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

3

The Adaptive Toolbox

Gerd Gigerenzer

Center for Adaptive Behavior and Cognition, Max Planck Institute for Human Development, Lentzeallee 94, 14195 Berlin, Germany

ABSTRACT

The notion of an adaptive toolbox provides a framework for nonoptimizing visions of bounded rationality, emphasizing psychological plausibility, domain specificity, and ecological rationality. Heuristics in the adaptive toolbox are modeled on the actual cognitive abilities a species has rather than on the imaginary powers of omniscient demons. They are designed for specific goals — domain specific rather than domain general — which enable them to make fast, frugal, and computationally cheap decisions. Heuristics are composed from building blocks that guide search, stop search, and make decisions. Heuristics that are matched to particular environmental structures allow agents to be ecologically rational. The study of ecological rationality involves analyzing the structure of environments, the structure of heuristics, and the match between them.

INTRODUCTION

Humans and animals make inferences about unknown features of their world under constraints of limited time, limited knowledge, and limited computational capacities. Models of rational decision making in economics, cognitive science, biology, and other fields, in contrast, tend to ignore these constraints and treat the mind as a Laplacean superintelligence equipped with unlimited resources of time, information, and computational might. Some forty years ago, Herbert Simon challenged this view with his notion of "bounded rationality," but with only limited success. Today, bounded rationality has become a diluted, fashionable term, used by the proponents of quite disparate visions of reasonableness: from optimization under constraints (e.g., Sargent 1993) to judgmental errors and human irrationality (e.g., Jolls et al. 1998).

The notion of an adaptive toolbox promotes a specific vision of bounded rationality based on three premises (Gigerenzer et al. 1999):

1. *Psychological Plausibility.* The goal of the program is to understand how actual humans (or ants, bees, chimpanzees, etc.) make decisions, as opposed to heavenly beings equipped with practically unlimited time, knowledge, memory, and other infinite resources. The challenge is to base models of bounded rationality on the cognitive, emotional, social, and behavioral repertoire that a species actually has.
2. *Domain Specificity.* The adaptive toolbox offers a collection of heuristics that are specialized rather than domain general as would be the case in subjective expected utility (SEU). These heuristics are composed of cognitive and emotional building blocks that can be part of more than one heuristic and allow the composition of new heuristics. The building blocks are more general than the heuristics.
3. *Ecological Rationality.* The "rationality" of domain-specific heuristics is not in optimization, omniscience, or consistency. Their success (and failure) is in their degree of adaptation to the structure of environments, both physical and social. The study of the match between heuristics and environmental structures is the study of ecological rationality.

VISIONS OF REASONABLENESS

The first premise, psychological plausibility, sets our vision of bounded rationality apart from the two species of "demons" in Figure 3.1. *Unbounded rationality* encompasses decision-making strategies that have little or no regard for the constraints in time, knowledge, and computational capacities that real humans face. For example, models that seek to maximize expected utility or perform Bayesian calculations often must assume demonic strength to tackle real-world problems. Real decision makers (as opposed to participants in an experiment in which all information is already conveniently laid out in front of them) need, first of all, to search for information. This search cannot go on indefinitely; it is constrained by limited time, money, attention, or other finite resources. The key difference between models of unbounded rationality and models of *optimization under constraints* is that the latter model *limited information search* that was terminated by a *stopping rule*. The assumption behind optimization under constraints is that the optimal stopping point can be calculated — the point at which the costs for further search begin to exceed the benefits that a continued search could bring (Stigler 1961). However, the rule "stop search when costs outweigh benefits" can require even more time, knowledge, and computation to calculate than models of unbounded rationality (Vriend 1996; Winter 1975). This leads to the paradoxical consequence that "limited minds" are assumed to have the knowledge and computational ability of sophisticated econometricians equipped with statistical software packages.

