers) to help with their accounting work. We have been interested in seeing how they went about selecting a particular computer from among the various ones that are available to them on the market. We have been equally interested in discovering why they got involved in this particular decision at the time that they did. It turns out in this case (it certainly is not true in the case of every fad and fancy) that almost all companies in the past five years, when they have gotten around to considering the computer decision, have ended up by deciding that they wanted a computer.

I do not want to argue here whether they should have or should not have decided that; whether it was a wise decision or not; whether it was because it would save them money, or because the vice president felt embarrassed with his golfing companions if he did not have one of these toys in his office. But, whatever the reason, the decisions have been nine to one — or, perhaps, nineteen to one — in favor of acquiring computers. From

a little contact with some naval installations, I gather that somewhat the same process is going on in the Navy at the present time: that computers are being acquired in large numbers and for a variety of purposes, Yet, there were great differences in the promptness with which companies looked at their own accounting and clerical problems and raised for themselves seriously the question of whether they ought to have such a device.

The real variation was in selecting the occasion for decision, and not in the content of the decision that was thereafter reached. One gets the impression in business organizations (and maybe this would be true of military installations as well) that the big difference between the organization that is just sort of getting along and the organization that is really on its toes does not lie in the care or skill with which they evaluate alternatives when somebody presents them with the alternatives and says, "Look! You decide!" — but that the big differences lies primarily in the promptness with which they detect new challenges in the environment and the promptness with which they detect new problems and new opportunities for innovation and for change — and respond to that challenge.

If one were to go into an organization and improve it on this score, one of the main things he would try to do would be to develop habits of systematically looking for occasions for decision, of systematically scanning the whole horizon and saying, "What new things are there in the world today — new technologies, new states of affairs, as far as the world situation is concerned — that this organization ought to be responding to, or ought at least be attending to — on which it ought at least have a study group?"

I know that in large organizations — military and otherwise — there are some built-in procedures of this kind. I am simply reflecting an over-all observation about organizations: that even though procedures of this sort are developed to a certain extent, they are still by far the weakest part of the decision process.

Organizing for Effective Decision Making. In addition to training to develop an orderly decision-making process, and to develop a process that will actively *search* for occasions to make decisions instead of them being presented, a third approach is to look at the organization structure itself as a major determinant of the way in which the individuals who have been put in that organization make decisions.

Let me first cite an industrial example, and then let me pose for you a probable example, from the military standpoint. If the shoe fits, you can put it on; if the shoe does not fit, you can throw it at me and I will try to dodge it up here.

The industrial example is this. I can tell an audience of business men a little story about a conference that involves the sales manager, the production manager, the product design man, the production-scheduling man, and two or three other characters that I can bring into the scene. This is a session in which these gentlemen are trying to straighten out some of their scheduling problems. I can construct a little dialogue in which the salesman says: "Well, if this factory had a little more flexibility in meeting customers' specifications and delivery dates, we would be able to build up our sales."

The production manager then says: "If you salesmen were not always promising day-before-yesterday deliveries to our customers, we could maintain an orderly schedule in the factory."

I won't bore you with the rest of the recital, but those are just characteristic complaints that one can hear in almost any business concern.

The usual reaction of a business group to this story is: "Gee, have you been planting microphones around my company?"

A sales manager reacts like a characteristic sales manager because different telephone calls and different letters come to his

desk than come to the desk of the production manager. He is praised and blamed for different things than is the production manager. The whole world of company decision making looks to him a somewhat different world than the world of the production manager. If you asked them about it, both of them would say: "Well, what we are after is company profits." But, company profits to a sales manager primarily means satisfied customers — customers who get what they want when they want it. Company profits to a production manager largely means low costs and large-volume production in the factory. It's like the story of the blind men who were looking at the elephant: one grabbed it by the trunk; the other grabbed it by the tail; they saw two entirely different elephants. Similarly, there are two entirely different views about how one can make profits in the company.

Executives do this not because they are stubborn people, not even because they are power-hungry — although the desire to get ahead in the world motivates people in business as it does in other organizations. They do this in considerable part simply because they are responding to the part of the environment that they see. The sales manager is shielded from the problems of the production manager because those problems are handled in the factory. The production manager never has the angry customer on the other end of the 'phone.

