Annu. Rev. Polit. Sci. 1999. 2:297–321

BOUNDED RATIONALITY

Bryan D. Jones
Department of Political Science, University of Washington, Seattle, Washington 98195; e-mail: bdjones@u.washington.edu

KEY WORDS: decision making, behavioral organization theory, behavioral decision theory, political psychology

ABSTRACT

Findings from behavioral organization theory, behavioral decision theory, survey research, and experimental economics leave no doubt about the failure of rational choice as a descriptive model of human behavior. But this does not mean that people and their politics are irrational. Bounded rationality asserts that decision makers are intendedly rational; that is, they are goal-oriented and adaptive, but because of human cognitive and emotional architecture, they sometimes fail, occasionally in important decisions. Limits on rational adaptation are of two types: procedural limits, which limit how we go about making decisions, and substantive limits, which affect particular choices directly. Rational analysis in institutional contexts can serve as a standard for adaptive, goal-oriented human behavior. In relatively fixed task environments, such as asset markets or elections, we should be able to divide behavior into adaptive, goal-oriented behavior (that is, rational action) and behavior that is a consequence of processing limits, and we should then be able to measure the deviation. The extent of deviation is an empirical issue. These classes are mutually exclusive and exhaustive, and they may be examined empirically in situations in which actors make repeated similar choices.

INTRODUCTION

Do people make rational decisions in politics and economics? Not if by "rational" we mean that they demonstrate conformity to the classic expected-utility model. There is no longer any doubt about the weight of the scientific evidence; the expected-utility model of economic and political decision making is not sustainable empirically. From the laboratory comes failure after failure of rational expected utility to account for human behavior. From systematic observation in organizational settings, scant evidence of behavior based on the expected-utility model emerges.

1094-2939/99/0616-0297$08.00

Does this mean that people (and therefore their politics) are irrational? Not at all. People making choices are intendedly rational. They want to make rational decisions, but they cannot always do so.

The implication for politics is that rational responses to the environment characterize decision making generally, but at points—often important points—rationality fails, and as a consequence there is a mismatch between the decision-making environment and the choices of the decision maker. We refer to this mismatch as "bounded rationality showing through" (Simon 1996b).

This conception has an important implication. In structured situations, at least, we may conceive of any decision as having two components: environmental demands (seen by the individual as incentives, positive or negative) and bounds on adaptability in the given decision-making situation. Ideally, an analysis based on rational choice should be able to specify what the environmental incentives are and to predict decisions based on those incentives. What cannot be explained is either random error (even the most rational of us may make an occasional mistake, but these are not systematic) or bounded rationality showing through. Standard statistical techniques give us the tools to distinguish systematic from random factors, so in principle it should be possible to distinguish the rational, adaptive portion of a decision from bounds on rationality.

One may think of any decision as arising from two sources. One is the external environment—how we respond to the incentives facing us. The other is the internal environment—those parts of our internal make-ups that cause us to deviate from the demands of the external environment (Simon 1996b).

We are not, however, thrown into a situation in which all residual systematic deviations from rational choices are treated prima facie as bounded rationality. A very limited set of facets of human cognitive architecture accounts for a very large proportion of the deviations from adaptation. These may be placed into two classes: procedural limits, which limit how we go about making decisions, and substantive limits, which affect particular choices directly. Of procedural limits, I cite two as being extraordinarily important in structured, institutional settings (such as voting in mass publics or in legislative bodies), attention and emotion. Of substantive limits, I cite but one—the tendency of humans to "overcooperate," that is, to cooperate more than strict adherence to rationality would dictate.

The primary argument in this essay is that most behavior in politics is adaptive and intendedly rational but that limits on adaptive behavior, imposed by human cognitive/emotional architecture, may be detected in even the most stable of environments. I advocate a research strategy that explicitly divides political action into the two categories of intended rationality and deviations from (or bounds on) intended rationality and explores empirically the implications for the outputs of institutions and the institutional processes responsible

for those outcomes. [The analysis presented here is further developed in my *Traces of Eve: Adaptive Behavior and Its Limits in Political and Economic Institutions* (manuscript in preparation).]

BOUNDED RATIONALITY: BIRTH AND DEVELOPMENT

Bounded rationality is a school of thought about decision making that developed from dissatisfaction with the "comprehensively rational" economic and decision theory models of choice. Those models assume that preferences are defined over outcomes, that those outcomes are known and fixed, and that decision makers maximize their net benefits, or utilities, by choosing the alternative that yields the highest level of benefits (discounted by costs). The subjective expected-utility variant of rational choice integrates risk and uncertainty into the model by associating a probability distribution, estimated by the decision maker, with outcomes. The decision maker maximizes expected utility. Choices among competing goals are handled by indifference curves—generally postulated to be smooth (twice differentiable)—that specify substitutability among goals.

A major implication of the approach is that behavior is determined by the mix of incentives facing the decision maker. A second implication is that adjustment to these incentives is instantaneous; true maximizers have no learning curves.

Like comprehensive rationality, bounded rationality assumes that actors are goal-oriented, but bounded rationality takes into account the cognitive limitations of decision makers in attempting to achieve those goals. Its scientific approach is different; rather than making assumptions about decision making and modeling the implications mathematically for aggregate behavior (as in markets or legislatures), bounded rationality adopts an explicitly behavioral stance. The behavior of decision makers must be examined, whether in the laboratory or in the field.

The Birth of Bounded Rationality

Simon (1999; see also Simon 1996a) reminds political scientists that the notion of bounded rationality and many of its ramifications originated in political science. Over his long career, Simon made major contributions not only to political science (as the founder of the behavioral study of organizations) but also to economics (as a Nobelist), psychology (as a founding father of cognitive psychology), and computer science (as an initiator of the field of artificial intelligence).

In the 1940s and 1950s, Simon developed a model of choice intended as a challenge to the comprehensive rationality assumptions used in economics.

The model first appeared in print in *Administrative Behavior* (1947), which critiqued existing theories of public administration and proposed a new approach for the study of organizational decision making. Simon gave great credit for the initiation of his innovative work to the behavioral revolution in political science at the University of Chicago, where he studied for all of his academic degrees. Although most political scientists are aware of Simon's contributions, many fail to appreciate that bounded rationality was the first, and because of its ripple effects in so many disciplines, the most important idea (even academic school of thought) that political science has ever exported.[1]

A brief retelling of the tale is in order. As an undergraduate at the University of Chicago, Simon returned to his native Milwaukee in 1935 to observe budgeting in the city's recreation department. He wrote:

> I came as a gift-bearing Greek, fresh from an intermediate price theory course taught by the grandfather of Chicago-School neoclassical laissez-faire economics, Henry Simons.... My economics training showed me how to budget rationally. Simply compare the marginal utility of a proposed expenditure with its marginal cost, and approve it only if the utility exceeds the cost. However, what I saw in Milwaukee didn't seem to be an application of this rule. I saw a lot of bargaining, of reference back to last year's budget, and incremental changes in it. If the word "marginal" was ever spoken, I missed it. Moreover, which participants would support which items was quite predictable.... I could see a clear connection between people's positions on budget matters and the values and beliefs that prevailed in their sub-organizations.
>
> I brought back to my friends and teachers in economics two gifts, which I ultimately called "organizational identification" and "bounded rationality." (Simon 1999)

In his autobiography, Simon noted the importance of these two notions for his later contributions to organization theory, economics, psychology, and computer science. "I would not object to having my whole scientific output described as largely a gloss—a rather elaborate gloss, to be sure—[on these two ideas]" (Simon 1996a:88).

