Administrative Decision Making

By HERBERT A. SIMON

Carnegie Institute of Technology

Amidst the general depreciation of progress in the social sciences, the author surveys an area where "we can speak without blushing about our present knowledge"—the examination of administration through decision-making concepts and language.

THERE is no need, at this late date, to justify the study of organization and administration in terms of the decision-making process, for decision-making concepts and language have become highly popular in writing about administration.[1] This paper will describe some of the progress that has been made over the past quarter century, employing this approach, toward deepening our scientific knowledge—what new facts have been learned about human behavior in organizations, what new scientific procedures for ascertaining facts, what new concepts for describing them, and what new generalizations for explaining them. This progress extends both to descriptive and normative matters: to the pure science of administration, and its application to the practical business of managing.

To satisfy limits on this journal's space, your patience and my time, the account will be highly selective. Only a few notable and significant advances have been selected; others for which equally plausible claims might be made are ignored. A frequent practice in the social sciences is to bemoan our present ignorance while making optimistic predictions about future knowledge. It is a pleasure to survey an area of social science where, by contrast, we can speak without blushing about our present knowledge—indeed, where only a small sample of the gains in knowledge that have been achieved in the past quarter century can be presented.

Operations Research and Management Science

One obvious answer to the question "What's new?" is the spectacular development in the normative theory of decision making that goes under the labels of "operations research" and "management science." Through these activities, many classes of administrative decisions have been formalized, mathematics has been applied to determine the characteristics of the "best" or "good" decisions, and myriads of arithmetic calculations are carried out routinely in many business and governmental organizations to reach the actual decisions from day to day. A number of sophisticated mathematical tools—linear programming, queuing theory, dynamic programming, combinatorial mathematics, and others—have been invented or developed to this end.

Like all scientific developments, this one has a long intellectual history, and did not spring, full-grown, from the brow of Zeus. Nevertheless, the state of the art today is so remarkably advanced beyond its position before World War II that the difference of degree becomes one of kind.[2]

The quantitative decision-making tools of operations research have perhaps had more extensive application in business than in governmental organizations. It is worth

[1] The term "decision-making" occurred three times in the titles of articles in the first fifteen volumes of the *Public Administration Review*—that is, through 1955; it occurred ten times in the next eight volumes, or about six times as often per annum as in the earlier period.

[2] Some notion of the state of proto-operations-research just before World War II, as it applied to municipal administration, can be obtained from Ridley and Simon, *Measuring Municipal Activities*, (Chicago: International City Managers' Association, first edition, 1938).

Reprinted from Public Administration Review, Vol. XXV, No. 1, March, 1965
quarterly journal of the American Society for Public Administration
1329 Eighteenth Street, N.W., Washington, D.C. 20036

recalling, however, that many of these tools underwent their early development in the American and British military services during and just after the Second World War (where the terms "operations research" and "operations analysis" were coined). Among the inventors of linear programming, for example, were Tjalling Koopmans, seeking, as statistician with the Combined Shipping Adjustment Board, a means for scheduling tanker operations efficiently; and George B. Dantzig and Marshall K. Wood, in the Office of the Air Force Controller, who used as one of their first (hypothetical) programming problems the scheduling of the Berlin Airlift.

Operations research, particularly in its governmental applications, has retained close intellectual ties with classical economic theory, and has sought to find effective ways of applying that theory to public budgeting and expenditure decisions. This has been a central preoccupation of the RAND Corporation effort, as exemplified by such works as Charles J. Hitch and Roland N. McKean, *The Economics of Defense in the Nuclear Age*.[3] In the past several years, Hitch, as Controller of the Department of Defense, and a number of his former RAND associates have played major roles in bringing the new tools to bear on Defense Department budget decisions. Thus, while the quarter century begins with V. O. Key's plaint about "The Lack of a Budgetary Theory,"[4] it ends with a distinct revitalization of the whole field of public expenditure theory, and with a burgeoning of new analytic tools to assist in allocating public resources.

Optimality and All That

In many ways the contributions of operations research and management science to decision-making theory have been very pragmatic in flavor. The goal, after all, is to devise tools that will help management make better decisions. One example of a pragmatic technique that has proved itself very useful, and has been rapidly and widely adopted over the past five years, is the scheduling procedure variously called PERT, or critical path scheduling. This technique does not use any very deep or sophisticated mathematics (which may account partly for the speed of its adoption), but is mainly an improvement of the common sense underlying the traditional Gantt Chart.