When we look at some of the problems that arise in the Department of Defense in connection with Armed Forces' Unification, irrespective of whether we look at these from the standpoint of one of the Services or whether we stand outside as a layman and grumble about all of the politics going on, we should realize that what is going on here is not merely (although surely there are elements of this involved) jockeying for a position in a large and complex organization, each unit trying to get its part of the organization into a position of power. Really, there are a number of people looking at an exceedingly complex situation — a situation that is far too large for the computational powers

that an individual human mind has available. Therefore, they are seeing enlarged the particular part of the situation for which they have particular responsibility, the part that is brought to their attention in their day-to-day operation — and they are perhaps seeing as rather small those parts of the situation that are other people's responsibilities and in other people's environments.

In these complex situations that organizations have to deal with — and certainly it would be difficult to find one more complex than the one which I am using for an illustration here — human beings tend to fix on subgoals. They tend to deal with a part of the problem rather than the whole problem — partly because they are operating against the very limits of their own thinking and computational abilities. A large part of the difficulty we get into when we try to put the parts of such a complex organization together stems not from failings of human motivation, not from stubbornness, not from desires for power or prestige or advancement — but from the fact that the people who are in the situation are simply seeing different worlds.

I think that some recognition in organizations, and some understanding by the decision makers, of this fact — and that it is a fact we are not going to do away with, for we are not going to make these problems very much simpler — might make them a little more tolerant and understanding when they try to deal rationally — or, what seems to them rationally — with persons who see the situation from quite a different standpoint and who are exposed to quite different parts of the problem. Related to this, we would see the problem of putting together a complex organization — putting together the Department of Defense, putting together the Navy, or even putting together one part of the Navy — as a problem of designing the environment of the individuals who have to make the decisions in that organization. We would try to design the environment so that the various specialized considerations that have to contribute to decisions are not lost sight of (we do have to have specialists), but, at the same time,

the differences are not exaggerated by the partial views that individuals have of the problem.

Military organizations long ago developed at least one important device for accomplishing this, or for helping to accomplish it — and that is rotation of duty. I know there are always mixed feelings about this in any organization that adopts the procedure — and certainly in an organization that adopts it as rigorously as does the Navy — feelings that one never really has a chance to acquire specialized knowledge about any one subject. But, against this, there are very great advantages.

When you sit across the table from an officer who has different responsibilities from yours and try to reach a decision with him, you can at least appreciate that if he views the problem a little differently than you view it, and if he comes up with somewhat different answers, it is not because he is full of human cussedness. You can appreciate, from other assignments you have had in the same organization, his point of view and the problems he is facing. Industrial organizations are increasingly adopting the technique of horizontal transfer of people — that is, rotation — as a device for giving executives this broader interdepartmental point of view.

These comments on how to make people in an organization more effective decision makers in set-piece decisions are, to be sure, generalizations. But, perhaps they also provide a check list against which to measure an organization or organizational unit. One can look at an organization and ask himself these questions:

First: "What are we doing to develop the decision-making skills of the members of this organization? what are we doing to get people to look at problems in an orderly fashion? what are we doing to internalize as a part of their own habit structure, or their own personality, the kinds of steps that are involved in the estimate of the situation?"

Second: "What are we doing to building into them habits of looking for occasions of decision, or scanning the horizon?"

Third: "What are we doing to arrange this organization so that people will be exposed to the kinds of experiences and contacts with the relevant parts of the world that will, in itself, bring to them the considerations that are important and relevant when they come to a decision-making situation?"

Toward a Science of Decision Making

Let me pass from the problem of training in decision making to one other topic; then, I will close. We often make the distinction between science and art in terms of the difference between those things that can be stated and taught through quite explicit principles and those things that are handled by people without knowing exactly how they handle them, without being able to state very explicitly what they are doing when they are handling them, and without being able to be explicit in training other people to handle them. I do not suppose we have any doubt that decision making has been largely an art rather than a science. I do not think that anything I have said this morning about the process would take it out of the category of art and put it in the category of science. Nevertheless, I have a deep conviction (and I will try to give you some of the basis for that conviction) that this state of affairs is about to change — is, in fact, changing already.

