Bounded Rationality

The Adaptive Toolbox

Edited by

G. Gigerenzer and R. Selten

Program Advisory Committee:
G. Gigerenzer and R. Selten, Chairpersons
R. Boyd, P. Hammerstein

The MIT Press

Cambridge, Massachusetts
London, England
2001

1

Rethinking Rationality

Gerd Gigerenzer[1] and Reinhard Selten[2]

[1]Center for Adaptive Behavior and Cognition, Max Planck Institute for Human Development, Lentzeallee 94, 14195 Berlin, Germany
[2]Juridicum, University of Bonn, Adenauerallee 24–42, 53113 Bonn, Germany

Visions of rationality do not respect disciplinary boundaries. Economics, psychology, animal biology, artificial intelligence, anthropology, and philosophy struggle with models of sound judgment, inference, and decision making. These models evolve over time, just as the idea of rationality has a history, a present, and a future (Daston 1988). Over the last centuries, models of rationality have changed when they conflicted with actual behavior, yet, at the same time, they provided prescriptions for behavior. This double role — to describe and to prescribe — does not map easily onto a sharp divide between descriptive and normative models, which plays down the actual exchange between the psychological and the rational (Gigerenzer et al. 1989). Herbert Simon's notion of bounded rationality was proposed in the mid-1950s to connect, rather than to oppose, the rational and the psychological (Simon 1956). The aim of this book is to contribute to this process of coevolution, by inserting more psychology into rationality, and vice versa.

This book, however, cannot and will not provide a unified theory of bounded rationality. Rather, its goals are (a) to provide a framework of bounded rationality in terms of the metaphor of the *adaptive toolbox*, (b) to provide an understanding about why and when the simple heuristics in the adaptive toolbox work, (c) to extend the notion of bounded rationality from cognitive tools to emotions, and (d) to extend the notion of bounded rationality to include social norms, imitation, and other cultural tools.

To reach these goals, this book adopts a broad interdisciplinary perspective. Bounded rationality needs to be, but it is not yet, understood. Before we look into the future of models of rationality, however, let us have a glance back into the past.

THE RATIONAL AND THE PSYCHOLOGICAL

The Aristotelian distinction between the realm of demonstrative proof (that is, the things people were absolutely certain about, such as matters of mathematics and religion) and that of mere probable knowledge has shaped millennia of Western thought. The empire of demonstrative proof and certainty, however, narrowed in Europe after the Reformation and the Counter-Reformation. In the mid-17th century, a new and more modest standard of reasonableness emerged that acknowledged the irreducible uncertainty of human life. God, as John Locke (1690/1959) asserted, "has afforded us only the twilight of probability." The theory of probability that emerged at that time became the major guide for reasonableness. Its birth-year has been dated at 1654, when Blaise Pascal and Pierre Fermat exchanged letters about gambling problems (Hacking 1975). The initial definition of reasonableness was to choose the alternative that maximizes expected value (in modern notation, $\Sigma p_i v_i$, where p_i and v_i are the probability and value, respectively, of the i^{th} consequence of a given alternative). Pascal's *wager* illustrates that this newly conceived rationality was not only a cognitive revolution, but a new form of morality as well. Pascal pondered on whether to believe in God or not. He reasoned that even if the probability that God exists were small, the expectation is infinite: infinite bliss for the saved and infinite misery for the damned. Therefore, Pascal argued that rational self-interest dictates that we forego our certain but short-lived worldly pleasures and bet on the uncertain but infinite prospect of eternal salvation. The new brand of calculating rationality replaced faith with moral expectation.