Contrasting with this pragmatic flavor, advances in operations research have been paralleled by developments in the pure theory of rational choice—a theory that has reached a very high level of mathematical and logical elegance and rigor. Among these developments perhaps the most important are: (1) rigorous, formal axiom systems for defining the concept of utility in operational terms, (2) extension of the theory of rational choice to encompass the maximization of expected utility under conditions of uncertainty, (3) extension of the theory to repeated choices over time—dynamic optimization, and (4) extension of the theory to competitive "gaming" situations. These formal advances have had an important influence, in turn, on directions of work in theoretical statistics (statistical decision theory, Bayesian statistics), and on the kinds of models that are preferred by operations researchers—or at least by the theorists among their number.[5]

An evaluation of these contributions on the pure theory of rational choice would return a mixed verdict. On the positive side, they have provided enormous conceptual clarification for discussions of "rationality." For

---

[3] Cambridge: Havard University Press, 1960.

[4] *American Political Science Review*, December 1940, p. 1142. Labels have an unfortunate tendency to compartmentalize knowledge. Thus, the literature of "budgeting" has been only partly informed by the literature on "decision making," and vice versa, and both of these have sometimes been isolated from the economics literature on resources allocation and public expenditure theory. Variants on the same basic sets of ideas are rediscovered each generation: "measurement of public services," "program budgeting," "performance budgeting," "engineering economy," "cost-benefit analysis," "operations analysis." What is genuinely new in this area in the past decade is the power and sophistication of the analytic and computational tools. Some impression of these tools may be gained from the Hitch and McKean book previously mentioned; from Roland N. McKean, *Efficiency in Government Through Systems Analysis* (Wiley, 1958); Arthur Maass, *et. al., Design of Water Resource Systems* (Harvard U. Press, 1962); or Allen V. Kneese, *The Economics of Regional Water Quality Management* (Johns Hopkins U. Press, 1964), and the references cited therein.

[5] Since I have discoursed at length on these matters elsewhere, I shall be brief here. See "Theories of Decision Making in Economics and Behavioral Science," 49 *American Economic Review* 253-283. (June 1959), and Part IV of *Models of Man* (Wiley, 1957).

example, it has always been unclear what rationality meant in a pure outwitting or bargaining situation, where each party is trying to outguess, and perhaps bluff, the other. If the theory of games, due to von Neumann and Morgenstern, did not solve this problem for all situations, it at least made painfully clear exactly what the problem is.

On the negative side, fascination with the pure theory of rational choice has sometimes distracted attention from the problems of decision makers who possess modest calculating powers in the face of a world of enormous complexity. (In the real world, the calculating powers of electronic computers as well as men must be described as "modest.") A normative theory, to be useful, must call only for information that can be obtained and only for calculations that can be performed. The classical theory of rational choice has generally ignored these information-processing limitations. It has assumed that rationality was concerned with choice among alternatives that were already specified, and whose consequences were known or were readily calculable. It has assumed, also, comparability of consequences—that is, a practically measurable utility index.

Since these conditions, on which the classical theory rests, are so seldom satisfied in the real world, great interest attaches to procedures that make less heroic assumptions about the "givens" and the knows; and there is considerable progress in devising less-than-optimal decision procedures for situations where the optimum is unknown and practically undiscoverable. These procedures, often called *heuristic methods,* are distinguishable from optimizing techniques in three respects: they grapple, as most optimizing techniques do not, with the problems of designing and discovering alternatives, as well as with choosing among given alternatives; they frequently "satisfice," or settle for good-enough answers in despair at finding best answers; they commonly do not guarantee the qualities of the solutions they provide, and often do not even guarantee they will find a solution. The second and third of these characteristics are, of course, not virtues, but are the price that must be paid for extending our theory and tools for decision making to the wide range of real-world situations not encompassed by the classical models.

By way of illustration, a common problem of business and governmental management involves locating a system of warehouses over a country so that products can be distributed from production points to ultimate users as economically as possible. Attempts to formulate the warehousing problem so that the optimizing methods known as linear programming can be used have failed because the computations become too lengthy. However, heuristic techniques have been applied successfully to find "good" solutions to the problem where "best" solutions are unattainable.[6]

It is traditional to observe, in any discussion of the modern decision-making tools, that knowledge of these tools runs far in advance of application, and that the domain of application has been limited largely to decisions that are well-structured or "programmed," and quantitative in character. The warehousing problem described above has both of these characteristics. Whether this limitation on applications is inherent or temporary is a more controversial question. One of the important tasks before us now is to see how far we can go in extending the applicability of the new decision-making tools to areas that are ill-structured, and qualitative, calling for "judgment," "experience," and even "creativity." To do this, we shall presumably have to understand what "judgment," "experience," and "creativity" are, a topic discussed later.

