Environments That Make Us Smart

Ecological Rationality

Peter M. Todd[1,2] and Gerd Gigerenzer[1]

[1]Center for Adaptive Behavior and Cognition, Max Planck Institute for Human Development, Berlin, Germany, and [2]Program in Cognitive Science, Indiana University

ABSTRACT—*Traditional views of rationality posit general-purpose decision mechanisms based on logic or optimization. The study of ecological rationality focuses on uncovering the "adaptive toolbox" of domain-specific simple heuristics that real, computationally bounded minds employ, and explaining how these heuristics produce accurate decisions by exploiting the structures of information in the environments in which they are applied. Knowing when and how people use particular heuristics can facilitate the shaping of environments to engender better decisions.*

KEYWORDS—*ecological rationality; decision making; adaptive toolbox; simple heuristics; environment structure; bounded rationality*

The importance of looking at the world to understand the mind has long been appreciated by a few prominent thinkers. Charles Darwin held that environmental forces had shaped human behavior through natural selection, leading to the modern call by evolutionary psychologists to look to our ancestral world for the problems our minds are designed to solve. More than 50 years ago, Egon Brunswik urged psychologists to study the texture of natural environments and the corresponding structure of cues the mind relies on to infer the state of its surroundings. Roger Shepard spoke of the mind as a mirror, reflecting long-standing physical aspects of the world such as the 24-hour light–dark cycle. Herbert Simon proposed the metaphor of the mind and world fitting together like the blades of a pair of scissors—the two must be well matched for effective behavior to be produced, and just looking at the cognitive blade will not explain how the scissors cut. In each case, the world is a key for understanding the workings of the mind (Todd & Gigerenzer, 2001). However, prevailing explanations of behavior are still expressed most often in terms of personality traits, cognitive styles, brain-region activation patterns, preferences and utilities, and other assumed entities "inside" the mind.

Address correspondence to Peter M. Todd, Program in Cognitive Science, Indiana University, 1101 E. 10th Street, Bloomington, IN 47405; e-mail: pmtodd@indiana.edu.

The research program on *ecological rationality* aims to explicate the mind–world interactions underlying good decision making. We build on the foundations from Darwin, Brunswik, Simon, and others to create a framework for understanding how environment structure—in the form of useful patterns of available information in the world—can be exploited by heuristics in the head to produce adaptive behavior (Gigerenzer, Todd, & the ABC Research Group, 1999; Todd, Gigerenzer, & the ABC Research Group, in press). Heuristics—simple decision algorithms that can work well in appropriate environments—generate both routine behavior and important decisions, for inference, choice, group deliberations, and even moral issues. For instance, consider the puzzling observation that only 28% of Americans become potential organ donors but 99.9% of French people do. To find out why, one might administer personality tests, measure moral attitudes, add a knowledge exam, and then perform multiple regression on the lot to find some significant predictors—but not the answer. Rather, most Americans and French seem to rely on the same simple *default heuristic*: "If there's a default choice, stick with it." The difference is in the external (here institutional) default-setting environment in each country: In most U.S. states there is a no-organ-donation default, so one has to actively opt in to become a donor, while in France one has to opt out to *not* be a donor (Johnson & Goldstein, 2003). It is the interaction between a heuristic and its social, institutional, or physical environment that explains behavior. This adaptive view has policy consequences: It explains why pro-donor information campaigns have had only limited success and indicates that changing the legal default should be more effective.

Or consider the inference task of predicting which of two tennis players will win an upcoming Wimbledon match. This

decision can be made on the basis of pieces of information, or cues, that could be looked up about each player, such as whether they are past champions, how many games they have won this season, what their seeding by the Wimbledon experts is, and so on. More simply, one could ignore all of this information and just rely on the *recognition heuristic*: "If you recognize one player and not the other, then predict the recognized one will win" (Gigerenzer, 2007). Tennis novices make predictions in line with this heuristic often (90% of the time). More surprisingly, their collective recognition can be even more accurate (e.g., correct on 72% of men's 2003 matches) than the Wimbledon experts' ratings (69%). But the recognition heuristic will only perform so well in environments that it is suited to—namely, those where the "biggest" objects (like the biggest winners in sports) are frequently discussed and hence likely to be recognized.