The definition of reasonableness in terms of expected value soon ran into problems: it conflicted with the intuition of educated people. The first in a series of conflicts was the so-called St. Petersburg paradox, based on the following monetary gamble. Pierre offers Paul a gamble in which a fair coin is tossed. If the coin comes up heads on the first toss, Pierre agrees to pay Paul \$1; if heads do not turn up until the second toss, Paul receives \$2; if not until the third toss, \$4, and so on. What is the fair price Paul should pay to play this game? The fair price is defined by the expected value, $\Sigma p_i v_i$, which is the sum of all possible outcomes (values) times their probability. This sum is $1/2 \times \$1 + 1/4 \times \$2 + 1/8 \times \$4 + \ldots + 1/2^n \times \$n + \ldots$, which is infinitely large. However, reasonable people offer only some \$5 to \$10 to play this game rather than a very large amount of money. This discrepancy was labeled a paradox, the "St. Petersburg paradox," not because it is a logical paradox, but because the mathematical theory was so at odds with the dictates of good sense. The theory of rationality was thought of as a description of human behavior as well as a prescription for it. To resolve the discrepancy between theory and good sense, Daniel Bernoulli proposed changing the theory: from maximizing expected value to maximizing expected utility, that is, through incorporating the psychological facts that money has diminishing returns (which he modeled by a logarithmic function between monetary value and utility) and that this utility depends on the amount of money a person already has (Bernoulli 1738/1954).

The St. Petersburg paradox was the first in a series of monetary gambles — such as the Allais paradox and the Ellsberg paradox in the 20th century — that led to changes, or at least to proposed changes, in the theory of rationality. By the early 19th century, it was realized that there are several definitions of expectation, not just one (e.g., arithmetic mean, geometric mean, and median; as well as several measures of variability around the expected value to model risk). This ambiguity was perceived as a defeat for the program to arrive at the *one* mathematical definition of reasonableness, and became one of the reasons why the program was abandoned by most mathematicians around 1840 (Daston 1988). The program, however, was revived in the mid-20th century. Like Bernoulli, many contemporary researchers attempted to resolve discrepancies between description and prescription by tinkering with the utility or probability function, while at the same time retaining the ideal of maximization or optimization.

In this book, we pursue a more radical alternative to connect the rational with the psychological. The theory of bounded rationality, as we understand it, dispenses with optimization, and, for the most part, with calculations of probabilities and utilities as well.

BOUNDED RATIONALITY

In the 1950s and 60s, the Enlightenment notion of reasonableness reemerged in the form of the concept of "rationality" in economics, psychology, and other fields, usually referring to the optimization (maximization or minimization) of some function. Before that time, economists assumed that people were motivated by "self-interest," whereas the term "rationality" was used only occasionally (Arrow 1986). Similarly, psychologists, before 1950, assumed that thinking can be understood by processes such as "restructuring" and "insight," whereas the calculation of probabilities and the notion of optimization played little, if any, role. This situation changed after 1950. In psychology, for instance, tools for statistical inference became institutionalized in the 1950s and, once entrenched in the daily routine of the laboratory, were proposed by the researchers as models of cognitive processes. Fisher's analysis of variance turned into a mechanism of causal attribution, and Neyman-Pearson decision theory became a model of stimulus detection and discrimination (Gigerenzer and Murray 1987). By this tools-to-theories heuristic, cognitive processes came to be seen as involving the calculation of probabilities, utilities, and optimal decisions. But this is not the only route by which the ideal of optimization entered the behavioral sciences. Animal biologists developed optimal foraging theory, and the emerging field of artificial intelligence bet on designing rational agents similar to the "optimal" animals, that is, pursuing the ideal of providing artificial agents with a complete representation of their environment and extensive computing capacities for calculating optimal behavior.

About the same time as the notion of rationality as optimization became entrenched in theory and research across many disciplines, the competing notion of "bounded rationality" emerged (Simon 1956; Sauermann and Selten 1962). Herbert Simon, who coined the term "bounded rationality," used the metaphor of a pair of scissors, where one blade is the "cognitive limitations" of actual humans and the other the "structure of the environment." Minds with limited time, knowledge, and other resources can be nevertheless successful by exploiting structures in their environments. In Simon's (1956) words, "a great deal can be learned about rational decision making ... by taking account of the fact that the environments to which it must adapt possess properties that permit further simplification of its choice mechanisms" (p. 129). Studying only one blade is not enough; it takes both for the scissors to cut.