Experiments on Decision Making

A second area of significant advance has been in applying the experimental method to the investigation of decision making. This has been done both by arranging for experiments on live real-world organizations—on the model of the Hawthorne experiments—and/or by bringing organizations, or organizationoid systems into the laboratory. For obvious reasons, the latter has been done more often than the former.

The first volume of the *Public Administration Review* contained a report of a large-scale field experiment on the decision-making

---

[6] Alfred A. Kuehn and Michael J. Hamberger, "A Heuristic Program for Locating Warehouses," *Management Science,* July 1963.

processes of social workers,[7] but similar experiments have been exceedingly rare in the succeeding twenty-five years. One of the few other examples to which I can refer is the study done in the Prudential Life Insurance Company by the Survey Research Center of the University of Michigan.[8] Either researchers on organizations decided that the information attainable from field experiments was not worth the trouble and cost of carrying out such experiments, or they found it difficult to secure the cooperation of business and governmental organizations in arranging such experiments—or both. Whatever the reason, field experiments have not been an important procedure for learning about organizational decision making.

In a few cases researchers have tried to import relatively sizeable organizations into the laboratory—hence, their studies lie on the boundary line between field and laboratory experiments. The Systems Research Laboratory of the RAND Corporation, for example, studied decision making by simulating, under controlled conditions, an entire air defense control center and associated early warning stations, manned on a full-time basis over a period of several months by a staff of some thirty subjects. While the studies conducted by the Systems Research Laboratory had as their direct outgrowth a major Air Force training program, the laboratory proved less tractable as a setting for obtaining data for testing theories of the decision-making process, and there has been no subsequent rash of studies of this kind.[9]

In contrast to the dearth of field experiments and large-scale laboratory experiments, laboratory experimentation with relatively small groups has been a thriving enterprise. Several examples of methodological advances in the art of small-group experimentation can be mentioned. Fred Bales, with his interaction process analysis, developed a scheme of data processing useful for studying the interaction of task-oriented and social-system oriented behavior in small problem-solving groups. Alex Bavelas devised a small-group task that permitted the experimenter to alter the decision-making process by opening or closing particular channels of communication between members of the group. In succeeding years, the Bales coding scheme and the Bavelas small-group task have both been used in a substantial number of studies, manipulating a great many different independent variables. Both have proved exceedingly valuable in permitting the cumulation of comparable knowledge from a whole series of experiments carried out by different investigators in different laboratories.

It is impossible to summarize here, or even to reference, the numerous contributions to the substantive knowledge of decision making that have been contributed by the small-group experiments. A single example will convey the flavor of such work. Cyert and March were able to produce bias in the estimates of members of a simulated organization by creating partial conflict of interest among them, but showed that under certain circumstances this bias did not affect organizational performance.[10]

New knowledge about organizational decision making can be obtained from appropriately planned experiments on individuals as well as from small-group experiments. Andrew Stedry, for example, has tested in this way theories about how budget controls affect behavior in organizations.[11] The series of studies of influence processes carried out at Yale by the late Carl Hovland and his associates belong in the same category.[12]

Persuasion and Evocation

Mention of the Yale research on influence processes marks a good point in our discussion to turn to several substantive developments in the theory of decision making. The notion that a decision is like a conclusion derived from a set of premises has been a useful

---

[7] Herbert A. Simon and William R. Divine, "Human Factors in an Administrative Experiment," 1 *Public Administration Review* 485-492. (Autumn 1941).

[8] N. C. Morse and E. Reimer, "Experimental Change of a Major Organizational Variable," 52 *Journal of Abnormal and Social Psychology* 120-129. (1955).

[9] Robert L. Chapman, *et al.*, "The System Research Laboratory's Air Defense Experiments," 5 *Management Science* 250-269. (April 1959).

[10] Richard M. Cyert and James G. March, *The Behavioral Theory of the Firm* (Prentice-Hall, 1963), pp. 67-77.

[11] *Budget Control and Cost Behavior* (Prentice-Hall, 1960), Chapter 4.

