

KNOWLEDGE AND ACTION

C. WEST CHURCHMAN

Center for Research in Management Science

University of California, Berkeley

Those who have thought and written about management often refer to management "know-how." By this they apparently mean the ability of the manager to understand what is going on in his environment. It is the chief concern of the management scientist to translate these vague stipulations about management knowledge into precise and verifiable assertions about how decisions are made and how decisions ought to be made.

One of the central problems of management science is the understanding of the role of information in decision making. The reason why this problem is so difficult is that we have failed to pay enough attention to the very subtle concept of the *use* of information. The much maligned, classical economic man was supposed to act in accordance with his self-interest, "given" complete information. But what does "given" mean in this context? What the economist meant was that information was automatically fed into the rational decision process, and the correct answer was thereby derived. If the manager of a firm can be adequately represented as a computer, this concept of "given" may be adequate, for then we merely mean by "given" information the inputs to a computer program.

Presented at TIMS International meeting, August 21-24, 1963, Tokyo.

However, when we consider real managers in their real environments, the concept of "given" is not clear at all. Suppose we picture a manager busily at work making a number of decisions. Suppose by some lucky chance that sufficient information is "available" for him to make a perfect decision in each case. We need some operational definition of "available," and the one that comes most readily to mind is that information is "available" if it is stored in the form of retrievable symbols in the environment of the manager. More precisely, a price of information is available to a manager if he can retrieve it at virtually zero cost (e.g., by asking someone, or looking it up, or retrieving it from computer memory). Even this more precise meaning of "available" leaves much to be desired, but the definition will suffice for the purpose of stating the problem of this essay.

The problem is this: there is sufficient evidence to show that a manager may have perfect information "available" to him and yet not make the correct decision; what is the explanation of this phenomenon? Why don't managers act on the information stored in their environments?

Before proceeding to discuss this problem in depth, something more needs to be said about the concept of a "rightchoice," i.e., about knowledge of the correct action.

A man may choose the right action by accident, so to speak, when in his clumsy way he stumbles upon the correct way to act. Thus, an archer who hits the bull's eye does not necessarily know how to shoot. What makes the expert archer is the ability to adjust his method of controlling the arrow that no matter what the motion of his target, within limits, he can perform successfully.

Hence, we will want to say that knowledge about a decision consists of choosing the correct action even though the conditions change. Knowledge is a sensitivity to changes in one's environment.

The point can be illustrated very well by referring to inventory control within industry. Most companies do reasonably well with their inventories during periods of stable demand. This is because the mana-

gers responsible for controlling inventory learn from experience that ordering too large or too small quantities is bad policy, and they have come to learn how to prepare for normal kinds of shifts in demand in terms of buffer stocks. But we would not say that the management had knowledge of the control of inventory because it has no way of adjusting to large shifts in demand or in costs.

Hence, when an operations research team comes in and tries to develop models, it often finds that its recommendations do not deviate very far from the current practice of the management. This does not imply, however, that the work of the operations research team has been wasted. On the contrary, once a model is developed the parameters of the model can be adjusted to take care of variations in the environment, such as the upswing or downswing of demand, the costs of carrying inventory or of placing orders or of shortages. Thus, the great advantage of modeling industrial operations is that the manager is then in an excellent position to adjust for changes. In other words, according to the definition of "knowledge" introduced above, the manager is provided with knowledge of the right action rather than more display of the right action.

The point cannot be overemphasized. Often in companies we find that men with years of experience have arrived at methods of managing operations which cannot be matched by analysis. I recall a vice president of sales at one of the companies with which we worked at Case Institute of Technology. This gentleman would forecast sales of the company's products based on reports received from the field, economical data and other types of information, coupled with his own judgment. He seemed to have an uncanny way of coming up with accurate estimates that we could not duplicate by any statistical methods known to us. In this case the company could be said to have arrived at the correct action, but again I doubt if we would want to say that the company had knowledge of forecasting. Once the vice president leaves, the environment is changed and the company has no obvious way to adapt to the change. Thus, our statistical methods, though perhaps slightly inferior to those of the

vice president under current conditions, constituted something more like knowledge than the vice president's decisions.

