

1

Preliminaries



Ralph stared at the figures in disbelief; he felt a headache coming on. What happened? After all the planning and work, all the careful analysis and getting his people to buy in, all the sleepless nights and worry, what went wrong? When he took this job he was sure he knew how to deal with the situation—he was the guy who had previously saved three other businesses that had the same problem. Had he been too optimistic? Had his past experience hindered rather than helped? Betty Miller, the company president, clearly believed he was on the wrong track. He hated to prove her right. After all, he was supposed to be the Knight in Shining Armor who would save the firm. Clearly, this Knight was going to have to think things through again.

He flipped back over the figures from previous months, comparing them carefully with the latest report. What were they trying to tell him? Had something changed that he had failed to notice? Was he focusing on the wrong things? Slowly he began to see a pattern. Although not broken out in a way that would make it obvious, the figures suggested that the problem was not with the products themselves but in sales. Some salespeople appeared to be doing a great job, but many more were not. That might be it! He had always read the reports from the viewpoint of fixing the product. Maybe that mind-set blinded him to the possibility that some salespeople simply do not do a good job. He felt like a fool. Most managers made just the opposite mistake, blaming the people first

2 THE PSYCHOLOGY OF DECISION MAKING

and the product last. Okay, he was more into product innovation than human resources, everyone knew that. But now that he began to look at the figures this new way, everything began to make sense.

Ralph felt a surge of relief and excitement. All was not lost—not yet, anyway. He was going to have to get a lot more information to make the details clear, but he thought he understood the problem. True, it was not something he had ever dealt with before, but he and Betty ought to be able to figure it out. Then it was a matter of coming up with plausible options, deciding on the right one, working hard, and watching closely so he didn't get misled again. It could be done. He almost called Betty to tell her. She probably would be so relieved that she wouldn't gloat. But, on second thought, maybe he would tell her tomorrow so he could enjoy this feeling of relief a little while longer.

To better understand the problems Ralph faces, we need way of structuring our thinking. The usual way begins with the diagnosis of the decision problem, moves on to selection of an action that will solve the problem, and ends with implementation of the selected action until the problem is solved:

Diagnosis → Action Selection → Implementation

Diagnosis

The need for a decision arises when anomalous events occur. Often these events stem from internal changes (wants) or external changes (demands), but they also can stem from the realization that an earlier decision was wrong and its implementation is not producing the desired results. In Ralph's case, the anomalous events were the unexpected figures in the monthly report, resulting in the realization that what he was doing was not making things better.

To make sense of the anomalous events, the decision maker must mentally put the events in the proper context in order to give them meaning (*framing*), which allows him or her to draw upon previous experience to decide what to do. If this is a situation that is very similar to a situation that has been encountered before, he or she can use that experience to deal with the events. If the situation is substantially different from previously encountered situations, he or she can set about

formulating an action plan that deals with its uniqueness. In Ralph's case, the company's problem had been interpreted as a product problem. But the figures in the reports did not make sense in this interpretation, suggesting that the frame was incorrect. When he began to view the situation as a sales problem, the figures suddenly began to make sense.

Whether the decision maker's framing of the events results in use of previously acquired knowledge about what to do or in formulation of a new action plan, he or she must use the events to guide the fine-tuning of the response. Even past experience usually provides only a general strategy for dealing with the situation. (When one can do exactly what one did before, it is called a conditioned response, which seldom is an effective way to deal with complex situations and almost never is effective for unique situations.) Therefore, the decision maker must *diagnose* the situation by evaluating the states of its most salient features. In Ralph's case, his past experience with product problems provided little help for dealing with the present sales problem, which probably is really a human resource (HR) problem. Therefore, he anticipated having to obtain more information in order to make judgments about precisely what was wrong and how serious it was before he could decide what to do about it.

Action Selection

We often talk as though the decision maker has some set of potential plans of action (options), and the decision consists of choosing the best from among them. However, consider your own decision making—you frequently have no idea what might be reasonable plans of action. Often you start off in one direction, only to change your mind when things go awry. In fact, decisions seldom are made at a single point. Rather, the process seems to feel its way along, changing in the light of feedback and often leading in directions that never were conceived of when it all began. In Ralph's case, he thinks that with Betty's help he can come up with options from which to choose, but the fact is that once he understands the nature of his sales/HR problem, he will start to do the most obvious things and then adjust his actions in light of how well they work.