Thus, models of bounded rationality describe how a judgment or decision is reached (that is, the heuristic processes or proximal mechanisms) rather than merely the outcome of the decision, and they describe the class of environments in which these heuristics will succeed or fail. These models dispense with the fiction of optimization, which in many real-world situations demands unrealistic assumptions about the knowledge, time, attention, and other resources available to humans. Note that dispensing with optimization (as a model of cognitive processes) does not imply that the outcome of a nonoptimizing strategy is bad. For instance, optimization is often based on uncertain assumptions, which are themselves guesswork, and as a consequence, there may be about as many different outcomes of optimizing strategies as there are sets of assumptions. In these real-world cases, it is possible that simple and robust heuristics can match or even outperform a specific optimizing strategy.

Simon's original vision, however, has met with only limited success. In particular, bounded rationality has become a fashionable label for almost every model of human behavior. In the following, we sketch our vision of bounded rationality, which is an elaboration of Simon's original concept. We begin with two prominent interpretations that are both inconsistent with our vision of bounded rationality.

WHAT BOUNDED RATIONALITY IS NOT

Bounded rationality is neither optimization nor irrationality. Nevertheless, a class of models known as *optimization under constraints* is referred to in the literature as "bounded rationality," and a class of empirical demonstrations of so-called errors and fallacies in judgment and decision making has also been labeled "bounded rationality." The fact that these two classes of models have little if anything in common reveals the distortion the concept of bounded rationality has suffered.

Optimization Under Constraints

A key process in bounded rationality is limited search. Whereas in models of unbounded rationality all relevant information is assumed to be available already, real humans and animals need to search for information first. Search can be for two kinds of information: alternatives (such as for houses and spouses) and cues (that is, for reasons and predictors when deciding between given alternatives). Search can be performed inside the human mind (memory) or outside it (e.g., library, internet, other minds). Internal search costs time and attention, and external search may cost even further resources, such as money. Limited resources constrain institutions, humans, animals, and artificial agents, and these limitations usually conflict with the ideal of finding a procedure to arrive at the optimal decision.

One way to model limited search without giving up the ideal of optimization is known as optimization with decision costs taken into account, also referred to as *optimization under constraints*. Stigler (1961) used the example of a person who wants to buy a used car, and stops searching when the costs of further search would exceed the benefit of further search. This is known as an optimal stopping rule. The problems with this form of optimization as a model of actual cognitive processes are well known. First, reliable estimates of benefits and costs, such as opportunity costs, can demand large degrees of knowledge; second, there is an infinite regression problem (the cost-benefit computations themselves are costly, and demand a meta-level cost-benefit computation, and so on); and finally, the knowledge and the computations involved can be so massive that one is forced to assume that ordinary people have the computational capabilities and statistical software of econometricians (Sargent 1993).

All in all, this attempt to model limited search leads to the paradoxical result that the models become even less psychologically plausible. The reason is that the desire for optimization is retained. Because optimization under constraints is often referred to as "bounded rationality" (e.g., Sargent 1993), this confusion has led to the dismissal of bounded rationality as a hidden form of optimization — one that just makes everything more complicated. As we will see, this interpretation is inappropriate and misleading. Models of bounded rationality use fast and frugal stopping rules for search that do *not* involve optimization.

Irrationality

Since the 1970s, researchers have documented discrepancies between a "norm" (e.g., a law of probability or logic) and human judgment. Unlike in Bernoulli's proposal, the blame was put on the human mind rather than on the norm. The discrepancies were labeled "fallacies," such as the base-rate fallacy and the conjunction fallacy, and attributed to humans' "bounded rationality," in the sense of limitations on rationality. This interpretation was put forward in psychology (e.g., Kahneman et al. 1982, p. xii), experimental economics (Thaler 1991, p. 4), and the law (Jolls et al. 1998, pp.1548–1549).