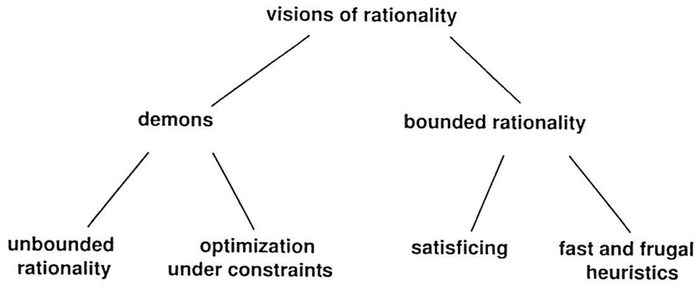

Figure 3.1 Models of bounded rationality consist of search for alternatives, such as houses and spouses (satisficing, Simon 1955; see also Selten, this volume) and search for cues, such as reasons for preferring one alternative to another (fast and frugal heuristics, Gigerenzer et al. 1999).

Unfortunately, optimization under constraints has often been labeled bounded rationality (e.g., Sargent 1993), which has (mis-)led many to conclude that there is, in the end, no difference between bounded and unbounded rationality.

The notion of ecological rationality puts models of bounded rationality — of which Figure 3.1 lists two general classes — in a functional, environmental perspective. In his 1956 article entitled "Rational Choice and the Structure of the Environment," Herbert Simon pointed out that there are two sides to bounded rationality: the "cognitive limitations" and the "structure of environments." His example was about an environment in which food is distributed randomly, and another environment in which cues for food distribution exist. An organism in the first environment can get away with very simple cognitive and behavioral strategies of search; an organism in an environment where cues exist, however, can benefit from cognitive abilities for learning cue–goal relations and planning search. The term environment, here, does not refer to a description of the total physical and biological environment, but only to that part important to an organism, given its needs and goals.

The notions of psychological plausibility and ecological rationality suggest two routes to the study of bounded rationality. The quest for psychological plausibility suggests looking into the mind, that is, taking account of what we know about cognition and emotion in order to understand decisions and behavior. Ecological rationality, in contrast, suggests looking outside the mind, at the structure of environments, to understand what is inside the mind. These research strategies are complementary, like digging a tunnel from two sides. However, the parties who dig from each side should meet at some point, and this has been a rare event, so far. On one hand, psychological plausibility has often been reduced to pointing out cognitive biases — typically by means of text problems, with little concern about environmental structures in which these purported biases could make sense (see Gigerenzer 1991; Gigerenzer and Murray 1987). On the other hand, the analysis of the structure of natural environments has often been paired with a behavioristic anxiety about opening the black box of the mind — from the psychologists Egon Brunswik and J.J. Gibson to ecologists and biologists. The neglect of one of the two sides of bounded rationality can even be traced in Simon's writings. For example, in his New Palgrave article, he explains that bounded rationality "takes into account the cognitive limitations of the decision maker — limitations of both knowledge and computational capacity" (Simon 1987, p. 266). The structure of environments he had emphasized repeatedly is not even mentioned.

I prefer to use the neutral term "cognitive abilities" over "cognitive limitations," since so-called limitations are relative to the fiction of optimizing demons, as in Figure 3.1. However, a serious program of bounded rationality needs to emancipate itself from the Christian ideal of an omniscient and omnipotent God, or its secularized version, Laplace's superintelligence. Bounded rationality in economics, cognitive science, and biology is about humans and animals, not about how they compare with demons and gods.

WHAT IS THE FUNCTION OF THE ADAPTIVE TOOLBOX?

The ultimate goal of organisms, according to evolutionary views, is reproduction, either of genes or some other units. The adaptive toolbox is designed to achieve proximal goals, such as finding prey, avoiding predators, finding a mate, and if a species is social or cultural, exchanging goods, making profits, and negotiating status. The tools are means to achieve proximal goals and include learning mechanisms that allow an adjustment of the tools when environments change. There can be a multitude of goals that attract the attention of an agent at any point in time, and these need not be compensatory, that is, a common psychological currency may not exist. Lack of a common currency is a psychological reality, and strategies of bounded rationality should be able to handle these situations in which no single goal variable exists that could be optimized. An example is aspiration adaptation theory (Selten, this volume).

