General Introduction

All students like to believe that their particular subject is the center of the universe. Doubtless, students of judgment and decision making are no different, but they may have a good argument for their view. After all, they can claim that the great moments of history all turned on someone's judgment as to what should be done and someone's decision to do it. Moreover, they will claim that although their subject is as old as civilization it has been studied in a scientific, empirical way only within the very last few decades. Indeed, most of the pioneers in this field are still alive and contributing to it. The fact that we are now able to study judgment and decision making in a scientific manner is, these students can claim, an exciting new discovery in and of itself.

Of course, the editors of this book and the authors of the chapters in it firmly believe in this view; judgment and decision making *are* of critical importance, and the fact that it is possible to study them in a scientific, empirical manner *is* a new and exciting event in the recent history of science.

Despite its central importance and long history, however, the field is still so new that it will be useful to turn to the dictionary to discover how these terms have been defined for common use. *Webster's Third New International Dictionary* says that *judgment* is "the mental or intellectual process of forming an opinion or evaluation by discerning and comparing," and the *capacity for judging* is "the power or ability to decide on the basis of evidence." Although the dictionary quotes E. L. Godkin as saying that "judgment is the highest of the human faculties," it also notes that Oliver Wendell Holmes said, "some of the sharpest men in argument are notoriously unsound in judgment." Apparently, we are to understand that the capacity to make sound judgments requires not only intelligence but wisdom and that the former does not guarantee the latter. *Webster's* definition of *decision*, "the act of settling or terminating . . . by giving judgment," suggests that there is little difference between *judgment* and *decision making* in ordinary discourse, so we shall not make a distinction here, although more advanced treatments of the topic do (see, e.g., Goldstein & Hogarth, 1997).

Not only the sources but the nature of sound judgment have fascinated scholars since the beginning of self-reflection; the Greek intellectuals apparently mused about those topics every day. And the discussion continues today among philosophers, psychologists, political scientists, lawyers, management scientists, and others inside and outside of academia, because sound judgment is, of course, of great practical as well as academic concern. No question will be of greater importance to the board of directors of the industrial firm that evaluates candidates for the position of chief executive officer than the soundness of each candidate's judgment. And while members of all the above academic departments debate theories of rational choice (each group in happy ignorance of the activities of the others), members of the board of regents will be exercising their judgment as they select the new president of the university. Indeed, the capacity for sound judgment of every person who desires a high (or even not so high) place in almost every segment of society will be judged by those responsible for selecting them. That is because within both government and industry there is a strong correlation between the prominence and power of one's position, the amount of time that one spends on problems requiring judgment, and the salary one receives. At the other end of the scale, inability to make the simple judgments required in the ordinary circumstances of day-to-day living leads to the diagnosis of mental illness.

In short, judgment and decision making are pervasive, important intellectual activities engaged in by all of us in academic, professional, and social pursuits throughout every day. The ability to form good judgments and make wise and effective decisions generally is considered the mark of a successful person in the smaller as well as the larger matters of living. Apparently, the same has been true of every human society.

What do we know about this salient feature of our lives? This book will not try to answer that question completely, but it will provide a general introduction to our knowledge of judgment and decision making and provide guideposts for those who may wish to pursue their inquiry further. Although the study of judgment and decision making is a field in its own right, it finds application in virtually every known human endeavor. (A recent survey conducted by one of us showed that articles related to judgment and decision making appeared in more than 500 different professional journals.) Therefore, we have chosen to group studies of judgment and decision making within those major fields in which studies of judgment and decision making are currently being conducted. These include judgment and social policy, economics, law, medicine and other fields indicated in the Table of Contents.

It is easy to find examples of the importance of the "intellectual process of forming an opinion or evaluation by discerning or comparing" or "decid[ing] on the basis of evidence." The decision to drop the atomic bomb on Hiroshima without warning is perhaps the most dramatic example of an act of judgment in the 20th century. Other examples include changes in health policy (e.g., the

decision to institute a National Health Service in Britain), economic policy (e.g., the deregulation of airlines in the United States), legal policy (e.g., the use of plea bargaining), and environmental policy (e.g., the protection of wildlife and pristine areas, the control of toxic waste), and the reduction of risk (e.g., the nationwide 55-mph speed limit in the United States); all provide examples of the attempt to exercise sound judgment. And in what follows we provide examples of efforts to study such judgments, both in the controlled conditions of the laboratory and in the world outside.

