filling in missing pieces of the value chain, or by strategic considerations to gain access to new knowledge, network ties become admission tickets to high-velocity races. Connectivity to an interorganizational network and competence at managing collaborations have become key drivers of the new logic of organizing.

The growth of new organizational forms is driven by divergent factors and pursued in a different manner by a wide array of organizations. Larger organizations are making their boundaries more permeable in order to procure key components or critical R&D. Subcontracting and outsourcing are steps taken to reduce fixed overheads. Organizations cooperate with ostensible competitors in order to take on projects too risky or challenging for one entity to pursue alone. Clusters of small organizations collaborate, cohering into a production network to create what no single small entity could on its own. In sum, organizations are coming to resemble a network of treaties because these multistranded relationships encourage learning from a broad array of collaborators and promote experimentation with new methods, while at the same time reducing the cost of expensive commitments. These developments do not mean that competition is rendered moot, instead the success of organizations is linked to the nature and depth of their ties to organizations in diverse fields.

See also: Authority: Delegation; Bureaucracy and Bureaucratization; Management: General; Network Analysis; Organization: Overview; Organizational Decision Making; Organizations: Authority and Power; Organizations, Sociology of

Bibliography

Chandler A D 1990 *Scale and Scope: The Dynamics of Industrial Capitalism*. Harvard University Press, Cambridge, MA

Coase R 1937 The nature of the firm. *Economica* **4**: 386–405

Cohen W, Levinthal D 1990 Absorptive capacity: a new perspective on learning and innovation. *Administrative Science Quarterly* **35**: 128–52

Doz Y L, Hamel G 1998 *Alliance Advantage: The Art of Creating Value Through Partnering*. Harvard Business School Press, Boston

Dyer J 1996 Specialized supplier networks as a source of competitive advantage: Evidence from the auto industry. *Strategic Management Journal* **17**(4): 271–92

Gomes-Casseres B 1996 *The Alliance Revolution: The New Shape of Business Rivalry*. Harvard University Press, Cambridge, MA

Helper S, MacDuffie J P, Sabel C 2000 Pragmatic collaborations: advancing knowledge while controlling opportunism. *Industrial and Corporate Change* **9**(3): 443–88

Hounshell D A 1984 *From the American System to Mass Production 1800–1932: The Development of Manufacturing Technology in the US*. Johns Hopkins University Press, Baltimore, MD

Merrill S A, Cooper R S 1999 Trends in industrial research and development: evidence from national data sources. In: *Securing America's Industrial Strength*. National Research Council Board on Science, Technology and Economic Policy. National Academy Press, Washington, DC, pp. 99–116

Mowery D C, Nelson R R (eds.) 1999 *Sources of Industrial Leadership: Studies of Seven Industries*. Cambridge University Press, New York

Pfeffer J, Salancik G R 1978 *The External Control of Organizations: A Resource Dependence Perspective*. Harper and Row Publishers, New York

Porter M E 1980 *Competitive Strategy: Techniques for Analyzing Industries and Competitors*. Free Press, New York

Powell W W 1996 Inter-organizational collaboration in the biotechnology industry. *Journal of Institutional and Theoretical Economics* **152**: 197–215

Powell W W, Koput K, Smith-Doerr L 1996 Interorganizational collaboration and the locus of innovation: Networks of learning in biotechnology. *Administrative Science Quarterly* **41**: 116–45

Powell W W, Smith-Doerr L 1994 Networks and economic life. In: Smelser N J, Swedberg R (eds.) *Handbook of Economic Sociology*. Princeton University Press, Princeton, NJ, pp. 368–402

Rosenbloom R S, Spencer W J 1996 The transformation of industrial research. *Issues in Science and Technology* **12**(3): 68–74

Sabel C F 1994 Learning by monitoring: the institutions of economic development. In: Smelser N J, Swedberg R (eds.) *Handbook of Economic Sociology*. Russell Sage Foundation, New York, pp. 137–65