In the discussion above, it is clear that I have been thinking of knowledge as a matter of ability. It is the ability of a decision-maker to adjust his decisions to changes in the situations that confront him. We could go on to describe deeper forms of knowledge. For example, we might want to talk about the way in which knowledge of decision-making in one area assists decision-making in another area or the way in which knowledge of decision-making can be refined, so that finer and finer distinctions can be made between types of decisions. I think these discussions would lead us to a consideration of "understanding" as opposed to mere knowledge, because understanding is essentially the process by which the decision maker knows why his method of decision-making is correct. To know "why" is to go beyond the present situation to a larger world and to understand why the decisions in the present situation are justified by considerations of the larger world. Essentially, understanding represents the highway from knowledge to wisdom: wisdom is the most general form of understanding.

Assuming that enough has been said about knowledge, suppose we return to the problem stated at the outset. We have pictured a man in an environment where sufficient information is "available" to provide him with knowledge of how to make decisions. Our problem is why doesn't knowledge occur?

One ready answer is that the information is too costly to obtain. Indeed, we have hidden in the definition of "available" the very difficult concept of "cost of information," and we could devote many pages to discussing this concept as is done, for example, in the work of Marschak and Rander (1964) on the theory of teams. Instead, in order to put the problem of this essay into sharp relief, we will assume for the moment that we are talking about a manager who can retrieve information readily at no cost. Why does manager fail to use the information so readily available to him?

The obvious answer is that he does not know how to transform the available information into a knowledge of action. In other words, he does not have a program built into his mind that will transform the inputs into correct outputs. Thus, operations researchers have suggested to the manager the wisdom of using personnel who do know how to do just exactly this: they know how to take the information available and calculate a decision system which, in effect, provides the manager with the knowledge of how to behave. We generally refer to this type of activity as analysis. Hence, the suggestion is that the missing gap for many managers is the lack of someone or some device capable of performing the necessary analysis.

Unfortunately, we have overwhelming evidence that available information plus analysis does not lead to knowledge. A management science team can properly analyze a situation and present recommendations to the manager, but no change occurs. The situation is so familiar to those of us who try to practice management science that I need hardly describe the cases. Some of my graduate students undertook to write to the authors of cases reported in *Operations Research* over the first six years of its publication, to determine to what extent the recommendations of the studies had been carried out by management. In no case was there sufficient evidence that the recommendations had been accepted.

At the University of California, Professor Ratoosh and I have been conducting some experiments in which we have been able to establish in controlled situations that analysis plus available information does not lead to knowledge. In our experiments, five subjects run a small firm. The instructions given to the subjects provide complete "available" information about the firm's operations. That is, it is possible so to analyze the information that the managers of our laboratory firm can adopt the right action. Furthermore, by analysis it is possible to generate a model which will provide the managers with a method of making decisions, even though demands and costs and other factors change. In other words, the model can provide the managers with a *knowledge* of how to act. In some

of our groups we have placed a person who has gone through the analysis beforehand. He is then given the task of persuading the rest of the management team. Therefore, in our laboratory we have constructed a situation in which there is available information plus analysis. The result has almost uniformly been one in which knowledge does not occur. The managers do not even accept the recommendations, and are far from gaining any knowledge from the analysis. See Churchman and Ratoosh (1960 and 1961).

Of course, one could say that the difficulty here is that the available information is in one location (the instructions given to the subjects), the analysis is in a second (in the head of one person) and the decision-making is in a third. One might feel that if all three of these components could be combined into one unit, the problem could be solved.

This suggestion is a reasonable one, and makes us return to the original definition of "available" information with the purpose of revising it somewhat. The original definition made "available" a very passive kind of thing, and ignored the very important concept of transmittal of information. How is the information to be retrieved? If the channels of retrieval are obscured by linguistic ambiguities and other noises, the information may not be "available," no matter how precisely it is stored.