Our research program has two components that correspond to the two blades of Simon's scissors:

(a) The study of the "adaptive toolbox" of decision mechanisms in the mind (see Table 1 for examples). The goal is to uncover and understand heuristics for inference and preference (e.g., tasks of categorization, estimation, and choice), their building blocks, the few pieces of information they use, and the evolved abilities they exploit. The methods employed are theoretical and experimental.
(b) The study of the ecological rationality of decision mechanisms. The goal is to determine what environmental structures enable a given heuristic to be successful, and where it will fail. The methods also include computer simulation and mathematical analysis.

The first component is explanatory, asking, for instance, how people make decisions about organ donation or winning sports competitors. The second is normative, determining what environment structures will increase organ-donation rates or help the recognition heuristic to predict match outcomes accurately. The study of ecological rationality, our focus here, requires clear definitions of both heuristics and environments.

STUDYING ECOLOGICAL RATIONALITY IN DECISION MAKING

The modern study of decision making began with the normative ideal that good decisions follow the mathematical prescriptions of Bayes's rule, or the maximization of expected utility. In these views, there is only one mental tool, and the question of this tool's ecological rationality—its fit to different environments—does not arise. But now there is an impressive body of experimental evidence showing that people often make decisions in an entirely different way: Humans rely on multiple simple decision heuristics, not one general-purpose calculus of rationality (e.g., Bröder & Schiffer, 2003; Gigerenzer, 2007). Individuals can certainly be led to use particular heuristics in inappropriate environments and consequently make errors, as the heuristics-and-biases research tradition emphasized (Kahneman, Slovic, & Tversky, 1982). The study of ecological rationality goes beyond this beginning: By (a) designing computational models of heuristics, it (b) specifies which environments are appropriate for which heuristics, and vice versa.

People often rely on a single reason to make decisions; but can this particular heuristic approach—restricting information

TABLE 1
Sample Heuristics, Environmental Structures That Make Them Ecologically Rational, and Surprising Predictions

Heuristic	Definition[1]	Ecologically rational if	Surprising predictions
Recognition	To decide which of two options is greater on some criterion, if only one option is recognized, choose that one.	Recognition is a valid cue (i.e., leads to correct decisions over half of the time)	Contradicting information about recognized object is ignored; recognizing fewer options can lead to greater accuracy.
Take The Best (see Fig. 1)	As above, but if both options are recognized, (1) search through cues in order of validity (2) stop search on first discriminating cue (3) choose option favored by this cue	Cue validities vary highly; moderate to high redundancy between cues	Can decide more accurately than multiple regression, neural networks, and exemplar models when generalizing to new data
Tallying (unit-weight linear model)	To estimate criterion for some object, count number of favoring cues.	Cue validities vary little; low cue redundancy (Hogarth & Karelaia, 2006).	Can decide as accurately as multiple regression
Try-a-dozen (satisficing)	To select a high-valued option from an unknown sequence, set an aspiration level at highest value seen in first 12 options, then choose next option that exceeds aspiration.	Unknown distribution of option values; no returning to previously seen options	Near-optimal performance over a wide range of sequence lengths (i.e., number of available options matters little)

Note. [1]For further details, see Gigerenzer et al. (1999).

use as much as possible—ever be reasonable? To answer, we must first define specific models of heuristics that use a single reason and then compare them with traditional rational approaches, which use several. We have proposed a class of "fast and frugal" heuristics that rely on the principle of one-reason decision making, such as the *Take The Best* heuristic (Gigerenzer & Goldstein, 1996; Gigerenzer et al., 1999; see Fig. 1). Take The Best searches for cues in order of their validity—that is, their rate of making correct decisions—and stops searching as soon as it finds a cue that distinguishes between the alternatives (which may be the first cue, or the second, or one further down the list); this is the single "best" available cue that the heuristic takes to make a decision.