The "feeling along" nature of decision making is very difficult to capture in a readable narration, and even more difficult to describe in a tractable theory. It is far easier, for both the theorist and the reader, to talk as though decision options actually exist in full-blown clarity and

4 THE PSYCHOLOGY OF DECISION MAKING

the choice among them occurs at a single point in time. But this is merely a useful fiction. In Ralph's case, the fiction is useful because we can think of him making choices among different options—even if he makes a sequence of choices in light of feedback about how things are progressing in his search for a solution to his sales/HR problem. The fiction also is useful to Ralph because it allows him to think in terms of rather specific plans of action—even though he knows that anything he decides to do will surely be modified as he goes along.

Implementation

After a plan of action has been decided upon, it must be implemented. That is, the plan must be used to guide behavior, and the decision maker must monitor its progress toward resolving the anomaly that started everything in the first place. Again, even though Ralph will feel his way along just like anyone else, he may talk as though he has a well-formulated plan—if only because it is easier to communicate with his colleagues and to help his HR people understand what is required of them. He may even believe it himself. But, after the dust settles and he has time to look back over what finally happened, the odds are good that Ralph will see that the route he took does not look much like the original plan. Indeed, if he laid a history of what happened side by side with the original plan, the only things they might share is a common starting point and, if he had successfully solved his sales/HR problem, a common ending point.

❖ SOME BACKGROUND

We will return to the “Diagnosis → Action Selection → Implementation” concept because it is the blueprint for the remainder of this book. For the moment, let us leave Ralph, who is busy enjoying his new insight into how to frame his problem, and turn to consideration of some of the larger issues in studying decision making.

Prescriptive Theory

First of all, decision making is studied by many different disciplines, each of which seems to regard it as uniquely its own. Economists address decision making by constructing axiomatic models that describe the

market forces at work in particular circumstances and that prescribe appropriate actions in light of the assumptions underlying the models. Operations researchers follow a similar logic, except that their models tend to be limited to specific problems encountered in specific enterprises. Applied statisticians also model decisions, either prescriptively (e.g., hypothesis testing) or descriptively (e.g., structural modeling).

With the possible exception of structural modeling, the emphasis in most work on decision making has been on prescribing what should be done, not on describing what decision makers actually do, and certainly not on diagnosis or implementation. This is because the models are designed solely to address the tasks involved in choice, which, as we shall see, is but one kind of decision making. Moreover, the models' logic need not mimic or even parallel the cognitive processes of the decision maker, even though their authors often use language that implies that it does. As a result, until quite recently, decision behavior was evaluated in light of how well it conformed to the prescriptive models rather than the other way around. That is, because the models followed logically from what usually were regarded as very attractive axioms, it did not seem reasonable to evaluate them by comparing them to what decision makers actually do. If such a comparison ever was made it was the behavior that was evaluated: Behavior that conformed to the models was judged to be rational and behavior that did not conform to the models was judged to be irrational. Prescriptive theory almost always focuses on choices among options and seldom focuses on either diagnosis or implementation.

Because psychologists, particularly students of organizational behavior, are more interested in describing what people do than in prescribing what they should do, we approach decision making in quite a different way. We are interested in the interplay of group and institutional dynamics and their effects upon the decisions made within and on the behalf of organizations. As a result, we are interested in the ways in which decision makers use information to arrive at decisions, and this is where our story begins.

Behavioral Theory

Behavioral decision theory began as the study of the degree to which unaided human decision making conforms to the processes and outputs of prescriptive decision theory. Almost from the beginning, however, it has gone beyond this rather narrow mandate by "psychologizing"

6 THE PSYCHOLOGY OF DECISION MAKING

prescriptive theory in order to make it more descriptive of what decision makers actually do.

Diagnosis

Behavioral research on diagnosis is based on the work of Egon Brunswik (1947), who studied perception. In doing so, he developed a model that describes how observers use sensory cues to infer (diagnose) the nature of the world around them. Hammond (1955) generalized this model, called the Lens Model, to more macro-level inferences. The result is that diagnosis has come to be thought of in terms of *policies* that decision makers use to evaluate the characteristics of decision situations prior to making decisions. The Lens Model has provided a powerful way of thinking about policies and has contributed greatly to an understanding of how decision makers use policies in situations as diverse as evaluating the quality of hogs (Phelps & Shanteau, 1978) and evaluating the quality of applicants for jobs as insurance salesmen (Roose & Doherty, 1976). These studies form the foundation of an empirical tradition in the study of policy use that has proved to be of enormous theoretical and applied value.