The concept we seem to require is "communication." Communication is a device for taking several minds and making them act more like one mind.

Thus, many people have suggested that the missing variable in the equation is proper communication: they say that there exists some form of communication such that available information *plus* analysis *plus* communication leads to knowledge. They say that if only operations researchers could learn to talk in a language the managers can understand, there would be no further difficulties in implementing operations research recommendations. Thus, articles on the subject of implementation of recommendation speak over and over again about the need for better communication between the scientist and the manager. These articles

point to the fact that the scientist talks in a semi-formal language (mathematics) and the manager talks in his own management terms, which themselves are not exactly those of normal discourse in his society. The problem, say these articles, is to make the proper translation.

Of course, to prove the invalidity of the new formula (information + analysis + communication leads to knowledge) is very difficult because there are so many modes of communication, but in our experiments we have found no adequate mode. We have presented the solution in simple graphical form, in simple arithmetical form, in ordinary discourse, as well as in the form of more complicated mathematical expressions. We can find no evidence that the mode of communication makes any difference. We have presented the solution in pieces and in a total report and, again, we have found no difference. I am aware of the fact that the sociological literature, and especially the work of Hovland (1954 and 1957) has talked a great deal about the importance of modes of communication with respect to persuasion, but we have not been able to translate these findings into our own work.

Of course, I am not saying that the mode of communication has no effect whatsoever. It is not very difficult, even in our experiments, to prepare messages which are so incomprehensible that they have no effect whatsoever, except to annoy the managers. What I am saying is that the formula will not do.

More precisely, the prescription for "better" communication ends up by being no more than a restatement of the problem: namely, to find some set of activities such that with "available" complete information, the manager will come to know the correct action.

It is the attempt to find such a set of activities that makes the problem become very difficult because by now we have exhausted all the obvious possibilities.

The next suggestion is one with which I am sure most managers and scientists are familiar. It points out that one must bear in mind the distinction between the personal goals of the manager and the organiza-

tional goals. Even if the corporate goals are explicit and clear, the manager's own personal ambitions may be at variance at times with the interests of the corporation. Therefore, if the scientist generates a method of making decisions that best serves the corporate aims, he may encounter resistance because a particular manager feels threatened or because the recommendations do not fit into the personal goals of the individual managers.

In other words, it is possible that a manager who really wanted to accomplish the organizational goals, would come to know the correct action.

Thus, it is a well known fact that managers feel threatened by modern techniques of analysis. They have reason to believe that analysis, with the help of high speed computers, may take over their roles. In order to preserve their status, they resist the recommendations that analysis provides. At this point one usually distinguishes between a perceived threat and a real threat. The perceived threat is one that may be incorrect; for example, it may happen that sophisticated analysis strengthens rather than weakens the manager's role. The task of the researcher in this case is one of clarifying the situation for the manager, though how this is to be done is often very obscure. A real threat is one that does in fact threaten the manager's role, either making it obsolete or less important.

A more general way to describe the problem is as follows. When a change is suggested to a manager, he reacts in ways that are typical of him, i.e., of his personality. Part of him will resist the suggestion, and often the resistance is so strong that he will reject the suggestion altogether. Nevertheless, the persons responsible for making the suggestion may learn enough about the manager's personality so that they know what to say to him to break down his resistance. Thus, they recognize that there is a part of the manager's personality that will adjust to change. On the other hand, there may be a part that remains invariant, no matter what the environment. Indeed, psychological literature leads us to believe that most men display invariant characteristics over most of

their lives: they rarely change from an extrovert to an introvert, or from a thinking type to a feeling type, for example.

These reflections imply that the missing ingredient in the process of implementation is the understanding of the manager. Any research team that fails to study the manager and his personality may very well be made up of persons, books, journals, and other communication devices of the manager's environment.

Thus, in the context of managerial politics, the researcher is apt to find that his recommendations are viewed from the point of view of their effect on a coalition, and not from the point of view of the whole organization. Since the researcher usually doesn't know who belongs to what coalition, and is far from understanding what holds the coalition together, he cannot determine how to overcome coalition biases. He finds himself in a confusing welter of contradictory reactions of managers.

