



ELSEVIER

Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

Journal of Economic Psychology 24 (2003) 143–165

JOURNAL OF
**Economic
Psychology**

www.elsevier.com/locate/joep

Bounding rationality to the world

Peter M. Todd *, Gerd Gigerenzer

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

Abstract

Simon proposed that human rationality is bounded by both internal (mental) and external (environmental) constraints. Traditionally, these constraints have been seen as independent, leading to a notion of bounded rationality that is either the attempt to do as well as possible given the demands of the world – the notion of optimization under constraints – or as the sub-optimal outcome of the limited cognitive system – the realm of cognitive illusions. But there is a third possibility, following Simon's original conception: rather than being unrelated, the two sets of bounds may fit together like the blades in a pair of scissors. The mind can take advantage of this fit to make good decisions, by using mental mechanisms whose internal structure exploits the external information structures available in the environment. In this paper we lay out a research program for studying simple decision heuristics of this sort that expands on Simon's own search for mechanisms of bounded rationality. We then illustrate how these heuristics can make accurate decisions in appropriate environments, and present detailed examples of two heuristics inspired by Simon's ideas on recognition-based processing and satisficing in sequential search.

© 2003 Elsevier Science B.V. All rights reserved.

PsycINFO classification: 2140; 2340

JEL classification: B31; D81

Keywords: Bounded rationality; Cognitive limits; Heuristics; Satisficing; Recognition heuristic; Sequential search

* Corresponding author. Tel.: +49-30-82-406-347; fax: +49-30-82-406-394.

E-mail address: ptodd@mpib-berlin.mpg.de (P.M. Todd).

1. Introduction

What does it mean to be perfectly rational? For Simon, the dream of rationality as embodied in the *Homo economicus* of economics and psychology required a decision maker to be a “supremely skillful actor, whose behavior could reveal something of the requirements the environment placed on him but nothing about his own cognitive make-up” (Simon, 1981, pp. x–xi). That is, the perfectly rational actor’s cognitive machinery is so powerful as to be able to follow any script set by the environment, with nary a forgotten line or missed cue. Of course, such perfection is a fiction, for humans and for artificial systems – our rationality is bounded, and our acting is constrained. The question then becomes, what is the nature of the bounds on our bounded rationality?

The traditional view has been that we are hemmed in by two unrelated sets of bounds: purely external ones, such as the costs of searching for information in the world, and, independently, purely internal constraints, such as limits on the speed with which we can process information and limits on the amount of information we can hold in working memory (e.g., Simon, 1981, Chapter 3). Given these constraints, bounded rationality can be seen either as the attempt to do as well as possible given the demands of the world – the notion of optimization under constraints – or as the suboptimal outcome of the limited cognitive system – the realm of irrationality and cognitive illusions.

But there is another possibility regarding the bounds, external and internal, that surround our rationality: rather than being separate and unrelated, the two sets of bounds may be intimately linked. As Simon put it, “Human rational behavior . . . is shaped by a scissors whose two blades are the structure of the task environments and the computational capabilities of the actor” (Simon, 1990, p. 7). These two blades – the two sources of bounds on our rationality – must fit together closely for rationality to cut. While the external bounds may be more or less immutable from the actor’s standpoint, the internal bounds comprising the capacities of the cognitive system can be shaped, for instance by evolution or development, to take advantage of the structure of the external environment (Todd, 2001). From this perspective, then, we can see bounded rationality as the positive outcome of the two types of bounds fitting together.¹ In other words, humans exhibit *ecological rationality* (Gigerenzer, Todd, & the ABC Research Group, 1999; Todd & Gigerenzer, 2000; Todd, Fiddick, & Krauss, 2000) – making good decisions with mental mechanisms whose internal structure can exploit the external information structures available in the environment.

In this paper we discuss the legacy of Simon’s concept of bounded rationality in the form of these three interpretations of the idea. We begin with the version that Simon decried: optimization under constraints. The second interpretation is one that

¹ See Shepard (2001) for another view of how the internal structures of the mind are matched to the enduring external structures of the world, and Todd and Gigerenzer (2001) for the relationship between this view and Simon’s.

Simon sometimes endorsed: erroneous deviations from a rational standard. The third is the most positive vision, indicating a new way to conceive of rationality: as a fit between the structures of the mind and the environment. We go on to present our development of this third way into the research program of ecological rationality, in which we study the particular simple, bounded, heuristics that fill up the mind's adaptive toolbox. Two examples, recognition-based inference and satisficing heuristics for sequential search, show how these simple heuristics can exploit the structure of information in the decision maker's environment to make good choices. Moreover, bounded decision mechanisms are not merely the only option available given the constraints of the agent and the world; less information and processing can actually enable greater accuracy than more in some cases. Finally, we indicate some of the research directions that remain to be explored within the program of ecological rationality.

2. Three faces of bounded rationality

The seeds of Simon's notion of bounded rationality appeared in his 1947 book *Administrative Behavior*, and grew to encompass satisficing a few years later (Simon, 1955, 1956). It was clear what Simon was opposing at that point: the concepts of full or substantive rationality, maximization of expected utility, the ideal of *Homo economicus*, and just plain optimization. But exactly what he was proposing was broad (and unknown) enough to be understood in different ways by different people. The three major interpretations have seen bounded rationality as meaning optimization under constraints, human irrationality and consequent cognitive illusions, or ecological rationality. We consider each in turn (see Gigerenzer, in press, for an extended and more personal account).

2.1. *Focusing on the world: Bounded rationality as optimization under constraints*

In models of full rationality, all relevant information is assumed to be available to *Homo economicus* at no cost. Real humans, however, need to search for information first. In an attempt to render economic theory more realistic, Stigler (1961) introduced constraints on full rationality, such as information not being free and humans having limited time and money to search for it. The idea of optimization under constraints proposes some set of these limitations, which usually stem from external factors in the world like information costs and search times, while retaining the ideal of optimization. In this common doctrine, the bounds in bounded rationality are just another name for constraints, and bounded rationality is merely a case of optimizing under constraints. Following this view, for instance, a person who wants to buy a used car of a certain brand stops looking at further alternatives and goes with the best seen so far as soon as the costs of further search – both direct costs and opportunity costs – exceed those of its benefits.

Introducing real constraints does make this approach more realistic, but maintaining the ideal of optimization, that is, calculating an optimal stopping point, does

not. What is lost is psychological plausibility, because such an ideal of optimization invokes new kinds of omniscience, being able to foresee what additional information further search would bring, what it would cost, and what opportunities one would forgo during that search. Thus, retaining the ideal of optimization can make models of optimization under constraints more demanding than models of full rationality, both mathematically and psychologically (Sargent, 1993). Furthermore, the possibility of performing *any* kind of optimization is problematic in many natural situations, as Simon has often pointed out. As a consequence, this interpretation of his idea of bounded rationality is one that he spoke against (Gigerenzer, in press).