Both the military services and our civilian economy have gone through a tremendous technological revolution in the last fifty years, which I do not need to detail to you. One of its consequences has been that the practical arts — the arts of making things and doing things in the everyday world — now rest for the first time on a foundation of fundamental research in the physical sciences. It becomes less and less possible to be effective in the practical arts without a bridge back to underlying knowledge of physics, chemistry and the other physical sciences. In one respect the most important aspect of this revolution is that it has

created the college-trained engineer as an essential intermediary between the basic physical sciences and day-to-day practice.

I believe that the change now coming about in our knowledge of the decision-making process is going to produce a corresponding change in the relations of the practical, everyday affairs of running, constructing, and designing organizations to the underlying human — behavioral or social — sciences. We will shortly be in a situation in which the men responsible for designing and operating organizations will have to build a bridge back to the behavioral sciences. It is entirely possible to make good decisions according to present-day standards without knowing anything in a formal way about psychology. This is a matter of considerable regret and chagrin to psychologists, and to others of us who are in the behavioral sciences and who would like to think that the world of practical affairs had some dependence upon us. What I am asserting — and I hope to produce a little evidence in a moment — is that this state of affairs is going to change rather drastically.

Let me give you a few of the signs and portents of this change. All of these developments belong to the post-World War II era, and most to the last ten years or less.

Theory of Games and Operations Research

We have had the invention of something called *Theory of Games and Economics Behavior* by the late John von Neumann, one of the most brilliant mathematicians of our generation. There was a period just after the war when some people — mostly scientists on the fringe of the military — thought that the *Theory of Games* was really the clue to fighting wars in a scientific fashion; that one formulated the military problem up as a game problem. Then one found the best strategy by mathematical means, according to the techniques of von Neumann and Morgenstern, one issued some directives based on this strategy, and then the battle was won. If there were any such illusions, I think that these have departed.

But the mathematical *Theory of Games* has, in a very important way, illuminated the concept of strategy and the concept of surprise. Whether any of you have been exposed to the theory in a formal way or not, a great deal of what goes on now in military instruction and in military war gaming in the armed services has been influenced and affected by some of the basic concepts that were developed in the *Theory of Games*.

Second, there is something abroad called "Operations Research," which I won't try to define because you can find about as many definitions as you can find operations researchers. One view is that it is common sense and mathematics applied to problems to which they had not been much applied before. By this definition, the Theory of Games is a part of operations research. Since the war, techniques have been developed that allow one — by the use of fairly high-powered mathematics and computing machines — to make good decisions in complex situations; in fact, to make decisions that are at least as good in some cases as those made previously by men of considerable experience using rule-of-thumb techniques.

The greatest successes of operations research techniques to date have been in such areas as inventory and production control; the scheduling of oil refineries; the scheduling of shipments from warehouses, and similar problems. Some of these techniques are now being used in military installations. In fact, this development actually first started (as many of you know) in the military. Operations research techniques have been applied to a number of naval and military problems, like the submarine search problem. You are better able than I to evaluate with what effectiveness they have been applied.

But I think the significance is this: in some areas of decision making one can now substitute, for what the man of judgment used to do, formal techniques making use of mathematics and computers. The area in which one can do this is still relatively

limited. The kinds of industrial problems I have mentioned are those that occupy the attention of a factory manager or a factory scheduler rather than a vice president. We have not displaced very many vice presidents yet — but, the development and the trend is nevertheless significant.

Routine, Well-Structured and Heuristic Decisions

What about the vice presidents? What about the whole host of decision-making situations where the decision is not a very precise thing — where it is not a question of how many “widgets” are to be kept in stock or how many are to be shipped out this month, but where the decision maker is faced with a rather ill-defined problem: where there are not any known automatic ways of cranking out a solution, and where the traditional elements of judgment and intuition have their field day? Is there really any prospect that we are going to understand these processes in the near future or that we are going to be able to do better than the human mind in performing them? I think there is such a prospect. Let me clarify what I am saying by making a couple of distinctions.