[12] See the Yale Studies in Attitude and Communication, edited by Hovland and Rosenberg, and published by the Yale University Press.

methaphor for analyzing the decision-making process. Following the metaphor a step further, we can view each member of an organization as "inputting" certain premises, and "outputting" certain conclusions, or decisions. But each member's conclusions become, in turn, the inputs, that is to say, the premises, for other members. For one person to influence another involves inducing him to use appropriate premises in his decision making.

What happens in an organization, or in any kind of social system, when there are conflicting premises pushing a particular decision in different directions? Much of the research on influence processes has been aimed at answering this question. In much of this research, influence has been conceived as a kind of "force," so that when several influences are brought to bear simultaneously, the outcome is interpreted as a "resultant" of the impinging forces. Persuasion is then a process of exerting such a force.

An important advance in understanding decision making has been to complement the notion of persuasion with the notion of evocation. When we want someone to carry out a particular action, we may think of our task as one of inducing him to *accept* latent decision premises favorable to the action that he already possesses. Thus, writing about food will often make a reader hungry, but we would hardly say that we had "persuaded" him that he was hungry; it would be better to say that we had "reminded" him.

Processes of persuasion play their largest role in decision making in conflict situations —where the issue is already posed, and the alternatives present. This is the framework within which most of the Yale studies on attitude change were carried out. It is also the framework for the important and well-known study of *Voting* by Berelson, Lazarsfeld, and McPhee.[13]

On the other hand, in studies of decision making where the focus of attention of the participants is one of the main independent variables, the evoking processes take on larger importance. The recent study of the Trade Agreements Act renewal, by Raymond Bauer, Ithiel Pool, and Lewis Dexter indicates that these processes played a major role in deciding the issue.[14] The authors describe the setting of their study thus (p. 5): "We are interested in the sources of information for each of these populations, the bases of its attitudes on the trade issue, *and the circumstances which lead some individuals to take active roles in the making of policy.*" (Emphasis supplied.) They demonstrate convincingly that the behavior of particular Congressmen on the trade issue depended as much on the alternative claims on their time and attention as on the distribution of interests of their constituents.

To the extent that the mechanism of evocation is important for decision making, many new ways arise in which organizational arrangements may affect behavior. As example, one of the findings of the study just mentioned (p. 229) can be cited:

> In summary, we would suggest that most significant of all to an understanding of what communication went out from business on foreign trade was neither self-interest nor ideology, but the institutional structure which facilitated or blocked the production of messages. Whether a letter to a congressman would get written depended on whether organization facilitated it, whether the writer's round of daily conversations would lead up to it, whether a staff was set up to produce it, and whether the writer conceived writing this letter to be part of his job.

Evoking mechanisms take on special prominence wherever dynamic change is occurring. Studies of the diffusion of innovations show that the timing of adoption of an innovation depends critically on the means for getting people to attend to it.[15] From every point of view, the new knowledge gained about evoking and attention-directing processes is a major substantive advance in our understanding of organizational decision making.

The Structure of Decisions

A decision is not a simple, unitary event, but the product of a complex social process generally extending over a considerable period of time. As noted, decision making includes attention-directing or intelligence processes that determine the occasions of decision, processes for discovering and design-

---

[13] University of Chicago Press, 1954.

[14] *American Business and Public Policy: The Politics of Foreign Trade* (Atherton Press, 1963).

[15] See J. Coleman, E. Katz, and H. Menzel, "Diffusion of an Innovation Among Physicians," 20 *Sociometry* 253-270. (1957); also, H. A. Simon and J. G. March, *Organizations* (Wiley, 1957), Chapter 7.

ing possible courses of action, and processes for evaluating alternatives and choosing among them. The complexity of decision making has posed grave difficulties in its study and description, difficulties only now being overcome by recent methodological innovations.

Traditionally, a decision-making process was captured and recorded by the commonsense tools of the historian using everyday language. The notion that a decision might be viewed as a conclusion drawn from premises—a notion mentioned earlier—introduced a modicum of system into the description of decision making. According to this view, in order to record a decision-making process it was necessary to discover the sources of the decision premises, and the channels of communication they followed through the organization to the point where they became the raw materials of decision.

Studies that adopted this general approach to the description of decisions, while remaining within the traditional case-study framework, became increasingly frequent during the period under discussion. One example is Herbert Kaufman's excellent study of *The Forest Ranger,* aimed at analyzing "the way their decisions and behavior are influenced within and by the Service."[16] Another is the study by the Carnegie Tech group of the influence of accounting information on operating decisions in large companies.[17]

The method of these studies is best described as "systematized common sense." The decision premise concept provides an ordering and organizing principle; it reduces somewhat the subjectivity of the description and the dangers of observer bias; but it falls far short of allowing complete formalization of the description. And it cannot, of course, solve the problem of how to validate generalizations with data from single cases.