The rather offbeat Phelps and Shanteau (1978) hog study illustrates how this kind of research is done. In the first of two experiments, seven members of the Kansas State University senior livestock judging team were each given a set of 64 index cards. On each was written a description of a hypothetical female breeding hog. The descriptions consisted of 11 characteristics of hogs, and each hog was said to be high or low on the characteristic (e.g., high body weight, low quality of nipples). The high and low values of the 11 cues across the 64 hypothetical hogs formed a partially replicated factorial design. Each livestock judge was asked to read each description and rate the hog for breeding quality. Thus the 11 cues were the independent variables, each judge's 64 ratings was the dependent variable, and the analysis was done using analysis of variance. The results showed that all of the judges used nine or more of the cues to arrive at their ratings, but it was not clear whether some cues counted for more than others. In the second study, the judges were shown pictures of hogs and rated each hog. This time results showed that most of the cues were used but that they were used in clusters. That is, various cues contributed to an inference about size, others contributed to meat quality, and others to breeding quality. This was interpreted to mean that the judges' policies involved a multi-stage

strategy in which cues are integrated into intermediate inferences, which then are combined to arrive at a final decision.

Choice

Behavioral research on choice began by comparing decision behavior with the dictates of prescriptive decision theory. The basic concept underlying prescriptive theory is that a decision consists of selecting one course of action, called an *option*, from an array of options. Because most options contain an element of risk, selecting one can be thought of as making a bet, and the goal is to select the option that is the best bet. In short, decision making is seen as analogous to gambling: The decision maker is a gambler who places a bet by selecting one of the options and then waiting to see whether the outcome does or does not materialize. This way of thinking about decision making is called the *gamble analogy*.

The key to choosing the best bet is conceptually very simple. The player should select the option that offers the most attractive payoff package. A simple way to summarize the attractiveness of an option's payoff package is to add up the values of the various payoffs it offers. However, this fails to take into account the possibility that those payoffs may not materialize if the option is chosen. To fully reflect the structure of the situation, the attractiveness of the payoffs must be tempered by the probability of acquiring them. To include the probabilities, the attractiveness of a bet can be summarized by first discounting each payoff by the probability that it will be acquired and then summing the discounted values of the payoffs. Because a probability is a decimal number, discounting is accomplished by multiplying the payoff's value by the probability. The sum of all possible discounted payoffs is called the option's *expected value*. It is assumed that the player should select the option that has the greatest expected value.

❖ BEHAVIORAL THEORY AND RESEARCH

The foregoing is the essence of formal, prescriptive decision theory; it is prescriptive because it is how a "rational" decision maker would make choices if he or she accepted the axioms and assumptions underlying the theory. About 50 years ago (some authorities trace things back 300 years, but 50 will do for us), psychologists became interested in decision behavior. Work by Ward Edwards (1954, 1955) set the

8 THE PSYCHOLOGY OF DECISION MAKING

paradigm for investigation: How well does the behavior of unaided decision makers compare to what a trained economist or statistician would do using prescriptive theory to make a specified decision?

We will examine a prototypical study in the Edwards tradition in a later chapter. However, a study by Gray (1975) provides a quaint, if atypical, illustration of behavioral research based on the presumption that decision makers' behavior is (or should be) reflected in prescriptive theory. Gray presented third-grade students with an array of six stacks of 10 cards, face down. Each stack contained arithmetic problems of roughly equal difficulty for the children, and the stacks ranged from very difficult at one end of the array to very easy at the other end. The student looked at a problem from each stack and stated how many of the 10 problems he or she would solve correctly if asked to work the problems in that stack. Dividing the estimate by 10 provides the student's assessment of his or her probability of success for each stack. The students' task was to select a stack from which a problem would be randomly selected for them to solve. They were told that if they selected the most difficult stack and got the correct answer, the payoff would be six red poker chips, the next most difficult paid five chips, and so on down to one chip for the easiest problems. If they got the wrong answer, they had to pay the experimenter the corresponding number of poker chips; they were loaned 10 chips to start out, which were deducted from their winnings at the end of the experiment.

The students' choices of problems to solve can be viewed as gambles—they could try a high-payoff difficult problem, but the odds were against getting the right answer. They could try a low-payoff easy problem, but the payoff was so low it hardly was worth the trouble. The students understood the dilemma and, in fact, selected the problems that prescriptive theory would dictate—the appropriate balance between the risk of failure and the magnitude of the payoff. It turned out that these third graders were reasonably efficient gamblers—with one exception. Like real gamblers, many of the students would play conservatively for a while and then, after having built up a supply of poker chips, would take a flyer on the really difficult stack. They almost invariably lost, whereupon they would select easy problems until they had regained their fortune; then they took another flyer.