Bounded rationality is, however, not simply a discrepancy between human reasoning and the laws of probability or some form of optimization. Bounded rationality dispenses with the notion of optimization and, usually, with probabilities and utilities as well. It provides an alternative to current norms, not an account that accepts current norms and studies when humans deviate from these norms. Bounded rationality means rethinking the norms as well as studying the actual behavior of minds and institutions.

The interpretation of fallacies as bounded rationality focuses on one blade of Simon's scissors (the cognitive limitations), but neglects the other blade (the structure of environments). Several of the so-called fallacies are based on norms that have been put forth without analyzing the structure of environments. For instance, what has been interpreted as base-rate neglect turns out to be a reasonable strategy under plausible assumptions about the environment (e.g., Birnbaum 1983; Mueser et al. 1999). Moreover, when information is represented in natural frequencies rather than probabilities, base rate neglect is perfectly rational (Gigerenzer and Hoffrage 1995). To sum up, bounded rationality is not an inferior form of rationality; it is not a deviation from norms that do not reflect the structure and representation of information in environments. Theories of bounded rationality should not be confused with theories of irrational decision making. The label "bounded rationality" signifies a type of theory, not outcomes.

MODELING OMNISCIENCE VERSUS HEURISTICS: AN ILLUSTRATION

In cricket, baseball, and soccer, players need to catch balls that come in high. Our thought experiment is to build a robot that can catch the ball. (No such robot exists as yet.) For the sake of simplicity, we consider only cases where a ball comes in front of or behind a player, but not to his left or right. One team of engineers, which we call the optimizing team, proceeds by programming the family of parabolas into the robot's mind (in theory, balls fly in parabolas). To select the proper parabola, the robot needs to be equipped with instruments that can measure the distance from which the ball was thrown or shot, as well as its initial velocity and projection angle. Yet in the real world, balls do not fly in parabolas due to air resistance and wind. Thus, the robot would need further measurement instruments that can measure the speed and direction of the wind at each point of the ball's flight, and compute the resulting path. Yet in a real game, there are myriad further factors, such as spin, that affect the flight of the ball, and the robot would need instruments to measure the initial direction and strength of the spin, and knowledge of how the various factors interact. Thus, the optimization team's program is to develop reliable measurement instruments that supply the

robot with all the relevant knowledge, and powerful computer software that can compute from these measurements where the ball will land. All this would have to be calculated within one or two seconds — the usual time a ball is in the air. Then the robot would run to this point and catch the ball.

A second team of engineers, which we call the boundedly rational team, makes a different start. They first study what experienced players actually do. (The optimizing team had discussed this option, but dismissed it because the visual measurements and computations players are assumed to perform are unobservable and unconscious; thus, observing and interviewing players would mean wasting time. Have you ever interviewed a soccer player?) Based on these observations, the heuristic team programs the robot not to move during the first half second or so but to make a crude estimate of whether the ball is coming down in front of or behind it, and then start running in this direction while fixating its eye on the ball. The heuristic the robot uses is to adjust its running speed so that the angle of gaze — the angle between the eye and the ball — remains constant (or within a certain range; see McLeod and Dienes 1996). By using this simple gaze heuristic, the robot will catch the ball while running. Note that this boundedly rational robot pays attention to only one cue, the angle of gaze, and does not attempt to acquire information concerning wind, spin, or the myriad of other causal variables, nor perform complex computations on these estimates. Note also that the gaze heuristic does not allow the robot to compute the point where the ball will land, run there, and wait for the ball. However, the robot does not need to make this difficult calculation; it will be there where the ball lands — just as actual players do not first calculate the landing point and then run there and wait for the ball. They catch the ball while running.