Stark D 1999 Heterarchy: distributing authority and organizing diversity. In: Clippinger J H III (ed.) *The Biology of Business*. Jossey-Bass, San Francisco, pp. 153–80

Williamson O E 1985 *The Economic Institutions of Capitalism*. Collier Macmillan, London

W. W. Powell

Bounded and Costly Rationality

Some kind of model of rational decision making is at the base of most current economic analysis, especially in microeconomics and the economics of organization. Such models typically assume that economic decision makers have sufficient cognitive capacities to solve the problems they face, and have preferences and beliefs that are consistent in a rather strong sense. In particular, it is typically assumed that decision makers (a) do not make logical errors, and can solve any relevant mathematical problems; (b) can process all available information in the time required (including computation, storage, and retrieval); and (c) have an adequate understanding of the decision problems they face, which includes having precise beliefs about the relevant uncertainties and precise preferences among the various consequences of their actions. However, this model has been criticized as inadequate from both normative and descriptive viewpoints. The various strands of this critical movement form the topic known as 'bounded rationality.' This article sketches the historical roots and some current developments of this movement, distinguishing between attempts to extend the standard models and the need for more radical departures.

Unease with mainstream models of *homo economicus* had already been voiced by J. M. Clark at the beginning of the twentieth century, but the work of Jacob Marschak and Herbert Simon provided the stimuli for a more intense level of activity. The term 'bounded rationality' was coined by Simon: 'Theories that incorporate constraints on the information-processing capabilities of the actor may be called *theories of bounded rationality*' (Simon 1972, p. 162). (For a review of empirical evidence of bounded rationality, see Conlisk 1996.)

The current mainstream theory of rational individual decision making in the face of uncertainty was elaborated by Savage (1954). This will here be called the 'Savage Paradigm,' and will be the main starting point of this article. Extensions of this theory to describe rational strategic behavior in multiperson situations are the subject of the theory of games (see below). The notion of optimizing is central to all these models of rationality.

The general concept of bounded rationality covers the two rather different approaches of Marschak and Simon. Marschak emphasized that the Savage Paradigm could be extended to take account of costs and constraints associated with information acquisition and processing in organizations, without abandoning the notion of optimizing behavior. This approach will here be called 'costly rationality,' and is elaborated in Sect. 4.

Simon was more concerned with behavior that could not so readily be interpreted, if at all, as optimizing. However, in some of his publications he apparently considered bounded rationality to be a broader concept, subsuming costly rationality as a particular case. The narrower concept will here be called 'truly bounded rationality' (Sect. 5).

1. Uncertainty

Discussions of rational decision making—unbounded and otherwise—have been closely tied to uncertainty. The very beginnings of formal probability theory were in part stimulated by questions of how to act rationally in playing games of chance with cards and dice. During the first half of the twentieth century a number of alternative views were developed concerning the nature of uncertainty, the possibly of different kinds of uncertainty, and whether, or in what circumstances, it could be measured (Arrow 1951, Savage 1954, Chap. 4). One could be uncertain about natural events, the consequences of action, the laws of nature, or the truth of mathematical propositions. There was general (but not universal) agreement that, if uncertainty could be measured (quantified), then that quantification should obey the mathematical laws of probability. However, the 'frequentist school' reserved the legitimacy of probabilistic reasoning for 'experiments' (planned or naturally occurring) that were repeated indefinitely under identical conditions. The 'personalist school,' which included a diverse set of methodologies, argued that the concept of probability was applicable to events outside the frequentist realm. Some personalists went so far as to deny that the frequentist view could be applied meaningfully to any events at all, i.e., all probability judgments were 'personal.' (See the accounts of Arrow 1951 and Savage 1954 of the work of Ramsey, Keynes, Jeffries, Carnap, and De Finetti. In some sense, the personalist view might also be ascribed to earlier authors, such as Bayes and Laplace.) The personalist view was given a solid foundation by Savage (1954).