The reader will notice that all of these examples of studies of judgment and decision making are recent. The systematic empirical study of judgment and decision making began to emerge as a discipline in its own right only in the 1960s. This occurred together with a strong surge of interest in the larger, more general field of cognitive psychology, which includes the study of memory, thinking, problem solving, mental imagery, and language. The explosion of research in cognitive psychology marked a sharp shift in interest from the concentration on motivation in psychological research to a concentration on "mental activity." There are two main reasons for this. First, something dropped out; by 1960 strict stimulus – response behaviorism lost credibility among many laboratory scientists, and Freudian psychology based largely on unconscious motivation lost credibility with almost everyone. Second, something dropped in, namely, the electronic computer, which immediately provided a credible metaphor for mental activity. Thus, within a decade of the introduction of the computer, psychologists were talking about and studying "human information processing." As one psychologist (George Miller) put it, "the mind came in on the back of the machine." The arrival of the computer made it possible to carry out research on human information processing (including judgment and decision making) in new and powerful ways. For example, those interested in problem solving were able to build computer models that simulated human information-processing activity, and this led rapidly to the creation of the new field of artificial intelligence. Those interested in constructing mathematical models of the judgment and decision-making process could rapidly test a variety of such models for their ability to represent and/or evaluate the rationality of human judgment and decision making. By the 1980s, work that would have been utterly impossible prior to the computer became commonplace.

Because two types of mathematical representations of judgment and decision-making behavior are frequently used, we present the basic ideas that underlie them in this General Introduction. The mathematical operations of both approaches are easy to grasp; a knowledge of simple algebra is all that is required. We first describe the approach known as *decision analysis* – which involves an a priori decomposition of the decision process – and, second, we describe the approach known as *judgment analysis* – which involves an a posteriori decomposition of the judgment process. Although the distinction between decision and judgment is somewhat arbitrary and need

not concern us here, the distinction between a priori and a posteriori decomposition is important and should be kept in mind, for it will be illustrated often in the chapters that follow.

Decision Analysis: A Priori Decomposition

A priori decomposition refers to separating the decision process into its components *before* the decision is made. Such components include (a) the *probabilities* or likelihood of occurrence of each alternative considered and (b) the *utility* attached to each alternative. The decision process is greatly aided when these concepts are used in the context of a *decision tree*.

Construction of a decision tree prior to making the decision is an easy way of guiding and simplifying the decision process because it *diagrams* the decomposition of the decision process into probabilities and utilities and thus provides a clear picture of the process and its components.

The decision maker needs only four types of information to construct a decision tree:

1. What are my possible courses of action? (Alternatives)
2. What are the events that might follow from those actions? (Outcomes)
3. What is the likelihood of each event?
4. What is the value of each event to me?

Here is a decision-making situation similar to one that actually confronted an elderly man known to us. The man had a very serious medical problem. His physician said that a difficult operation was necessary to remedy the situation. The physician added, however, that, given the man's very advanced age and the nature of the operation, there was a 40% *probability* that the patient would not recover from the operation. If the operation were not performed, the serious medical problem would linger, causing the patient discomfort and impairing his mobility. There was no chance that the problem would "go away," and there was a 20% probability that without the operation the man would die within the next 6 months. What should the man do? Should he have the operation?

Figure I.1 depicts the decision tree for this situation. The box represents a "decision node." The two possible courses of action emanate from this node. They are "operate" and "don't operate," and thus they comprise the first of the four types of information needed to construct a decision tree. The circles are chance nodes. (They are called chance nodes because no decision can be made to cause one of the outcomes to occur rather than the other. Their occurrence is therefore left to chance.) The events emanating from these circles are the possible events that might occur following the courses of action. This is the second type of information needed. Preceding each

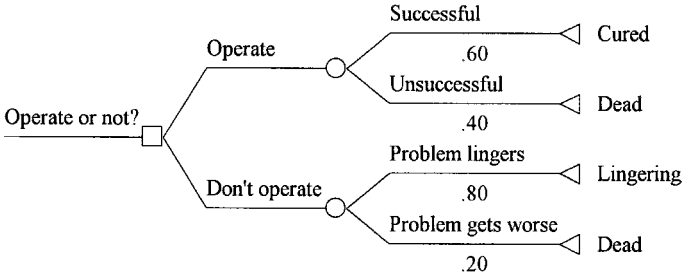

Figure I.1. Decision tree showing probabilities and outcomes associated with alternative actions.

possible event is its probability of occurrence, based on the physician's best estimate. This is the third essential type of information. Finally, we need to know what value the patient places on each of these outcomes.

Because the value of any commodity is judged differently by everyone, the term *utility* rather than *value* is used. This term captures the subjective nature of the evaluation; a particular amount of money may have different utility for me than for you. Even health may have different utilities for different people. To calculate the utility of each outcome, we shall call the worst outcome 0.0 on a utility scale and the best outcome 1.0.