Bounded rationality and organizational identification (now considered a consequence of bounded rationality) won ready acceptance in political science, with its emerging empiricist orientation, but they were largely ignored in the more theoretical discipline of economics. Or, as Simon (1999) puts it,

[1]Two recent incidents convinced me of the need to remind political scientists that Simon's "tribal allegiance" (1999) is to our discipline. A well-regarded political scientist recently commented, "I didn't know that Simon was a political scientist." In a written review, a cognitive psychologist somewhat haughtily informed me that Simon's work on organizations, and in particular March & Simon's *Organizations* (1958), was intended to extend his work on problem solving to organizational behavior. Of course, the intellectual path was the other way around.

economists "mostly ignored [bounded rationality] and went on counting the angels on the heads of neoclassical pins."

Procedural Rationality

Simon spent a great deal of time and energy attacking the abstract and rarefied economic decision-making models. Much of his attack was negative—showing how the model did not comport with how people really made decisions. But Simon also developed what he termed a procedural model of rationality, based on the psychological process of reasoning—in particular his explanation of how people conduct incomplete searches and make tradeoffs between values.

> Since the organism, like those of the real world, has neither the senses nor the wits to discover an "optimal" path—even assuming the concept of optimal to be clearly defined—we are concerned only with finding a choice mechanism that will lead it to pursue a "satisficing" path that will permit satisfaction at some specified level of all of its needs. (Simon 1957:270–71)

Simon elaborated on his "satisficing" organism over the years, but its fundamental characteristics did not change. They include the following:

1. Limitation on the organism's ability to plan long behavior sequences, a limitation imposed by the bounded cognitive ability of the organism as well as the complexity of the environment in which it operates.

2. The tendency to set aspiration levels for each of the multiple goals that the organism faces.

3. The tendency to operate on goals sequentially rather than simultaneously because of the "bottleneck of short-term memory."

4. Satisficing rather than optimizing search behavior.

An alternative satisfices if it meets aspirations along all dimensions (attributes). If no such alternative is found, a search is undertaken for new alternatives.

Meanwhile, aspirations along one or more dimensions drift down gradually until a satisfactory new alternative is found or some existing alternative satisfices (Simon 1996b:30).

In detailing the general requirements of an organism operating under bounded (as contrasted with comprehensive) rationality, Simon (1983:20–22; see also Simon 1995) notes the following requisites: (*a*) "Some way of focusing attention," (*b*) "a mechanism for generating alternatives," (*c*) "a capacity for acquiring facts about the environment," and (*d*) "a modest capacity for drawing inferences from these facts."

I cannot do justice to the importance for other disciplines of Simon's "gloss" on bounded rationality. Just one note: The study of problem solving is grounded in the intended rationality of problem solvers, as is the study of judgment (Newell 1968, 1990). By imposing a task environment, experimenters can examine that part of the problem solver's behavior that may be explained objectively, via the nature of the task environment, and compare it with that part that can be explained only with reference to failures to overcome systematic internal limitations—bounded rationality showing through (Newell & Simon 1972, Simon 1996b).

The principle that rationality is intended but not always achieved, that what "shows through" from the inner environment of the problem solver can be systematically studied, is a principle that I consider extraordinarily useful in the study of human behavior in relatively set institutional task environments.

BOUNDED RATIONALITY IN POLITICAL SCIENCE

Bounded rationality has been a key component since the 1950s in public-administration and public-policy studies. In more recent times, partly in reaction to the attitudinal model of voting behavior, the approach has been used to understand political reasoning (Iyengar 1990, Sniderman et al 1991, Marcus & McKuen 1993). Nevertheless, bounded rationality, born in organization theory (Simon 1947), has had its greatest impact in political science in the study of governmental organizations.

The fundamental premise underlying organizational studies in political science is that the behavior of organizations mimics the bounded rationality of the actors that inhabit them (March 1994). This correspondence is not simply an analogy among phenomena at different levels; the relationship is causal. This premise characterized behavioral organization theory generally, along with the insistence that organizational science be grounded in the observation of behavior in (and analysis of data from) organizational settings. The most important components of the political theory of organizations were the concepts of limited attention spans, habituation and routine, and organizational identification. Behavioral organization theory, unlike the subjective expected-utility approach, viewed uncertainty not as simple probabilities attached to specified outcomes, but as infecting the very specification of outcomes themselves.

Over and over again, students of the behavior of public organizations reported findings that did not comport with the demands of "objective rationality" (Simon 1985:294). Search was incomplete, selective, and nonoptimal (Simon 1985, Jones & Bachelor 1994). Decision makers did not need simply to choose among alternatives; they had to generate the alternatives in the first place (Simon 1983, 1996b; Chisholm 1995). Problems were not givens; they

had to be defined (Rochefort & Cobb 1994). Solutions did not automatically follow problems; sometimes actors had set solutions ready to apply to problems that could occur (Cohen et al 1972, Kingdon 1996, Jones & Bachelor 1994). Choice was based on incommensurate goals, which were ill-integrated (March 1978; Simon 1983, 1995; Jones 1994). Organizations seemed to have limited attention spans and, at least in major policy changes, serial processing capacity (Simon 1983, Jones 1994, Cobb & Elder 1972, Kingdon 1996).

The three most important strands of research stemming from behavioral organizational theory in political science focused on incremental budgeting, on the impacts of organizational routine on policy outputs, and on policy agendas.

Incremental Budgeting

Incremental decision making was developed not only as a descriptive model of decisions by bounded actors but as a normative mechanism for use in an uncertain world (Lindblom 1959). If people are handicapped by limited cognition, and if the world is fundamentally complex and ambiguous, then it made sense for a decision maker to (*a*) move away from problems, rather than toward solutions; (*b*) make only small moves away from the problem; and (*c*) be willing to reverse direction based on feedback from the environment. Wildavsky (1964; see also Fenno 1966, Meltsner 1971), in his classic observational studies of federal budgeting, noted that such incremental budgeting was governed by decision rules based on two norms: base and fair share. What was the agency's base, and what was a fair share given changes in the agency's environment since last year's budget meeting? Incrementalism was even criticized as too rational a characterization of budget processes, because of the adoption of roles by budget decision makers (Anton 1966, Crecine 1969). Incrementalism, in effect a small-step hill-climbing algorithm, implied adjustment to local optima rather than global ones.

Incrementalism in decision making implied incrementalism in organizational outcomes—so long as one also modeled exogenous "shocks" (Davis et al 1966, 1974). Students of the budgetary process concluded that incrementalism did not fit even endogenous decision processes (Wanat 1974, Gist 1982). Pure incrementalism did not seem to characterize governing organizations. In essence, there were too many large changes in budget processes. But it was realized that attentional processes are selective (as the incremental model recognized) and subject to occasional radical shifts. Incorporating this aspect of attentional processes better accounts for the distribution of budget outcomes (Padgett 1980, 1981; Jones et al 1996, 1997, 1998).

Organizational Habits and Routines

Cognitive limits of human decision makers imposed limits on the ability of the organization to adjust to its environment. Rather than maximizing, organiza-

tions tended to adopt task performance rules, which routinized even the most important decisions of the organization (March & Simon 1958). Firms routinized price and output decisions (Cyert & March 1963). Learning in organizations seemed to be a slow, evolutionary, conflictual process (Sabatier & Jenkins-Smith 1993, Lounamaa & March 1985, Ostrom 1990) rather than the instantaneous adjustment process that rational organization theory would imply. Participants identified with the rules of the organization, adhering to them even in the face of evidence of problems (Jones 1980, 1985). This could cause disjoint "lurches" as organizations were finally forced to adjust to changes in their environments (Dodd 1994).