2.2. Focusing on the mind: Bounded rationality as cognitive illusions

Another conception of bounded rationality held by many psychologists and some economists is that it means internal cognitive limitations and the systematic errors – also called irrationality, biases, and cognitive illusions – that the mind's constraints lead to in judgment and decision making. For instance, in his article “Bounded rationality in individual decision making,” Camerer (1997, p. 179) summarizes anomalies in decisions and errors in judgments and calls this the “exploration of procedural (bounded) rationality of individuals.” This view has spread from psychology into economics and law, shaping new research areas such as behavioral economics (e.g., Camerer, 1995) and law and economics (e.g., Jolls, Sunstein, & Thaler, 1998). In Camerer's (1995, p. 588) words, “The goal is to test whether normative rules are *systematically* violated and to propose alternative theories to explain any observed violations.” Conlisk (1996, p. 672) shows the strong connection between this vision of bounded rationality and economic thinking when he ends a section of his paper on “Evidence of bounds on rationality” by saying “the bias evidence suggests that people are capable of a wide variety of substantial and systematic reasoning errors relevant to economic decisions.” The evidence in question has led to a list of well-known cognitive biases such as base rate neglect, overconfidence bias, and the sunk-cost effect (Kahneman, Slovic, & Tversky, 1982).

The irrationality interpretation is seldom spelled out in much detail, but its underlying idea can be summarized like this: The meaning of bounded rationality is that humans have internal cognitive limitations, which express themselves in errors in judgment and decision making; therefore, the study of errors is the study of bounded rationality (see Conlisk, 1996, for an argument along these lines). Simon's own early writing reflects this idea: “. . . behavior is determined by the irrational and nonrational elements that bound the area of rationality” (Simon, 1947, p. 241). But the problem is that demonstrations of cognitive illusions in fact accept as normative the very optimization theories that Simon fought against. A true theory of bounded rationality need not rely on optimization theories, neither as descriptions nor as norms of behavior, as we will see in the following section.

Why is bounded rationality not the same as irrationality? Simon (1990) gave the answer in the form of the analogy mentioned earlier: Bounded rationality is like a pair of scissors, with the mind as one blade, and the structure of the environment as the other. To understand behavior, one has to look at both and how they fit to-

gether. In other words, to evaluate cognitive strategies as rational or irrational, one needs to also analyze the environment, because a strategy is rational or irrational only with respect to a particular (physical or social) environment (Simon, 1955, 1956). The study of cognitive illusions and errors, however, focuses on the first blade, and compares it (the workings of the mind) with laws of probability and logic rather than with the structure of the environment.² One blade alone does not work as well as two; by introducing a properly fit second blade (crucial aspects of environment structure), apparently stable cognitive illusions can be made to disappear.

As an example of such scissors-repair, Gigerenzer, Hoffrage, and Kleinbölting (1991) theoretically derived and experimentally demonstrated that two well-studied cognitive illusions, the overconfidence bias and the hard–easy effect, disappear when the underlying questions are randomly sampled from an environment rather than systematically selected – that is, when people face an appropriate environmental structure. (Juslin, Winman, & Olsson, 2000, confirmed this initial demonstration in a quantitative review of over one hundred extant studies.) Note that this is different from the “debiasing” approach often taken by followers of the error-oriented perspective of bounded rationality. Generally in debiasing studies, rather than changing the information structure of the environment as above, broader methods such as providing greater incentives and greater learning opportunities are used in an attempt to help people overcome their reasoning biases. These efforts though have typically met with limited success, leaving the negative impression that “biases are not fragile effects which easily disappear, but rather substantial and important behavioral regularities” (Conlisk, 1996, p. 671). Taking both blades of Simon’s scissors into account can help to undermine this impression.

2.3. Putting mind and world together: Bounded rationality as ecological rationality

The two preceding perspectives see bounded rationality in a rather negative light, as something that usually keeps us from being truly optimal or properly rational. This arises in part because the internal and external constraints on decision makers are seen as being unaligned and hence often at odds – accordingly, if we had greater memory or better computational ability, perhaps, then we might not need to pay the costs for searching for extra information. But this opposition between internal and external bounds need not be the case (and, for systems that have adapted their internal structure to the external world over time, typically would not be the case). Instead, as Simon’s scissors metaphor indicates, bounded rationality can be seen as emerging from the joint effect of two interlocking components: the internal limitations of the (human) mind, and the structure of the external environments in which the mind operates. This fit between the internal cognitive structure and the external information structure underlies the perspective of bounded rationality as ecological

² There are some exceptions; for instance, Tversky and Kahneman (1973) wrote of availability being an “ecologically valid cue” for frequency judgments (p. 209), reflecting the environment’s structure.

rationality – making good (enough) decisions by exploiting the structure of the environment.

The first blade of Simon's scissors implies that humans "must use approximate methods to handle most tasks" (Simon, 1990, p. 6). These methods include recognition processes that largely obviate the need for further information search, heuristics that guide search and determine when it should end, and simple decision rules that make use of the information found. Simon's second blade is also of crucial importance because it can explain when and why simple decision mechanisms perform well: if the structure of the mechanism is adapted to the structure of the information in the environment. Heuristics that are matched to particular environments allow agents to be ecologically rational, making adaptive decisions that combine accuracy with speed and frugality. (We call the heuristics "fast and frugal" because they process information in a relatively simple way, and they search for little information.) The study of ecological rationality thus involves analyzing the structure of environments, the structure of heuristics, and the match between them.

3. The ecological rationality research program

The goal of the study of ecological rationality is an understanding of the particular decision mechanisms people and other animals use to make good decisions given particular structures of information in the environment. By letting the world do some of the work – by relying on the presence of particular useful information patterns – the decision mechanisms themselves can be simpler; hence our focus on simple, fast and frugal heuristics. (This is certainly not to say that all of cognition is handled by such heuristics, but we expect that a large proportion of the inferences and choices people make can be explained through their use.) The overall approach we take to studying heuristics draws on methodologies from a variety of disciplines, including experimental psychology, computer modeling, mathematical analysis, evolutionary biology, and cognitive anthropology. These are combined in a multi-stage process that begins with identifying important decision tasks (e.g., on psychological or evolutionary grounds – see Todd, 2000) and specifying the structure of information in the environment that can be used to make those decisions. After this initial groundwork, we proceed by (1) proposing computational models of candidate heuristics that are realistically based on human competences, and testing whether they work via simulation; (2) mathematically analyzing when and how the heuristics work with particular environmental structures; and (3) experimentally testing when people use these heuristics. We now turn to a more detailed consideration of these three steps.

3.1. *Computational modeling of specific heuristics*

To study particular heuristics in detail, it is crucial to develop computational models that specify the precise steps of information gathering and processing involved in generating a decision, so that the heuristic can be instantiated as a testable

computer program. This allows us to be sure we have identified the important components that will make a particular mechanism actually work in a particular environment, and to discover the implications of our assumptions more readily (Simon, 1981, pp. 18–19). In particular, we have modeled heuristics as being made up of building blocks that guide the search for alternatives to decide between or information about the available alternatives, that stop that search, and that make a decision based on the search results. In more detail, these building blocks can be described as follows.

Search. A central concept in Simon's notion of bounded rationality is search (Simon, 1955, 1956). Because most decisions outside of laboratory tasks are not made with the useful options or information at hand or in mind, we must seek this information in the world or in memory. When knowledge that can guide search is limited or absent, search can be simply random. With more available knowledge, options or information about options can be sought in some order determined by a criterion related to usefulness, or in an order based on a recollection about which options or cues worked previously when making the same type of judgment.

Stopping. To enable fast inferences and decisions, search for alternatives or information must be terminated at some point in a readily computable way. For example, one simple stopping rule is to cease searching for information about available alternatives and make a decision as soon as the first cue or reason that favors one alternative is found. This stopping rule does not need to compute an optimal cost-benefit trade-off as in optimization under constraints; in fact, it need not compute any costs or benefits at all. For search among sequentially encountered alternatives, aspiration-level stopping rules (Simon, 1955) can be used.