Take The Best and other one-reason decision heuristics are frugal in that they do not look for any more information than is needed to make an inference. Such heuristics are also fast because they do not involve any complex computation—not even the multiplication and addition required by weighted additive mechanisms (standard decision models that add up all cue values multiplied by cue importance weights to find an overall value for each choice option). Furthermore, they work well to make correct inferences in particular types of environments. For instance, Take The Best will not do well compared to weighted additive or tallying mechanisms in environments where the distribution of cue importance is uniform (i.e., the available pieces of information are roughly equal in their usefulness). But many environments, such as those of consumer choice or mate choice, are characterized instead by a distribution of cue importance that falls off rapidly (a "J-shaped" distribution), so that the most influential cue is considerably more important than the second, which is considerably more important than the third, and so on. In such environments, Take The Best can outperform the weighted additive model, particularly when generalizing to new decisions (Gigerenzer et al., 1999; Hogarth & Karelaia, 2006). When environments are moderately unpredictable and learning samples are small, as with many social and economic situations, Take The Best also tends to make inferences as accurately as or better than multiple regression, neural networks, and other machine-learning models (Chater, Oaksford, Nakisa, & Redington, 2003). Thus, in challenging environments with high variability, low predictability, and little opportunity for learning, good decisions may nonetheless be made more often by simple mechanisms than by complex ones.

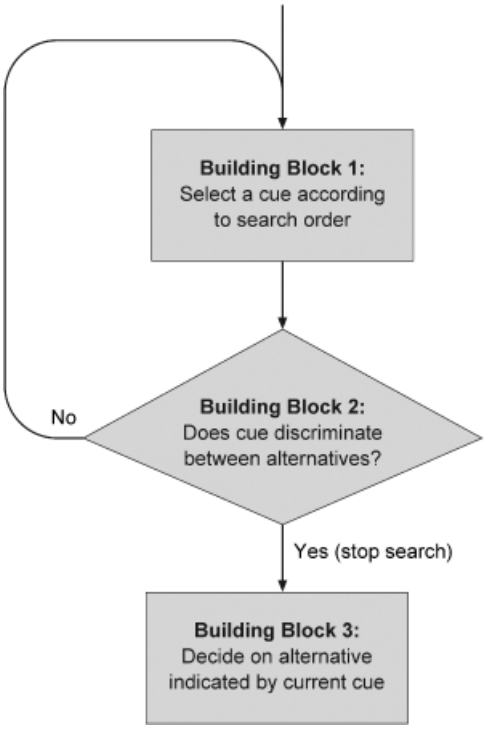

Fig. 1. Information processing steps in heuristics in the class of *one-reason decision mechanisms*. Many classes of heuristics consist of building blocks, including: (1) a *search rule* that determines what information (cues or features, in memory or external environment) is searched for, in what order; (2) a *stopping rule* that terminates search for cues; and (3) a *decision rule* that makes a decision using the cues found. In the class of heuristics shown here, the third building block uses only one cue—one reason—to make the decision. Different types of one-reason heuristics can be built by selecting building blocks that are adapted to particular environmental structure. For instance, if one knows which cues are more or less valid (individually accurate) in a particular domain, the validity-ordered search used as the first building block in the *Take The Best* heuristic will work well; while if one has little knowledge about an environment, random search through the cues can be adaptive (as embodied in the *Minimalist* heuristic, which selects cues to check in a random order).

Furthermore, because the human mind has been shaped by the adaptive processes of evolution and learning, we predict that people will tend to be ecologically rational themselves, often using simple decision heuristics that confer the twin advantages of speed and accuracy in particular environments. This prediction has been supported in numerous studies; for instance, one-reason decision mechanisms have been reported to be used when people must pay for information (Newell, Weston, & Shanks, 2003), must search for information in memory (Bröder & Schiffer, 2003), or are under time pressure—all situations in which it is advantageous to limit information search. Moreover, people are sensitive to the distribution of cues in an environment, appropriately applying either Take The Best or a weighted additive mechanism, depending on which will be more accurate (Rieskamp & Otto, 2006). Exactly how people are able to determine which type of environment they are in, and then which heuristics will be appropriate to apply, remains an open question.