The agenda that Edwards (1954, 1955) set was to derive a psychology of decision making using prescriptive theory as a starting point, mapping the points at which observed behavior deviated from prescription, and then modifying prescriptive theory to make it descriptive of

what decision makers actually do. In short, to transform predictive theory into descriptive theory and then, if necessary, move on to a broader psychology of decision making. In large part, this is exactly what happened, and the story of how this happened is the underlying theme of this book. We begin with prescriptive theory and end having reviewed two generations of descriptive theories with a third generation in the making. The first generation of descriptive theories closely resembles prescriptive theory in that it attempts to retain the general logic and much of the mathematical rigor in its formulation, and it retains the idea that decisions are analogous to gambles. First generation theory has been extensively researched and is broadly accepted as a valid description of human decision making by disciplines outside of psychology, primarily economics. However, from a psychological viewpoint, its resemblance to prescriptive theory is both a strength and a weakness. It gains strength from its legacy of centuries of thinking about decision making that lead up to prescriptive theory, but it is weakened by the constraints this legacy places on its breadth and flexibility. In an attempt to go beyond these constraints, researchers began to explore new ideas about decision making—ideas that bore little or no relation to predictive theory. The result was a second generation of research and theory that is still in the process of coalescing into a unitary body of work.

Second generation research and theory (also called naturalistic research and theory) has its origins in first generation work but grew beyond it in an attempt to more realistically describe how decisions actually are made, notably by professional decision makers such as managers. Second generation work abandons the gamble analogy and views decision making as a form of problem solving. It has been very much influenced by the need for practical knowledge about decision making for systems planning, management, and other real-world applications.

To a large degree, second generation research and theory is based on the work of Herbert Simon. In his classic *Administrative Behavior* (1945), Simon emphasizes that the behavior of a person in an organization is constrained by the position he or she holds in that organization—something so obviously true that it is easy to overlook both its insightfulness and its very broad implications. This means that decision making in organizations is strongly influenced by the structure and norms of the organization, and that decision makers do not entertain the full array of options that an outsider might consider available. Another feature of Simon's analysis is the emphasis on the individual as a decision making agent for the organization, emphasizing that organizations do

10 THE PSYCHOLOGY OF DECISION MAKING

not themselves make decisions. As pointed out by Davis (1992), after all these years this small but important point is still overlooked by many people who study organizational and group decision making.

Simon is a Nobel Laureate in economics who has produced a stream of useful insights that have inspired large amounts of research by many people. Two of his most famous insights, based on his observations of behavior in organizations, are *bounded rationality* and *satisficing*. Because decision makers' cognitive capacity is rather limited, they must reduce information processing demands by simplifying the problems they encounter. To do so, Simon contends, they construct "small worlds" that are limited representations of the problem at hand. The representation contains only the most salient information, and the decision maker proceeds to make his or her decision based solely upon that "bounded" representation. The decision may in fact be "rational" in that it conforms to the prescriptions of the appropriate prescriptive theory, but the decision maker uses only the information contained in the bounded representation.

Satisficing is another way of reducing the information processing load. This is not a "rational" decision strategy in the sense of following the prescriptions of prescriptive theory, but it allows the decision maker to arrive at a decision without all the computational effort required by prescriptive theory. Here it is assumed that the decision maker has some set of standards that an option must meet for it to be at least minimally satisfactory. The idea is that the first option that comes along that meets all of the standards is the one that is selected. Satisficing is not rational in the prescriptive sense because the decision maker has no assurance that on one or more standards an as-yet-unseen option might not be superior to the option that has been selected, and he or she therefore would fail to select the best possible option. However, the simplicity of this decision strategy presumably makes it worthwhile to risk missing the best option in favor of choosing one that is at least sufficient.

The second generation viewpoint is strongly influenced by decision making in organizations, particularly work-related decisions. Organizational theorists have tended to examine decisions at a higher, more macro, level than do most behavioral decision theorists. The former often treat the organization as the unit and talk as though it makes decisions, presuming something like an organizational mind. The latter focus on individuals, often without specifying the particular context,

presuming that an organization's decisions are a function of the collaborative decisions made by its individual members. Things often are confused, however, by a tendency for both organizational and behavioral theorists to switch back and forth between levels of analysis, which muddles things considerably.