Decision. Once search has been stopped, an inference or decision must be made. Many cognitive process models have focused exclusively on this stage (i.e., without explicitly modeling search and stopping rules) and proposed that all pieces of information are weighted and integrated in some linear and nonlinear form. In contrast, fast and frugal heuristics can make use of simple decision strategies, such as one-reason decision making and elimination, that follow naturally from the concepts of limited search and stopping. For instance, a decision or inference could be based on only one cue or reason, whatever the total number of cues found during search. Such one-reason decision making does not need to weight or combine cues, and so no common currency between cues need be assumed.

Different types of these building blocks can be put together to form a variety of fast and frugal heuristics. In turn, these heuristics can be thought of as making up part of the *adaptive toolbox*: the collection of specialized cognitive mechanisms that evolution and learning has built into the human mind. Four main classes of heuristics in the adaptive toolbox have received the most attention so far (Gigerenzer et al., 1999).

Ignorance-based decision making. Some very simple heuristics can actually rely on a lack of knowledge to make appropriate decisions. The *recognition heuristic* (Goldstein & Gigerenzer, 1999, 2002), for instance, based on the fundamental adaptation in humans and animals to recognize faces, voices, or names, exploits the information carried by the recognition, and lack thereof, of particular objects in certain

environments. A decision maker who recognizes the name of one object (e.g., a city) but not of a second one (e.g., another city) can use that pattern of recognition to infer that the recognized object has a higher value on a criterion (e.g., population). In laboratory experiments, 90% or more of participants rely on the recognition heuristic in situations where it is ecologically rational, that is, when recognition (e.g., of brand names, stocks, sports teams) is correlated with the criterion (e.g., quality, market performance, winning a game). (This heuristic is covered in more detail in a later section.)

One-reason decision making. When information beyond mere recognition is available, so that the recognition heuristic cannot be used, a second class of simple heuristics becomes applicable: those that rely on just a single cue to make a decision (Gigerenzer & Goldstein, 1996, 1999). To perform the task of inferring which of two options or objects has a higher value on a criterion, these one-reason decision heuristics proceed through the following steps: (1) select a cue dimension as directed by the search building block and look for the corresponding cue values of each option – the *Take The Best* heuristic, for instance, selects cues in order of their validity, while the *Minimalist* heuristic checks them in random order, and *Take The Last* checks them in order of most recent usefulness; (2) compare the two options on their values for that cue dimension; (3) if the options do not differ, then return to the beginning of this loop (step 1) to look for another cue dimension; (4) if the options do differ, then stop (this is the stop-search building block); and then (5) choose the option with the cue value indicating a greater value on the choice criterion (the decision building block). Here, search can involve many cues, but only one will be used to determine the choice, minimizing the amount of computation that must be done.

Elimination heuristics. A third class of heuristics uses elimination, which is particularly useful when larger numbers of objects are involved in categorization or estimation tasks. *Categorization by Elimination* (Berretty, Todd, & Martignon, 1999), which is similar to Tversky's (1972) Elimination by Aspects model of preference-based choices, assigns a stimulus to a particular category by using one cue after another in a specified order to narrow down the set of remaining possible categories until only a single one remains. The *QuickEst* heuristic (Hertwig, Hoffrage, & Martignon, 1999) is designed to estimate the quantitative values of objects along some criterion while using as little information as possible through another elimination procedure.

Satisficing heuristics for sequential search. When alternatives are encountered sequentially (such as houses and potential spouses), Simon's (1955) concept of *satisficing* by means of aspiration levels provides a tool for choice. An aspiration level is set, and search for alternatives is stopped as soon as the aspiration is met. Simple mechanisms for setting the aspiration level (and revising it), such as checking the first few alternatives and taking the best as the aspiration level for further search, can prove ecologically rational both in individual and mutual search (e.g., for a mating or trading partner, see Simão & Todd, in press; Todd & Miller, 1999). (Satisficing search heuristics are also covered in more detail later.)

After specifying particular heuristics in the form of computable algorithms, the performance of these heuristics can be assessed by testing them on data from real-

world environments. While there may be no optimal strategies in many real-world environments, we can establish reasonable performance criteria. As a measure of the success of a heuristic, we compare its performance with the actual requirements of its environment, which can include making accurate decisions, in a minimal amount of time, and using a minimal amount of information. We have thus replaced the multiple coherence criteria – such as following the laws of logic and probability – with multiple correspondence criteria relating to real-world decision performance (Gigerenzer & Todd, 1999). These multiple criteria might be weighted differently in particular decision environments; in some situations, for instance, it may be more important to make a decision quickly rather than focusing on accuracy. However, one of the surprising empirical results to emerge from testing simple heuristics is that they need not always make such trade-offs. For instance, in comparing the performance of one-reason decision making heuristics against more traditionally rational mechanisms (including multiple regression and cue tallying) on a set of 20 real-world data sets, it was found that the fast and frugal heuristics used less information (about a third the number of cues on average) and still came close to – and, in the case of generalizing to new data, even beat – the performance of the more information-hungry and computationally expensive standard benchmark algorithms (Czerlinski, Gigerenzer, & Goldstein, 1999). Thus, in appropriately structured environments, simple heuristics can be faster, more frugal, and more accurate at the same time – no trade-off need be considered.

3.2. Analysis of the ecological rationality of simple heuristics

How is such surprising performance possible? To find out, the next step in the ecological rationality research program is to analyze the structure of environments, the corresponding structure of the heuristics, and how the two fit together. Standard statistical models, and standard theories of rationality, aim to be as general as possible, so they make broad and mathematically convenient assumptions about the data to which they will be applied. But the way information is structured in real-world environments often does not follow convenient simplifying assumptions. For instance, whereas most statistical models are designed to operate on datasets where means and variances are independent, Pearson (1897) noted that in natural situations these two measures tend to be correlated, and thus each can be used as a cue to infer the other (Einhorn & Hogarth, 1981, p. 66). While general statistical methods strive to ignore such factors that could limit their applicability, evolution would seize upon informative environmental dependencies like this one and exploit them with specific heuristics if they would give a decision-making organism an adaptive edge.

Finding out when and how structures of information in environments can be used to good advantage by simple heuristics is thus a central component of the ecological rationality research program. Some of the types of information structures that have been explored so far, and their impact on the performance of particular heuristics, are listed below (see Gigerenzer et al., 1999, for more details).

Noncompensatory information. The Take The Best heuristic equals or outperforms any linear decision strategy when information is noncompensatory, that is, when the

potential contribution of each new cue falls off rapidly so that combinations of later cues cannot outweigh earlier ones (Martignon & Hoffrage, 1999). Such environments seem fairly commonplace, at least in an approximately noncompensatory form.

Scarce information. Take The Best outperforms linear models on average when few cues are known relative to the number of objects (Martignon & Hoffrage, 1999).

J-shaped distributions. Many environments are characterized by a distribution of criterion values in which small values are common and big values are rare (a so-called “J-shaped” distribution). Such distributions can be formed through natural phenomena such as accretionary growth (e.g., creating a J-shaped distribution of city populations or business sizes). The QuickEst heuristic (Hertwig et al., 1999) estimates quantities about as accurately as, and with considerably less information than, more standard complex estimation strategies when the criterion to be estimated follows a J-shaped distribution.