Routines in service organizations invariably generated unintended consequences, many of which went unrecognized or unaddressed. For example, distributional consequences of supposedly neutral rules were often ignored (Levy 1974, Mladenka 1978).

In other cases, an organization might have contradictory demands on it. Such contradictory demands are handled in economics via indifference curves, which specify a decision maker's preferences under all combinations of the demands. Instead of a rational process for handling tradeoffs, public service organizations tended to develop task performance rules for each demand. The response of the organization depended on which set of rules was activated. A study of Chicago's Building Department revealed that two sets of task performance rules were in effect. One set directed resources in accordance with the severity of the problem. These rules embodied the classic administrative norm of neutral competence. A second set of rules, less explicit but just as important, directed resources based on responsiveness to political forces. The distribution of organizational outputs to neighborhoods depended on an attention rule, activated by middle management, that governed which set of rules was to be put in force. Neutral competence was the default; response to political forces required an override of standard operating procedures, but the attention rule override happened so often that it could easily be detected in organizational outputs (Jones 1985).

Policy Agendas

If individuals have limited attention spans, so must organizations. The notion of policy agendas recognizes the "bottleneck" that exists in the agenda that any policy-making body addresses (Cobb & Elder 1972). These attention processes are not simply related to task environments—problems can go for long periods of time without attracting the attention of policy makers (Rochefort & Cobb 1994). A whole style of politics emerges as actors must strive to cope with the limits in the attentiveness of policy makers—basically trying to attract allies to their favored problems and solutions. This style of politics depends on connections driven by time-dependent and often emotional attention processes

rather than a deliberate search for solutions (Cohen et al 1972, March & Olsen 1989, Kingdon 1996, Baumgartner & Jones 1993).

Because attention processes are time dependent and policy contexts change temporally, connections between problems and solutions have time dependency built into them. As an important consequence, policy systems dominated by boundedly rational decision makers will at best reach local rather than global optima. Because of the time dependence of attentional processes, all policy processes will display considerable path dependence (March 1994).

OBJECTIONS TO THE EXPECTED-UTILITY MODEL: BEHAVIORAL DECISION THEORY

The expected-utility model incorporates risk and uncertainty into models of rational choice. Instead of maximizing utility, decision makers maximize expected utility in choice situations in which the consequences of choice are risky (may be characterized by known probabilities) or uncertain (are characterized by unspecified probabilities).

Numerous empirical studies of human decision making, from experiments in the laboratory to large-scale social surveys to observational studies in the field, have demonstrated that humans often do not conform to the strictures of choice theory (Slovak 1990). This study of how people actually behave in choice situations is known as behavioral decision theory. Even defenders of choice theory have retreated in the face of the onslaught of empirical findings. Expected-utility theory is no longer seriously entertained as an accurate descriptive theory (Halpern & Stern 1998b).

Again, this does not imply that people are irrational, nor that people interacting in large-scale institutions make large-scale mistakes. Intendedly rational actors in large-scale institutions may respond collectively to the tasks they face adaptively. Wittman (1995:16) notes that "even if some individuals make incorrect choices, the law of large numbers is likely to yield the correct majority choice."

Many of these objections are quite fundamental—so much so that it seems impossible to develop a serious empirical theory of choice without taking them into consideration. They address both (*a*) the limitations of humans to comprehend and act on inputs from the environment and (*b*) the fundamental complexity of the environment, which is vastly underestimated in standard rational choice theories.

The Nature of the Decision Maker

Empirical objections to rational choice are so voluminous that they are, in effect, a laundry list of problems. The first set has to do with the nature of the decision maker.

SEARCH BEHAVIOR In general, people do not consider all aspects of a decision facing them. They must factor the decision to make it manageable, examining only relevant aspects. They do not undertake complete searches for information, and they ignore available information—especially if it is not relevant to the factors they have determined to characterize the structure of the problem.

SEARCH MUST INCLUDE BOTH ALTERNATIVES AND ATTRIBUTES Different physiological and psychological mechanisms probably underlie the search for attributes (which is equated in ordinary language with understanding a problem) and the search for alternatives (which involves the choice under a given decisional structure, design, or understanding) (Jones 1996).

CALCULATIONS People generally cannot perform the calculations necessary even for a reduced set of options in a decision-making situation. This is actually the least problematic limitation in decision making. They can, given time, write down and manipulate the numbers.

COGNITIVE ILLUSIONS AND FRAMING When identical options are described in different terms, people often shift their choices. For example, if a choice is described in terms of gains, it is often treated differently than if it is described in terms of losses. This shift demonstrates the concept of framing, developed by psychologists Daniel Kahneman and Amos Tversky. They claim that this tendency violates a major, if often unstated, assumption of rational choice—namely the axiom of invariance, which states that the "preference order between prospects should not depend on the manner in which they are described" (Kahneman & Tversky 1983:343). They bolster their claim with numerous convincing experiments indicating that decision makers tend to choose different alternatives when they are described in positive terms (for example, in terms of the number of lives saved with a vaccine) than when they are described in negative terms (the number of people who will die). Kahneman & Tversky (1983:343) state, "In their stubborn appeal, framing effects resemble perceptual illusions more than computational errors."

SELF-CONTROL People often seem to need to bind themselves in some way to establish self-control over their behavior in the future. A major mechanism for dealing with likely future lapses in self-control is to establish binding rules that prohibit the unwanted behavior. For example, Thaler (1991) has developed the notion of mental accounting to explain the tendency of people to separate categories of income and impose more constraints on some (investment income) than on others (a Christmas bonus). People also tend to treat gains differently from losses, applying different risk functions to them, essentially being more risk-adverse for gains than for losses (Kahneman & Tversky 1983, 1985).

INCOMMENSURATE ATTRIBUTES In multi-attribute situations, people often have severe difficulties in making the tradeoffs that look so simple in consumer choice theories. They tend to use a variety of shortcuts that avoid making the direct tradeoff.

DESIGN People have trouble figuring what factors are relevant to a given decision-making situation, and these framings are subject to radical shifts in a short period of time (Jones 1994).

UPDATING People are "incomplete Bayesians." In uncertain situations, they do not update their choices in light of incoming information about the probability of outcomes in the manner predicted by calculations from probability theory (Bayes' rule is the relevant yardstick) (Edwards 1968; Kahneman & Tversky 1983, 1985; Piattelli-Palmarini 1994). Some literature in political science suggests that voters update partisan attachments in the aggregate in a Bayesian fashion (Gerber & Green 1998). Intendedly rational voters would update (if not strictly according to Bayes' rule). Whether the law of large numbers acts to push incomplete Bayesian voters toward a closer approximation to Bayes' rule in the aggregate remains an open question.

IDENTIFICATION WITH MEANS In situations of repeated decision making, people often come to identify both cognitively and emotionally with the means, or subgoals, of a decision-making process. If they do, they are likely to become too conservative in shifting to a more effective means for solving a problem (March 1994). A scientist may, for example, become expert in a mode of analysis and apply it to all sorts of problems, even if the approach yields suboptimal results. The rational choice debate in political science has aptly illustrated the tendency to identify emotionally as well as cognitively with means (Green & Shapiro 1994).

A Note on Experimental Economics

In recent years, a vigorous experimental movement in economics has emerged. The methodology is direct: Derive a result from theoretical economics, set up a laboratory situation that is analogous to the real-world economic situation, and compare the behavior of subjects to the predicted behavior. These experiments have been criticized by practitioners of disciplines with much longer traditions of experimentation, such as psychology and biology; these criticisms are substantial. Perhaps most importantly, the economics experiments typically fail to study control groups (Green & Shapiro 1994:125–27). The justification is that the theoretical prediction serves as the comparison, or control.