Central to the development of thinking about uncertainty was the simple idea that uncertainty about the consequences of an action could (or should) be traced to uncertainty about the 'state of the world' in which the action would be taken. An essential feature of the concept of the 'state of the world' is that it is beyond the control of the decision maker in question. A further clarification was provided by the theory of games (put forward by J. von Neumann and O. Morgenstern in 1944, and further elaborated by J. Nash, J. Harsanyi, R. Selten, and others). In multiperson decision-making situations in which the participants have conflicting goals, this theory distinguishes between two aspects of the state of the world from the point of view of any single decision maker, namely (a) the 'state of Nature,' which is beyond the control of any of the persons involved; and (b) the actions of the other persons, the latter being called 'strategic uncertainty.' (For material on the theory of games, especially noncooperative games, see *Game Theory*; *Game Theory: Noncooperative Games*; *Game Theory and its Relation to Bayesian Theory*)

The remainder of this article concentrates on a critique of the theory of rational behavior as it is applied to single-person decision making, or to multiperson situations in which the persons do not have conflicting goals. (For discussions of bounded rationality in a game-theory context, see Rubinstein 1998 and Radner 1997.)

2. The Savage Paradigm

A sketch of the Savage Paradigm is needed here in order to understand the notions of costly and bounded rationality. (For a systematic treatment, see *Utility and Subjective Probability: Contemporary Theories*; *Utility and Subjective Probability: Empirical Studies*.) The essential building blocks of the model are (a) a set of alternative states of the world, or simply states, which are beyond the decision-maker's control; (b) a set of alternative actions available to the decision maker, or as Savage calls them, 'acts'; and (c) a set of alternative consequences. An act determines which consequence will be realized in each state (of the world). Hence a parsimonious way to think about an act is that it is a function from states to consequences.

The decision maker (DM) is assumed to have preferences among acts. These preferences reflect both the DM's beliefs about the relative likelihood of the different states, and the DM's tastes with regard to consequences. A few axioms about the independence of beliefs from consequences enable one to impute to the DM two scales: (a) a probability measure on the set of states, reflecting the DM's beliefs; and (b) a utility scale on the set of consequences, reflecting the DM's tastes. Using these two scales, one can calculate an expected utility for each act, in the usual way, since an act associates a consequence with each state. Thus, for each state, one calculates the product of its probability times the utility of the associated consequence, and then adds all of the products to obtain the expected utility of the act. One proves that expected utility represents the DM's preferences among acts in the following sense: *the DM prefers one act to another if and only if the first act has a higher expected utility*. (This theorem is sometimes called the expected utility hypothesis.) The rational DM is assumed (or advised) to choose an act that is most preferred among the available ones, i.e., has the highest expected utility; this is the assumption of optimization.

The simplicity of this formulation hides a wealth of possible interpretations and potential complexities. First, the axioms of the theory enable the DM to infer preferences among complicated acts from those among simpler ones. Nevertheless, the required computations may be quite onerous, even with the aid of a computer.

Second, if the decision problem has any dynamic aspects, then states can be quite complex. In fact, a full description of any particular state will typically require a full description of the entire history of those features of the DM's environment that are relevant to the decision problem at hand.

Third, the description of the set of available acts reveals—if only implicitly—the opportunities for the DM to acquire information about the state of the world and react to it. The laws of conditional probability then determine how the DM should learn from observation and experience. In fact, this is what gives the Savage Paradigm its real 'bite.' An act that describes how the DM acquires information and reacts to it dynamically is sometimes called a strategy (plan, policy). In a model of a sequential decision problem, the space of available strategies can, of course, be enormous and complex. This observation will be a dominant motif in what follows.

(The formula for learning from observation is sometimes called 'Bayes's theorem,' after the eighteenth-century author Reverend Thomas Bayes. Hence, the method of inference prescribed by the Savage Paradigm is called 'Bayesian learning.')

