MARCH 1994
Pp. 1-23

CHAPTER ONE

Limited Rationality

By far the most common portrayal of decision making is one that interprets action as rational choice. The idea is as old as thought about human behavior, and its durability attests not only to its usefulness but also to its consistency with human aspirations. Theories of rational choice, although often elaborated in formal and mathematical ways, draw on everyday language used in understanding and communicating about choices. In fact, the embedding of formal theories of rationality in ordinary language is one of their distinctive features. Among other things, it makes them deceptively comprehensible and self-evident. This chapter examines the idea of rational choice and some ways in which theories of limited rationality have made that idea more consistent with observations of how decisions actually happen.

1.1 The Idea of Rational Choice

Like many other commonly used words, "rationality" has come to mean many things. In many of its uses, "rational" is approximately equivalent to "intelligent" or "successful." It is used to

describe actions that have desirable outcomes. In other uses, "rational" means "coldly materialistic," referring to the spirit or values in terms of which an action is taken. In still other uses, "rational" means "sane," reflecting a judgment about the mental health displayed by an action or a procedure for taking action. Heterogeneous meanings of rationality are also characteristic of the literature on decision making. The term is used rather loosely or inconsistently.

In this book, "rationality" has a narrow and fairly precise meaning linked to processes of choice. Rationality is defined as a particular and very familiar class of procedures for making choices. In this procedural meaning of "rational," a rational procedure may or may not lead to good outcomes. The possibility of a link between the rationality of a process (sometimes called "procedural rationality") and the intelligence of its outcomes (sometimes called "substantive rationality") is treated as a result to be demonstrated rather than an axiom.

1.1.1 The Logic of Consequence

Rational theories of choice assume decision processes that are consequential and preference-based. They are *consequential* in the sense that action depends on anticipations of the future effects of current actions. Alternatives are interpreted in terms of their expected consequences. They are *preference-based* in the sense that consequences are evaluated in terms of personal preferences. Alternatives are compared in terms of the extent to which their expected consequences are thought to serve the preferences of the decision maker.

A rational procedure is one that pursues a logic of consequence. It makes a choice conditional on the answers to four basic questions:

1. The question of *alternatives:* What actions are possible?
2. The question of *expectations:* What future consequences might follow from each alternative? How likely is each possible consequence, assuming that alternative is chosen?
3. The question of *preferences:* How valuable (to the decision maker) are the consequences associated with each of the alternatives?
4. The question of the *decision rule:* How is a choice to be made among the alternatives in terms of the values of their consequences?

When decision making is studied within this framework, each of these questions is explored: What determines which alternatives are considered? What determines the expectations about consequences? How are decision maker preferences created and evoked? What is the decision rule that is used?

This general framework is the basis for standard explanations of behavior. When asked to explain behavior, most people "rationalize" it. That is, they explain their own actions in terms of their alternatives and the consequences of those alternatives for their preferences. Similarly, they explain the actions of others by imagining a set of expectations and preferences that would make the action rational.

A rational framework is also endemic to theories of human behavior. It is used to understand the actions of firms, marriage partners, and criminals. It underlies many theories of bargaining, exchange, and voting, as well as theories of language and social structure. Rational choice processes are the fundamentals of microeconomic models of resource allocation, political theories of coalition formation, statistical decision theories, and many other theories and models throughout the social sciences.

1.1.2 Rational Theories of Choice

Within rational processes, choice depends on what alternatives are considered and on two guesses about the future: The first guess is a guess about future states of the world, conditional on the choice. The second guess is a guess about how the decision maker will feel about that future world when it is experienced.

PURE THEORIES OF RATIONAL CHOICE

Some versions of rational choice theory assume that all decision makers share a common set of (basic) preferences, that alterna-

tives and their consequences are defined by the environment, and that decision makers have perfect knowledge of those alternatives and their consequences. Other versions recognize greater inter-actor subjectivity but nevertheless assume perfect knowledge for any particular decision—that all alternatives are known, that all consequences of all alternatives are known with certainty, and that all preferences relevant to the choice are known, precise, consistent, and stable.

These pure versions of rational choice have well-established positions in the prediction of aggregate behavior, where they are sometimes able to capture a rational "signal" within the subjective "noise" of individual choice. They are sources of predictions of considerable generality, for example the prediction that an increase in price will lead (usually) to an aggregate decrease in demand (although some individuals may be willing to buy more at a higher price than at a lower one).