The most significant outputs of the hidden managerial coalition are the importance of issues and the ways in which important issues should be considered. Managers control a very scarce commodity: their own time and attention. Their most conscious problem is one of determining what they should pay attention to. A researcher who claims he can save them a few thousands or millions of dollars may immediately lose their attention, because the managers believe that a new market, or a threat to the corporation's existence, are far more demanding of their attention than cost savings, especially if the method of analysis is so alien to them.

In other words, one aspect of the formula suggested earlier was quite faulty. It may be true that information+analysis ideally leads to knowledge, but analysis takes up a significant portion of the manager's time and energy. One measure of the cost of analysis *for him* is the distance of the analytic method from his typical way of thinking about problems. His typical way of thinking comes from his coalition. A very striking example of a manager's reaction to "scientism," i.e., alien thinking, is to be found in a recent article of Lilienthal's (1963).

Thus, to discover what a manager thinks is important, and how he believes he should think about important issues, one must determine to whom he listens, and to determine to whom he listens one must understand the coalitions to which he belongs. Some of these coalitions are external to the company; they are the other members of the managerial fail to bring about a recommended change. Furthermore, any research team that believes it can implement a recommendation by the same process, regardless of who is managing, is simply naive. For a further discussion of this point, see Cyert and March (1963).

Because he often lacks a methodology of understanding people, the researcher may give up any attempt to implement broad changes of policy. Instead, he may be satisfied to work in areas where the status and role of the manager remains invariant no matter what is changed. A manager may not feel threatened if the equipment in his shop is redesigned, or if physical sequences of actions are changed. But if someone asks whether his shop should manufacture the items it does, in the quantities it customarily makes, he cannot help but regard this question as one that is directed to his own role. If he is a reflective type, he may enjoy the question; if he is an anxious type, he may not. Even if he enjoys the question, he may believe that a solution arrived at by analysis stifles his imagination. His personality may be such that he must act instinctively and creatively, or not act at all.

Perhaps the most important invariances of personality occur in the formation of coalitions in organizations. There are the so-called political aspects of management in firms. A coalition may arise because the members of the coalition recognize certain common economic advantages. But it is well known that coalitions also come about because of the personalities of their members; some mix of attitude, trait, and opinion creates a loyalty that is hard to dissolve. The loyalty is strengthened by opposition. Thus, various obscure and complicated coalition frameworks occur among managers; they are obscure because no one ever writes down their bylaws and articles of confederation, or ever announces them publicly.

Indeed, most managers are not clearly aware of them.

These coalitions of managers are like the “invisible colleges” of scientists; see for example, Derek Price (1961). They are the sources the manager consults: they are the basis of the language he uses; they provide the criteria of what is important and what is unimportant. They are the community whom the manager respects. These large coalitions of managers in specific industries account for the managerial styles and fads. If these coalitions come to believe that science, and especially mathematics and computers, are important, managers will pay attention to scientists. If not, they won't pay attention, no matter how elegant the recommendations the scientists make.

I have purposefully been vague about the meaning of a coalition, in order to emphasize a point. But the definition is really quite simple: a coalition of a manager is the group of people who influence what the manager attends to. I do not say that they are the group that influences what he does, because this is too narrow a concept. The manager of a competing firm may hire an operations research group. If this influences a manager to learn about operations research, the first manager belongs to the second's coalition.

There is little doubt that in the experiments we have run at Berkeley, the problem is one of attention. The subjects quickly form one or more coalitions, and then become much too busy to want to attend to the recommendations that are made to them, even though their recommendations are correct. The analysis offered is foreign to the way in which they have taught themselves to think about their task. In some mysterious way, the subjects agree that the way they are organized, “not going into the red,” and various other secondary aspects of the task are the most important.