Decreasing populations. In situations where the set of alternatives to choose from shrinks over time, such as in a seasonal mating pool, a satisficing heuristic that commits to an aspiration level quickly will outperform rules that sample many alternatives before setting an aspiration (Dudey & Todd, 2002; Todd & Miller, 1999).

While simple heuristics are designed to apply to specific environments, they typically do not contain enough detail to match any one environment precisely. General strategies that can be made to conform to a broad range of environments, on the other hand, can end up being too highly focused to be of much real use – having a large number of free parameters to adjust can be a hindrance. This is because of the important difference between the two typical applications of a strategy, *fitting* (modeling a given set of data) and *generalization* (predicting new data). In fitting, it is usually true that the more parameters a model has, and the more information (cues) it uses, the better it will fit given data. In generalization, on the other hand, more is not necessarily better – a computationally simple strategy that uses only some of the available information can be more robust, making more accurate predictions for new data, than a computationally complex, information-hungry strategy. This failure of generalization, a phenomenon known as *overfitting*, strikes models with many free parameters, from multiple linear regression to neural networks, when they try to make sense of every piece of information they encounter. As Simon (1981, p. 44) put it, “Although uncertainty does not . . . make intelligent choice impossible, it places a premium on robust adaptive procedures instead of strategies that work well only when finely tuned to precisely known environments.”

Robustness goes hand in hand with speed, accuracy, and especially information frugality. Fast and frugal heuristics can reduce overfitting by ignoring the noise inherent in many cues and looking instead for the “swamping forces” reflected in the most important cues. Thus, simply using only one or a few of the most useful cues can automatically yield robustness. Furthermore, important cues are likely to remain important – the informative relationships in the environment are likely to hold true even when the environment changes. In contrast, the random fluctuations of noise and even the effects of smaller systematic factors may well frequently vary. Because of this difference, fast and frugal heuristics that pay attention to systematic informa-

tive cues while overlooking more variable uninformative cues can ride out environmental change without suffering much decrement in performance.

Of course, this immediately raises the question of how we know which cues *are* the most informative ones to use in the first place. Heuristics do not work well merely by using little information – they must use *appropriate* (little) information. For some decisions, evolution has equipped us with knowledge of the most important cues to pay attention to (e.g., a food's taste, a parent's face). In other settings, we have mechanisms for learning useful cues, individually or from our culture (much research remains to be done to uncover these mechanisms). We also design and construct our own environments, explicitly or implicitly, to include salient useful cues that reduce our need to search for information; for instance, a wedding ring is a good cue of a married state, and a turn signal is a good cue of a driver's upcoming behavior. (Note that the problem of selecting which cues to use is related to the problem of selecting which heuristic to use, an open question discussed again at the end of this paper.)

In sum, by matching the structure of information in the environment with the structure implicit in their building blocks, heuristics can be accurate without being too complex. In addition, by being simple, these heuristics can avoid being *too* closely matched to any particular environment, and hence can escape the curse of overfitting. This marriage of structure with simplicity produces the counterintuitive situations in which there is little trade-off between being fast and frugal and being accurate.

3.3. *Empirical investigations of simple heuristics*

The research program described so far encompasses three big questions: (1) What are reasonable heuristic principles for guiding information or alternative search, stopping search, and making a decision using the results of that search? (2) When and why do these heuristics perform well, that is, how can they be ecologically rational? (3) How well do fast and frugal heuristics actually perform in real-world environments? By precisely modeling cognitive processes that work with limited resources and fit to particular environment structures, answering these questions helps us to show how ecological rationality emerges from real psychology. Exploring just these three questions is sufficient if we are interested in investigating new heuristics for various applied settings – the realms of artificial intelligence and decision-support systems, for instance. But because we are also concerned to discover the principles that guide natural human and animal behavior, we must add a fourth question to our research program: What is the evidence that humans or animals use specific fast and frugal heuristics?

As Simon pointed out, experimental tests of particular cognitive mechanisms are a challenge, because effectively adaptive behavior “will reflect characteristics largely of the outer environment . . . and will reveal only a few limiting properties of the inner environment” (Simon, 1981, p. 66). Thus, gathering data about the outcomes of decision making is typically insufficient to determine the underlying mechanisms; moreover, many mechanisms can generate the same outcomes in common situations. One

approach is then to present people with uncommon situations and observe how their decision mechanisms “break” – this is the method usually adopted in the generation and study of cognitive illusions (e.g., Piattelli-Palmarini, 1994), and has been proposed by Simon as well.³ But such a method is inappropriate for understanding how heuristics and environment structures fit together well. Instead, experiments that focus on elucidating the process of decision making – again, as also proposed by Simon (Ericsson & Simon, 1993) – are likely to be more profitable (see, e.g., Rieskamp & Hoffrage, 1999).

Experimental evidence is growing that humans do indeed use simple heuristics to make decisions in an ecologically rational manner, using as little information as possible and tailoring their information and option search to the structure available in the environment. Experts have been shown to base their judgments on surprisingly few pieces of information (Shanteau, 1992). In other studies, people have been found to trade off the effort involved in making a choice against the accuracy of that choice, and choose a quick and simple decision strategy that would achieve the desired balance (Payne, Bettman, & Johnson, 1993). Situations in which people use only a single piece of information to make a choice have also been studied, both in the case of the recognition heuristic (Goldstein & Gigerenzer, 1999, 2002), and the Take The Best and other one-reason decision heuristics (Bröder, 2000; Newell & Shanks, in press; Rieskamp & Hoffrage, 1999).

The evidence for the use of simple heuristics can also take the form of demonstrations of their impact on other aspects of human cognition. For instance, by exploring the effects of the Take The Best heuristic when it is used in memory updating, the first process model of the well-known hindsight effect in memory was proposed and experimentally supported (Hoffrage, Hertwig, & Gigerenzer, 2000). This model predicts about 75% of individual occurrences of hindsight bias and reversed hindsight bias, whereas earlier verbal accounts were unable to make any predictions at the level of individual items and participants.

In addition, it is important to look for the presence of simple heuristics in settings outside of the laboratory. Discovering how people process information to make judgments and decisions in their everyday lives – as well as how to structure information so that better decisions can be made (see Gigerenzer & Hoffrage, 1995) – can have important real-world implications in domains ranging from deciding on guilt in the courtroom to choosing whether or not to have a risky medical test. Empirical studies have shown that taking the ecologically rational decision heuristics of humans into account when communicating medical or legal information can help people to make better decisions in such domains (Hoffrage, Lindsey, Hertwig, & Gigerenzer, 2000).

³ “In a benign environment we would learn from the [system] only what it had been called upon to do; in a taxing environment we would learn something about its internal structure – specifically about those aspects of the internal structure that were chiefly instrumental in limiting performance.” (Simon, 1981, p. 16)

4. Two examples of simple heuristics

To give a further concrete indication of how ecological rationality can be studied, we next present in more detail two types of bounded decision mechanisms that Simon considered crucial (Simon, 1990; see also Simon, personal communication quoted in Gigerenzer, in press): recognition-based inference, and satisficing search keeping track of the best item seen so far. In both cases, we describe specific heuristics that have been proposed and the tests that have been used to explore when and how they can be ecologically rational.