Beyond Optimization

The strategies in the adaptive toolbox do not try to optimize, that is, they do not try to compute the maximum of some function, and for several good reasons. The general reason is that optimization is feasible in only a restricted set of problems, typically on the basis of simplifying assumptions. For example, even for well-defined games such as chess and go — where the rules and goals are simple, unambiguous, and not subject to change — no optimal strategy is known. Specific reasons are numerous, including (a) if multiple competing goals exist, optimization can become a heavy and unbearable computational burden, (b) if incommensurable goals exist, optimization can be impossible, (c) if incommensurable reasons (or cues) exist, optimization can be impossible, (d) if the alternatives are unknown and need to be generated in a lengthy process of search, optimization models assuming a finite, known choice set do not apply, (e) if the

cues or reasons are unknown and need to be generated in a lengthy process of search, optimization models that assume a finite, known set of predictors do not apply, (f) if future consequences of actions and events are unknown, optimization models assuming a known, finite set of consequences do not apply, and (g) if optimization is attempted for the multitude of decisions an organism faces in real time, this can lead to paralysis by computational explosion (see Klein, this volume).

Note that optimization is always relative to a number of assumptions about the environment, social and physical, that are typically uncertain — except in textbook examples. The heuristics in the adaptive toolbox just "bet" on the environment on the basis of past experience or a little probing, without attempting a complete analysis and subsequent optimization (Gigerenzer and Todd 1999).

Beyond Consistency

The heuristics in the adaptive toolbox also do not take consistency as a *sine qua non* for rational behavior. In contrast, axioms and rules of internal consistency, such as property alpha, transitivity, and additivity of probabilities, are often seen as the cornerstones of human rationality, and of animal rationality as well (e.g., McGonigle and Chalmers 1992). Similarly, violations of consistency are typically seen as instances of irrationality, such as preference reversals (see, however, Sen's [1993] argument that consistency is an ill-defined concept unless the social goals and objectives of people are specified).

From a functional view, however, consistency in choice and judgment is not a general norm to follow blindly, but rather a tool for achieving certain proximal goals. For a given goal, consistent behavior can be an advantage, a disadvantage, or unimportant. For example, in cooperative relationships within families and businesses, some forms of consistent behaviors seem to be indispensable. They contribute to producing and maintaining a social climate of trust, fairness, and commitment. In competitive relationships, however, strategies with built-in inconsistencies can be an advantage. Protean behavior provides one such example: the erratic escape movements of prey, the "crazy dances" of predators, the sequential moves of two competitors who mutually try to make their future choices unpredictable, and the biochemistry of viruses (Driver and Humphries 1988). When unpredictability is adaptive, consistent preferences or behaviors may be deadly. There is also a class of situations in which neither consistency nor inconsistency seems functional, but other features of decision making are in the foreground: to make a fast decision, to make a frugal decision (with only meager information), or to make a decision at all — not to get it right, but to get a business or social interaction going.

Domain Specificity: Making Choices Fast and Frugal

Domain-specific heuristics allow faster reaction with less information than a general-purpose algorithm. A domain is a subset G' of the set G of goals. I consider two complementary ways to define domains: *cognitive tasks* and *adaptive problems*. Cognitive tasks include classification, estimation, and two-alternative choice. Here, the subset G' consists of tasks that afford the same type of result, such as classifying an object into one of several categories or estimating the value of an object on a criterion. Adaptive problems include mate choice, habitat choice, food choice, social exchange, and their modern equivalents such as dieting and stock markets. Adaptive problems are characterized by their common content rather than a common cognitive task. Cognitive psychology has almost exclusively focused on studying mechanisms for cognitive tasks rather than for adaptive problems. There is a large literature on classification, estimation, and choice — most of which ignores search and stopping and proposes demon-like models of decision, such as exemplar models of classification, multiple regression models of estimation, and SEU models of choice.