In any case, two sets of findings have emerged. The first is that, in many situations that mimic real markets, in both animal and human experiments,

market incentives have a major effect on behavior (Kagel et al 1995). On the other hand, the maximization models do not predict behavior very well—and they fail to predict behavior just where bounded rationality should show through in decision making. To cite but one very important result, the laboratory studies of Kagel & Levin (1986) show that auctions, and particularly auctions with numerous bidders, produce aggressive bidding that results in negative profits—the "winner's curse." Overbidding afflicts experienced as well as inexperienced bidders.

The Nature of the Environment

In addition to the objections based on the nature of the decision maker, there are objections to rational choice theory that involve the nature of the environment.

AMBIGUITY AND UNCERTAINTY Proponents of limited rationality suggest that the environment is fundamentally more uncertain than is understood in prevailing choice models. Uncertainty, in rational choice models, means not knowing the probability of decisional consequences. In limited-rationality models, uncertainty also involves lack of knowledge of the attributes that characterize the problem (these are termed ill-structured problems). It can also involve ambiguity, which itself has two connotations. The first refers to situations in which the attributes are clear, but their relative importance is not (Jones 1996). The second, more fundamental ambiguity is one in which "alternative states are hazily defined or in which they have multiple meanings, simultaneously opposing interpretations" (March 1994:178).

Ambiguity and uncertainty in the environment feed back into characteristics of the decision maker. Preferences are desires about end states. In the rational choice model, people maximize the probabilities of achieving a desired state. But if end states are ambiguous, then our preferences must be ambiguous! If our preferences are ambiguous, then that mainstay of rational choice—fixed, transitive preferences—cannot hold (see March 1994:ch. 5, for an extended discussion).

REPEATED DECISIONS AND ENDS-MEANS CAUSAL CHAINS People never make decisions in isolation. They interact with others, who themselves have decision strategies. They must modify their goals in light of the social milieu in which they find themselves. Indeed, some analysts have argued that preferences should be viewed as fluid, not fixed, because of the necessity to be flexible in the face of changing circumstances. It is common for decisions to exist in complex ends-means causal chains (Simon 1983). In many problems, as we take one step down the path toward solution, we preclude other options and we open new opportunities. That is, problem solving is an ongoing process involv-

ing interaction with the environment, which changes the set of constraints and opportunities we face.

THE COST OF INFORMATION

A major (he says *the* major) contribution of Downs' *An Economic Theory of Democracy* was to introduce the notion that search behavior is subject to a rational calculus (Downs 1957, 1993). The more valuable the likely outcome of a decision, the more extensive one's search should be. Where the decision maker possesses limited information about alternatives, it would be rational for that decision maker to use shortcuts, such as ideology or party identification, as cues to action in order to save expensive search time. The addition of a cost-of-search function to the model of rationality, along with the understanding of the role of risk and uncertainty, are the major additions to our understanding of rational choice (see Becker 1976). Downs' notion has been developed extensively in political science, especially in the study of voter turnout (Ferejohn & Kuklinski 1990). Lupia & McCubbins (1998) have explored the notion of source effects in cuing voting direction as a rational action, and they provide experimental evidence.

In my opinion, information cost functions cannot save comprehensive rationality. First, no studies have yet been directed at the search process in realistic political situations. And, make no mistake about it, this is an important process assumption. One will need to show that decision makers explicitly and consciously substitute such considerations as party and ideology for seeking information on public policy proposals. If the shortcut is buried in the backwaters of habit and routine, only bounded rationality can be used to understand the phenomenon.

Second, a model of rationality including search costs fails to incorporate the tendency of people to identify with means (organizational identification). If people act out of organizational loyalty rather than self-centered calculation, then the model fails. This is, in principle, testable, but again the proponents of information-cost functions have not done the empirical work.

It is possible to include organizational loyalty as part of the utility function—that is, as a separate goal. Then it will be necessary to map the trade-offs made by the decision maker between, for example, party loyalty and a policy goal—just the kind of trade-offs that laboratory studies show that people accomplish poorly.

Third, information costs cannot explain many of the observations of organizational behavior and laboratory results. These include at least the following: (*a*) the tendency of people to think of risk differently when they are losing than when they are winning, (*b*) the "winner's curse" in auctions, and (*c*) the tendency of people to fail to act according to Bayes' rule in updating information.

INFORMATION THEORY AND INFORMATION PROCESSING

Nowhere do comprehensive and bounded rationality differ more than in the treatment of information. The transmission of information has always been an important component of politics, but it has received renewed emphasis in recent years via signaling theory. Rational-actor theories of decision making require no theory of decision makers, because all behavior is explained in terms of incentives incoming from the environment (Simon 1979). Similarly, rational-actor theories of information need a theory of signals and a theory of senders but have no need of a theory of receivers. In modern signaling theory, information is costly and noisy; the receiver wants the information and the sender may or may not have incentive to supply correct information. If the sender does transmit the information, the signal will reduce the variance (noise) affecting the receiver's view of the world.

In information processing, the receiver must attend to and interpret incoming information. Often, the problem for the receiver is not a lack of information but rather an overload. The scarce resource is not information; it is attention (Simon 1996b). In essence, one needs a theory of the receiver to understand his or her response to a signal. If the receiver's frame of reference is multidimensional, then the concept of noise reduction is not enough to explain the receiver's response. The sender of the information may also try to influence the relative importance of the attributes that structure the multidimensional frame of reference held by the receiver (Jones 1996). Framing effects in political communication stem primarily from the limited attention spans (short-term memories) of decision makers and the necessity to retrieve coded patterns from long-term memory (Iyengar 1990).

WHAT SHOULD REPLACE RATIONAL CHOICE?

The response of social scientists to the onslaught of empirical findings showing the failure of the rational model may be divided into two camps. The first camp continues to do business as usual, ignoring the demonstrated weaknesses of the underlying assumptions—for example, denying that assumptions ought to be subject to empirical test. Then there has been a tendency to "discover" incentives in the environment that must have been there to account for any observed deviations. Green & Shapiro (1994) refer to this as post-hoc theorizing and offer numerous examples.

The second camp has begun a research program of incorporating elements of bounded rationality into models of political and economic decision making. Behavioral economist Colin Camerer (1998:56) recommends replacement assumptions that allow economic agents

> to be impulsive and myopic, to lack self-control, to keep track of earning and spending in separate "mental accounting" categories, to care about the outcomes of others (both enviously and altruistically), to construct preferences from experience and observation, to sometimes misjudge probabilities, to take pleasure and pain from the difference between their economic state and some set of reference points, and so forth.

This challenge has been taken seriously. In international relations, Kahneman & Tversky's prospect theory has been used to understand foreign policy decision making (Farnham 1994, Levy 1997, Quattrone & Tversky 1988). I have shown elsewhere that shifts of attention among the attributes that structure a situation can yield discontinuous behavior in political institutions—even when actors maximize utility (Jones 1994). In the study of voting behavior, Hinich & Munger (1994) predict that "rigorous formal models may someday account for emotion, history, and the idiosyncrasies of human cognition." Economists have directly modeled economic phenomena, using selected assumptions based on bounded rationality (Sargent 1993). Financial economists have incorporated decision-making models based on heuristic shortcuts, emotion, and contagion to understand large jumps in the behavior of asset markets (Lux 1995, 1998). This approach, though promising, is feasible in the long run only if rigorous empirical tests of the new models are undertaken.