If we are to learn more about the implementation of recommendations, we must learn more about how people decide where to direct their attention. It is for this reason that at System Development Corporation, Herbert Eisenberg, Martin Shubik and I have started a few very simple

experiments on deliberation and its role in decision making. Deliberation is the process by which the mind in reaching a decision scans various aspects of the problem. In our experiments, we present one subject with alternative points of view to which he may pay attention, and other subjects try to influence him towards one viewpoint or another. We are attempting to learn more about the way in which managers come to pay attention to issues.

The story of this essay is a mystery novel, with the added frustration that the culprit remains unidentified even at the end. Perhaps the advantage of such a devilish novel is that it may suggest a better plot for the next one to be written. With this in mind, let me end by introducing some broad philosophical generalities, not supported by the "available" evidence, but nonetheless helpful in future research.

We started by looking for the ingredients that a research staff would seek to supply in order to bring about a recommended change: available information, analysis, and communication. We argued in the end that none of these ingredients matters at all unless the manager pays attention to the problem, and that paying attention is an obscure process of the managerial mind, little understood by management scientists.

The obscurity we face in this regard is simply the obscurity of the concept of decision making itself. We do not yet understand how to describe a human decision. The descriptions usually offered reflect the psychological traits of the describers. A thinking type believes that the mathematical theory of optimization and specifically of games will provide all the concepts necessary to define a decision clearly, as well as a correct decision. He is arrogant enough to label his efforts "decision theory," without any uneasy pang of conscience. The feeling type asserts that a decision is essentially a unique expression of human values, and that the meaning of a decision cannot be captured by generalized mathematical expressions. An intuitive type believes that decisions are insights, quick-flashes of understanding how to solve a problem. They frequently assert that the manager leaps to his conclusion without benefit of or even need

for analysis. Finally, there are those who assume that the whole business of decision making is contained in available information: what decision is made depends solely on what facts are known.

Philosophy attempts to take one step back from the issues that divide scientists into intellectual camps. All the points of view mentioned above are valid. They amount to saying that decision making can be conceptualized in many ways. What seems to be common to these ways of describing decisions is the concept of a focus. (See Cowan, 1961). A specific decision is the focus of a mathematical model, a general value structure, of insightful behavior, of masses of data. The focusing that leads to decision making takes the manager's whole world and displays a sub-area where he must seek a solution. Within this subarea the coalitions that influence his behavior lead him to confine his attention to certain aspects. Eventually he is led to one alternative, to *the* choice.

The rational mind of the scientist would like to remove all irrationality from this focusing of attention of the manager. The trouble is that in order to do this, we scientists must understand the world of the manager: not a piece of it, but the whole world. If we only understand a piece of the manager's world, we have no justification for asserting that he should pay attention to the piece that we present to him.

Thus, in order to recommend important changes to a manager, we must understand the process by which this whole world becomes focused on certain issues and aspects of his environment. Any decision is a snapshot of the universe of the manager. An optimal decision is a snapshot of the rational universe.

REFERENCES

1. Cowan, Thomas A.: An address delivered at the American Association for the Advancement of Science meetings in Denver, December, (1961).
2. Cyert, Richard M. and March, James G.: "Behavioral Theory of the Firm," Prentice-Hall, Inc., (1963).
3. Hovland, Carl I., Janis, Irving L. and Keeley, Harold H.: "Communication and Persuasion," Yale University Press, (1954).

4. Hovland, Carl I., *et al.*: "The Order of Presentation on Persuasion," Yale University Press, (1957).
5. Lilienthal, D. E.: "Skeptical Look at 'Scientific Experts'." New York Times, September 29, (1963).
6. Price, Derek J. de Solla: "Science since Babylon," Yale University Press, (1961).
7. Radner, Roy and Marschak, J.: "Economic Theory of Teams" to appear in 1964.
8. Ratoosh, Philburn and Churchman, C. W.: "Innovation in Group Behavior," Working Paper No. 10, Center for Research in Management Science, University of California, Berkeley, January, (1960).
9. -----: "Report on Further Implementation Experiments," Working Paper No. 26, Center for Research in Management Science, University of California, Berkeley, March, (1961).