In spite of their utility for these qualitative aggregate predictions, pure versions of rational choice are hard to accept as credible portraits of actual individual or organizational actors. Consider the problem of assigning people to jobs in an organization. If it were to satisfy the expectations of pure rationality, this decision would start by specifying an array of tasks to be performed and characterizing each by the skills and knowledge required to perform them, taking into account the effects of their interrelationships. The decision maker would consider all possible individuals, characterized by relevant attributes (their skills, attitudes, and price). Finally the decision maker would consider each possible assignment of individuals to tasks, evaluating each possible array of assignments with respect to the preferences of the organization.

Preferences would be defined to include such things as (1) profits, sales, and stock value (tomorrow, next year, and ten years from now); (2) contributions to social policy goals (e.g. affirmative action, quality of life goals, and the impact of the assignment on the family); and (3) contributions to the reputation of the organization among all possible stakeholders—shareholders, potential shareholders, the employees themselves, customers, and citizens in the community. The tradeoffs among these various objectives would have to be known and specified in advance, and all possible task definitions, all possible sets of employees, and all possible assignments of people to jobs would have to be considered. In the end, the decision maker would be expected to choose the one combination that maximizes expected return.

A considerably less glorious version of rationality—but still heroic—would assume that a structure of tasks and a wage structure are given, and that the decision maker assigns persons to jobs in a way that maximizes the return to the organization. Another version would assume that a decision maker calculates the benefits to be obtained by gathering any of these kinds of data, and their costs.

Virtually no one believes that anything approximating such a procedure is observed in any individual or organization, either for the job assignment task or for any number of other decision tasks that confront them. Although some people have speculated that competition forces the outcomes of actual decision processes to converge to the outcomes predicted from a purely rational process, even that speculation has been found to be severely restricted in its applicability. Pure rationality strains credulity as a description of how decisions actually happen. As a result, there have been numerous efforts to modify theories of rational choice, keeping the basic structure but revising the key assumptions to reflect observed behavior more adequately.

RATIONAL DECISION MAKING AND UNCERTAINTY ABOUT CONSEQUENCES

The most common and best-established elaboration of pure theories of rational choice is one that recognizes the uncertainty surrounding future consequences of present action. Decision makers are assumed to choose among alternatives on the basis of their expected consequences, but those consequences are not known with certainty. Rather, decision makers know the likelihoods of various possible outcomes, conditional on the actions taken.

Uncertainty may be imagined to exist either because some processes are uncertain at their most fundamental levels or be-

cause decision makers' ignorance about the mechanisms driving the process make outcomes look uncertain to them. The food vendor at a football game, for example, knows that the return from various alternative food-stocking strategies depends on the weather, something that cannot be predicted with certainty at the time a decision must be made.

Since a decision maker does not know with certainty what will happen if a particular action is chosen, it is unlikely that the results of an action will confirm expectations about it. Postdecision surprise, sometimes pleasant sometimes unpleasant, is characteristic of decision making. So also is postdecision regret. It is almost certain that after the consequences are known (no matter how favorable they are) a decision maker will suffer regret—awareness that a better choice could have been made if the outcomes could have been predicted precisely in advance. In such a spirit, investors occasionally rue the gains they could have realized in the stock market with perfect foresight of the market.

The most commonly considered situations involving uncertainty are those of decision making under "risk," where the precise consequences are uncertain but their probabilities are known. In such situations, the most conventional approach to predicting decision making is to assume a decision maker will choose the alternative that maximizes expected value, that is, the alternative that would, on average, produce the best outcome if this particular choice were to be made many times. The analog is gambling and the choice of the best gamble. An expected-value analysis of choice involves imagining a decision tree in which each branch represents either a choice to be made or an "act of nature" that cannot be predicted with certainty. Procedures for constructing and analyzing such trees constitute a large fraction of modern decision science.

In more elaborate rational theories of choice in the face of risk, an alternative is assessed not only by its expected value but also by its uncertainty. The value attached to a potential alternative depends not only on the average expected return but also on the degree of uncertainty, or risk, involved. For risk-averse

decision makers, riskiness decreases the value of a particular alternative. For risk-seeking decision makers, riskiness increases the value.