3. *The Simon Critique*

As Herbert Simon emphasized in his work, the cognitive activities required by the Savage Paradigm (and its related precursors) are far beyond the capabilities of human decision makers, or even modern human/computer systems, except with regard to the simplest decision problems. This led Simon and his colleagues (especially at Carnegie-Mellon University) to investigate models of human decision making that are more realistic from the point of view of cognitive demands, and yet do not entirely abandon the notion of rationality. This research also had an impact on the emerging field of artificial intelligence (see Simon 1972, 1981, and references therein). Savage himself was aware of the problem of bounded rationality, but he nevertheless felt that his model was a useful one for thinking about rational decision making (Savage 1954, pp. 16, 17).

4. *Costly Rationality and the Extended Savage Paradigm*

As just sketched in the previous section, the Savage Paradigm does not appear to take account explicitly of the costs of decision making. However, nothing prevents the DM from incorporating into the description of the consequences of an act the costs—in terms of resources used—of implementing the corresponding actions. The costly activities involved in decision making include:

observation and experimentation;
information processing, i.e., computation;
memory; and
communication.

The last category may be important when the decision-making process is undertaken by a team of individuals.

If the resources used by these decision-making activities are limited, then those limits may impose binding constraints on the activities themselves—constraints that must be taken into account in the DM's optimization problem. If the constraints are on the rate of resource use per unit time, then more extensive decision-making activities may cause delays in the implementation of the eventual decisions. To the extent that a delay lowers the effectiveness of a decision (e.g., by making it more obsolete), one may think of delay as an 'indirect cost.' Extending the Savage Paradigm to incorporate the costs of decision making may in some cases be natural, and in other cases problematic. The first class of cases is here called 'costly rationality.'

The notion that observation is costly was implicit in the Neyman–Pearson theory of hypothesis testing, and was made explicit by Abraham Wald in his pioneering studies of sequential statistical procedures (see Wald 1950 for an influential codification of his general approach). The cost of observation also figures in more classical (nonsequential) statistical problems such as the design of sample surveys and agricultural experiments. Given some model of the costs of observation, the DM chooses the kind and amount of observation, optimally balancing the expected benefits of additional observations against their costs. Such decision problems fit naturally into the Savage Paradigm, although taking account of these costs typically complicates the analysis. For example, in the case of clinical trials and similar problems, the calculation of optimal policies quickly becomes computationally intractable for many problems of realistic size.

Even after the information has been collected, it still must be further processed to produce the required decisions. This information-processing task may be quite demanding. Examples include (a) computing a weekly payroll; (b) scheduling many jobs on many machines; (c) managing multiproduct inventories at many locations; and (d) project selection and capital budgeting in a large firm. Such tasks are typically too complex to be handled by a single person, even with the aid of modern computers. In such circumstances the required processing of the information is decentralized among many persons in the organization. The theoretical study of decentralized decision making in an organization whose members have identical goals was introduced by J. Marschak in the theory of teams (Marschak and Radner 1972).

Computer science has provided a number of useful models of information processing by both computers and humans, and the decentralization of information processing in human organizations finds its counterpart in the theories of parallel and distributed processing in computer systems. T. A. Marschak and C. B. McGuire, in 1971, were probably the first to suggest the use of a particular model of a computer (the finite automaton) to represent the limited information-processing capabilities of humans in economic organizations. S. Reiter and K. R. Mount were early contributors to this line of research, and went further in analyzing economic organizations as networks of computers. (For more recent developments, see Radner 1997, Van Zandt 1999, and references therein.) One conclusion from this literature is the iron law of delay for networks of processors of bounded individual capacity. This 'law' can be paraphrased in the following way: as the size of the information-processing task increases, the minimum delay must also increase unboundedly, even for efficient networks, and even if the number of available processors is unlimited (Radner 1997, and references therein).