4.1. Ignorance-based decision mechanisms

One of the simplest forms of decision that can be made is to select one option from two possibilities, according to some criterion on which the two can be compared. What simple cognitive mechanisms can be used to make this type of decision? This will depend on the amount and type of information that is available in the environment. If the only information available is whether or not each possibility has ever been encountered before, then the decision maker can do little better than rely on his or her own partial ignorance, choosing recognized options over unrecognized ones. This kind of “ignorance-based reasoning” is embodied in the recognition heuristic mentioned earlier (Goldstein & Gigerenzer, 1999, 2002): When choosing between two objects (according to some criterion), if one is recognized and the other is not, then select the former. For instance, Norway rats have evolved to behave according to a rule of this type, preferring to eat things they recognize through past experience with other rats (e.g., items they have smelled on the breath of others) over novel items (Galef, 1987).

Following the recognition heuristic will be ecologically rational – that is, will yield correct responses more often than would random choice – in those environments in which exposure to different possibilities is positively correlated with their ranking along the decision criterion being used. Thus, the rats’ food preference copying presumably evolved because the things that other rats have eaten (i.e., recognized items) are more often palatable than are random (unrecognized) items sampled from the environment. Such useable correlations are likely to be present in species with social information exchange where important environmental objects are communicated and unimportant ones are ignored (see Todd & Kirby, 2001, for a simulation testing this idea). But could an organism faced with a recognized and an unrecognized alternative use more than just recognition, for instance, use what else it knows or can find out about the recognized object, to make a better choice between the two? In some cases yes – but often extra knowledge and the time taken to access and process that information will not yield a much improved answer.

Goldstein and Gigerenzer (1999) investigated this question in an environment consisting of pairs of all German cities with over 100,000 inhabitants (i.e., the 83 largest German cities), where the task was to decide which of the two cities in each pair was larger. Because we hear about large cities more often than small cities, using recognition to decide which of two cities is larger will often yield the correct answer

(in those cases where one city is recognized and the other is not). If only recognition could be used to make this decision (with average recognition rates for each city assessed from a survey of American students, and around half of the cities assumed to be recognized), accuracy across all pairs could reach 65% – well above chance performance. When nine informative cues were added for the recognized cities and multiple regression used over these cues to determine the larger city, accuracy only went up by about another 7%, and most of this increase was for more accurate choices in pairs where both cities were recognized. Thus in this case, extra time and knowledge provide little benefit over the recognition heuristic in those decisions where one alternative is recognized and the other is not.

In fact, adding more knowledge for the recognition heuristic to use, by increasing the proportion of recognized objects in an environment, can even decrease decision accuracy. This *less-is-more effect*, in which an intermediate amount of (recognition) knowledge about a set of objects can yield the highest proportion of correct answers, is straightforward from an information theory perspective, but surprising from a cognitive one. Knowing more is not usually thought to *decrease* decision-making performance, but when using simple heuristics that rely on little knowledge, this is exactly what is theoretically predicted, and can be found experimentally as well (Goldstein & Gigerenzer, 1999, 2002).

How widely these findings apply in other environments and tasks must still be determined. There is already some evidence that the recognition heuristic can outperform strategies relying on much more knowledge and computation in a very complex and dynamic environment: the stock market. When deciding which companies to invest in from among those trading in a particular exchange, the recognition heuristic would lead us to choose just those that we have heard of before. Such a choice can be profitable assuming that more-often-recognized companies will typically have better-performing stocks. This is a reasonable assumption, given that many companies that are large and successful now spend vast amounts to build their brands, working to get consumers to recognize their name (and associate it with as many aspects of daily life as possible) much more than striving to develop their products (Klein, 2000).

This assumption has also been experimentally tested by asking several sets of people what companies they recognized and forming investment portfolios based on the most familiar firms (Borges, Goldstein, Ortman, & Gigerenzer, 1999). Nearly 500 people in the United States and Germany were asked which of 500 American and 298 German traded companies they recognized. To form portfolios based on very highly recognized companies, we used the American participants' responses to select their top 10 most-recognized German companies, and the German responses to choose the top 10 most-recognized American firms. In this trial performed during 1996–1997, the simple ignorance-driven recognition heuristic beat highly trained fund managers using all the information available to them, as well as randomly chosen portfolios (which fund managers themselves do not always outperform). This does not prove that people use the recognition heuristic when making such choices (though common investment advice suggests this is so), but it does show that heuristics designed to be fast and use as little information as possible can perform well even in a complex environment. Furthermore, here the measure of performance is not

merely decision accuracy measured in terms of the correct choices, but rather the more realistic yardstick of resources accrued. Some “correct” stock choices could be worth much more than others, and it is important to compare decision heuristics where possible in a manner that takes the value or significance of their individual choices into account.

4.2. *Satisficing heuristics for sequential choice*

The recognition heuristic can help agents make good choices with a minimum of information when confronted with multiple options to choose between simultaneously. But many of the important decisions we face are not structured this way – rather, we often must choose between a set of options that we do not see all at once, but one after another, sequentially. This happens for instance when we are trying to find the best price on tomatoes as we drive from store to store, or deciding between potential mates met one at a time, or finding a new house among those that come on the market over some extended period. These situations are typically characterized by low (or zero) probability of being able to recall, or return to and choose, previously seen options once they have been passed by (e.g., a nice house on the market one month will probably be sold by the next). The problem then becomes one of deciding when to stop searching and select the currently available option.

How can we make a good choice in such a sequential setting? This depends on how much knowledge we have. For instance, we might know the distribution of possible values of the options we might encounter, and we might have a limited number of possible options that we can check, in which case we can compute the optimal point to stop search so as to balance the gain from the particular chosen option against the risk of missing a better option later on (DeGroot, 1970). But what is the *minimal* amount of information that could be used for these sequential choices? While this is a difficult question to answer, one possible approach is to use only a single bit of information registered for each option, corresponding to whether (bit = 1) or not (0) the current option is the best one seen so far. Then, each sequence of possible options would be conceived of as a binary string of 1’s and 0’s, and the decision maker would just have to decide where to stop (usually on a 1).

Note that a fair amount of processing may have to go into determining the value of this *best-so-far* cue for each option, comparing its criterion value to that of the previous best-so-far candidate (which must therefore be stored somehow); but the same is true of the recognition heuristic, which makes use of a simple binary cue – is the object *recognized* or not? – that must itself be computed with further (potentially complex) machinery such as a face recognition mechanism. We are interested here in the simpler higher-order cognitive decision mechanisms that can be built upon the outputs of possibly more complex lower-level mechanisms including perception and categorization (see Todd, 1999).

Given this minimal-information framework, we now want to ask, how can a decision maker go about finding a good option, and how well will the decision maker be able to do? The answers all depend on what our definition of “good” is (see

Dudey & Todd, 2002, for a more detailed exploration). We start with perhaps the most widely studied version of this sequential search problem, in which “good” is defined as having a high probability of finding and choosing the single best candidate in the available set of options. This version is known in probability theory as the *secretary (or dowry) problem*. In the secretary problem it is the searcher’s aim to find the (one) best of N applicants for a job as a secretary, with the assumption that they can all be arranged on a single dimension of overall quality. The applicants are presented to the searcher sequentially in a random order and the searcher has no knowledge about the distribution of the applicants’ quality values. With each new applicant, the searcher learns the quality (or current rank in the strictest version) of this applicant, and then must choose between stopping the search and thus hiring the current applicant or continuing the search to look for a better applicant. If the searcher continues she cannot go back and choose an earlier applicant – that is, there is no ability to “recall” past applicants in this search.