The riskiness of an alternative is defined in different ways in different theories, but most definitions are intended to reflect a measure of the variation in potential outcomes. This variation has a natural intuitive measure in the variance of the probability distribution over outcome values. For various technical reasons, such a measure is not always used in studies of choice, but for our purposes it will suffice. When risk is taken into account, a decision is seen as a joint function of the expected value (or mean) and the riskiness (or variance) of the probability distribution over outcomes conditional on choice of a particular alternative.

MODIFYING THE ASSUMPTIONS

The introduction of risk and the development of ways to deal with it were major contributions to understanding and improving decision making within a rational framework. Such developments were, however, just the first step in modifying the knowledge assumptions of rational choice. Most modern theories of rational choice involve additional modifications of the pure theory. They can be distinguished by their assumptions with respect to four dimensions:

1. *Knowledge:* What is assumed about the information decision makers have about the state of the world and about other actors?
2. *Actors:* What is assumed about the number of decision makers?
3. *Preferences:* What is assumed about the preferences by which consequences (and therefore alternatives) are evaluated?
4. *Decision rule:* What is assumed to be the decision rule by which decision makers choose an alternative?

Although most theories "relax" the assumptions of the pure theory on at least one of these dimensions, they tend to be con-

servative in their deviations from the assumptions underlying a pure conception of rationality. For example, most theories of limited knowledge are not simultaneously theories of multiple actors; most theories of multiple actors (for example, microeconomic versions of game theory) are not simultaneously theories of limited knowledge; and virtually none of the limited knowledge or multiple-actor theories introduce conceptions of ambiguous or unstable preferences. In that sense at least, the pure model still permeates the field—by providing an overall structure and significant (though different) parts for various different theories.

1.1.3 Enthusiasts and Skeptics

Enthusiasts for rational models of decision making notice the widespread use of assumptions of rationality and the successes of such models in predictions of aggregates of human actors. They easily see these symptoms of acceptance and usefulness as impressive support for the models. Skeptics, on the other hand, are less inclined to give credence to models based on their popularity, noting the historical fact that many currently rejected theories have enjoyed long periods of popularity. They are also less inclined to find the models particularly powerful, often emphasizing their less than perfect success in predicting individual behavior. They easily see these symptoms of conventionality and imperfection as making the models unattractive.

Both enthusiasts and skeptics endorse limited rationality, the former seeing limited rationality as a modest, natural extension of theories of pure rationality, and the latter seeing limited rationality as a fundamental challenge to pure rationality and a harbinger of much more behaviorally based conceptions of decision making.

1.2 Limited (or Bounded) Rationality

Studies of decision making in the real world suggest that not all alternatives are known, that not all consequences are considered, and that not all preferences are evoked at the same time. Instead of considering all alternatives, decision makers typically

appear to consider only a few and to look at them sequentially rather than simultaneously. Decision makers do not consider all consequences of their alternatives. They focus on some and ignore others. Relevant information about consequences is not sought, and available information is often not used. Instead of having a complete, consistent set of preferences, decision makers seem to have incomplete and inconsistent goals, not all of which are considered at the same time. The decision rules used by real decision makers seem to differ from the ones imagined by decision theory. Instead of considering "expected values" or "risk" as those terms are used in decision theory, they invent other criteria. Instead of calculating the "best possible" action, they search for an action that is "good enough."

As a result of such observations, doubts about the empirical validity and usefulness of the pure theory of rational choice have been characteristic of students of actual decision processes for many years. Rational choice theories have adapted to such observations gradually by introducing the idea that rationality is limited. The core notion of limited rationality is that individuals are intendedly rational. Although decision makers try to be rational, they are constrained by limited cognitive capabilities and incomplete information, and thus their actions may be less than completely rational in spite of their best intentions and efforts.

In recent years, ideas of limited (or bounded) rationality have become sufficiently integrated into conventional theories of rational choice to make limited rationality viewpoints generally accepted. They have come to dominate most theories of individual decision making. They have been used to develop behavioral and evolutionary theories of the firm. They have been used as part of the basis for theories of transaction cost economics and game theoretic, information, and organizational economics. They have been applied to decision making in political, educational, and military contexts.

