CHAPTER 1

Introduction

JONATHAN BENDOR

There are two main orientations toward bounded rationality (BR) in political science. The first orientation sees the glass as half full, emphasizing that decision makers often manage to do "reasonably well"—even in complex tasks—despite their cognitive limitations. Virtually all of Simon's work and also the theory of "muddling through" (Lindblom 1959; Braybrooke and Lindblom 1963) belong to this branch, which we can call the problem-solving approach. In the second orientation the glass is half empty: the emphasis is on how people make mistakes even in simple tasks. Most of the research on heuristics and biases, following Tversky and Kahneman's pioneering work (1974), belongs here.[1]

Prominent early use of the problem-solving approach can be found in Aaron Wildavsky's studies of budgeting. In, for example, *The Politics of the Budgetary Process*, he devotes much space to showing how and why making resource allocation decisions in the federal government is beset by complexities *and* how the professionals cope with their difficult tasks: "It [is] necessary to develop mechanisms, however imperfect, for helping men make decisions that are in some sense meaningful in a complicated world" (1964, p. 11). One might argue that his orientation was due simply to the time paths of these different intellectual currents: Simon and Lindblom had launched the problem-solving branch before Aaron wrote his pioneering book on budgeting, whereas the Tversky-Kahneman branch didn't get started until nearly a decade later. But there is a deeper explanation. Aaron did field research on federal budgeting, including 160 interviews with "agency heads, budget officers, Budget Bureau staff, appropriations committee staff, and Congressmen" (1964, p. v). He was not interested in how experimental subjects committed errors of judgment or choice in laboratory settings; he was interested in how real decision makers tackled problems of great complexity. Hence, he was intrigued by how they managed to do this extremely difficult task reasonably well.[2] One sees in the book a respect for the decision makers, arising in large measure from an appreciation of the difficulty of the tasks they confronted.

Indeed, I suspect that the seriousness with which Aaron thought about the tasks of budgetary officials was part of a long-standing theme of his professional life: a passionate interest in the real-world problems confronting government officials in a modern society. (Helping to found Berkeley's Graduate School of Public Policy was another reflection of this theme.)

This is more than biographical detail. It also illustrates an important—though neglected—part of the problem-solving approach to bounded rationality: a close examination of decision makers' tasks. In Simon's pioneering formulation, the focus was always on a *comparison* between a decision maker's mental abilities and the complexity of the problem he or she faces: for example, "the capacity of the human mind for formulating and solving complex problems is very small compared with the size of the problems whose solution is required for objectively rational behavior in the real world—or even for a reasonable approximations to such objective rationality" (1957, p. 198). Thus, for Simon, as for Wildavsky, the idea of bounded rationality is *not* a claim about the brilliance or stupidity of human beings, independent of their task environments. Many social scientists miss this central point and reify the idea of BR into an assertion about the *absolute* capacities of human beings.[3] The fundamental notion here is that of cognitive limits; and, as for any constraint, if cognitive constraints do not bind in a given choice situation, then they will not affect the outcome. *And whether they bind depends vitally on the information-processing demands placed on the decision makers by the problem at hand.* More vividly, Simon has called the joint effects of "the structure of task environments and the computational capacities of the actor . . . a scissors [with] two blades" (1990, p. 7): Theories of BR have cutting power—especially when compared to theories of (fully) rational choice—only when both blades operate. Thus, any analysis that purports to fall into this branch of the research program yet examines only the agent's properties is badly incomplete.

Thus, Wildavsky not only belonged squarely in the problem-solving branch of the BR program; his intellectual propensities—his interest in how real officials tackle real problems of great complexity—predisposed him to use *both* blades of Simon's scissors. That was unusual. It was also productive: many of his insights about budgeting flowed from his effective use of Simon's scissors.

Of course, every research method focuses our attention on some scholarly questions in the domain at hand *and* deemphasizes others in that same domain. (Lindblom's warnings [1959] about the utopian folly of trying to be comprehensive apply to academics as well as to government officials.) So it is not surprising that Wildavky's research methods led him to ignore certain topics. In particular, his interest in *applying* the basic ideas of bounded rationality to the study of real-world budgeting steered him away from analyzing the foundations of BR theory. That simply was not part of his intellectual agenda. But a serious focus on those foundations is long overdue. Brilliant as they were, neither Simon nor Lindblom said it all. We political scientists—particularly those of us who work on the behavioral (bounded rationality) side—have done too much quoting and too little reworking. I believe that we will see vigorous scientific competition between rational choice (RC) theories of policy making and behavioral theories only if behavioralists take the foundations of their theories as seriously as RC theorists take theirs. Further, I think that this entails transforming verbal theories into mathematical models. (For an argument on this point in the context of incrementalism, see chapter 4.)