In our own work on fast and frugal heuristics (Gigerenzer et al. 1999), we decided on a complementary research strategy that studies both types of domain-specific heuristics for cognitive tasks and for adaptive problems. For example, the "QuickEst" heuristic (Hertwig et al. 1999) and the "Take The Best" heuristic (Gigerenzer and Goldstein 1996) are specific to certain cognitive tasks of estimation and two-alternative choice, respectively. Mate choice, in contrast, is an adaptive problem that demands boundedly rational heuristics designed for situations of mutual choice, sequential encounter of alternatives, and stopping rules that can include emotions such as love (Todd and Miller 1999). We found that studying heuristics for both types of domains has proved helpful. One might assume that the adaptive toolbox employs the same dual approach for handling new problems. Heuristics that are task specific can easily be generalized to new tasks of the same type, and strategies designed for specific adaptive problems can be generalized to new cultural versions of the strategy.

Emotions are prime examples of domain-specific tools for bounded rationality, in particular for solving adaptive problems. Emotions can help to limit the number of decisions to be made. Disgust, for example, limits the choice set of items to eat and serves the adaptive function of preventing food poisoning. However, disgust is of little help for other adaptive problems, such as social exchange, where anger and, conversely, feelings of guilt can keep people from cheating on contracts.

The Mind as a Backwoods Mechanic and Used Parts Dealer

Leibniz had a beautiful dream of discovering the universal logical language in which God had written the book of nature. This language, the Universal Characteristic, would replace all reasoning with one calculus (Leibniz 1677/1951). It would put an end to scholarly bickering and clamorous controversy — if there is a problem, just sit down and calculate. For some time, Enlightenment thinkers hoped that the calculus of probability would make Leibniz' dream a reality, but by the 1840s, most mathematicians had given up the task of reducing rational reasoning to a general calculus as a thankless and even antimathematical one (Daston 1988). However, the dream has survived in some quarters, where hearts still beat for a unified formula of rationality, be it some version of SEU or Bayesianism.

The notion of an adaptive toolbox full of specialized devices lacks the beauty of Leibniz' dream or that of SEU. It invokes the more modest abilities of a "backwoods mechanic and used part dealer," as Wimsatt (1999) describes nature. The backwoods mechanic has no general-purpose tool nor are all spare parts available to him. He must fiddle with various imperfect and short-range tools, a process known as vicarious functioning (Brunswik 1955). He will have to try one thing, and if it does not work, another one, and with step-by-step adjustments will produce serviceable solutions to almost any problem with just the things at hand.

The design of domain-specific mechanisms can be a viable alternative to building intelligent machines. Since its inception, artificial intelligence has relied upon perfect rationality — a complete representation of the world and some form of general optimization algorithm — as the desired property of intelligent systems, resulting in a wide gap between theory and practice hindering progress in the field (Russell and Subramanian 1995). This "Good Old Fashioned AI," or GOFAI, tends to produce robots that move a few feet and then sit there for a long time until they have computed the optimal direction in which to move the next few feet, and so on. In contrast, the argument has been made that smart robots need to be modeled after humans, equipped with special-purpose abilities without a centralized representation and computation control system, and an intelligence that emerges from their interactions with the environment (e.g., Brooks 1993).

The function of the adaptive toolbox is, thus, to provide strategies — cognitive, emotional, and social — that help to handle a multitude of goals by making decisions quickly, frugally, accurately, or, if possible, not at all. The function of the adaptive toolbox is not to guarantee consistency or solve differential equations to optimize some function. Clearly, I believe that the importance of optimization and consistency has been overestimated in theories of rational choice.

WHAT ARE THE TOOLS IN THE ADAPTIVE TOOLBOX?

The adaptive toolbox provides heuristics, and these are composed of building blocks. I describe three functions these building blocks have: they give search a direction, stop search, and make a decision.