Still, observations of individual decision makers in the context of their work provide invaluable information about how decision making takes place. The result has been a number of mini-theories—theories designed to cover some limited kind of decisions but that contribute to an emerging new view about how decision making takes place, particularly in organizational contexts. These mini-theories are beginning to coalesce into more general theories. It is this imminent coalescence that constitutes the second generation of behavioral research and theory.

Paradigms, Theories, and Models

It is useful to differentiate between paradigms, theories, and models. Although the terms often are used interchangeably, they are not the same thing. In fact they are hierarchically related: paradigms are the most general—rather like a philosophical or ideological framework. Theories are more specific, based on the paradigm and designed to describe what happens in one of the many realms of events encompassed by the paradigm. Models are even more specific, providing the mechanisms by which events occur in a particular part of the theory's realm. Of the three, models are most affected by empirical data—models come and go, theories only give way when evidence is overwhelmingly against them (and sometimes not even then), and paradigms stay put until a radically better idea comes along.

Prescriptive (Normative) Models

Now to tie all this about paradigms, theories, and models to the study of decision making. As we said above, prescriptive decision theory views the decision maker as a maximizer of expected value. First generation research translated this into the assumption that decision performance will be best if the decision makers' behavior reflects economics' utilitarian paradigm, which means that they strive to acquire desirable payoffs, and if their decision processes correspond to the processes dictated by utility theory and probability theory. That is,

12 THE PSYCHOLOGY OF DECISION MAKING

decision making, in the sense of acquisition of desired payoffs, will be most successful if the attractiveness of each option is summarized as the sum of the probability-discounted utilities corresponding to its potential payoffs, and if the decision maker chooses the option that offers the greatest sum. If performance falls short of what an economist could achieve by following these prescriptions, the fault must lie in the failure of the decision maker to behave according to the prescriptions. Therefore, by examining the non-correspondence between what the decision maker "should" do and what he or she actually does, first generation researchers and theorists strove to (1) better understand human decision processes and (2) help decision makers perform better. In first generation research, those aspects of utility theory and probability theory that pertain to the decision at hand are referred to as *normative models* because they are seen as defining the norms for correct processing of relevant information—not because it is a social norm to actually behave this way. First generation researchers interpret systematic differences between the prescriptions of the normative models and what decision makers actually do as reflections of people's cognitive limitations and their systematic processing errors. Techniques for helping decision makers overcome these limitations and avoid these errors are called *decision aids*.

In contrast, second generation research is influenced by observations of professional decision makers (managers) and attempts to discover the cognitive processes in which they engage while making various kinds of decisions. This viewpoint lacks the organizing framework provided by prescriptive theory, so it often looks (and is) less systematic than traditional research. At the moment, second generation research and theory consist of a collection of mini-theories, each focusing on some aspect of real-world decision making. There have been a few attempts to construct more encompassing theories, and research on them is presently under way.

Before moving on, it is important to note that by and large, the second generation viewpoint grew out of the first generation, which in turn grew out of the notion of using prescriptive theory as a foundation for behavioral research, so it is not entirely correct to think of the first and second generations as competing with one another. On the other hand, the second generation often has defined itself by how it differs from the first generation, which often makes it look like they all are in competition.

❖ THE PLAN

We begin with the basics of prescriptive theory (Chapters 2 through 4). Then we describe the work that led to “psychologizing” some of the central concepts of prescriptive theory and to first generation behavioral theory of decision making (Chapters 5 through 7). This is followed by an examination of work on interpersonal, organizational, and group decisions (Chapter 8) and descriptions of some interesting mini-theories (Chapter 9), with an eye to the implications for second generation behavioral theory. Then we describe a second generation theory that attempts to build on first generation theory and on subsequent research results to produce a descriptive psychological theory of decision making (Chapter 10). The book ends with a short recap, specification of needed research, and some predictions about the future (Chapter 11).

❖ SUMMARY

We have characterized decision making as a sequence of events: diagnosis, action selection, and implementation. We have introduced the idea of prescriptive theory and its gambling analogy. And we have discussed two generations in the construction of a psychology of decision making: the first generation, which adheres to the structure and spirit of prescriptive theory, and the second generation, which grows out of a broader appreciation of cognitive and organizational psychology. In the course of this we have described how parts of prescriptive theory were used as normative models by researchers, how differences between normative prescriptions and observed decision behavior were used to construct the first generation behavioral theories, and how dissatisfaction with the breadth of these theories prompted research leading to a second generation of behavioral theories of decision making.