The next section surveys a family of theories that has been central to the problem-solving branch of the BR program: those that use the idea of *aspiration levels* as a major concept.

THEORIES OF ASPIRATION-BASED PROBLEM REPRESENTATION AND CHOICE

The main claim I offer in this section is that the idea of aspiration-based choice constitutes a major family of theories in the bounded rationality research program. The word *family* matters: I think it is a serious mistake to view satisficing per se as an alternative to theories of optimization. As careful scholars working in the optimization tradition have often pointed out, there is no single RC theory of (e.g.) electoral competition (see, e.g., Roemer's comparison [2001] of Downsian theory to Wittman's), much less just one RC theory of politics. Similarly, satisficing is a theory of search. It is *not* the Behavioral Theory of Everything. Moreover, a key part of satisficing—the idea of aspiration levels—is shared by several other important behavioral theories: theories of reinforcement learning (Bush and Mosteller 1955) and prospect theory (Kahneman and Tversky 1979). So I first argue that a "family" of theories is a significant grouping that fits into the more conventional hierarchy of research program, theories, and models, and that substantively we can gain some insight by focusing our attention on this common feature of aspirations.

(A reason that it is methodologically important to identify this family of theories is that the size of this set is probably indefinite. That is, an indefinitely long list of choice problems may be representable via aspiration levels. I see no reason why this should not hold. All that is required of the choice problem is that there be more than two feasible payoffs, but the agent simplifies the problem by reducing that complex set into two simple equivalence classes. Importantly, there is *no* restriction on the substantive type of problem, or the task of the decision maker, which can be represented this way.)[4]

Technically, an aspiration level is a threshold in an agent's set of feasible payoffs. This threshold partitions all possible payoffs into two disjoint sets: those below the threshold and those that are greater than or equal to the threshold.[5] In all aspiration-based theories, this dichotomous coding *matters*: that is, important further implications flow from this representation of the choice problem.[6] The details of these implications vary, for they naturally depend on the substantive nature of the theory at hand, as we will see shortly. This parallels how the details of optimal strategies vary, depending on the type of problem confronting the decision maker. Strategies of optimal candidate location in a policy space look quite different from strategies of optimal nuclear deterrence. But they share a common core, that of optimal choice. Similarly, psychological theories of learning look quite different from prospect theory, but they too have a common core. Interestingly, we will see that early on, scholars working on certain members of this family of theories were not even aware that their particular theories *required*, as a necessary part of their conceptual apparatus, the idea of aspirations; they backed into this idea. Let me now briefly describe several important members of this family of theories: satisficing and search, reinforcement learning, and prospect theory.

Theories of Aspiration-Based Behavior

Search Search was Simon's original context for satisficing. The idea is simple. Simon posited that when an agent looked for, say, a new job or house, he or she had an *aspiration level* that partitioned candidates into satisfactory options and unsatisfactory ones.[7] As soon as the decision maker encountered a satisfactory one, search ended. The verbal theory suggests that the aspiration level adjusts to experience, falling in bad times (when one searches without success) and rising in good times (swiftly encountering something that far exceeds the aspiration level), but this was not represented by the formal model. (Cyert and March [1963] allow aspirations to adjust to experience in their computational model.)

There is, however, a rival formulation: optimal search theory. The basic idea is that agents assess both the expected marginal gains and the expected marginal costs from searching further and set an optimal stopping rule that equates the two. Some behavioralists (e.g., Schwartz et al. 2002) are completely unaware that this rival exists. This is a pity. The ignorance allows such behavioralists to have scholarly aspirations that are too low: they think satisficing easily beats RC theories because the latter predict that decision makers search all options exhaustively before making a choice. Since this is obviously false in many (most?) domains, satisficing (hence BR, etc.) comes out on top in this scientific competition. But this horse race was bogus: optimal search theory does *not*, in general, predict exhaustive search. Indeed, in many environments the optimal stopping rule takes the form of a cutoff: if the agent stumbles on an option worth at least $\overline{v}$, stop searching; otherwise, continue.[8] Significantly, the main qualitative features of this prediction are exactly the same as those of the satisficing theory: the set of feasible options is partitioned into two subsets, and the searcher stops looking upon finding something belonging to the better subset. Clearly, then, some additional cleverness is required in order to derive different predictions from RC and BR theories of search. It's doable, but we can't stop with the obvious predictions.

Now that the notion of optimal search has been spelled out, it is clear that in most search problems the stopping rule could be *either* suboptimally low *or* suboptimally high. If the costs of search are sufficiently low, then one should keep searching until one has found the highest-quality option, but this will rarely be the case in policy making, particularly not when decision makers are busy (Behn and Vaupel 1982).