Search Rules

One can think of search as an exploration of two dimensions: search for alternatives (the choice set) and cues (to evaluate the alternatives). On one hand, Simon's concept of satisficing involves search for alternatives, but not for cues

(Simon 1955). Cues can be thought of as implicit in his concept of an aspiration level. On the other hand, the fast and frugal heuristics studied by our research group (Gigerenzer et al. 1999) search for cues and are designed for situations in which the alternatives (such as job candidates or stocks) are already known. Thus, the division of bounded rationality into satisficing and fast and frugal heuristics, as shown in Figure 3.1, reflects the type of search: search for alternatives (satisficing) or search for cues (fast and frugal heuristics).

Building blocks for guiding search include random search, ordered search (e.g., looking up cues according to their validities), and search by imitation of conspecifics, such as stimulus enhancement, response facilitation, and priming. Imitation is an effective mechanism that allows humans (and the few other species that imitate behavior on some larger scale) to learn quickly where to look and what to look for.

When evaluating models of search, psychological plausibility should not be confused with irrationality. In noisy environments, search heuristics can be both psychologically plausible, simple, and successful. One reason is the robustness of simple search heuristics. For example, the Take The Best heuristic computes a simple order (of cue validities) to direct search for cues, which is suboptimal given the data it knows. Nevertheless, when it comes to making predictions about new data, Take The Best actually can make more accurate inferences using the simple order than the order that was optimal for the data available (Martignon and Hoffrage 1999).

The family of noncognitive tools for search has not been explored in depth. For example, emotions such as disgust can eliminate large numbers of alternatives from the search space. In general, emotions can narrow down choice sets more effectively and for a longer time than cognitive search tools.

Stopping Rules

Search for alternatives and cues must be stopped at some point. Strategies in the adaptive toolbox employ stopping rules that do not try to compute an optimal cost-benefit tradeoff as in optimization under constraints. Rather, building blocks for stopping involve simple criteria that are easily ascertained. In Simon's (1955) satisficing models, search is stopped when the first alternative is found that is as high as or better than the aspiration level; aspiration levels may go up or down depending on the time spent searching. Selten's (1998) aspiration adaptation theory provides a more general framework in which several goals, each with an aspiration level, exist, and the goals need not be commensurable. Simple rules for stopping search for cues are employed by Take The Best, Take The Last, and other heuristics, where search is stopped as soon as the first cue that favors one alternative is found (Gigerenzer and Goldstein 1996).

Building blocks for stopping need not be cognitive, as in these three examples. There are certain adaptive problems where cognitive stopping tools, such as comparison between an alternative and an adaptation level, are vulnerable to

instability. The moment a more attractive alternative comes into sight, the chosen alternative might be discarded and search taken up again. For adaptive problems such as rearing children, dispensing with one's wife or husband every time a more attractive partner comes in sight might not be a successful strategy. In these situations, emotions can function as effective tools for stopping search. For example, love can stop search for partners more effectively and for a longer time than an aspiration level comparison, and, in addition, strengthen commitment to the loved one. Similarly, feelings of parental love, triggered by one's infant's presence or smile, can prevent cost-benefit computations as to whether it really pays to stay with one's child. The question of whether or not it is worthwhile to endure all the sleepless nights and physical constraints associated with infant care simply never arises once the emotion is triggered. Emotions illustrate domain-specific rather than domain-general mechanisms of stopping search.

Decision Rules

Once search has been stopped, a decision or inference must be made. Models of judgment and decision making have ignored search and stopping rules traditionally and have focused exclusively on the decision rule: Are cue values combined by multiple linear regression? By Bayes's rule? Or in some other fashion? There is no evidence that humans could perform the extensive computations demanded by multiple regression or Bayesian networks to make judgments and decisions in situations with large numbers of cues. However, this is not to say that fewer computations and less information imply significantly less accuracy, not to mention irrationality. For example, simple linear models that use only unit weights (+1 or –1), and forego the matrix computations linear multiple regression demands, can make predictions about as well as regression (e.g., Dawes 1979).