1.2.1 Information Constraints

Decision makers face serious limitations in attention, memory, comprehension, and communication. Most students of individ-

ual decision making seem to allude to some more or less obvious biological constraints on human information processing, although the limits are rarely argued from a strict biological basis. In a similar way, students of organizational decision making assume some more or less obvious information constraints imposed by methods of organizing diverse individuals:

1. *Problems of attention.* Time and capabilities for attention are limited. Not everything can be attended to at once. Too many signals are received. Too many things are relevant to a decision. Because of those limitations, theories of decision making are often better described as theories of attention or search than as theories of choice. They are concerned with the way in which scarce attention is allocated.

2. *Problems of memory.* The capabilities of individuals and organizations to store information is limited. Memories are faulty. Records are not kept. Histories are not recorded. Even more limited are individual and organizational abilities to retrieve information that has been stored. Previously learned lessons are not reliably retrieved at appropriate times. Knowledge stored in one part of an organization cannot be used easily by another part.

3. *Problems of comprehension.* Decision makers have limited capacities for comprehension. They have difficulty organizing, summarizing, and using information to form inferences about the causal connections of events and about relevant features of the world. They often have relevant information but fail to see its relevance. They make unwarranted inferences from information, or fail to connect different parts of the information available to them to form a coherent interpretation.

4. *Problems of communication.* There are limited capacities for communicating information, for sharing complex and specialized information. Division of labor facilitates mobilization and utilization of specialized talents, but it also encourages differentiation of knowledge, competence, and language. It is difficult to communicate across cultures, across generations, or across professional specialties. Different groups of people use different frameworks for simplifying the world.

As decision makers struggle with these limitations, they develop procedures that maintain the basic framework of rational choice but modify it to accommodate the difficulties. Those procedures form the core of theories of limited rationality.

1.2.2 Coping with Information Constraints.

Decision makers use various information and decision strategies to cope with limitations in information and information-handling capabilities. Much of contemporary research on choice by individuals and organizations focuses on those coping strategies, the ways choices are made on the basis of expectations about the future but without the kind of complete information that is presumed in classical theories of rational choice.

THE PSYCHOLOGY OF LIMITED RATIONALITY

Psychological studies of individual decision making have identified numerous ways in which decision makers react to cognitive constraints. They use stereotypes in order to infer unobservables from observables. They form typologies of attitudes (liberal, conservative) and traits (dependent, extroverted, friendly) and categorize people in terms of the typologies. They attribute intent from observing behavior or the consequences of behavior. They abstract "central" parts of a problem and ignore other parts. They adopt understandings of the world in the form of socially developed theories, scripts, and schemas that fill in missing information and suppress discrepancies in their understandings.

The understandings adopted tend to stabilize interpretations of the world. For the most part, the world is interpreted and understood today in the way it was interpreted and understood yesterday. Decision makers look for information, but they see what they expect to see and overlook unexpected things. Their memories are less recollections of history than constructions based on what they thought might happen and reconstructions based on what they now think must have happened, given their present beliefs.

A comprehensive review of psychological studies of individual information processing and problem solving would require more space and more talent than are available here. The present intention is only to characterize briefly a few of the principal speculations developed as a result of thatresearch, in particular speculations about four fundamental simplification processes: editing, decomposition, heuristics, and framing.

Editing. Decision makers tend to edit and simplify problems before entering into a choice process, using a relatively small number of cues and combining them in a simple manner. Complex problems or situations are simplified. Search may be simplified by discarding some available information or by reducing the amount of processing done on the information. For example, decision makers may attend to choice dimensions sequentially, eliminating all alternatives that are not up to standards on the first dimension before considering information from other dimensions. In other situations, they may consider all information for all alternatives, but weight the dimensions equally rather than weight them according to their importance.

Decomposition. Decision makers attempt to decompose problems, to reduce large problems into their component parts. The presumption is that problem elements can be defined in such a way that solving the various components of a problem individually will result in an acceptable solution to the global problem. For example, a decision maker might approach the problem of allocating resources to advertising projects by first decomposing the global advertising problem of a firm into subproblems associated with each of the products, then decomposing the product subproblems into problems associated with particular geographic regions.

One form of decomposition is working backward. Some problems are easier to solve backward than forward because, like mazes, they have only a few last steps but many first steps. Working backward is particularly attractive to decision makers who accept a "can do" decision making ideology, because it matches an activist role. Working backward encourages a perspective in which decision makers decide what they want to have happen and try to make it happen.