The patient assigned "death" the former value and "complete cure" the latter value. Using this range (0.0–1.0), the patient felt that living in discomfort and having decreased mobility was worth .6 to him. It was a state closer to "complete cure" than to "death," but not by much.

We now have all the information needed to make a decision tree. First, it is necessary to examine each outcome. The utility of each outcome needs to be weighted according to its likelihood. An outcome of 1.0 ("wonderful") that has a high probability of occurring should definitely be preferred to another outcome of utility 0.0 that has very little likelihood of occurring. To accomplish this mathematically, the utility of each outcome is multiplied by its probability of occurrence. This product is the *expected utility* of each outcome. Figure I.2 contains these calculations, which are located at the right edge of each branch of the decision tree.

All that remains now is the process called "folding back," which consists of pruning all but the most preferred course of action at each decision node. There are two possible courses of action: operate and don't operate. For each of these two options we add together all of the expected utilities associated with that option. For example, the option "don't operate" has associated with it two expected utilities of .48 and 0. Their sum, .48, is the expected utility of the "don't operate" course of action. Because the "operate" course of action is higher (.6), the preferred course of action is to operate.

One immense virtue of a decision tree is that it is a wonderfully general decision aid. As long as the four types of information are available, any decision

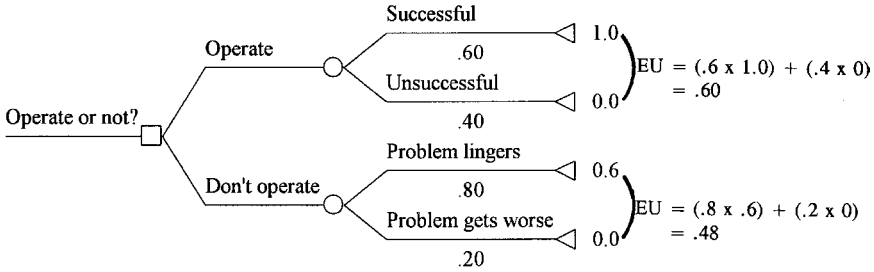

Figure I.2. Calculation of expected utility for each alternative.

can be analyzed by use of the tree. Should I move to Minneapolis to accept this new job? Should I switch insurance policies from X to Y? Should we get a loan to get a new car now, or should we limp along with the one we have?

One difficulty often encountered in constructing a decision tree is that likelihoods and utilities are often not easy to assess. When an expert opinion is available, as in the case of our medical example, reasonable likelihoods can often be provided. Because the helpfulness of a decision tree is based largely on the accuracy of the likelihoods and utilities used, every effort should be made to obtain good estimates.

Occasionally the decision maker is uneasy about the "verdict" of the decision tree. In our medical example, the "operate" option was only .12 superior to the "don't operate" option. "What if I later decide that living with discomfort and decreased mobility isn't so bad? Maybe it's worth a .7 and not a measly .6," thinks the elderly man as the morning of the operation approaches. A quick calculation will reveal that "operate" is *still* the preferred choice, even if .7 is deemed the utility of an uncomfortable and immobile existence. Modifying the probabilities and utilities in this way is called a "sensitivity analysis," because such manipulating of the probabilities and utilities tests how sensitive the final choice is to the numbers initially assigned. Reasonable modifications of the probabilities and utilities often leave the decision unchanged. The decision maker can then rest comfortably with the decision that has been reached.

Concern about the assignment of accurate probabilities, the calculation of expected utilities, and the performance of various arithmetic tasks should not obscure what may be the greatest virtue of a decision tree: It forces the decision maker to make explicit all the bases for the decision. In the tree are contained all the courses of action, all the probabilities, all the utilities, and all the outcomes of which the decision maker is aware – or, at least, those that he or she plans to consider. Every analysis works on a simplified version of the real situation, but the tree at least makes it explicit what is being included, what excluded. It also makes it explicit how far into the future the decision maker is thinking about the consequences of this action. The time-frame issue is important when the utility of an outcome shifts over time. In exercising, for

example, a short time-frame stresses negative outcomes like feeling tired, but a longer time-frame stresses more positive ones like good health and vigor. Merely having to generate the information necessary to draw the tree may force the decision maker to confront the situation in a much more organized and thoughtful way than would otherwise be the case.

Judgment Analysis: A Posteriori Decomposition

If a priori decomposition implies decomposing the decision process *prior* to its occurrence, then a posteriori decomposition obviously implies that decomposition will take place *after* a series of judgments have been made. As we shall see, a person's *judgment policy* can be "captured" after judgments are made regarding hypothetical cases; the policy may then be applied to real cases.