Models of rationality rely on weighting and summing. Simple linear models dispense with optimal weighting; heuristics that use *one-reason decision making* dispense with summing. For example, Take The Best and other lexicographic heuristics rely only on one cue to make the decision and ignore all others. An apparently paradoxical result is that Take The Best — which uses less information and fewer computations than multiple regression does — can make the *more* accurate predictions. This result has been obtained for two-alternative choice problems involving predictions about biological, demographic, economic, ecological, and psychological variables, including homelessness rates in U.S. cities, Chicago's inner city high-school dropout rates, fertility of individual charr fishes, professors' salaries, and sales prices of houses (Czerlinski et al. 1999).

Thus, simple, psychologically plausible decision tools need not be inferior to complex combination schemes; there are situations where there is no trade-off between simplicity and accuracy. The study of these conditions is part of the study of ecological rationality.

Incommensurability between goals is a psychological phenomenon that optimization models postulating a common currency cannot handle. Models of bounded rationality can (e.g., Selten 1998), as mentioned before. Incommensurability between cues or reasons is a second psychological phenomenon that prevents optimization. For example, when Darwin pondered whether to marry or not, the pro and contra reasons he wrote down on a piece of paper included having children and having conversations with clever friends. Children and conversations are not of a common currency for many of us. The question of how many conversations with clever friends equal having one child will be rejected, based on the invalid assumption that everything has a price tag. Moral institutions are built on the principle that some things have no price: doctorates, military honors, and love (Elster 1979). One-reason decision making accepts the possibility of incommensurable reasons and does not impose a summing-up of values. A heuristic such as Take The Best goes with the best reason and ignores the rest.

ECOLOGICAL RATIONALITY

Ecological rationality is possibly the most important idea for understanding why and when bounded rationality works. Consider two classes of strategies: simple lexicographic ordering and multiple regression. Traditional definitions of rationality are concerned with the internal order of beliefs and inference, such as consistency. By internal criteria, one might conclude that lexicographic strategies are poor strategies because, unlike multiple regression, they employ one-reason decision making, ignore much of the information available, and some lexicographic strategies even produce intransitive judgments. For example, when Keeney and Raiffa (1993) discuss lexicographic ordering, they declare that it "is more widely adopted in practice than it deserves to be" because "it is naively simple" and "will rarely pass a test of 'reasonableness' "(pp. 77–78).

Environmental Structure

The notion of ecological rationality, by contrast, is not concerned with internal criteria. The question of ecological rationality concerns the match between a strategy and an environment. A match concerns structural properties. For example, consider a lexicographic strategy and the task of inferring which of two alternatives scores higher on a criterion. Lexicographic strategies are *noncompensatory*, that is, the first cue on which two alternatives differ determines the choice, and no further cue or combination of cues can reverse this decision. Consider an environment consisting of M binary cues, $C_1, \ldots, C_M$. These cues are noncompensatory for a given strategy if every cue C_j outweighs any possible combination of cues after C_j, that is, C_{j+1} to C_M. In the special case of a weighted linear model with a set of weights, $W = \{w_1, w_2, w_3, \ldots, w_M\}$, a strategy is noncompensatory if for every $1 \leq j \leq M$ we have

$$w_j > \sum_{k>j} w_k \,. \qquad (3.1)$$

An example is the set of weights {1, 1/2, 1/4, 1/8, 1/16}. For a lexicographic strategy called Take The Best, which uses the cue order $C_1, \ldots, C_M$, the following result can be proven: If an environment consists of cues that are noncompensatory for a given linear model, this model cannot be more accurate than the faster and more frugal Take The Best (Martignon and Hoffrage 1999).

This example illustrates a match between a strategy and an environment with respect to a property, noncompensatoriness. Heuristics that are matched to particular environments allow agents to be *ecologically rational*, making adaptive decisions that combine accuracy with speed and frugality. The degree to which a match exists (e.g., the cues in the environment may only be approximately noncompensatory) determines how accurate a heuristic is. In the present example, a match makes the simple Take The Best as accurate as multiple regression, that is, a judgment based on one reason is as good as one based on many reasons. Given that simple heuristics tend to be more robust when environments are noisy and information is scarce (see below), one-reason decision making can actually become more accurate than regression.