Decomposition is closely connected to such key components of organizing as division of labor, specialization, decentralization, and hierarchy. An important reason for the effectiveness of modern organization is the possibility of decomposing large complex tasks into small independently manageable ones. In order for decomposition to work as a problem solving strategy, the problem world must not be tightly interconnected. For example, if actions taken on one advertising project heavily affect the results of action on others, deciding on the projects independently will produce complications. The generality of decomposition strategies suggests that the world is, in fact, often only loosely interconnected, so subproblems can be solved independently. But that very generality makes it likely that decomposition will also be attempted in situations in which it does not work.

Heuristics. Decision makers recognize patterns in the situations they face and apply rules of appropriate behavior to those situations. Studies of expertise, for example, generally reveal that experts substitute recognition of familiar situations and rule following for calculation. Good chess players generally do more subtle calculations than novices, but their great advantage lies less in the depth of their analysis than in their ability to recognize a variety of situations and in their store of appropriate rules associated with situations. Although the problem solving of expert salespersons has been subjected to less research, it appears to be similar.

As another example, people seem not to be proficient at calculating the probability of future events by listing an elaborate decision tree of possible outcomes. However, they are reasonably good at using the output of memory to tell them how frequently similar events have occurred in the past. They use the results of memory as a proxy for the projection of future probability.

Such procedures are known to the literature of problem solving and decision making as "heuristics." Heuristics are rules-of-

thumb for calculating certain kinds of numbers or solving certain kinds of problems. Although psychological heuristics for problem solving are normally folded into a discussion of limited rationality because they can be interpreted as responses to cognitive limitations, they might as easily be interpreted as versions of rule-following behavior that follows a logic quite different from a logic of consequence (see Chapter 2).

Framing. Decisions are framed by beliefs that define the problem to be addressed, the information that must be collected, and the dimensions that must be evaluated. Decision makers adopt paradigms to tell themselves what perspective to take on a problem, what questions should be asked, and what technologies should be used to ask the questions. Such frames focus attention and simplify analysis. They direct attention to different options and different preferences. A decision will be made in one way if it is framed as a problem of maintaining profits and in a different way if it is framed as a problem of maintaining market share. A situation will lead to different decisions if it is seen as being about "the value of innovation" rather than "the importance of not losing face."

Decision makers typically frame problems narrowly rather than broadly. They decide about local options and local preferences, without considering all tradeoffs or all alternatives. They are normally content to find a set of sufficient conditions for solving a problem, not the most efficient set of conditions. Assigning proper weights to things in the spatial, temporal, and causal neighborhood of current activity as opposed to things that are more distant spatially, temporally, or causally is a major problem in assuring decision intelligence (see Chapter 6). It is reflected in the tension between the frames of decision makers, who often seem to have relatively short horizons, and the frames of historians, who (at least retrospectively) often have somewhat longer horizons.

The frames used by decision makers are part of their conscious and unconscious repertoires. In part they are encased in early individual experiences that shape individual approaches to problems. In part they are responsive to the particular sequences of decision situations that arise. There is a tendency for frames to persist over a sequence of situations. Recently used frames hold a privileged position, in part because they are more or less automatically evoked in a subsequent situation. In addition, past attention strengthens both a decision maker's skills in using a frame and the ease of justifying action to others within the frame.

These internal processes of developing frames and using them is supplemented by an active market in frames. Decision makers adopt frames that are proposed by consultants, writers, or friends. They copy frames used by others, particularly others in the same profession, association, or organization. Consequential decision making itself is, of course, one such frame. Prescriptive theories of decision making seek to legitimize a consequential frame for considering decisions, one that asks what the alternatives are, what their expected consequences are, and what the decision maker's preferences are.

THE STATISTICS OF LIMITED RATIONALITY

Faced with a world more complicated than they can hope to understand, decision makers develop ways of monitoring and comprehending that complexity. One standard approach is to deal with summary numerical representations of reality, for example income statements and cost-of-living indexes. The numbers are intended to represent phenomena in an organization or its environment: accounting profits, aptitude scores, occupancy rates, costs of production. The phenomena themselves are elusive—real but difficult to characterize and measure. For example, income statements confront a number of uncertainties. How quickly do resources lose their value (depreciate or spoil)? How should joint costs be allocated to various users? How should inventory be counted and valued? How can the quality of debts be assessed? What is the value of a contract? Of a good name? There is ambiguity about the facts and much potential for conflict over them. As a result, the numbers are easily described as inventions, subject to both debate and ridicule. They have elements of magic about them, pulled mysteriously

from a statistician's or a manager's hat. For example, estimates of U.S. government subsidies to nuclear power went from $40 billion under one administration to $12.8 billion under another with no change in actual programs.