Memory storage and communication among humans and computers are also resource-using activities, and cause further delays in decision making. Both the storage and transmission of information and the results of information processing seem to be relatively 'cheap' compared with observation and processing, at least if we consider computer-supported activities. The proliferation of large data banks, and the flood of junk mail, telephone calls, and e-mail, lend support to this impression. It appears that today it is much cheaper, in some sense, to send, receive, and store memos and papers than it is to process them. (For game-theoretic models of players with limited memory, see Rubinstein 1998. For models of costly communication in organizations, and some implications for organizational structure, see Marschak and Reichelstein 1998.)

5. *Truly Bounded Rationality*

Many real decision problems present difficulties that prevent the DM from usefully treating them as optimization problems. Among these difficulties are:

inconsistency;
ambiguity;
vagueness;
unawareness; and
failure of logical omniscience.

As will be seen, these difficulties are somewhat related and overlapping. In particular, it is difficult to distinguish in practice between 'ambiguity' and 'vagueness.'

Regarding inconsistency, Savage (1954, p. 57) wrote:

> According to the personalistic view, the role of the mathematical theory of probability is to enable the person using it to detect inconsistencies in his own real or envisaged behavior. It is also understood that, having detected an inconsistency, he will remove it. An inconsistency is typically removable in many different ways, and the theory gives no guidance for choosing.

Some 'inconsistencies' have been observed so frequently, and have been so 'appealing,' that they have been used to criticize the Savage axioms, and to form a basis for a somewhat different set of axioms (e.g., the so-called 'Allais paradox' and 'Ellsberg paradox' (see *Utility and Subjective Probability: Contemporary Theories*; *Measurement Theory: Conjoint*). In other cases, it has been argued that inconsistent preferences arise because the DM is forced to articulate preferences about which they are not 'sure.' (This explanation is related to 'vagueness'; see below.) In particular, this unsureness may be related to uncertainty about what are the states of the world, a circumstance that can lead to a preference for 'flexibility' (see below). Finally, it has been observed in experiments that inconsistencies are more frequent the closer the alternatives are in terms of preference. This observation led J. Marschak and others to the elaboration of models of 'stochastic choice' (see *Decision and Choice: Random Utility Models of Choice and Response Time*).

Allusion has already been made to the DM's possible vagueness about his/her preferences. However, he/she could also be vague about any aspect of his/her model of the decision problem, and is likely to be so if the problem is at all complex. Vagueness can be about the *interpretation* of a feature of the model, or about its *scope*, or both. Unfortunately, there has been little if any formal theorizing about problems of vagueness.

At a moment of time, the DM may be *unaware* of some aspect of the problem: for example, he/she may be unaware of some future contingencies that could arise, or of some actions that are available to him/her. This phenomenon is particularly interesting if the DM is aware of the possibility that he/she may be unaware of something. For example, if the DM anticipates that in the future he/she will become aware of new acts of which he/she is currently unaware, then he/she may prefer present actions that allow for 'flexibility' of choice in the future. (This idea was formalized by T. C. Koopmans, D. Kreps, and others: Kreps 1992. For other theoretical treatments of unforeseen contingencies, see Dekel et al. 1998.)

As a consequence of the preceding considerations, decision theorists recognize that it is impossible for a DM to construct a complete model of his/her 'grand decision problem,' i.e., for their whole life! A common research strategy is to suppose that the DM can break up the grand decision problem into subproblems that can be solved independently without (much) loss of utility. Savage called this the device of constructing 'small worlds,' but showed that the conditions for this to be done without loss are unrealistically stringent (Savage 1954, pp. 82–91).