Robustness

Simple heuristics can be successful for two reasons: they can exploit *environmental structure*, as the example above illustrates, and they can be *robust*, that is, generalize well to new problems and environments. If there is uncertainty in an environment, in the sense of some degree of unpredictability and changing environments, robustness becomes an issue. A model with many free parameters can achieve a good fit to a given body of data but may not generalize well to new data if it *overfitted* the old data. Overfitting occurs when a model with more parameters fits a sample of data better than a model with fewer parameters but makes less-accurate predictions for a new data sample than the simpler model. Complex models with many free parameters, such as multiple regression or Bayesian methods, tend to overfit in environments where information is noisy or fluctuating, particularly when forced to make predictions from small samples.

Akaike (1973) discovered a way of estimating the degree of overfitting, which becomes larger as the number of parameters increases, resulting in higher rates of error and a larger sum of squares (see also Forster and Sober 1994). There are statistical techniques that expend considerable computational power and time trying to determine the point at which a model maximizes its predictive accuracy without overfitting. Fast and frugal heuristics sidestep this expenditure. Their simplicity helps them to avoid overfitting and to make robust predictions without doing these computations.

To summarize, the reasonableness of models of bounded rationality derives from their ecological rationality, not from coherence or an internal consistency of choices. A strategy is ecologically rational to the degree that it is adapted to the information in an environment, whether the environment is physical or social.

Social Rationality

The study of social rationality is a special case of ecological rationality when environments consist of other agents with which to interact. Humans are one of the few species that cooperate with genetically unrelated conspecifics for mutual benefit, and economic markets and educational institutions are products of this reciprocal altruism.

Social rationality adds a further class of goals to decision making: social goals that are important for creating and maintaining social structure and cooperation. These goals include transparency (i.e., making decisions that are understandable and predictable by the group with which one associates), fairness (i.e., making decisions that do not violate the expectations between people of equal social standing), and accountability (i.e., making decisions that can be justified and defended [Tetlock 1983]).

Social imitation can help make decisions with limited time and knowledge. Heuristics such as "eat what your peers eat" and "prefer mates picked by others" can speed up decision making by reducing the time spent on information search. Forms of social rationality can be found in the animal world, as well. For instance, female guppies tend to copy the mate choices of other female guppies — a tendency strong enough to reverse their prior preferences for one male over another (Dugatkin 1996). Female quail use a related form of mate copying (Galef and White 1998). In humans, mate copying is enhanced by the media; even academic hiring often seems to be under the spell of social imitation.

To summarize, the adaptive toolbox contains boundedly rational strategies that employ social norms, social imitation, and social emotions in addition to the cognitive building blocks outlined earlier. These additional heuristic tools are particularly important in the realm of social rationality.

REFERENCES

Akaike, H. 1973. Information theory and an extension of the maximum likelihood principle. In: Second International Symposium on Information Theory, ed. B.N. Petrov and F. Csaki, pp. 267–281. Budapest: Akademiai Kiado.

Brooks, R.A., and L.A. Stein. 1993. Building Brains for Bodies. A.I. Memo No. 1439, August. Cambridge, MA: MIT, Artificial Intelligence Lab.

Brunswik, E. 1955. Representative design and probabilistic theory in a functional psychology. *Psych. Rev.* **62**:193–217.

Czerlinski, J., G. Gigerenzer, and D.G. Goldstein. 1999. How good are simple heuristics? In: Simple Heuristics That Make Us Smart, G. Gigerenzer, P.M. Todd, and the ABC Research Group, pp. 97–118. New York: Oxford Univ. Press.

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

Dawes, R.M. 1979. The robust beauty of improper linear models in decision making. *Am. Psych.* **34**:571–582.

Driver, P.M., and D.A. Humphries. 1988. Protean Behavior: The Biology of Unpredictability. Oxford: Oxford Univ. Press.

Dugatkin, L.A. 1996. Interface between culturally based preferences and genetic preferences: Female mate choice in *Poecilia reticulata. Proc. Natl. Acad. Sci. USA* **93**:2770–2773.