First, there are decisions that are of quite a routine sort: cut-and-dried decisions. One can write them out in an SOP in fair detail. Then, having written out an SOP, one can get a clerk to do them, or, sometimes, one can even get a machine to do them. Many of the tactical decisions of air defense are of this kind: that is, the decisions involved in operating an air defense installation or a surveillance operation. As long as one’s patience lasts, he can set down in considerable detail what a person ought to do, and he can reduce a large part of the process very nearly to a clerical routine.

Second, there are decisions that one might call “well-structured.” One knows what the objective is; one can even measure the payoff in dollars or some other unit. One knows what the alternatives are: as in the inventory or production control decisions,

or many of your logistic decisions, at least in peacetime. Operations research techniques are gradually giving a broader and deeper understanding of these well-structured decisions, and gradually showing how to supplant human decision-making processes with machine processes.

Now I will talk about a third class of decisions. Let me give them a name. It is not a very familiar word, but it is the closest thing that I can find in the dictionary — let's call these "heuristic" decisions. These are the unstructured decisions, or those that are not very well structured. They are the ones that involve intuition and judgment, if anything does.

Research on Heuristic Decisions

Within the past several years, there have been several successful attempts to program an ordinary digital computer — a large one, but no larger than those now on the market — to make decisions of this kind, and to make them essentially by imitating the processes that humans use in making them. I can outline just briefly the furthest developed example.

If you think back to your high school geometry, you will perhaps recall that getting a *proof* of a theorem was kind of a high art. There was not any systematic way of cranking out answers to problems, or of rearranging things in an orderly way as there was in algebra. In high school geometry, you just had to have a good idea to find a proof for a theorem. If you did not get the proof, and a fellow classmate or the teacher did, it always had an element of surprise and trickery in it. "How did he ever think to do that?"

In research that is going on cooperatively between Carnegie Tech and The RAND Corporation (which is an Air Force contractor), we picked a problem like the geometry problem (it happens to be a problem in symbolic logic, but it has very much the flavor of Euclidean geometry), and we asked ourselves whether we could induce a computing machine to discover proofs for theorems in

much the same way as in high school geometry. We now know how to do this — the machine can prove most of the theorems in Chapter II of *Principia Mathematica* — which is not going very far, but perhaps as far as one could go with a class of college sophomores in a semester.

In the course of doing this, we think we have learned a good deal about how human beings solve problems heuristically. We did it not by using the brute force of the machine — not just by making the wheels turn faster. We did it by imitating as closely as we could the processes we thought we saw humans going through.

This research answers one of the questions I posed in my opening remarks: Are there some mysterious things hidden in those terms of “intuition” and “judgment?” Are there some processes going on that are over and above the common ones that have been observed in the problem-solving process? I think the answer is quite generally “No.” Judgment and intuition — at least, insofar as our imitation of this particular area of human activity is concerned — turn out to be just more check lists of things to look at, things to think about, and things to try first. Successful problem solving in this unstructured, heuristic area is largely a matter of judicious balance between willingness to search a range of possibilities (even though one is not sure that any of them is going to work) and having good facilities for evaluation of those possibilities so that one does not waste all his time in trying them out.

The monkeys in the British museum did not succeed (you remember they were trying to type out the works of Shakespeare by sitting in front of typewriters and pecking at random) simply because the number of possible sequences that one can type on the typewriter is so large that they never got around to the particular sequences which correspond to the works of Shakespeare.

The reason why a computing machine, programmed as machines usually are for operations research or for scientific computing problems, would not be able to type out the works of Shakespeare is that it would not do *enough* random searching; it would not do *enough* exploring. It would be tied down to the basic computing program that one gave it.

What humans do down in this area is to search selectively — not to search completely at random, but to be very sensitive to the kinds of cues that arise as they go along. If you have children at the Easter egg-hunting age, you will have observed that it is very easy to hide an Easter egg so that a child will take an hour or two hours to find it. He will even get angry with you before he finds it. But if you tell him when he is getting warmer and when he is getting cooler, he can find the same egg in the matter of a few minutes.