The numbers are magical, but they also become quite real. Numbers such as those involved in a cost-of-living index or an income (profit and loss) statement come to be treated as though they were the things they represent. If the cost-of-living index goes down, decision makers act as though the cost of living has gone down—even though they are well aware of the many ways in which, for many people, the cost of living may actually have gone up. Indeed, the whole concept of "cost of living" moves from being an abstract hypothetical figure to being a tangible reality.

Three main types of such numbers can be distinguished:

1. Representations of *external reality* are numbers purporting to describe the environment in which decision makers exist. Measures of external reality include such numbers as the balance of payments with another country, the number of five-year olds in a school district, the number of poor in a country, the cost of living, the unemployment rate, and the number of people watching a particular television program on a given night.

2. Representations of *processes* are numbers purporting to measure "work" performed. They include the fraction of the time of a machinist or lawyer that is allocated to a particular product or client, the total number of hours worked, and the length of time taken to produce a product. They also include records of how resources were allocated—for example, how much was spent on administration, on pure versus applied research, and on graduate versus undergraduate education.

3. Representations of *outcomes* are numbers purporting to report the outcomes of decisions or activities. In a business firm, this includes outcomes such as sales or profits. In a school, student achievement is represented by a number. Numbers are also constructed to measure such outcomes as number of enemy killed, changes in crime rates, and budget deficits.

The construction of these magic numbers is partly problem solving. Decision makers and professionals try to find the right

answer, often in the face of substantial conceptual and technical difficulties. Numbers presuppose a concept of what should be measured and a way of translating that concept into things that can be measured. Unemployment numbers require a specification of when a person is "seeking employment" and "not employed". The concepts and their measurement are sufficiently ambiguous to make the creation of unemployment statistics a difficult technical exercise. Similarly, the definition and measurement of corporate profits, gross national product (GNP), or individual intelligence are by no means simple matters. They involve professional skills of a high order.

The construction of magic numbers is also partly political. Decision makers and others try to find an answer that serves their own interests. Unemployment levels, profits, GNP, individual intelligence, and other numbers are negotiated among contending interests. If the cost-of-living index affects prices or wages, affected groups are likely to organize to seek a favorable number. If managers are evaluated in terms of their profits, they will seek to influence transfer prices, depreciation rates, and the application of accounting rules and conventions that affect the "bottom line." If political leaders care about GNP, they will involve themselves in the negotiation of those numbers. Management involves account and number management as much as it involves management of the things that the numbers represent.

These simultaneous searches for truth and personal advantage often confound both participants and observers. Realist cynics portray the pursuit of truth as a sham, noticing the many ways in which individuals, experts, and decision makers find it possible to "discover" a truth that happens to be consistent with their own interests. Idealist professionals portray the pursuit of personal advantage as a perversion, noticing the many ways in which serious statisticians struggle to improve the technical quality of the numbers without regard for policy consequences. Both groups have difficulty recognizing the ways in which the process subtly interweaves truth seeking and advantage seeking, leaving each somewhat compromised by the other, even as each somewhat serves the other.

The tenuousness and political basis of many key numbers is

well-known to decision makers. They regularly seek to improve and influence the numbers. At the same time, however, decision makers and others have an interest in stabilizing the numbers, securing agreement about them, and developing shared confidence in them as a basis for joint decision making and communication. The validity of a number may be less important than its acceptance, and decision makers may be willing to forgo insisting on either technical correctness or immediate political advantage in order to sustain social agreement.

1.2.3 Satisficing and Maximizing

Most standard treatments of rational decision making assume that decision makers choose among alternatives by considering their consequences and selecting the alternative with the largest expected return. Behavioral students of decision rules, on the other hand, have observed that decision makers often seem to *satisfice* rather than *maximize.* Maximizing involves choosing the best alternative. Satisficing involves choosing an alternative that exceeds some criterion or target.