There is a third possibility, as yet unexplored in political science. That is to use the rational model to estimate what fully rational actors would do given the external situation. This is a possibility only when one understands the structure of the situation and the frame that would be used by rational actors—as is the case when cognitive psychologists study problem solving in set task environments. This approach causes us to consider explicitly the conditions under which bounded rationality will show through in structured decision-making situations. It is not as useful in fluid, nonstructured environments.

THE PROBLEM OF PROBLEM SPACES

Results from problem-solving experiments in psychology laboratories and from studies in artificial intelligence suggest that decision makers process information by applying operators in a problem space constructed to search for solutions (Newell 1990). In these experiments, the task environment is tightly specified, so that the investigator knows exactly the preferences (goals) of the subject—to solve the problem. Process-tracing methods allow the study of the steps that subjects take to solve the problem (Newell & Simon 1972). Results suggest that, as the time allocated to solve the problem increases, the demands of the task environment overwhelm the limitations imposed by human cognitive architecture. However, some facets of the underlying architecture continue to show through even in tightly specified task environments.

What if task environments are uncertain, ambiguous, or contradictory? Then the direct representation of the task environment in the problem space of the decision maker is not so evident. Considerable work in political science has been directed at the study of how policy makers understand the problems they face (Rochefort & Cobb 1994). Perhaps the major problem with the use of rational choice in political science has been the confident postulation of a problem definition that may or may not fit the problem definition held by the decision maker. Even when the goals of decision makers are clear and unambiguous (such as the postulate that legislators wish to be reelected), subgoals may not be at all clear.

BOUNDED RATIONALITY SHOWING THROUGH

If, however, the task environment can be specified tightly enough to predict rational responses from decision makers, then it becomes possible to compare observed behavior with that expected from rational predictions (Jones 1999). We have already seen that this is the technique commonly used by experimental economists. (Unfortunately, the economists have no alternate hypothesis to accept when the null expectation is rejected.) But it is also possible to extend the approach to structured institutional decision making, in which the incentives generated by the institution are understood well enough to model quantitatively.

The "efficient market thesis" provides a powerful example. In what economists call an informationally efficient market, the price of a stock tomorrow cannot be predicted from the price of a stock today. The reason, as Samuelson (1965:41) put the enigma, is this: "In competitive markets, there is a buyer for every seller. If one could be sure that a price will rise, it would have already risen."

There will be plenty of price movement in the stock market. But because all participants are fully rational, they will use up all of the available systematic information. That means they will fully value the stock. The next move of the stock cannot be predicted—it could be up or down. But, because the factors that will affect the price of the stock (after it has been bid up or down by investors based on systematic information) are random, the distribution of the changes in a stock's prices—hourly, daily, or yearly—follow a random walk through time (Fama 1965).

The simplest form of the random-walk hypothesis may be written as follows:

$$P_t = \mu + P_{t\text{-}1} + \varepsilon_t, \qquad \varepsilon_t \sim \text{IID}\ (0, \sigma^2). \qquad 1.$$

Here the price at t is a function of $t - 1$, a term μ (which assesses the long-run expected change, or drift, in the price series), and an error term that is assumed

to be independently and identically distributed with finite variance. In this circumstance, markets would follow a random walk with drift, and prices would be Normally distributed because of the central limit theorem.

The implication of the efficient-market thesis is that a stock market (or other asset market, such as bonds) will follow a random walk.[2] In a random walk, we cannot predict the next step from previous steps. If we define the (daily) returns in a stock price as the price on day two minus the price on day one, and we make a frequency distribution of these daily returns over a long period of time, this frequency distribution will approximate a Normal. The many factors that could affect the price of a stock (or a whole market), when added up, mostly cluster around the average return, with very few changes a long way (either up or down) from the average return.

So we have a clear prediction for the behavior of market returns. Unfortunately, the evidence is not supportive. Asset market returns are invariably leptokurtic; they have slender peaks and fat tails in comparison to the Normal (see Figure 1). They are, in effect, subject to bubbles and crashes. Bubbles and crashes are not to be expected in the efficient-market thesis because sophisticated traders will be able to make money on the underlying dependence in the error-generating process (Fama 1965:38). This, in effect, would restore Normality to the price series.

Note that the observed leptokurtic distributions are not bizarre deviations from the Normal distribution, and that we have standard techniques for estimating the "excess kurtosis problem" (Lux 1998:149). Intendedly rational actors may deviate from fully rational actors, but the deviation will be attenuated in well-functioning institutions. The fat-tailed, slender-peaked distribution is what we would expect if we thought that market participants were intendedly rational. Markets would be affected by such factors as (*a*) non-Bayesian updating behavior in the face of incoming information (basically underreacting to some information and overreacting to other information, depending on the context), (*b*) contagion, and (*c*) emotion (Lux 1998). This disproportionate reaction to information would not be fully compensated by sophisticated traders, since they also would be subject to the same cognitive and emotional forces.

It may be objected that external shocks surprising to everyone are responsible for the observed fat tails. Experimental economists, however, have directly observed bubbles and crashes (the fat tails) in their toy economies (Smith et al 1988). The experimentalists have also directly observed leptokurtic distributions in simulated markets (Plott & Sunder 1982). It is unlikely that what occurs endogenously in the lab would be explained exogenously in real markets.

[2]The initial random-walk hypothesis has been supplemented by more sophisticated models of random processes. The theoretical justification and implications remain similar. See Campbell et al (1997) for a discussion.

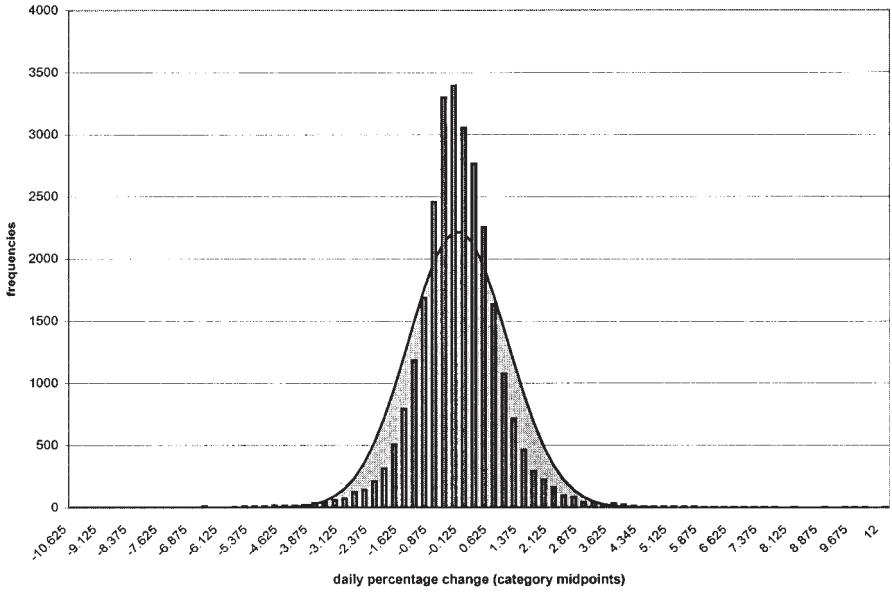

Figure 1 Frequency distribution of percentage changes in the Dow-Jones Industrial Average, 1896–1996, compared with a normal distribution with similar mean and variance.