Elster, J. 1979. Ulysses and the Sirens: Studies in Rationality and Irrationality. Cambridge: Cambridge Univ. Press.

Forster, M., and E. Sober. 1994. How to tell when simpler, more unified, or less ad hoc theories will provide more accurate predictions. *Brit. J. Phil. Sci.* **45**:1–35.

Galef, B.G., and D.J. White. 1998. Mate-choice copying in Japanese quail, *Coturnix coturnix japonica. Anim. Behav.* **55**:545–552.

Gigerenzer, G. 1991. How to make cognitive illusions disappear. Beyond heuristics and biases. *Eur. Rev. Soc. Psych.* **2**:83–115.

Gigerenzer, G., and D.G. Goldstein. 1996. Reasoning the fast and frugal way: Models of bounded rationality. *Psych. Rev.* **103**:650–669.

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

Gigerenzer, G., and P.M. Todd. 1999. Fast and frugal heuristics: The adaptive toolbox. In: Simple Heuristics That Make Us Smart, G. Gigerenzer, P.M. Todd, and the ABC Research Group, pp. 3–34. New York: Oxford Univ. Press.

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

Hertwig, R., U. Hoffrage, and L. Martignon. 1999. Quick estimation: Letting the environment do the work. In: Simple Heuristics That Make Us Smart, G. Gigerenzer, P.M. Todd, and the ABC Research Group, pp. 209–234. New York: Oxford Univ. Press.

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

Keeney, R.L., and H. Raiffa. 1993. Decisions with Multiple Objectives. Cambridge: Cambridge Univ. Press.

Leibnitz, G.W. 1951. Preface to the universal science. In: Selections, ed. P.P. Wiener. New York: Charles Scribner's Sons. (Original published 1677).

Martignon, L., and U. Hoffrage. 1999. Why does one-reason decision making work? A case study in ecological rationality. In: Simple Heuristics That Make Us Smart, G. Gigerenzer, P.M. Todd, and the ABC Research Group, pp. 119–140. New York: Oxford Univ. Press.

McGonigle, B., and M. Chalmers. 1992. Monkeys are rational. *Qtly. J. Exp. Psych.* **45B**:189–228.

Russell, S.J., and D. Subramanian. 1995. Provably bounded-optimal agents. *J. Artif. Intel. Res.* **1**:1–36.

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

Selten, R. 1998. Aspiration adaptation theory. *J. Math. Psych.* **42**:191–214.
Sen, A. 1993. Internal consistency of choice. *Econometrica* **61**:495–521.
Simon, H.A. 1955. A behavioral model of rational choice. *Qtly. J. Econ.* **69**:99–118.
Simon, H.A. 1956. Rational choice and the structure of environments. *Psych. Rev.* **63**:129–138.
Simon, H.A. 1987. Bounded rationality. In: The New Palgrave. A Dictionary of Economics, ed. J. Eatwell et al., pp. 266–268. London: Macmillan.
Stigler, G.J. 1961. The economics of information. *J. Pol. Econ.* **69**:213–225.
Tetlock, P.E. 1983. Accountability and complexity of thought. *J. Pers. Soc. Psych.* **45**:74–83.
Todd, P.M., and G.F. Miller. 1999. From pride and prejudice to persuasion: Satisficing in mate search. In: Simple Heuristics That Make Us Smart, G. Gigerenzer, P.M. Todd, and the ABC Research Group, pp. 287–308. New York: Oxford Univ. Press.
Vriend, N.J. 1996. Rational behavior and economic theory. *J. Econ. Behav. Org.* **29**:263–285.
Wimsatt, W.C. 1999. Re-engineering Philosophy for Limited Beings: Piecewise Approximations to Reality. Cambridge: Harvard Univ. Press.
Winter, S.G. 1975. Optimization and evolution in the theory of the firm. In: Adaptive Economic Models, ed. R.H. Day and T. Groves. New York: Academic.