The shopkeeper in a small retail store could determine price by assessing information about the complete demand of the relevant population at a set of various prices and selecting the price that best serves her or his preferences. Alternatively he or she could use a simple mark-up over cost in order to ensure an acceptable profit margin on each item. A maximizing procedure for choosing equipment at a new manufacturing facility would involve finding the best combination of prices and features available. A satisficing strategy would find equipment that fits specifications and falls within budget. A marketing manager could seek to find the best possible combination of products, pricing, advertising expenses, and distribution channels; or he or she could create a portfolio of products that meets some sales, market share, or profit target.

DO DECISION MAKERS SATISFICE OR MAXIMIZE?

Neither satisficing nor maximizing is likely to be observed in pure form. Maximizing requires that all possible alternatives be

compared and the best one chosen. Satisficing requires only a comparison of alternatives with a target until one that is good enough is found. Maximizing requires that preferences among alternatives meet strong consistency requirements, essentially requiring that all dimensions of preferences be reducible to a single scale—although that scale need not exist in conscious form. Satisficing specifies a target for each dimension and treats the targets as independent constraints. Under satisficing, a bundle that is better on each criterion will not be chosen over another bundle that is good enough on each criterion if the latter bundle is considered first. Satisficing also makes it possible that no bundle will satisfy all criteria, in which case a decision will not be made.

In personnel decisions, a maximizing procedure would involve finding the best possible combination of persons and tasks. A satisficing procedure, on the other hand, would involve finding a person good enough to do the job. A decision maker would define a set of tasks adequate to accomplish the job, and would set targets (performance standards, job requirements) for performance on the job. A decision maker would consider candidates sequentially, perhaps by looking at the current job holder or an immediate subordinate, and would ask whether that person is good enough. When universities consider granting tenure to professors, or when individuals consider mates, for example, they can choose among a host of decision rules varying from relatively pure satisficing rules ("Does this person meet the standards set for satisfactory performance as a tenured professor or spouse?") to relatively pure maximizing rules ("Is this person the best possible person likely to be found—and available—for tenure or marriage in the indefinite future?").

There are problems with using empirical data to tell whether (or when) decision makers maximize or satisfice. The usual difficulties of linking empirical observations to theoretical statements are compounded by the ease with which either vision can be made tautologically "true." True believers in maximization can easily use circular definitions of preferences to account for many apparent deviations from maximizing. True believers in satisficing can easily use circular definitions of targets to account for many apparent deviations from satisficing.

Assessing whether organizations satisfice or maximize involves inferring decision rules from one or more of three kinds of data: (1) data drawn from listening to participants as they talk about the process, (2) data drawn from observing decision processes, and (3) data drawn from observing decision outcomes. The different kinds of data lead to different impressions.

When participants talk about the process, they seem generally to accept the ideology of maximization, but their descriptions sound a lot like satisficing. There is a strong tendency for participants to talk about targets as critical to the process of decision. Although there are frequent efforts to reduce a few separate goals to a common measure (e.g. profit), separate targets are treated as substantially independent constraints unless a solution satisfying them all cannot be found. In addition, alternatives are considered semisequentially. It may not be true that only one alternative is considered at a time (as in the pure form of satisficing), but only a few seem to be considered at a time.

In observations of the process of decision making, targets frequently appear as components of both official and unofficial practices. It is common to specify goals as constraints, at least at first. There is a tendency for only a few alternatives to be considered at a time, but consideration often continues for some more or less predetermined time, rather than strictly until the first satisfactory alternative is found. Decision makers sometimes seem to maximize on some dimensions of the problem and satisfice on others. Sometimes they seem to try to maximize the chance of achieving a target. Targets seem to be especially important when they are defined in terms of surviving until the next period, meeting a deadline, or fulfilling a contract. The pure maximization model seems not to fit the data, although in some situations people might be described as maximizing within a much-edited choice set.

When decision outcomes are observed, it is difficult to differentiate maximizing from satisficing. Most decisions are interpretable in either way, so it is necessary to find situations in which the two yield distinctively different outcomes. Maximization emphasizes the relative position of alternatives. A maximizing procedure is sensitive to nonhomogeneous shifts in alternatives, when one alternative improves relative to another. A maximizing search is sensitive to changes in expected return and costs. Satisficing, on the other hand, emphasizes the position of alternatives relative to a target. A satisficing procedure is sensitive to a change in the absolute value of the current choice, and thus to homogeneous downward shifts in alternatives if they include the chosen one. A satisficing search is sensitive to current position relative to the target.