RELEVANT STRUCTURE IN ENVIRONMENTS

The patterns of information that decision mechanisms may (or may not) be matched to can arise from a variety of environmental processes, including physical, biological, social, and cultural sources. First, brains on this planet have likely evolved longest to deal with patterns in the physical environment (e.g., how visual cues such as binocular disparity, size of projection on the retina, or occlusion can be used to decide which of two objects is closer). Many of these patterns can be characterized in terms of

cue validities (how often a cue indicates a correct decision), discrimination rates (how often a cue distinguishes between decision alternatives, regardless of its correctness), and redundancies (correlations between cue values across alternatives). Patterns in the occurrence of events in time and space also matter: People use *satisficing* heuristics that set aspiration levels (for instance, a minimal acceptable job salary or a maximal acceptable house-buying price) for searching through sequences of options when it is difficult to return to previously seen options and to tell what options lie ahead (as on job, housing, or even mating markets—Gigerenzer et al., 1999). And humans and other animals searching for resources that occur in patches (such as fruits on bushes or information on Web sites) do best using *incremental* or *decremental* rules, which increase or decrease the tendency to stay in a patch with each resource item found, depending on whether patches have widely varying or roughly equal numbers of resources (Hutchinson & Gigerenzer, 2005).

Probably most important for humans are structures in and from social environments. We can use satisficing heuristics to make ecologically rational decisions about other people as potential mates based on the sequential pattern of people we have previously encountered (Gigerenzer et al., 1999). We can use the *majority rule* (which simplifies computation by just tallying yes/no judgments rather than, for instance, weighting them according to importance of different judges) and other heuristics to make group decisions about people as potential employees, based on the distribution of information within our group (Todd et al., in press). Patterns of recognition knowledge gained by individuals via social exchange can be successfully exploited by the recognition heuristic mentioned earlier (Gigerenzer et al., 1999) to decide which of two items is larger (e.g., cities) or more successful (e.g., tennis players or stocks), and such recognition patterns are also given prominence in group decision making (Reimer & Katsikopoulos, 2004). But people do not use the recognition heuristic blindly—rather they seem to evaluate intuitively its ecological rationality for a given situation (via mechanisms that have been associated with particular neural structures—see Gigerenzer, 2007). For instance, few people rely on name recognition when they know that they have heard of a city (e.g., Chernobyl) for a reason that is uncorrelated with its size.

Environment structures are also specifically created by cultures or institutions to influence the behavior of others, as in the organ-donor example earlier. Sometimes this is felicitous, as when traffic laws determining intersection right-of-way are designed hierarchically to match our one-reason decision mechanisms (Todd et al., in press). In other cases, institutions create environment structures that do not fit well with people's decision mechanisms and instead cloud minds and lead to poor choices. For instance, information about medical treatments is often represented in terms of misleading relative risks (e.g., stating that mammography screening reduces mortality risk by 25%) whereas absolute risks would lead to greater understanding (e.g., explaining that 25% mortality reduction means going from 4 cancer deaths out of 1,000 without screening to 3 out of 1,000 with screening; see Todd et al., in press).

Finally, some forms of environment structure emerge without design through the social interactions of multiple decision makers. For instance, people choosing a city to move to are often attracted by large, vibrant metropolises, so that "the big get bigger," which can result in a J-shaped (or power-law) distribution of city populations (i.e., a few teeming burgs, a number of medium-sized ones, and numerous smaller towns). Such an emergent distribution can be seen in many domains where people make active choices, such as buying books (leading to a few bestsellers and many languishing titles) or visiting Web sites (a few Yahoos versus countless low-traffic personal Web pages). This structure can in turn be exploited by heuristics for choice or estimation that, for instance, rely on the fact that most objects are not big (Gigerenzer et al., 1999). Similarly, drivers may choose a parking space using a particular satisficing heuristic that creates a pattern of taken and still-available spots that serves as the environment for future drivers to search through with their own heuristic (Hutchinson & Gigerenzer, 2005). In these cases, individuals are, through the effects of their own choices, essentially shaping the environment in which they and others must make further choices, creating the possibility of a coadapting loop between mind and world.

LOOKING AHEAD

Studying ecological rationality as the fit between structures of information-processing mechanisms in the mind and structures of information in the world gives us three things to focus on: the mind (decision heuristics), the world (information patterns), and how they can match. So far, only a modest set of the simple heuristics contained in the mind's adaptive toolbox have been explored. To uncover more of its contents, we can proceed in two directions. Delving downwards, we need to expand our understanding of the set of building blocks and deeper evolved abilities (e.g., the capacities for recognition or for trust) that can combine to create decision mechanisms. Connecting upwards, it is necessary to consider how the adaptive toolbox of heuristics for inference and preference ties in with other cognitive, memory, perceptual, and motor systems to produce adaptive behavior (as has been done in implementing the recognition heuristic within a broader cognitive modeling framework—see Schooler & Hertwig, 2005).