The invention of the modern digital computer radically changed the situation. As gradually became apparent to those who came into contact with computers, the computer is a device that is capable of making decisions. (One demonstration of this is its use to implement the analytic decision-making schemes introduced by operations research.) Hence, a language suitable for describing the processes going on in computers might well be appropriate for describing decision making in organizations. At least the notion appeared to be worth a trial: to equate "decision premise" with the concepts of data input and program of instructions in a computer, and to equate the concept of a conclusion with the concept of the output of a computer program.

An early, and rather primitive, attempt to describe an organization decision-making process in computer programming terms appeared in 1956.[18] In this study the authors recounted the steps taken by a business firm to reach a decision about the installation of an electronic computer. They then showed how this sequence of events could be explained by a program composed of an organized system of relatively simple and general information-gathering, searching, problem-solving, and evaluating processes. Of particular interest was the fact that the decision examined in this study was not a highly structured, quantitative one, but one that called for large amounts of professional and administrative judgment.

Encouraging results from early studies of this kind raised hopes that it might be possible to use computer programming languages formally as well as informally to construct theories of organizational decision making, and to test those theories by simulating the decision process on the computer. Computer programs seeking to explain several kinds of organizational decision-making situations have, in fact, been constructed, and have shown themselves adequate to simulate important aspects of the human behavior in these situations. The decisions that have been simulated in this way to date are still relatively simple ones, but they encompass behavior that would generally be regarded as professional, and as involving judgment. Two of the best-developed examples are a simula-

---

[16] Johns Hopkins U. Press, 1960, p. 4.

[17] *Centralization versus Decentralization in Organizing the Controller's Department* (New York: The Controllership Foundation, 1954). The study is summarized in John M. Pfiffner and Frank P. Sherwood, *Administrative Organization* (Prentice-Hall, 1960), Chapter 21.

[18] R. M. Cyert, H. A. Simon, and D. B. Trow, "Observation of a Business Decision," 29 *Journal of Business* 237-248. (1956).

tion of a department store buyer and a simulation of a bank trust investment officer.[19]

I am not aware that any single comparable simulation of a decision-making process in the area of public administration has yet been carried out, but it appears that several are under way in current research. Perhaps the most likely target for initial attempts is public budgeting. If we examine the strategies described in recent empirical studies, like those of Wildavsky,[20] we will see that they can be rather directly translated into components of computer programs.

Parallel with these simulations of administrative decision making there has been a considerable exploration of individual thinking and problem solving processes, also using computer simulation as the tool of theory formulation and theory testing.[21] Today, we have a considerable specific knowledge on how human beings accomplish complex cognitive tasks. We have reasons for optimism, too, that this body of knowledge will increase rapidly, for in the digital computer language we have an analytic tool and a means for accurate expression whose powers are commensurate with the complexity of the phenomena we wish to describe and understand.

[19] Descriptions of these two simulations may be found in Chapters 7 and 10, respectively, of Cyert and March, *Behavioral Theory of the Firm, op cit.*

[20] Aaron Wildavsky, *The Politics of the Budgetary Process* (Little, Brown and Company, 1964).

[21] For a survey, and numerous examples, see Edward Feigenbaum and Julian Feldman, *Computers and Thought* (McGraw-Hill, 1963).

Landmarks and New Roads

These, then, are some of the more prominent landmarks along the road of decision-making research over the past twenty-five years. On the normative side, the analytic tools of modern operations research have secured an important place in the practical work of management. Their role in everyday decision making promises to be much enlarged as present techniques are supplemented by new heuristic approaches.

On the side of the pure science of administration, there have been equally fruitful developments. The experimental method, in the small-group laboratory, can now be used to study a wide range of decision-making behaviors that are relevant to organizations. We have introduced the concept of evocation into our theories of influence, and have used it to gain new understanding of the decision-making process in changing environments. Finally, the modern digital computer, a powerful new tool, has provided both a language for expressing our theories of decision making and an engine for calculating their empirical implications. Theories can now be compared with data of the real world of organizations.

The attention-directing mechanisms so important in decision making also have played their part in determining the particular developments sampled in this paper. Another scientist, with a different set of research concerns, would chose a different sample. The fact that even one such sample exists shows how far we have come during the past twenty-five years toward understanding human behavior in organizations.