This thought experiment can illustrate several more general points. First, contrary to conventional wisdom, limitations of knowledge and computational capability need *not* be a disadvantage. The heuristic tools of humans, animals, and institutions can be simple, but nevertheless effective in a given environment. The optimizing robot that needs a complete representation of the environment (some AI researchers have fed their programs more than a million facts to approximate this ideal) and bets on massive computation, may never finish its analysis before the ball has hit the ground. Simplicity, by contrast, can enable fast, frugal, and accurate decisions. Second, a simple heuristic can exploit a regularity in the environment. In the present case, the regularity is that a constant angle of gaze will cause a collision between the ball and the player. Third, boundedly rational heuristics are, to some degree, domain-specific rather than universal strategies. These heuristics are middle-ranged, that is, they work in a class of situations (e.g., pilots are taught a similar heuristic to avoid collisions with other planes), but they are not general-purpose tools such as the ideal of an all-purpose optimization calculus. What we call the "adaptive toolbox" contains a number of these middle-range tools, not a single hammer for all purposes.

MODELS OF BOUNDED RATIONALITY

We mentioned at the beginning that, so far, there is no complete theory of bounded rationality. Nevertheless, we can specify three classes of processes that models of bounded rationality typically specify:

1. *Simple search rules*. The process of search is modeled on step-by-step procedures, where a piece of information is acquired, or an adjustment is made (such as to increase running speed to keep the angle of gaze constant), and then the process is repeated until it is stopped.
2. *Simple stopping rules*. Search is terminated by simple stopping rules, such as to choose the first object that satisfies an aspiration level. The stopping rule can change as a consequence of the length of search or other information, as in aspiration adaptation theory (Selten 1998). Simple stopping rules do not involve optimization calculations, such as computations of utilities and probabilities to determine the optimal stopping point.
3. *Simple decision rules*. After search is stopped and a limited amount of information has been acquired, a simple decision rule is applied, like choosing the object that is favored by the most important reason — rather than trying to compute the optimal weights for all reasons, and integrating these reasons in a linear or nonlinear fashion, as is done in computing a Bayesian solution.

The process of search distinguishes two classes of models of bounded rationality: those that search for alternatives (e.g., aspiration level theories such as satisficing, see Chapters 2 and 4), and those that search for cues (e.g., fast and frugal heuristics, see Chapters 3, 4, and 9). The chapters in this book describe and discuss models of both kinds in detail. Here, we will address the four questions that, in our opinion, pose a challenge for research on bounded rationality and around which this book is organized.

FOUR BIG QUESTIONS

The goal of this book is to promote bounded rationality as the key to understanding how actual people make decisions without calculating utilities and probabilities — that is, without optimizing.

1. Is There Evidence for an Adaptive Toolbox?

As mentioned above, models of bounded rationality consist of simple step-by-step rules that function well under the constraints of limited search, knowledge, and time — whether or not an optimal procedure is available. The repertoire of these rules or heuristics, available to a species at a given point in its

evolution is called its "adaptive toolbox." The concept of an adaptive toolbox, as we see it, has the following characteristics: First, it refers to a collection of rules or heuristics rather than to a general-purpose decision-making algorithm (as was envisioned in Leibniz's dream of a universal calculus or in versions of subjective expected utility theory). Second, these heuristics are fast, frugal, and computationally cheap rather than consistent, coherent, and general. Third, these heuristics are adapted to particular environments, past or present, physical or social. This "ecological rationality"— the match between the structure of a heuristic and the structure of an environment — allows for the possibility that heuristics can be fast, frugal, and accurate all at the same time by exploiting the structure of information in natural environments (Gigerenzer et al. 1999). Fourth, the bundle of heuristics in the adaptive toolbox is orchestrated by some mechanism reflecting the importance of conflicting motivations and goals. This mechanism is not yet well understood.

2. Why and When Do Simple Heuristics Work?

Bounded rationality is sometimes interpreted as suboptimal or even irrational in comparison with nonbounded rationality. However, evidence exists showing that fast and frugal rules can be about as accurate as complex statistical models (e.g., multiple regression, Bayesian networks), while demanding less information and computational power (e.g., Martignon and Laskey 1999). One reason simple heuristics work is that they can exploit structures of information in the environment. That is, their rationality is a form of ecological rationality, rather than of consistency and coherence. A second reason is the robustness of simple strategies compared to models with large numbers of parameters, which risk overfitting. Third, there are real-world situations involving multiple goals (e.g., accuracy, speed, frugality, consistency, accountability) that have no known common denominator, which poses serious problems to optimization, but can be handled by models of bounded rationality (e.g., Chapter 2).