We may use this approach in the study of politics. In the study of elections, there has been considerable debate about realignments. Electoral realignments imply leptokurtic distributions. If we were to plot election-to-election changes over a long period of time, we would expect to see most elections cluster around the center, with very little change in the pattern of standing allegiances to the parties for most elections. Once in a while, however, a major change would occur, falling in the fat tails. Very few cases would fall in the shoulders, or wings, of the distribution.

There is, however, a second hypothesis. In this approach, parties are the creations of ambitious, election-driven politicians (Aldrich 1995). Politicians play the part of entrepreneurs in market economies, immediately responding to the preferences of voters for "packages" of public policies. This suggests a relatively efficient response to information because of the activities of entrepreneurial politicians. Elections, under the hypothesis that elections are relatively informationally efficient, would have output distributions similar to the stock market.

Nardulli (1995) has produced a phenomenally complete analysis of presidential elections since 1824 at the county level. He finds scant evidence of

national realignments but points to a series of "rolling realignments" that are regionally based.

Nardulli's data may be plotted in a frequency distribution, similar to the stock market data discussed above. Figure 2 plots 110,000 observations on election margin swings in US counties for presidential elections from 1824 through 1992. On the one hand, the distribution is leptokurtic—the fat tails and slender peaks are in evidence. On the other hand, the distribution is no more leptokurtic than the US stock market. It would seem that bounded rationality—not sweeping realignments or a fully rational interaction between voters and politicians—best characterizes the data.

Another way to look at distributional data is to subtract the observed relative frequency of categories from expected (based on the Normal). Figure 3 does this for Nardulli's election data. The graph makes clear how the election data deviates from what is expected based on the Normal. Specifically, the graph shows an excess of cases clustered around the central peak and in the tails of the distribution. There are too few cases in the shoulders of the distribution, in comparison to the Normal. This means that a great many elections are

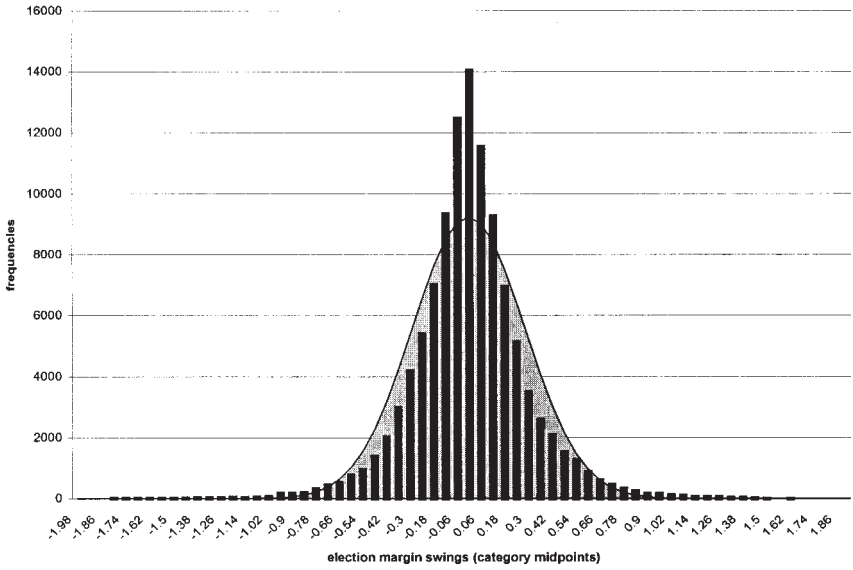

Figure 2 Frequency distribution of election margin changes in county-level presidential elections, 1824–1992, compared with a normal distribution with similar mean and variance. Election margin swing is the Democratic proportion of the two-party vote minus the Republican portion. Compiled from P Nardulli, personal communication.

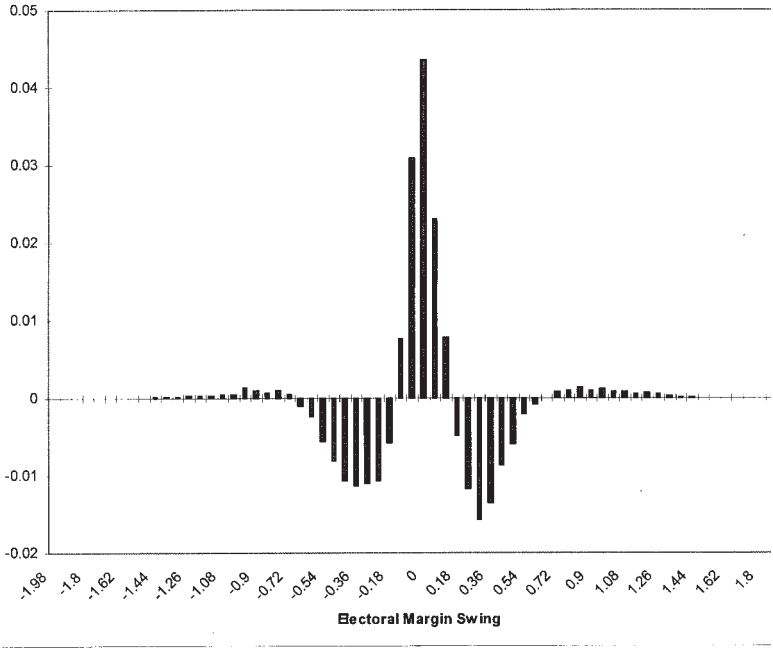

Figure 3 Observed-expected election data. Compiled from P Nardulli, personal communication.

incremental changes from the previous election, but a few exhibit considerable punctuations. There are too few modest changes in the distribution (again with the Normal as a standard).

Although I cannot make the full argument here, I note that this is the kind of distribution one would expect, given a set of intendedly rational actors subject to certain cognitive limits—particularly limits on attention—facing a world in which incoming information is approximately normally distributed.[3]

SUBSTANTIVE LIMITS ON RATIONALITY: OVERCOOPERATION

Riker's *The Theory of Political Coalitions* (1963) is an elegant analysis of political behavior in formal group decision-making situations, such as committees or legislative bodies. In many situations, committees must decide how to share a divisible good (one that can be broken up in any number of ways). The

[3]If decision makers are modeled based on a power function and if information is normally distributed, then the outcome distribution will be exponential. Substantial reasons exist for using a power function as a first approximation for bounded decision makers in institutional settings. Both elections and stock markets are exponential distributions (Jones 1999).

good could be streetlights that must be allocated to neighborhoods by a city council or highway projects to congressional districts or dollars to any number of worthy projects. Riker (1963) assumed that each member of the decision-making body had one vote and that a plurality was necessary for a proposal to win.

If decision makers are rational in such situations, they will form minimum winning coalitions, Riker reasoned. That is, rational legislators will find a way to divide the good up in such a way that the winning coalition will share all of the benefits while totally excluding the losers. In that way, each member of the winning coalition will maximize his or her part of the spoils. Any sharing with the minority will dilute the benefits gained by the majority.

Many political scientists spent a great amount of effort searching for the predicted minimum winning coalition, with mixed results at best. In the real world of democratic politics and government, there seemed to be too much cooperation. Moreover, participants often seemed to cooperate with the "wrong" others. In an early study of coalition formation in parliamentary democracies, Axelrod (1970; see also 1997) showed that parties tended to form coalitions based on ideological similarity rather than on the size of the governing coalition.

Why would rational actors seemingly overcooperate? Political scientists uncovered all sorts of reasons that minimum winning coalitions might not form on any given vote in a committee or legislature. The vote might not be on something divisible. It might be on something that was not excludable from the minority, as would be the case in voting for an increase in social security benefits in Congress. Social security recipients can live in any district. Political institutions are sometimes set up to require supermajorities, in effect forcing more cooperation than might be predicted from preferences alone.