Interestingly, however, the possibility that uncalculated aspirations may be *too high* has gone almost unnoticed by the behavioral literature. Indeed, an auxiliary premise has been smuggled into the concept of satisficing: it is virtually *defined* as search with a suboptimally low threshold (as in the phrase that probably most of us have heard, "*merely* satisfice").

The problem is not with the definition per se—one can stipulate a technical term as one pleases—but with its uses. This implicit smuggling of an important property into the heart of satisficing helps us to overlook the *possibility* of excessively high aspirations. It reinforces the mistake of equating "optimal" with "best quality," or worse, assuming that optimal equals perfect. (Chapters 3 and 5 analyze the implications of equating optimal with perfect.) We thereby neglect some of the empirical content of aspiration-based models of adaptation.

Learning In experiments on learning, psychologists often give subjects a set of options that they can repeatedly try. In so-called bandit problems, every option either pays off some fixed amount $v > 0$ or yields nothing. Some options are better than others—pay v with a higher probability—but subjects aren't told which. Instead, they must learn which options are better.

Originally, these experiments were part of the behaviorist research program in psychology, which eschewed mentalistic concepts such as aspirations.[9] However, learning theorists in psychology discovered (the hard way) that they needed this concept to explain the behavior of subjects.[10]

This issue becomes more pressing in choice situations where there are more than two payoffs. Given only two payoffs, it is quite natural to hypothesize that subjects will regard getting something as a success while getting nothing is a failure.[11] But many choice situations do not provide such an obvious coding. In, for example, the two-person prisoner's dilemma, is the payoff to mutual cooperation reinforcing? How about the payoff to mutual *defection*?

Prospect Theory Perhaps the best-known postulate of prospect theory (Kahneman and Tversky 1979) is that decision makers are risk averse regarding gains but risk seeking regarding losses. This is not, however, a good way to remember the theory. Its fundamental axiom—its most important departure from classical utility theory—is that people evaluate outcomes relative to a reference point.[12] Indeed, under the classical theory, the claim that people are, for example, risk averse about gains *makes no sense*: the idea of a "gain" has no place in the conceptual apparatus in standard utility theory.[13] Decision makers simply have preferences over baskets of consequences; that's all that matters. They do not compare baskets to an internal standard of goodness or acceptability.

A reference point is, in effect, an aspiration level.[14] Of course, aspirations in prospect theory have a different function than they do in satisficing-and-search theory. (Prospect theory has not been applied to search problems, as far as I know.) Rather than serving as a stopping rule, aspirations in the context divide the set of feasible outcomes into those coded as gains and those coded as losses. But again we see a dichotomizing of payoffs into two qualitatively different subsets.

An Important Problem: The Empirical Content of Aspiration-Based Theories

Although I think that aspiration-based theories form a tremendously important family in the BR research program, no set of theories in the social sciences is free of problems. (At least, I have not been lucky enough to encounter such a set!) And since I completely agree with Martin Landau's view that scientific progress depends tightly on criticism, I think it is vital for scholars who work on these theories to detect their weaknesses and work on them. I do not think that aspiration-based theories of choice are perfect. That would be a reflexively bizarre claim.

Here I want to identify just one problem—but, I think, a significant one. It concerns the empirical content of several kinds of aspiration theories: at a minimum, those of search and of learning. (Here I am relying entirely on the results of Bendor, Diermeier, and Ting [2003, 2007], hereafter BDT.) The problem is simple: many of these theories have very little empirical content. Indeed, some probably cannot be falsified. This is troubling.

These are strong claims; they should be demonstrated. I will sketch out one of BDT's results (which are all deductively established) to give a sense of the logic. Consider an agent who learns via classical trial and error: try an action; if it "works" (the current payoff, π_t, is at least as big as the current aspiration, a_t), then the agent becomes more disposed to use it again. If it fails ($\pi_t < a_t$), then the agent becomes less likely to try it in the future. In accord with the conventional wisdom, aspirations also change with experience, moving up in good times and down in bad. (These informal ideas can be made mathematically precise: see Bendor, Diermeier, and Ting (2003, p. 264) for a formal specification of these axioms of propensity adjustment and aspiration adjustment. Theorem 1, below, refers to these as axioms 1 and 3, respectively.) Then we get the following result.

> THEOREM 1. *Consider any repeated game with deterministic and stationary payoffs in which players adjust their action propensities by any arbitrary mix of adaptive rules that satisfy axiom 1 and adjust their aspirations by any arbitrary mix of rules that satisfy axiom 3. Then any outcome of the stage game can be sustained as a stable outcome by some self-replicating equilibrium. (BDT 2003, 2007)*

This is a "folk theorem" for reinforcement learning, in the same sense that there are folk theorems for repeated games in noncooperative game theory: if anything is stable, then the theory isn't predicting much.