3. What Role Do Emotions and Other Noncognitive Factors Play in Bounded Rationality?

Emotions are often seen as obstacles to rationality. However, emotions like disgust or parental love can provide effective stopping rules for search and a means for limiting search spaces. In particular, for important adaptive problems such as food avoidance, an emotion of disgust, which may be acquired through observation of conspecifics, can be more effective than cognitive decision making. Similarly, in social species, imitation and social learning can be seen as mechanisms that enable fast learning and obviate the need for individual calculations of expected utilities. For instance, a monkey may be prepared to acquire a fear of snakes the moment it sees another conspecific exhibit signs of fear in the presence of a snake (one-trial learning), or a child may make choices by imitating parents and peers.

Thus, many decisions seem to be made on the basis of factors other than cognitive ones, that is, factors other than estimations of probabilities, gains, costs, and the like. In general, the questions are: What noncognitive factors allow humans and animals to make fast and accurate decisions and to learn more efficiently than from mere experience? In what situations are noncognitive factors indispensable for adaptive decision making?

4. What Is the Role of Culture in Bounded Rationality?

Models of bounded rationality have been focused not only on cognitive factors but also on individual decision making. However, just as emotions can lead to sensible decisions and commitments, so can social norms and institutions. Social norms can be seen as fast and frugal behavioral mechanisms that dispense with individual cost-benefit computations and decision making. For instance, culture (as a system of values and beliefs) can help actual humans diminish the problem of combinatorial explosion and the related problem of how to make an infinite number of possible decisions in real time, which has tormented attempts at building intelligent machines and robots. Cultural systems of belief need not be "correct" to work. For instance, navigational systems were successful long before our present-day astronomical knowledge was available, and they worked often on the basis of false cultural beliefs about the movement of the earth and heavenly bodies (Hutchins 1995). The point is that humans do not need to wait until all knowledge is acquired and all truth is known (which probably will never be the case). Adaptive solutions can be found with little knowledge; the price for this is that they are not general, but do work in a specific environment, culture, or time.

INTERDISCIPLINARITY

Bounded rationality is a genuinely interdisciplinary topic: its subject matter is the proximal mechanisms (the heuristics) that animals, humans, institutions, and artificial agents use to achieve their goals. The common denominator is that decisions need to be made with limited time, knowledge, and other resources, and in a world that is uncertain and changing. Thus, the framework of bounded rationality — the building blocks of heuristics; the notion of ecological rationality; the cultural acceleration of learning by social norms and imitation — may help to integrate fields that, so far, have been dissociated and did not access relevant knowledge outside their disciplines.

The lack of information flow between disciplines can hardly be underestimated. A brilliant example is the sunk cost fallacy (e.g., an agent has a choice to

invest either in A or B, where B is the more promising option; however, because the agent had previously invested in A but not in B, he chooses A). Hundreds of papers were written in economics and psychology on the sunk cost fallacy, and hundreds of papers were written in evolutionary biology (by some of the most eminent biologists) on the Concorde fallacy — which is the same fallacy. There is not a single cross reference in these hundreds of papers, nor any awareness that both fields came to opposite conclusions: in economics and psychology, it is taken for granted that humans often commit the sunk cost fallacy, in animal biology, no conclusive evidence has been found that a single animal species would commit the sunk cost fallacy (Arkes and Ayton 1999). Bounded rationality and the study of the adaptive toolbox may help to shift discipline-oriented research toward more problem-focused research.