Even on divisible goods, rational legislators might not form minimum coalitions because of what Axelrod (1984) called "the shadow of the future." In the language of game theory, a vote is not a one-shot game. Axelrod began studying cooperative (or coalitional) behavior in computer simulations of games in which rational players repeated play. Cooperative behavior indeed emerged in such situations, whereas it did not in a single play. If people know that they will be interacting with others over a long period of play, then they are more likely to use strategies that involve offers of cooperation in the rational hope that such offers will be reciprocated, making both parties better off. This sort of cooperation is rationally based—it is a reasoned response to the task environment.

Divide the Dollar

In a laboratory game that mimics political coalition formation, a subject is asked to divide a set amount among several players. The players can reject the

offer, but they cannot modify it. This game, or a variant of it, has been played in laboratories many times. The inevitable result is that leaders give away too much of the spoils. They generally do not divide the results equally, but they do share too much. Remember that this is a one-shot game; there is no repeated play. The leader will not see these people again. (Repeated play and knowing the participants cause leaders to share even more.)

In contrast to the process limitations discussed earlier in this paper, however, the limitations here are not that the leader was unable to do the necessary calculations or attend to the relevant factors. In the simplest form of this game, the ultimatum game, consisting of only two players (proposer and responder), the same results hold. The typical offer to the responder is 30–40% of the total. Even when the responder cannot reject the offer (the "dictator" game), the proposer offers more than predicted.(See Camerer & Thaler 1995 for a nontechnical review of these findings.)

These games, which are not affected by the size of the reward but do seem to be affected somewhat by cultural differences, illustrate major deviations from strict self-centered rationality. Heuristic decision making took over, but this was not a heuristic that served as an informational shortcut. I term these and other similar heuristics substantive, because the heuristic directly affects the decisional outcome.

CONCLUSIONS: RATIONAL CHOICE AS TASK ENVIRONMENT

Political decision makers are invariably intendedly rational; they are goal-directed and intend to pursue those goals rationally. They do not always succeed. I have detailed an approach to political choice that has three components: (*a*) the task environment, (*b*) the problem space constructed by the decision maker, and (*c*) the limits imposed by the cognitive/emotional architecture of human decision makers.

The behavior of a fully rational decision maker would be completely determined by the task environment. If we know the environment and the goals of the decision maker, then we may deduce the decision maker's actions. If, however, the decision maker intends to be rational but may fail, then we will need to know something about the cognitive and emotional architecture of the decision maker.

This conception of decision making leads to two important hypotheses:

1. In relatively fixed task environments, such as asset markets and elections, observed behavior (B) of actors may be divided into two mutually exclusive and exhaustive categories: rational goal attainment (G) and limited rationality (L). This leads to the fundamental equation for fixed task environments, $B = G + L$.

2. In uncertain, ambiguous, or contradictory task environments, behavior is a function of goals, processing limits, and the connection between the decision maker's problem space and the task environment (objectively characterized). In this far more complex situation, problem-space representations may interact nonlinearly with goals and processing limits.

The strategy I have suggested here is to divide these two separate situations for analytical purposes and treat them separately. In relatively fixed task environments, we should be able to divide behavior into adaptive, goal-oriented behavior and behavior that is a consequence of processing limits, and we should be able to measure the deviation. I have offered a first cut at such a strategy above for outcome distributions from structured institutional settings. Having so divided outcome behavior, we might want to reexamine the internal workings of such institutions—in effect, to trace the processes that lead to the outcomes of interest.

Literature Cited

Aldrich J. 1995. *Why Parties?* Chicago: Univ. Chicago Press

Anton T. 1966. *The Politics of State Expenditures in Illinois.* Urbana: Univ. Ill. Press

Axelrod R. 1970. *Conflict of Interest.* Chicago: Markham

Axelrod R. 1984. *The Evolution of Cooperation.* New York: Basic Books

Axelrod R. 1997. *The Complexity of Cooperation.* Princeton, NJ: Princeton Univ. Press

Baumgartner FR, Jones BD. 1993. *Agendas and Instability in American Politics.* Chicago: Univ. Chicago Press

Becker GS. 1976. *The Economic Approach to Human Behavior.* Chicago: Univ. Chicago Press

Camerer CF. 1998. Behavioral economics and nonrational organizational decision making. See Halpern & Stern 1998a, pp. 53–77

Camerer CF, Thaler RF. 1995. Ultimatums, dictators and manners. *J. Econ. Perspect.* 9:209–19

Campbell JY, Lo AW, MacKinlay AC. 1997. *The Econometrics of Financial Markets*. Princeton, NJ: Princeton Univ. Press

Chisholm D. 1995. Problem-solving and institutional design. *J. Public Admin. Res. Theory* 5:451–91

Cobb R, Elder C. 1972. *Participation in American Politics.* Baltimore, MD: John Hopkins Univ. Press

Cohen M, March JG, Olsen J. 1972. A garbage can model of organizational choice. *Admin. Sci. Q.* 17:1–25

Crecine JP. 1969. *Government Problem-Solving: a Computer Simulation of Government Budgeting.* Chicago: Rand McNally

Cyert RM, March JG. 1963. *A Behavioral Theory of the Firm.* Englewood Cliffs, NJ: Prentice Hall

Davis OA, Dempster MAH, Wildavsky A. 1966. A theory of the budget process. *Am. Polit. Sci. Rev.* 60:529–47

Davis OA, Dempster MAH, Wildavsky A. 1974. Towards a predictive theory of government expenditure: US domestic appropriations. *Br. J. Polit. Sci.* 4:419–52

Dodd LC. 1994. Political learning and political change: understanding development