The principal concepts of a judgment analysis are best illustrated by reference to the model of the judgment situation presented in Figure I.3, which indicates that judgment is a cognitive process similar to inductive inference. That is, judgment is a cognitive or intellectual process in which a person draws a conclusion, or an inference (Y_s), about something (Y_e), which *cannot* be seen, on the basis of data (X_i), which *can* be seen. In other words, judgments are made from *tangible* data, which serve as *cues* to *intangible* events and circumstances. The wide-ranging arc connecting Y_s and Y_e (labeled r_a in

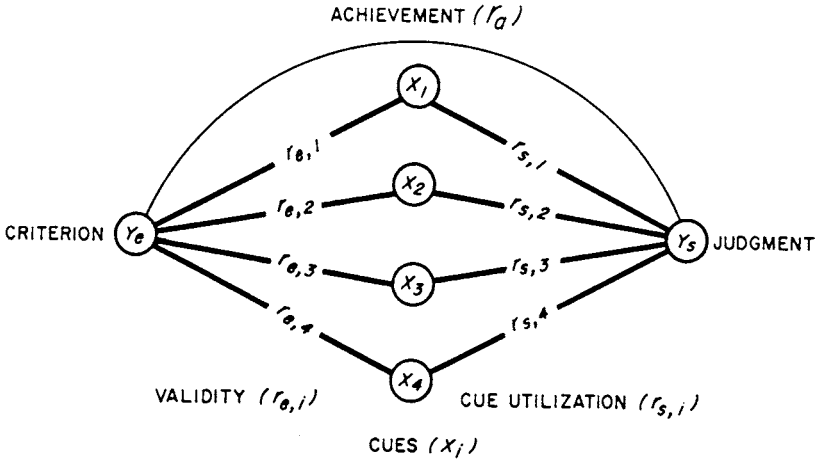

Figure I.3. The "Lens Model," a schematic illustration of how a subject froms a judgment (Y_s) of some criterion variable (Y_e) on the basis of a set of imperfect cues, X_i. Cue validity ($r_{e,i}$) measures the extent to which each cue reflects the value of the criterion; cue utilization ($r_{s,i}$) measures the relation between each cue and the subject's judgment; and achievement, r_a, measures how well the subject's judgments correspond to the actual criterion values.

Figure I.3) indicates the degree to which the judgment Y_s was correct, that is, the extent to which the judgment coincides with the actual circumstance to be judged. A rough example can be found in the judgments of the weather forecaster who looks at certain tangible cues (X_i) such as wind speed, temperature, and barometric pressure and makes a judgment (Y_s) about what tomorrow's weather (Y_e) will be. The arc, r_a, indicates the degree of accuracy over a series of judgments.

Throughout any ordinary day one frequently encounters similar circumstances. Tangible data (e.g., events in the news, activities of the stock market, actions of friends and neighbors) evoke judgments as to the unperceived events that gave rise to the events perceived. *Causes* (Y_e) are frequently being inferred (Y_s) from those cue events (X_i), or *effects*, that are being observed. And the ability to make correct inferences (indicated by r_a) is, of course, an ability in which persons are believed to differ widely. High judgmental accuracy is considered to be an essential attribute of persons with high responsibility; low accuracy indicates persons very likely to be in difficulty with their social or physical surroundings.

The model in Figure I.3 also indicates the concept of *differential weight*. Cues may have differential weight in that they are of differential value in making inferences about events. That is, if a cue has a very strong relation (a high degree of covariation) with an event to be inferred, it will be more useful than one that has a weak relation. Therefore, cues with high degrees of covariation with the event to be inferred have a large degree of *ecological validity*; their weight is greater than those with low degrees of covariation.

The counterpart to the ecological validity (r_e) of a cue is its utilization (r_s) by the subject (see Figure I.3). Cues also may be used or depended upon to a larger or smaller degree, therefore, with regard to their *subjective utilization*. Thus an observer may compare the differential weights of a set of cues ($r_{e,i}$) in the task with the weights implicitly assigned to them by the person making the inference. Mismatches between ecological validities and subjective utilization of cues are one source of inaccurate judgments. In other words, one source of poor judgment lies in the failure to attach the correct relative weights or importance to cues.

Not only do cues have different task weights, but they may be related to the variable to be inferred (Y_e) by means of different functional relations, or *function forms*. These may include positive linear function forms, negative linear function forms, or a variety of curvilinear function forms (see Figure I.4). Of course, cues may be related to judgment (Y_s) by means of various function forms also, and the comparison, or match, between task function form and subjective function form will also form the basis for accurate or inaccurate judgments.