It is necessary to find situations in which the position of the chosen alternative is changing relative to either other alternatives or the target, but not both. As an example, take the willingness of people to pursue energy conservation. Maximizers will be sensitive to shifts in relative prices but not to whether they reach a target or not (except secondarily). Satisficers will be sensitive to whether they are reaching a target but not to shifts in relative prices (except secondarily). Observations of actual decision making in such domains as new investments, energy conservation, and curricular decisions indicate that satisficing is an aspect of most decision making but that it is rarely found in pure form.

Beyond the evidence that such a portrayal seems to match many observations of decision making behavior, there are two broader theoretical reasons—one cognitive and one motivational—why behavioral students of decision making find satisficing a compelling notion. From a cognitive perspective, targets simplify a complex world. Instead of having to worry about an infinite number of gradations in the environment, individuals simplify the world into two parts—good enough and not good enough. From a motivational perspective, it appears to be true that the world of psychological sensation gives a privileged position to deviations from some status quo.

SATISFICING, ADAPTIVE ASPIRATIONS, AND THE STATUS QUO

In classical theories of rational choice, the importance of a potential consequence does not depend on whether it is portrayed as a "loss" or as a forgone "gain." The implicit aspiration level represented by the status quo is irrelevant. This posture of the

theory has long been resisted by students, and generations of economists have struggled to persuade students (and managers) to treat cash outlays and forgone gains as equivalent. The resistance of students has a natural satisficing explanation. Satisficing assumes that people are more concerned with success or failure relative to a target than they are with gradations of either success or failure. If out-of-pocket expenditures are treated as decrements from a current aspiration level (and thus as unacceptable) and forgone gains are not, the former are more likely to be avoided than the latter. A satisficing decision maker is likely to make a distinction between risking the "loss" of something that is not yet "possessed" and risking the loss of something that is already considered a possession.

The tendency to code alternatives as above or below an aspiration level or a status quo has important implications for decision making. Whether a glass is seen as half-empty or half-full depends on how the result is framed by aspiration levels and a decision maker's history. The history is important because aspiration levels—the dividing line between good enough and not good enough—are not stable. In particular, individuals adapt their aspirations (targets) to reflect their experience. Studies of aspiration level adjustment in situations in which information on the performance of others is lacking indicate that decision makers revise aspirations in the direction of past performance but retain a bit more optimism than is justified by that experience. Thus, current aspirations can be approximated by a positive constant plus an exponentially weighted moving average of past experience.

If aspirations adapt to experience, then success contains the seeds of failure, and failure contains the seeds of success. In a very general way, empirical data seem to support such a conception. Although there are some signs that chronically impoverished individuals are less happy than chronically rich individuals, studies of lottery winners reveal that they are no more happy than other people, and studies of paraplegics reveal that they are no less happy than others. This pattern of results has led some people to describe life as a "hedonic treadmill." As individuals adapt their aspirations to their experience, both their satisfactions and their dissatisfactions are short-lived.

The world is more complicated than such a simple model would suggest, of course. Aspirations adapt not only to one's own experience but also to the experience of others. They can become attached not just to the level of reward but to the rate of change of reward. They do not adapt instantaneously, and they appear to adapt upward more rapidly than downward. As a result, deviations in a negative direction seem to be more persistently noticed than positive deviations. This "predisposition to dissatisfaction" is, of course, a strong stimulus for search and change in situations where it exists.

1.3 Theories of Attention and Search

In theories of limited rationality, attention is a scarce resource. The evoked set, of alternatives, consequences, and preferences, and the process that produces the evoked set, take on an importance not found in models of infinitely rational decision makers. Not all alternatives are known, they must be sought; not all consequences are known, they must be investigated; not all preferences are known, they must be explored and evoked. The allocation of attention affects the information available and thus the decision.

Ideas that emphasize the importance of attention are found throughout the social and behavioral sciences. In psychology, the rationing of attention is central to notions of editing, framing, and problem solving "set"; in political science, it is central to the notion of agendas; in sociology, it is central to the notion that many things in life are "taken as given" and serve as constraints rather than as decision alternatives. In economics, theories of search are a central concern of the study of decisions. The study of decision making is, in many ways, the study of search and attention.

1.3.1 The Rationing of Attention

In contrast to traditional societies, which are ordinarily described as short of physical and human resources rather than short of time, the modern world is usually described as stimu-