Finally, I come to what is perhaps the most difficult aspect of truly bounded rationality. Up to this point it has been assumed—if only implicitly—that the DM has no difficulty performing mathematical calculations or other logical operations. In particular, having formulated a decision model, he/she will be able to infer what it implies for their optimal strategy. As has already been pointed out, this assumption is absurd, even for small-world models, except for 'Mickey Mouse' problems that are constructed for textbooks and academic articles. The crux of the matter is that, in any even semirealistic decision problem, *the DM does not know all of the relevant logical implications of what he or she knows*. This phenomenon is sometimes called 'the failure of logical omniscience.' Examples of the failure of logical omniscience are:

(a) A DM who knows the axioms of arithmetic is uncertain about whether they imply that 'the 123rd digit in the decimal expansion of pi is 3,' unless he/she have a long time to do the calculation and/or has a powerful computer with the appropriate software.

(b) Twenty years ago, a DM who knew the axioms of arithmetic was still uncertain about whether they imply Fermat's last theorem.

The following examples are closer to practical life, and possibly more intimidating:

(a) Given all that a DM knows about the old and new drugs for treating a particular disease, what is the optimal policy for conducting clinical trials on the new ones?

(b) Given all that is known, theoretically and empirically, about business organizations in general, and about telecommunications and AT&T in particular, should AT&T reorganize itself internally, and if so, how?

Savage (1954, p. 7fn) commented on this class of problems:

> The assumption that a person's behavior is logical is, of course, far from vacuous. In particular, such a person cannot be uncertain about decidable mathematical propositions. This suggests, at least to me, that the tempting program sketched by Polya of establishing a theory of the probability of mathematical conjectures cannot be fully successful in that it cannot lead to a truly formal theory …

(For further discussion and references, see Savage 1972.)

In spite of some interesting efforts (see, for example, Lipman 1995), it does not appear that there has been significant progress on what it means to be rational in the face of this kind of uncertainty.

6. *Satisficing, Heuristics, and Non-Bayesian Learning*

In view of the difficulties posed by the various manifestations of 'truly bounded rationality,' a number of authors have proposed and studied behavior that departs more or less radically from the Savage Paradigm. These will be discussed under three headings: satisficing, heuristics, and non-Bayesian learning.

The term 'satisficing' refers to behavior in which the DM searches for an act that yields a 'satisfactory,' as distinct from an optimal, level of expected utility. The target, or 'satisfactory,' level of expected utility is usually called the DM's 'aspiration level.' In the simplest model, the aspiration level is exogenous, i.e., a given parameter of the model. More ambitious models describe some process whereby the aspiration level is determined within the model, and may change with experience (Simon 1972, Radner 1975). Such aspiration levels are called 'endogenous.' In some problems even optimal behavior bears a resemblance to satisficing. One category is the 'secretary problem' (Radner 2000).

The term 'heuristics' refers generally to behavior that follows certain rules that appear to produce 'good' or 'satisfactory' results most of the time in some class of problems (Simon 1972, see *Heuristics for Decision and Choice*). For example, the calculation of an optimal schedule for assigning jobs to machines is typically intractable if the numbers of jobs and machines are even moderately large. Nevertheless, human schedulers routinely construct 'satisfactory' schedules with such numbers, using various rules of thumb that have been developed with experience. Heuristics are central to many artificial intelligence applications. Satisficing plays an important role in many heuristic methods, and also in the processes of their modification.

The discussion of heuristics leads naturally to the consideration of non-Bayesian learning (NBL). Bayesian learning (i.e., the application of the calculus of conditional probability) is of course part of the Savage Paradigm in any decision problem in which the DM conditions his/her action on information about the state of the world. Many standard statistical methods use NBL. For example, the use of the sample mean to estimate a population mean is typically inconsistent with the Savage Paradigm (although in some cases the latter can be shown to be a limit of Bayesian estimates, as some parameter of the problem goes to infinity). Most psychological theories of learning postulate some form of NBL. A central question in the theory of NBL is: under what conditions, if any, does a particular NBL procedure converge asymptotically to a procedure that is Savage-Paradigm optimal as the DM's experience increases? (Rustichini 1999).

Again, one must ask: is there any satisfactory meaning to the term 'rationality' when used in the phrase 'bounded rationality'? The convergence of NBL to optimal actions could provide one (weak) meaning. Nevertheless, the problems raised by the various phenomena grouped under 'truly bounded rationality' may eventually lead students of decision making to answer this last question in the negative.