To maximize the chance of selecting the best applicant in the secretary problem setting, the searcher should sample the first 37% of the applicants and the select the first candidate thereafter who is better than all previous applicants. (See Ferguson, 1989, for a review of the literature on the secretary problem and this optimal solution.) One simple way to do this is to set an aspiration level equal to the quality of the best applicant seen so far, and then after 37% of the applicants have been seen, remember this level and use it to stop the search with the next applicant seen who exceeds this threshold – an approach following Simon’s (1955) notion of aspiration-based satisficing search, though here with an optimal method for setting the aspiration level. Alternatively, using the best-so-far cue, one can just pass by the first 37% of the population and then take the next applicant encountered with a value of 1 on this cue. While this approach is optimal given the problem setting, it only yields a success rate of 37%, and on average requires searching through 74% of the population before stopping.

Seale and Rapoport (1997) studied the behavior of people engaged in a sequential search problem equivalent to the secretary problem. They regarded the optimal 37% rule as only a benchmark and put more emphasis on finding simple heuristics that would be a better explanation of the actual behavior of the participants in their studies. They proposed three such heuristics, namely a *cut-off rule*, a *candidate count rule*, and a *successive noncandidate rule*. The cut-off rule is a generalization of the optimal solution, where searchers simply pass by a certain number (not necessarily the optimal number) of applicants and then hire the next encountered top-ranked applicant. Each applicant that is top-ranked at the moment of being interviewed is termed a *candidate*. The candidate count rule then simply implies choosing the n th candidate seen. The successive noncandidate rule on the other hand chooses the first candidate that is interviewed after observing at least k consecutive noncandidates – that is, it stops searching after the gap between successive candidates has grown sufficiently large. All of these heuristics demand only minimal cognitive requirements (mainly counting and computing the best-so-far cue).

To test these heuristics, Seale and Rapoport (1997) first showed by means of Monte Carlo simulations that the rules can be very effective at finding the best ap-

plicant (in comparison to the benchmark 37% rule). With the best parameter choice, the cut-off rule is obviously identical to the optimal benchmark solution, while the successive noncandidate rule also matches the benchmark 95% of the time; the candidate count rule does rather more poorly. Next, Seale and Rapoport compared the predictions of the three search heuristics with the actual behavior of their participants when searching through sequences of 80 values (presented as relative ranks). The cut-off rule was most successful in this regard, being most consistent with observed search behavior (accounting for between 41% and 74% of the decisions made) for 21 out of 25 participants. The successive noncandidate rule gave the best fit for 8 of the 25 (matching the performance of the cut-off rule for some participants), while the candidate count rule best matched only one participant's decisions. Thus, some form of cut-off rule seems to work well and to describe what people are doing in sequential search situations of this type.

Cut-off rules using the minimal best-so-far cue also work well when the goals of search are changed. In the original secretary problem, the single best option is the only good one. But this strictest-possible payoff function would be found in few natural situations. In many species, most animals find some mate, some food source, and some place to live (and indeed most companies find some secretary), and thus receive some payoff, even if not the highest possible payoff. Thus in these cases, a payoff proportional to the quality of the alternative chosen (e.g., the typing speed of the hired secretary) is more appropriate than the all-or-none payoff of the standard secretary problem. Alternatively, the search payoff function could fully reward only choices made in the top 10% or 25% of all available alternatives (e.g., if only a quarter of available habitats have enough water for survival) and give zero payoff to all other choices. Todd and Miller (1999) examined the performance of cut-off rules with such different payoff structures, and found that little search was needed before a good aspiration level could be set. For instance, after checking only a dozen applicants and then taking the best seen so far, the chance of picking an applicant in the top 10% (out of 100–1,000 possible applicants) was near its maximum level of over 80%.

More generally, in environments where possible options are encountered independently, one after another, with no knowledge of their distribution and no recall available, searchers can be ecologically rational by simply using the best-so-far cue with the cut-off (or successive noncandidate count) rule and a low parameter value. This approach (e.g., Todd and Miller's "Try a dozen" rule) is a simple way to perform well, according to multiple criteria, in sequential search. However, the ecological rationality of such rules breaks down in slightly different search environments. When the population of potential options is decreasing in value over time, then using the best-so-far cue is a bad idea, because a new best-so-far option is unlikely ever to be seen (Dudey & Todd, 2002). Such decreasing-value environments could arise for instance in a pool of potential mates as a breeding season progresses and the best individuals are mated and taken out of the pool first. Rules that could stop search on applicants that are not candidates (i.e., not the best seen so far) would clearly fare better in these environments.

5. Benefits of cognitive limits

Studying ecological rationality enables us to go beyond the widespread fiction that basing decision making on more information and computation will always lead to more accurate inferences. There is a point at which increasing information and information processing can actually do harm, as we indicated earlier in our discussion of robustness through simplicity and the less-is-more effect when using the recognition heuristic. But more than this, there are cases where cognitive limitations actually seem to be beneficial, enabling new functions that would be absent without them, rather than constraining the possible behaviors of the system (see Hertwig & Todd, *in press*, for a detailed account). When Simon (1981, p. 66) spoke of the “limiting properties of the inner environment,” he downplayed this possibility.⁴

Here we mention just a pair of examples. First, limitations in human cognitive capacity may actually lay the groundwork for inferences of causality in terms of the early detection of covariation. Kareev (1995, 2000) has argued that limited working memory capacity forces people to rely on small samples of information drawn from real-world environments or long-term memory. Such small samples, however, may have an advantage: They can increase the chance of detecting a correlation in population. This is because a small sample has a greater likelihood than a large sample of exhibiting a magnified estimate of the real correlation in the sampled population. Kareev has amassed theoretical and empirical evidence for this idea, leading him to propose that cognitive limitations in working memory are not a liability but in fact enable important adaptive functions such as the detection of covariation. While his conclusions have proven controversial, they are important if correct: The ability to detect contingencies early – that is, after little sampling of the environment – would be particularly useful in domains where the costs of false alarms are outweighed by the benefits of exploiting discovered contingent relationships.

Stronger arguments about the benefits of cognitive limitations have been made for the domain of language learning. Newport (1990) has proposed that the constraints of the developing brain are a necessary hindrance for a first language to be learned fluently. This leads to another type of “less-is-more” phenomenon: Lesser ability to remember and process morphological and semantic components into form-meaning mappings in young children allows them more rapidly and accurately to learn those mappings that they do acquire, and then to build further upon these as language learning proceeds. Late language learners, in contrast, may falter when attempting to learn the full range of semantic mappings with their mature mental capacities all at once. Elman (1993) captured this phenomenon in a neural network model of language acquisition. He described the restrictions of the developing mind as enabling accurate early learning about a small portion of the environment, which then provides a scaffold to guide learning and hypothesizing about the rest of the

⁴ However, at other points Simon did emphasize the usefulness of constraints as opportunities and guides for creative processes, particularly in design (Simon, 1981). March (1978) follows this approach in considering how best to engineer decision systems that capitalize on people’s intelligent use of their own constraints in predicting their future preferences.

environment in fruitful, adaptive directions. Cognitive “constraints” are no longer a negative limitation of our (or our children’s) ability to behave adaptively in our environment—rather, “the early limitations on memory capacity assume a more positive character . . . they act as a filter on the input, and focus learning on just that subset of facts which lay the foundation for future success” (Elman, 1993, pp. 84–85).