Optimization is an attractive fiction; it is mathematically elegant, and one can draw on a well-developed calculus. Compared to the beauty of optimization, the actual proximal mechanisms of humans and animals resemble the tools of a backwoods mechanic. The pleasing ideal of a universal calculus may have distracted researchers in many fields from exploring the contents of the adaptive toolbox. However, there is also another sense of beauty: the aesthetics of simplicity. There is a sense of wonder in how simplicity can produce robustness and accuracy in an overwhelmingly complex world. The gaze heuristic, the adaptation of aspiration levels, Tit-for-Tat, the recognition heuristic, Take the Best, and many other smart heuristics described in this book can produce this sense of wonder.

REFERENCES

Arkes, H.R., and P. Ayton. 1999. The sunk cost and Concorde effects: Are humans less rational than lower animals? *Psych. Bull.* **125**:591–600.

Arrow, K.J. 1986. Rationality of self and others in an economic system. *J. Business* **59**:S385–S399.

Bernoulli, D. 1738/1954. Specimen theoriae novae de mensura sortis. *Commentarii academiae scientarum imperialis Petropolitanae* **5**:175–192. (Engl. transl. by L. Sommer, Exposition of a new theory on the measurement of risk. *Econometrica* **22**:23–26.)

Birnbaum, M.H. 1983. Base rates in Bayesian inference: Signal detection analysis of the cab problem. *Am. J. Psych.* **96**:85–94.

Daston, L.J. 1988. Classical Probability in the Enlightenment. Princeton, NJ: Princeton Univ. Press.

Gigerenzer, G., and U. Hoffrage. 1995. How to improve Bayesian reasoning without instructions. Frequency formats. *Psych. Rev.* **102**:684–704.

Gigerenzer, G., and D.J. Murray. 1987. Cognition as Intuitive Statistics. Hillsdale, NJ: Erlbaum.

Gigerenzer, G., Z. Swijtink, T. Porter, L. Daston, J. Beatty, and L. Krüger. 1989. The Empire of Chance. How Probability Changed Science and Everyday Life. Cambridge: Cambridge Univ. Press.

Gigerenzer, G., P.M. Todd, and the ABC Research Group. 1999. Simple Heuristics That Make Us Smart. New York: Oxford Univ. Press.

Hacking, I. 1975. The Emergence of Probability. Cambridge: Cambridge Univ. Press.

Hutchins, E. 1995. Cognition in the Wild. Cambridge, MA: MIT Press.

Jolls, C., C.S. Sunstein, and R. Thaler. 1998. A behavioral approach to law and economics. *Stanford Law Rev.* **50**:1471–1550.

Kahneman, D., P. Slovic, and A. Tversky, eds. 1982. Judgment under Uncertainty: Heuristics and Biases. Cambridge: Cambridge Univ. Press.

Locke, John. 1690/1959. An Essay Concerning Human Understanding, ed. Alexander Campbell Fraser. New York: Dover.

Martignon, L., and K.B. Laskey. 1999. Bayesian benchmarks for fast and frugal heuristics. In: Simple Heuristics That Make Us Smart, G. Gigerenzer, P.M. Todd, and the ABC Research Group, pp. 169–188. New York: Oxford Univ. Press.

McLeod, P., and Z. Dienes. 1996. Do fielders know where to go to catch the ball or only how to get there? *J. Exp. Psych.: Human Percep. Perf.* **22**:531–543.

Mueser, P.R., N. Cowan, and K.T. Mueser. 1999. A generalized signal detection model to predict rational variation in base rate use. *Cognition* **69**:267–312.

Sargent, T.J. 1993. Bounded Rationality in Macroeconomics. Oxford: Oxford Univ. Press.

Sauermann, H., and R. Selten. 1962. Anpassungstheorie der Unternehmung. *Z. ges. Staatswiss.* **118**:577–597.

Selten, R. 1998. Aspiration adaptation theory. *J. Math. Psychol.* **42**:191–214.

Simon, H.A. 1956. Rational choice and the structure of environments. *Psych. Rev.* **63**:129–138.

Stigler, G.J. 1961. The economics of information. *J. Pol. Econ.* **69**:213–225.

Thaler, R.H. 1991. Quasi Rational Economics. New York: Russell Sage.