Tasks that involve curvilinear functions are apt to be more difficult to learn than those with positive linear function forms, and people's judgments

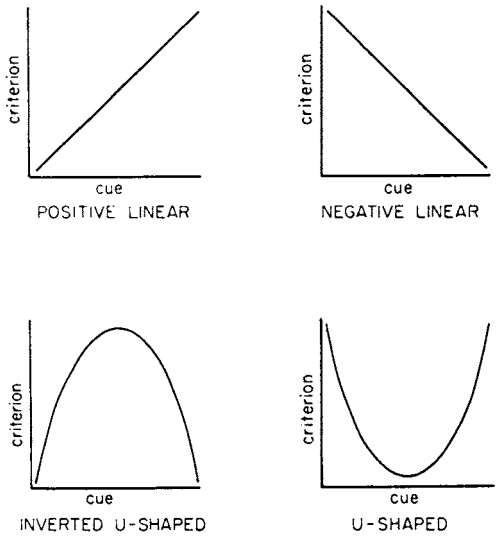

Figure I.4. Four commonly observed function forms between cues and a person's judgment.

related to cues in curvilinear forms are apt to be more difficult for other people to understand. Moreover, people seldom make explicit to others exactly which function form they are employing in a given situation, simply because most people are not aware of this concept. It may appear in colloquial form, however. For example, one way to indicate that another person is using a positive linear function form when he or she should be using a curvilinear one is to say that the person's error lies in believing that "if a little is good, more is better" – a rule that, if followed in the taking of medicine, is apt to lead to disaster. The reader will find it easy to think of other examples of an inappropriate use of linear function forms. One should also consider the difficulties of *changing* a function form to fit the function form of a task, or to fit one preferred by a friend, a teacher, or a therapist.

The principles by which the cue data are organized into a judgment are of considerable importance. Such data may be organized by adding them, $Y_s = X_1 + X_2 + X_3$; by averaging them, $Y_s = (X_1 + X_2 + X_3)/3$; or by making use of some configural or patterning principles, $Y_s = X_1 + X_2 X_3$.

When asked about how they organize information into a judgment, most persons are apt to report that they make use of a pattern or configuration of the data. Physicians who make a diagnosis, experts in investment, and others whose professional judgment is of great importance generally reply to questions about their judgment processes by referring more or less vaguely to their intuitive ability to recognize "patterns." Empirical research, however,

in general has not supported these contentions; simpler organizing principles have been found to account for, or at least to predict, judgments from data better than patterns. Although it hardly seems doubtful that human beings *can* organize data by means of patterns, the extent to which they do in fact is unknown; in any event, reports of the use of such principles certainly cannot be taken for granted.

Finally, it is important to consider the *consistency* with which the same judgment is made in response to the same data. Although everyone is apt to assume that they always make the same judgments when confronted with the same facts, it is virtually certain that they will not do so except under the very simplest circumstances. That is, perfect consistency in judgments is apt to occur only when there is no uncertainty whatever in the task situation. Such simple task situations, of course, require little in the way of judgment, inasmuch as a given cue always evokes the same judgment.

The simplest and best way to discover the cues, differential weights, function forms, and consistency of a person's judgment process is to use a computer to present a number of cases to the person making the judgment. After the judgments have been made, a computer program can readily decompose the judgment process into weights, function forms, and consistency. Because this information is extracted *after* a series of judgments have been made, the decomposition is obviously a posteriori. But it is important to observe that, because the judgments can be made with regard to *hypothetical* cases, the person's *judgment policy* (consisting of specific weights and function forms) can then be applied to a real case or a series of real cases. In short, even though the method extracts the various components of the judgment policy after the judgments have been made, the policy may be applied to new cases or to any one judgment problem, just as in the situation where a priori decomposition takes place.

Decision analysis and judgment analysis start at opposite ends of the spectrum of description versus prescription (or advice-giving), but each expands toward the other pole. If one starts with a decision tree, it isn't long before one starts to wonder where these alternatives came from, how the decision maker assessed these probabilities and utilities, chose a time frame, and so on – all interesting descriptive issues. Conversely, if one starts with a descriptive study of, say, a physician attempting to make a diagnosis, it isn't long before one starts to wonder why these cues are being used rather than these others, and how consistently they are being used, and how accurate the diagnoses turn out to be on average, and whether alternative ways of using the available information might be more accurate – all interesting prescriptive issues. Such interplay of description and prescription is deeply woven into JDM research and gives it much of its special interest and value. The central phenomenon, after all, is human efforts to choose actions with some purpose in mind. We are interested in judgment because the thing being judged is of