Bibliographic Notes

Many important works on bounded rationality have been omitted from the bibliography because of space limitations. The references cited in the body of the entry have historical interest, provide an overview of a topic discussed, and/or provide a key to other literature. The following provide additional references and information on the application of notions of bounded rationality in economics and management: Arrow 1974, Radner 2000, Shapira 1997, Van Zandt 1999.

See also: Bounded Rationality; Decision and Choice: Paradoxes of Choice; Decision Research: Behavioral; Game Theory; Heuristics for Decision and Choice; Intentionality and Rationality: A Continental-European Perspective; Intentionality and Rationality: An Analytic Perspective; Rational Choice Explanation: Philosophical Aspects

Bibliography

Arrow K J 1951 Alternative approaches to the theory of choice in risk-taking situations. *Econometrica* **19**: 404–37

Arrow K J 1974 *The Limits of Organization*. Norton, New York

Conlisk J 1996 Why bounded rationality. *Journal of Economic Literature* **34**: 669–700

Dekel E, Lipman B, Rustichini A 1998 Recent developments in modeling unforeseen contingencies. *European Economic Review* **42**: 523–42

Kreps D 1992 Static choice in the presence of unforeseen contingencies. In: Dasgupta P, Gale D, Hart O, Maskin E (eds.) *Economic Analysis of Markets and Games*. MIT Press, Cambridge, MA, pp. 258–81

Lipman B L 1995 Decision theory without logical omniscience: toward an axiomatic framework for bounded rationality. Unpublished thesis, Queens University, Canada

Marschak J, Radner R 1972 *Economic Theory of Teams*. Cowles Foundation, New Haven, CT

Marschak T A, Reichelstein S 1998 Network mechanisms, informational efficiency, and hierarchies. *Journal of Economic Theory* **79**: 106–41

Radner R 1975 Satisficing. *Journal of Mathematical Economics* **2**: 253–62

Radner R 1997 Bounded rationality, in determinacy, and the managerial theory of the firm. In: Shapira Z (ed.) *Organizational Decision Making*. Cambridge University Press, Cambridge, UK

Radner R 2000 Costly and bounded rationality in individual and team decision-making. *Industrial and Corporate Change* **9**: 623–58

Rubinstein A 1998 *Modeling Bounded Rationality*. MIT Press, Cambridge, MA

Rustichini A 1999 Optimal properties of stimulus-response learning models. *Games and Economic Behavior* **29**: 244–73

Savage L J 1954 *The Foundations of Statistics*. Wiley, New York. Dover, New York

Savage L J 1972 *The Foundations of Statistics*, 2nd edn. Dover, New York

Shapira Z (ed.) 1997 *Organizational Decision Making*. Cambridge University Press, Cambridge, UK

Simon H A 1972 Theories of bounded rationality. In: McGuire C B, Radner R (eds.) *Decision and Organization*. North-Holland, Amsterdam, pp. 161–76

Simon H A 1981 *The Sciences of the Artificial*. 2nd edn. MIT Press, Cambridge, MA

Van Zandt T 1998 Organizations with an endogenous number of information processing agents. In: Majumdar M K (ed.) *Organizations with Incomplete Information*. Cambridge University Press, Cambridge, UK, pp. 239–305

Van Zandt T 1999 Real-time decentralized information processing as a model of organizations with boundedly rational agents. *The Review of Economic Studies* **66**: 633–58

Wald A 1950 *Statistical Decision Functions*. Wiley, New York

R. Radner

Bounded Rationality

1. Introduction

The central ideas of bounded rationality (BR) are straightforward. First, humans are cognitively constrained in various ways, e.g., we can consciously attend to only one choice problem at a time. Second, these

International Encyclopedia of the Social & Behavioral Sciences ISBN: 0-08-043076-7