Such arguments, while still under debate (see, e.g., Rohde & Plaut, 1999), suggest an even bolder hypothesis. Rather than adhering to the usual assumption that simple decision-making strategies have evolved (or arose through learning) in response to the cognitive limitations of the human mind, we can reverse this causality and ask whether the constraints may have arisen in support of the evolution of simple strategies. That is, given that cognitive capacities are free parameters that have been adjusted in the course of evolution, the ability to use simple heuristics may have required the evolution of no more than a certain limited amount of cognitive capacity necessary to execute those heuristics. This argument requires that simple heuristics had a selective advantage over more complex cognitive strategies (which would have required more processing power). But we have already seen that such advantages do exist, in terms of speed and robustness and enabling new functionality. Thus, the benefits of simple limited decision mechanisms may actually partly underlie the emergence of bounded rationality itself.

6. Where to next

In this paper we have seen some of the progress being made in the study of ecological rationality as inspired by Simon’s conceptions of a more psychologically plausible understanding of the workings of the mind. Exploring the contents of the mind’s adaptive toolbox – finding the environmentally matched heuristics we use in particular situations, as well as the building blocks that compose them – remains the overarching goal. A number of questions in more specific areas also fall within this goal, including:

Selecting heuristics. How does the mind choose which heuristic to use? A fast and frugal mind would not employ a meta-level algorithm that makes optimal cost–benefit computations to select a heuristic. The fact that heuristics are designed for particular tasks, rather than being general-purpose strategies, solves part of the selection problem by reducing the choice set. In other cases, the associations between appropriate heuristics for particular situations may be learned – the environment may take care of selection for us. Can we identify other selection mechanisms?

Social norms and emotions. What simple heuristics do people use for navigating the complexities of social domains? How are heuristics acquired socially, for instance through imitation, word of mouth, or cultural heritage? Can social norms, cultural strictures, historical proverbs, and the like enable fast and frugal social reasoning by obviating cost–benefit calculations and extensive information search? Relatedly, do emotions facilitate rapid decision making by putting strong limits on the search for information or alternatives, as when falling in love stops partner search and facilitates commitment? The search for evidence for the decision-making functions of

these cultural and emotional processes is one of the most important areas to be mapped out (see, e.g., Hanoch, 2002, for an approach following Simon).

Describing environment structure. What is the appropriate set of concepts for describing those aspects of environment structure, whether physical or social, that shape the design and performance of decision heuristics? Here one can turn for inspiration to other fields, including ecology and statistics, that have analyzed environment structure from different perspectives. Evolutionary psychology reminds us to reflect on possible differences between present and past environments, by considering the important adaptive problems our ancestors faced, the information available in their environment to solve those problems, and how these inputs have changed in the modern world (Cosmides & Tooby, 1987, p. 302). Furthermore, we should consider how environments are constructed by the decision makers living and acting in them, and how for instance socially emergent patterns can interact with the heuristics of individuals in turn.

Most of these questions were anticipated by Simon during his long career, and they remain as challenges today. In promoting the importance of the structure of the environment for understanding how real organisms can behave sufficiently well given their bounded rationality, Simon opened up a new direction for psychological and economic research that has been all-too-little explored since. As he put it at the end of his groundbreaking paper on environment structure in 1956:

The analysis set forth here casts serious doubt on the usefulness of current economic and statistical theories of rational behavior as bases for explaining the characteristics of human and other organismic rationality. It suggests an alternative approach to the description of rational behavior that is more closely related to psychological theories of perception and cognition, and that is in closer agreement with the facts of behavior as observed in laboratory and field. (Simon, 1956/1979, p. 28)

Our concept of ecological rationality is intended to highlight both of Simon's stated bounds on cognition, the internal and the external, the mind and the world, and the way that human rationality emerges between them. By exploring the particular fast and frugal heuristics and other psychological mechanisms that fill the mind's adaptive toolbox, and explicating how these mechanisms work by exploiting particular environment structures, we hope to fulfill some of the promise of the new direction laid out by Simon nearly half a century ago.

References

- Berretty, P.M., Todd, P.M., & Martignon, L. (1999). Categorization by Elimination: Using few cues to choose. In G. Gigerenzer, P.M. Todd, and the ABC Research Group, *Simple heuristics that make us smart* (pp. 235-254). New York: Oxford University Press.
- Borges, B., Goldstein, D.G., Ortmann, A., & Gigerenzer, G. (1999). Can ignorance beat the stock market? Name recognition as a heuristic for investing. In G. Gigerenzer, P.M. Todd, and the ABC Research Group, *Simple heuristics that make us smart* (pp. 59-72). New York: Oxford University Press.

- Bröder, A. (2000). Assessing the empirical validity of the “Take The Best” heuristic as a model of human probabilistic inference. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26, 1332–1346.
- Camerer, C. F. (1995). Individual decision making. In J. H. Kagel & A. E. Roth (Eds.), *The handbook of experimental economics* (pp. 587–703). Princeton, NJ: Princeton University Press.
- Camerer, C. F. (1997). Progress in behavioral game theory. *Journal of Economic Perspectives*, 11, 167–188.
- Conlisk, J. (1996). Why bounded rationality? *Journal of Economic Literature*, 34, 669–700.
- Cosmides, L., & Tooby, J. (1987). From evolution to behavior: Evolutionary psychology as the missing link. In J. Dupré (Ed.), *The latest on the best: Essays on evolution and optimization* (pp. 277–306). Cambridge, MA: MIT Press/Bradford Books.
- Czerlinski, J., Gigerenzer, G., & Goldstein, D. G. (1999). How good are simple heuristics? In G. Gigerenzer, P.M. Todd, and the ABC Research Group, *Simple heuristics that make us smart* (pp. 97–118). New York: Oxford University Press.
- DeGroot, M. H. (1970). *Optimal statistical decisions*. New York: McGraw-Hill.
- Dudey, T., & Todd, P. M. (2002). Making good decisions with minimal information: Simultaneous and sequential choice. *Journal of Bioeconomics*, 3, 195–215.
- Einhorn, H. J., & Hogarth, R. M. (1981). Behavioral decision theory: Processes of judgment and choice. *Annual Review of Psychology*, 32, 53–88.
- Elman, J. (1993). Learning and development in neural networks: The importance of starting small. *Cognition*, 48, 71–99.
- Ericsson, K. A., & Simon, H. A. (1993). *Protocol analysis: Verbal reports as data* (rev. ed). Cambridge, MA: MIT Press.
- Ferguson, T. S. (1989). Who solved the secretary problem? *Statistical Science*, 4, 282–296.
- Galef, B. G., Jr. (1987). Social influences on the identification of toxic foods by Norway rats. *Animal Learning & Behavior*, 18, 199–205.
- Gigerenzer, G. (in press). Striking a blow for sanity in theories of rationality. In M. Agier & J.G. March (Eds.), *Essays in honor of Herbert Simon*. Cambridge, MA: MIT Press.
- Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103, 650–669.
- Gigerenzer, G., & Goldstein, D.G. (1999). Betting on one good reason: The Take The Best heuristic. G. Gigerenzer, P.M. Todd, & the ABC Research Group, *Simple heuristics that make us smart* (75–95). New York: Oxford University Press.
- Gigerenzer, G., & Hoffrage, U. (1995). How to improve Bayesian reasoning without instructions: Frequency formats. *Psychological Review*, 102, 684–704.
- Gigerenzer, G., Hoffrage, U., & Kleinbölting, H. (1991). Probabilistic mental models: A Brunswikian theory of confidence. *Psychological Review*, 98, 506–528.
- Gigerenzer, G., & Todd, P.M. (1999). Fast and frugal heuristics: The adaptive toolbox. In G. Gigerenzer, P.M. Todd, and the ABC Research Group, *Simple heuristics that make us smart* (pp. 3–34). New York: Oxford University Press.
- Gigerenzer, G., Todd, P. M., & the ABC Research Group (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.
- Goldstein, D. G., & Gigerenzer, G. (1999). The recognition heuristic: How ignorance makes us smart. In G. Gigerenzer, P.M. Todd, and the ABC Research Group, *Simple heuristics that make us smart* (pp. 37–58). New York: Oxford University Press.
- Goldstein, D. G., & Gigerenzer, G. (2002). Models of ecological rationality: The recognition heuristic. *Psychological Review*, 109(1), 75–90.
- Hanoch, Y. (2002). “Neither an angel nor an ant”: Emotion as an aid to bounded rationality. *Journal of Economic Psychology*, 23, 1–25.
- Hertwig, R., Hoffrage, U., & Martignon, L. (1999). Quick estimation: Letting the environment do the work. In G. Gigerenzer, P.M. Todd, and the ABC Research Group, *Simple heuristics that make us smart* (pp. 209–234). New York: Oxford University Press.