All that we are able to discover in the processes of heuristic or judgmental problem solving is that the problem solver, while he does not have any way of turning a crank and systematically reaching the conclusion, does know enough about the situation so that while he is exploring he can at the same time be evaluating the “warmer” and the “cooler” of his explorations. If we can program a machine to do this (as we now can — at least in some areas), we can get the machine to do judgmental problem solving.

I do not want to make a prediction as to how many years it will be before machines will be doing better than vice presidents do now, nor the corresponding statement about military ranks. Nor do I want to predict when machines will be more *economical* problem solvers than people.

You are perhaps familiar with the Air Force story (or, maybe it is a Naval Air Force story) about a pilot who was told about the marvels of the new pilotless aircraft, and all the things

they would do: the feed-back controls, and the servomechanisms that control these aircraft.

The officer's reply was: "Yes, but where can you get a non-linear servomechanism for \$600 a month that reproduces itself?"

So, I am going to be a little wary about my prediction. What I will be bold to predict is that we are now on the verge (or over the verge) of understanding the human problem-solving process in some of its vaguer and more mysterious manifestations. The line that has been drawn up to the present, limiting the things one can attack with systematic techniques and high-powered mathematical tools, is becoming obscured. The opportunities for these tools are moving directly into the area that has been traditionally called "judgmental decision making."

I do not want to make any very definite predictions, nor do I want to argue whether, when a machine is doing these things, it is "thinking" or not. You can decide how you want to define the word "thinking." I do discover (as also my colleagues and other people who have worked in this area) that when a machine has been programmed like this, after a few days it becomes very hard to refer to the machine as "it" instead of "he." But that is another question.

It would be my guess that these developments are going to have a major impact on the processes of running organizations — both military and civilian organizations — not in the distant future, but in our lifetimes. I do not feel a real hesitation in saying that it is going to have an impact within the next ten years.

What I said earlier about set-piece decisions would apply at the present time to the need to be alert to these impending changes. It is exceedingly important at the present time that organizations in our society have their scanning mechanisms

turned to this particular part of the horizon; that they begin to attend to and evaluate the possibilities of a revolution in our decision-making technology and in the whole man-machine relationship in organizations that will be a consequence of this revolution.

This recommendation will seem less strange perhaps to the organizations which you here represent than to almost any others because you have already had important changes in the man-machine relations — particularly in observational and surveillance techniques — in military operations. What I foresee is a similar shift on a much larger scale, affecting a much larger part of the total range of activities that are now carried on by human beings — by these non-linear self-reproducing servomechanisms. Research is going on in this area in the military at the present time, and there is some indication that research is also going on in Russia. On the whole, our development appears to be somewhat further along than the Russian, but there are probably people in this audience who are better able to evaluate that particular aspect than I am.

In Conclusion

I have talked already a few more minutes than I had intended to. Let me stop on this note and suggest that in the discussion period I would be glad to carry the discussion further on the problems of training people to become more effective decision makers in either of the two types of decisions I have mentioned; or, if you like, I would be glad to be more explicit about some of the things that I have said very vaguely and very briefly concerning the impending mechanization of decision making.

Thank you very much!

BIOGRAPHIC SKETCH

Professor Herbert A. Simon

Professor Simon was graduated from the University of Chicago with an A.B. degree in 1936, and received his Ph.D. from there in 1943.

From 1936 to 1939, he was associated with the University of Chicago as a research assistant and staff member of the International City Managers Association. After three years as Director of Administrative Measurement Studies, Bureau of Public Administration at the University of California, he returned to the Illinois Institute of Technology to become a Professor of Political Science. He became Chairman of the Department of Political Science in 1946.

Professor Simon has held positions with the United States Bureau of the Budget, the United States Census Bureau, the Cowles Commission for Research in Economics, and has served as Consultant and Acting Director of the Management Engineering Branch of the Economic Cooperation Administration. Since 1949, he has been a Professor of Administration and Head of the Department of Industrial Management at Carnegie Institute of Technology in Pittsburgh.

Professor Simon has been the author and co-author of several books relating to his field, including *Administrative Behavior* and *Public Administration*.