across time. In *The Dynamics of American Politics*, ed. LC Dodd, C Jilson, pp. 331–64. Boulder, CO: Westview
Downs A. 1957. *An Economic Theory of Democracy*. New York: Harper & Row
Downs A. 1993. The origins of an economic theory of democracy. In *Information, Participation, and Choice*, ed. B Groffman, pp. 197–99. Ann Arbor: Univ. Mich. Press
Edwards W. 1968. Conservatism in human information processing. See Kleinmuntz 1968, pp. 17–52
Fama E. 1965. The behavior of stock market prices. *J. Business* 12:34–105
Farnham B. 1994. *Avoiding Losses, Taking Risks*. Ann Arbor: Univ. Mich. Press
Fenno R. 1966. *The Power of the Purse: Appropriations Politics in Congress*. Boston: Little, Brown
Ferejohn J, Kuklinski J, eds. 1990. *Information and the Democratic Process*. Urbana: Univ. Ill. Press
Gerber A, Green DP. 1998. Rational learning and partisan attitudes. *Am. J. Polit. Sci.* 42:794–818
Gist JR. 1982. "Stability" and "competition" in budgetary theory. *Am. Polit. Sci. Rev.* 76:859–72
Green DP, Shapiro I. 1994. *Pathologies of Rational Choice Theory*. New Haven, CT: Yale Univ. Press
Halpern J, Stern R. 1998a. *Debating Rationality*. Ithaca, NY: Cornell Univ. Press
Halpern J, Stern R. 1998b. Beneath the social science debate: economic and social notions of rationality. See Halpern & Stern 1998, pp. 1–20
Hinich M, Munger M. 1994. *Ideology and the Theory of Political Choice*. Ann Arbor: Univ. Mich. Press
Iyengar S. 1990. Shortcuts to political knowledge: selective attention and the accessibility bias. In *Information and the Democratic Process*, ed. J Ferejohn, J Kuklinski, pp. 160–85. Urbana: Univ. Ill. Press
Jones BD. 1980. *Service Delivery in the City: Citizen Demand and Bureaucratic Rules*. New York: Longman
Jones BD. 1985. *Governing Buildings and Building Government. A New Perspective on the Old Party*. University: Univ. Ala. Press
Jones BD. 1994. *Reconceiving Decision-Making in Democratic Politics*. Chicago: Univ. Chicago Press
Jones BD. 1996. *Attributes, alternatives, and the flow of ideas: information processing in politics.* Presented at Annu. Meet. Am. Polit. Sci. Assoc., San Francisco
Jones BD. 1999. Bounded rationality, political institutions, and the analysis of outcomes. In *Competition and Cooperation: Conversations with Nobelists about Economics and Political Science*, ed. J Alt, M Levi, E Ostrom. New York: Russell Sage Found. In press
Jones BD, Bachelor L. 1994. *The Sustaining Hand*. Lawrence: Univ. Press Kans. 2nd ed.
Jones BD, Baumgartner FR, True JL. 1996. *The shape of change: punctuations and stability in U.S. budgeting, 1947–94.* Presented at Annu. Meet. Midwest Polit. Sci. Assoc., Chicago
Jones BD, Baumgartner FR, True JL. 1998. Policy punctuations: US budget authority, 1947–1995. *J. Polit.* 60:1–33
Jones BD, True JL, Baumgartner FR. 1997. Does incrementalism stem from political consensus or from international gridlock? *Am. J. Polit. Sci.* 41:1319–39
Kagel J, Battalio R, Green L. 1995. *Economic Choice Theory*. New York: Cambridge Univ. Press
Kagel J, Levin D. 1986. The winner's curse and public information in common value auctions. *Am. Econ. Rev.* 76:894–920
Kahneman D, Tversky A. 1983. Choices, values, and frames. *Am. Psychol.* 39:341–50
Kahneman D, Tversky A. 1985. Prospect theory: an analysis of decision-making under risk. *Econometrica* 47:263–91
Kingdon J. 1996. *Agendas, Alternatives, and Public Policies*. Boston: Little, Brown. 2nd ed.
Kleinmuntz B, ed. 1968. *Formal Representation of Human Judgment*. New York: Wiley
Levy F. 1974. *Urban Outcomes: Schools, Streets, and Libraries*. Berkeley: Univ. Calif. Press
Levy JS. 1997. Prospect theory, rational choice, and international relations. *Int. Stud. Q.* 41:87–112
Lindblom CE. 1959. The science of muddling through. *Public Admin. Rev.* 19:79–88
Lounamaa PH, March JG. 1985. Adaptive coordination of a learning team. *Manage. Sci.* 33:107–23
Lupia A, McCubbins M. 1998. *The Democratic Dilemma.* Cambridge, UK: Cambridge Univ. Press
Lux T. 1995. Herd behaviour, bubbles and crashes. *Econ. J.* 105:881–96
Lux T. 1998. The socio-economic dynamics of speculative markets. *J. Econ. Behav. Org.* 33:145–65
March J. 1978. Bounded rationality, ambiguity, and the engineering of choice. *Bell J. Econ.* 9:578–608
March JG. 1994. *A Primer on Decision-Making*. New York: Free

March JG, Olsen JP. 1989. *Rediscovering Institutions.* New York: Free
March JG, Simon HA. 1958. *Organizations.* New York: Wiley
Marcus GE, McKuen M. 1993. Anxiety, enthusiasm, and the vote. *Am. Polit. Sci. Rev.* 87:688–701
Meltsner A. 1971. *The Politics of City Revenue.* Berkeley: Univ. Calif. Press
Mladenka K. 1978. Rules, service equity, and distributional decisions. *Soc. Sci. Q.* 59: 192–202
Nardulli P. 1995. The concept of a critical realignment, electoral behavior, and political change. *Am. Polit. Sci. Rev.* 89:10–22
Newell A. 1968. Judgment and its representation. See Kleinmuntz 1968, pp. 1–16
Newell A. 1990. *Unified Theories of Cognition.* Cambridge, MA: Harvard Univ. Press
Newell A, Simon HA. 1972. *Human Problem Solving.* Englewood Cliffs, NJ: Prentice Hall
Ostrom E. 1990. *Governing the Commons.* Cambridge, UK: Cambridge Univ. Press
Padgett JF. 1980. Bounded rationality in budgetary research. *Am. Polit. Sci. Rev.* 74: 354–72
Padgett JF. 1981. Hierarchy and ecological control in federal budgetary decision making. *Am. J. Sociol.* 87:75–128
Piattelli-Palmarini M. 1994. *Inevitable Illusions.* New York: Wiley
Plott CR, Sunder S. 1982. Efficiency of experimental security markets with insider trading. *J. Polit. Econ.* 90:663–98
Quattrone GA, Tversky A. 1988. Contrasting rational and psychological analyses of political choice. *Am. Polit. Sci. Rev.* 3: 719–36
Riker W. 1963. *The Theory of Political Coalitions.* New York: Yale Univ. Press
Rochefort D, Cobb R. 1994. *The Politics of Problem Definition.* Lawrence: Univ. Press of Kansas
Sabatier P, Jenkins-Smith H. 1993. *Policy Change and Learning*. Boulder, CO: Westview
Samuelson P. 1965. Proof that properly anticipated prices fluctuate randomly. *Ind. Manage. Rev.* 6:41–49
Sargent T. 1993. *Bounded Rationality in Macroeconomics.* Oxford, UK: Oxford Univ. Press
Simon HA. 1947. *Administrative Behavior.* New York: Macmillan
Simon HA. 1957. *Models of Man.* New York: Wiley
Simon HA. 1979. Rational decision-making in business organizations. *Am. Econ. Rev.* 69: 495–501
Simon HA. 1983. *Reason in Human Affairs.* Stanford, CA: Stanford Univ. Press
Simon HA. 1985. Human nature in politics: the dialogue of psychology with political science. *Am. Polit. Sci. Rev.* 79:293–304
Simon HA. 1995. Rationality in political behavior. *Polit. Psychol.* 16:45–61
Simon HA. 1996a. *Models of My Life, MIT Edition*. Cambridge, MA: MIT Press
Simon HA. 1996b. *The Sciences of the Artificial*. Cambridge, MA: MIT Press. 3rd ed.
Simon HA. 1999. The potlatch between political science and economics. In *Competition and Cooperation: Conversations with Nobelists about Economics and Political Science*, ed. J Alt, M Levi, E Ostrom. Cambridge, UK: Cambridge Univ. Press
Slovak P. 1990. Choice. In *Thinking: An Invitation to Cognitive Science*, ed. DN Osherson, EE Smith, 3:89–116. Cambridge, MA: MIT Press
Smith VL, Suchanek GL, Williams AW. 1988. Bubbles, crashes, and endogenous expectations in experimental spot asset markets. *Econometrica* 56:1119–51
Sniderman PM, Brody RA, Tetlock PE. 1991. *Reasoning and Choice: Explorations in Political Psychology.* Cambridge, UK: Cambridge Univ. Press
Thaler RH. 1991. *Quasi Rational Economics*. New York: Russell Sage Found.
Wanat J. 1974. Bases of budgetary incrementalism. *Am. Polit. Sci. Rev.* 68:1221–28
Wildavsky A. 1964. *The Politics of the Budgetary Process.* Boston: Little, Brown
Wittman D. 1995. *The Myth of Democratic Failure.* Chicago: Univ. Chicago Press