Fortunately, there are ways of ameliorating this problem and restoring empirical content to models of aspiration-based search and learning. BDT identify two options: one can either assume that players randomly experiment (this is precluded by classical satisficing) or let them obtain *vicarious* experience by making their aspirations depend partly on other people's experience (payoffs). The latter method builds the venerable sociological idea of reference groups into these models. (Classical satisficing theory is asocial.)

CONCLUSION: SOME POLICY IMPLICATIONS

The contest between RC and BR theories will last for decades. But life and its problems go on. Given this difference in tempos between basic and applied science, it's important to put some pressure on basic scientists. After all, if the change is as big a deal as the proponents of the challenging research program claim, it should generate interesting policy implications. Below are two that I think are significant. The first implication concerns the normative analysis of public policies; the second is descriptive as well as prescriptive.

Happiness, Aspirations, and Policy Evaluation

Since Bentham, evaluations of public policies have had a strongly utilitarian cast. But modern applications of utilitarianism, such as cost-benefit analysis, are based not only on the classical theory's normative principles but also on descriptive axioms regarding utility—or, more informally, happiness. As noted earlier, these descriptive axioms do not include the concept of an aspiration level. Scholars working in the new field of hedonic psychology, however, have found that concept very useful indeed in explaining some intriguing empirical findings.[15] (In this sense the largely empirical field of hedonic psychology is linking hands with theories of aspiration-based adaptation discussed throughout this book.) For example, although per capita GNP rose dramatically from 1946 to 1990 in France, Japan, and the United States, "there was no increase in mean reports of [subjective well-being]" (Diener et al. 1999, p. 288). Endogenously rising aspirations explain this Faustian dynamic of doing well but feeling no better.

The implications of these empirical findings and their aspiration-based explanation for policy evaluation could be profound. To cut to the chase: if the subjective well-being of citizens affected by *every* public program is determined by their comparing objective payoffs to their subjective aspiration levels, then cost-benefit analysis is seriously flawed. As the data reported by Diener et al. suggests, it could produce systematic biases. In particular, a better-grounded evaluation of the collective impact of our public (and private) policies might show that though we're getting a lot richer (on average), we're not getting much happier.[16]

More specific normative implications require a specific model of aspiration-based choice. Consider prospect theory with endogenous aspirations.[17] Given the axiom of loss aversion, plus the conventional assumption that aspirations adjust to payoffs with a nontrivial lag, such a model implies that a utilitarian measure of social welfare would be enhanced if public policies cushioned people against the short-term effects of sudden catatrophic losses.

Expertise in the Policy Process: Discovering and Teaching Effective Heuristics

Some scholars in the problem solving branch of the BR program (e.g., Ericsson and Lehmann 1996; Simon 1999b) have studied expertise closely. A major finding in this area is that most experts labor under the same general mental constraints—such as the limited number of pieces of information that can be held in working memory—as do the rest of us. Thus, they do not overcome bounded rationality; instead, they *finesse* it: they've learned domain-specific methods that are procedurally rational (Simon 1990, 1996). One of these methods is heuristic search.

This assertion is, I believe, inconsistent with the views of many political scientists, who seem to believe in the proposition that *amateurs satisfice but experts optimize*. This belief is based on a misunderstanding of how satisficing in particular and heuristics in general fit into Simon's theory of problem solving. Simon realized long ago that chess is so difficult—the decision tree explodes after only a few moves—that even grand masters must search heuristically (Simon and Schaeffer 1992). Part of the difference—as Lindblom guessed—is that they use powerful heuristics; we duffers do not.[18]

Thus, if we are to teach our students to be genuine experts, in whatever policy field they go into, we must take heuristics seriously.[19] This would be a sharp break with the past. With a few exceptions, theoretically oriented political scientists and most policy scientists trained in economics have been oriented toward *strategies*—especially optimal ones—rather than toward heuristics. The two are quite different. Whereas a strategy (in its technical, game-theoretic sense) is a complete plan of action, heuristics can be incomplete. Heuristics—"get your middle pieces out early," "reciprocate" (in the iterated prisoner's dilemma), "try a proof by contradiction"—may be valuable pieces of advice even if they aren't complete solutions. Moreover, whereas we are accustomed to making very sharp evaluations of strategies—they are either optimal (or part of Nash equilibria, in strategic settings) or not—heuristics come in many shades of gray. Indeed, we use heuristics when we don't have the slightest idea of what *is* an optimal plan for the task at hand. Learning which are "pretty good" heuristics in a given domain and which are mediocre is an important part of becoming an expert.[20] I think we do our professional-degree students a profound disservice if we pretend that all they will need in the real world are strategies. Most of the time they will need heuristics; full-blown strategies will be beyond their (or anyone's) reach. Lindblom's cautions have lost none of their punch.