Researchers have also started to put together a vocabulary for describing environment structures, for instance, in terms of cue validities and distributions of objects. This effort is still woefully incomplete and incohesive, though. Useful ways to describe psychologically relevant aspects of spatial structure, temporal patterns, and social environments (among others) must be developed or imported from other disciplines.

However, the greatest challenge is tying the two types of structure, mental and environmental, together. Researchers can

explore the ecological rationality of particular decision mechanisms by predicting their performance based on how well their building blocks fit to certain information patterns, and then testing the match via experimentation, simulation, and mathematical analysis in different environments. But a deeper theoretical account of the reasons why certain environmental patterns fit with certain heuristic mechanisms must still be developed. Knowing both when and why structures in the mind fit to structures in the world is necessary for understanding the ecological rationality of our decision heuristics and the environments that make us smart.

Recommended Reading

Gigerenzer, G. (2002). *Calculated risks: How to know when numbers deceive you*. New York: Simon & Schuster.

Gigerenzer, G., & Engel, C. (2006). *Heuristics and the law*. Cambridge, MA: MIT Press.

Gigerenzer, G., & Selten, R. (Eds.). (2001). *Bounded rationality: The adaptive toolbox*. Cambridge, MA: MIT Press.

Payne, J.W., Bettman, J.R., & Johnson, E.J. (1993). *The adaptive decision maker*. Cambridge, UK: Cambridge University Press.

Todd, P.M., & Gigerenzer, G. (2000). Simple heuristics that make us smart (précis). *Behavioral and Brain Sciences*, *23*, 727–741.

REFERENCES

Bröder, A., & Schiffer, S. (2003). "Take The Best" versus simultaneous feature matching: Probabilistic inferences from memory and effects of representation format. *Journal of Experimental Psychology: General*, *132*, 277–293.

Chater, N., Oaksford, M., Nakisa, R., & Redington, M. (2003). Fast, frugal, and rational: How rational norms explain behavior. *Organizational Behavior and Human Decision Processes*, *90*, 63–86.

Gigerenzer, G. (2007). *Gut feelings: The intelligence of the unconscious*. New York: Viking Press.

Gigerenzer, G., & Goldstein, D.G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, *103*, 650–669.

Gigerenzer, G., Todd, P.M., & the ABC Research Group. (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.

Hogarth, R.M., & Karelaia, N. (2006). "Take-the-Best" and other simple strategies: Why and when they work "well" with binary cues. *Theory and Decision*, *61*, 205–249.

Hutchinson, J.M.C., & Gigerenzer, G. (2005). Simple heuristics and rules of thumb: Where psychologists and biologists might meet. *Behavioural Processes*, *69*, 97–124.

Johnson, E.J., & Goldstein, D.G. (2003). Do defaults save lives? *Science*, *302*, 1338–1339.

Kahneman, D., Slovic, P., & Tversky, A. (1982). *Judgment under uncertainty: Heuristics and biases*. Cambridge, UK: Cambridge University Press.

Newell, B.R., Weston, N., & Shanks, D.R. (2003). Empirical tests of a fast and frugal heuristic: Not everyone "takes-the-best". *Organizational Behavior and Human Decision Processes*, *91*, 82–96.

Reimer, T., & Katsikopoulos, K.V. (2004). The use of recognition in group decision-making. *Cognitive Science*, *28*, 1009–1029.

Rieskamp, J., & Otto, P.E. (2006). SSL: A theory of how people learn to select strategies. *Journal of Experimental Psychology: General*, *135*, 207–236.

Schooler, L., & Hertwig, R. (2005). How forgetting aids heuristic inference. *Psychological Review*, *112*, 610–628.

Todd, P.M., & Gigerenzer, G. (2001). Shepard's mirrors or Simon's scissors? *Behavioral and Brain Sciences*, *24*, 704–705.

Todd, P.M., Gigerenzer, G., & the ABC Research Group. (in press). *Ecological rationality: Intelligence in the world*. New York: Oxford University Press.