- Hertwig, R., & Todd, P.M. (in press). More is not always better: The benefits of cognitive limits. In L. Macchi & D. Hardman (Eds.), *The psychology of reasoning and decision making: A handbook*. Chichester, UK: Wiley.
- Hoffrage, U., Hertwig, R., & Gigerenzer, G. (2000). Hindsight bias: A by-product of knowledge updating? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26(3), 566–581.
- Hoffrage, U., Lindsey, S., Hertwig, R., & Gigerenzer, G. (2000). Communicating statistical information. *Science*, 290, 2261–2262.
- Jolls, C., Sunstein, C. R., & Thaler, R. (1998). A behavioral approach to law and economics. *Stanford Law Review*, 50, 1471–1550.
- Juslin, P., Winman, A., & Olsson, H. (2000). Naive empiricism and dogmatism in confidence research: A critical examination of the hard–easy effect. *Psychological Review*, 107, 384–396.
- Kahneman, D., Slovic, P., & Tversky, A. (Eds.). (1982). *Judgment under uncertainty: Heuristics and biases*. Cambridge, UK: Cambridge University Press.
- Kareev, Y. (1995). Through a narrow window – working memory capacity and the detection of covariation. *Cognition*, 56, 263–269.
- Kareev, Y. (2000). Seven (indeed, plus or minus two) and the detection of correlations. *Psychological Review*, 107, 397–402.
- Klein, N. (2000). *No logo*. London: Flamingo.
- March, J. G. (1978). Bounded rationality, ambiguity, and the engineering of choice. *The Bell Journal of Economics*, 9(2), 587–608.
- Martignon, L., & Hoffrage, U. (1999). Why does one-reason decision making work? A case study in ecological rationality. In G. Gigerenzer, P.M. Todd, and the ABC Research Group, *Simple heuristics that make us smart* (pp. 119–140). New York: Oxford University Press.
- Newell, B. R., & Shanks, D. R. (in press). Take-the-best or look at the rest? Factors influencing ‘one-reason’ decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*.
- Newport, E. L. (1990). Maturational constraints on language learning. *Cognitive Science*, 14, 11–28.
- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. New York: Cambridge University Press.
- Pearson, K. (1897). On the scientific measure of variability. *Natural Science*, 11, 115–118.
- Piattelli-Palmarini, M. (1994). *Inevitable illusions: How mistakes of reason rule our minds*. New York: Wiley.
- Rieskamp, J., & Hoffrage, U. (1999). When do people use simple heuristics and how do we know this? In G. Gigerenzer, P.M. Todd, & the ABC Research Group, *Simple heuristics that make us smart* (pp. 141–167). New York: Oxford University Press.
- Rohde, D. L. T., & Plaut, D. C. (1999). Language acquisition in the absence of explicit negative evidence: How important is starting small? *Cognition*, 72(1), 67–109.
- Sargent, T. J. (1993). *Bounded rationality in macroeconomics*. Oxford, UK: Oxford University Press.
- Seale, D. A., & Rapoport, A. (1997). Sequential decision making with relative ranks: An experimental investigation of the “secretary problem”. *Organizational Behavior and Human Decision Processes*, 69(3), 221–236.
- Shanteau, J. (1992). How much information does an expert use? Is it relevant? *Acta Psychologica*, 81, 75–86.
- Shepard, R. N. (2001). Perceptual-cognitive universals as reflections of the world. *Behavioral and Brain Sciences*, 24(4), 581–601.
- Simão, J., & Todd, P. M. (in press). Modeling mate choice in monogamous mating systems with courtship. *Adaptive Behavior*.
- Simon, H. A. (1947). *Administrative behavior*. New York: Macmillan.
- Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69, 99–118. Reprinted in Simon, H.A. (1979), *Models of thought*, vol. I (pp. 7–19). New Haven, CT: Yale University Press.
- Simon, H. A. (1956). Rational choice and the structure of environments. *Psychological Review*, 63, 129–138. Reprinted in Simon, H.A. (1979), *Models of thought*, vol. I (pp. 20–28). New Haven, CT: Yale University Press.

- Simon, H. A. (1981). *The sciences of the artificial* (2nd ed). Cambridge, MA: MIT Press.
- Simon, H. A. (1990). Invariants of human behavior. *Annual Review of Psychology*, 41, 1–19.
- Stigler, G. J. (1961). The economics of information. *Journal of Political Economy*, 69, 213–225.
- Todd, P. M. (1999). Simple inference heuristics versus complex decision machines. *Minds and Machines*, 9(4), 461–477.
- Todd, P. M. (2000). The ecological rationality of mechanisms evolved to make up minds. *American Behavioral Scientist*, 43(6), 940–956.
- Todd, P. M. (2001). Fast and frugal heuristics for environmentally bounded minds. In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality: The adaptive toolbox (Dahlem Workshop Report)* (pp. 51–70). Cambridge, MA: MIT Press.
- Todd, P. M., Fiddick, L., & Krauss, S. (2000). Ecological rationality and its contents. *Thinking and Reasoning*, 6(4), 375–384.
- Todd, P. M., & Gigerenzer, G. (2000). Simple heuristics that make us smart. *Behavioral and Brain Sciences*, 23(5), 727–741.
- Todd, P. M., & Gigerenzer, G. (2001). Shepard's mirrors or Simon's scissors? Commentary on Shepard, R., Perceptual-cognitive universals as reflections of the world. *Behavioral and Brain Sciences*, 24(4), 704–705.
- Todd, P. M., & Kirby, S. (2001). I like what I know: How recognition-based decisions can structure the environment. In J. Kelemen & P. Sosik (Eds.), *Advances in artificial life: 6th European conference proceedings (ECAL 2001)* (pp. 166–175). Berlin: Springer-Verlag.
- Todd, P.M., & Miller, G.F. (1999). From pride and prejudice to persuasion: Satisficing in mate search. In G. Gigerenzer, P.M. Todd, and the ABC Research Group, *Simple heuristics that make us smart* (pp. 287–308). New York: Oxford University Press.
- Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review*, 79(4), 281–299.
- Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5, 207–232.