THINKING AND REASONING, 2000, *6* (4), 375–384

Discussion

Ecological rationality and its contents

Peter M. Todd, Laurence Fiddick, and Stefan Krauss
Max Planck Institute for Human Development, Berlin, Germany

Our research group's book, *Simple Heuristics That Make Us Smart* (Gigerenzer, Todd, & the ABC Research Group, 1999; hereafter *Simple Heuristics*), reports on investigations of simple, fast and frugal heuristics that can effectively solve practical decision problems accurately and economically. We are happy that this book has sparked interest and debate in the fields of psychology and philosophy (and further in economics and animal behaviour), as evidenced by David Over's recent review in this journal (Over, 2000; see also the 30 commentaries discussed in Todd, Gigerenzer, & the ABC Research Group, 2000, and Mysterud, 2000). Over's review revealed certain confusions about our research programme, and we are grateful for the opportunity to expand on those issues here.

The central concept of *Simple Heuristics* is that of *ecological rationality*: how decision mechanisms can produce useful inferences by exploiting the structure of information in their environment. As the book's title indicates, we focus on a class of simple heuristics that are fast and frugal, using little time and little information, and which work well in part *because* of these limitations, not in spite of them as one might expect. Simple heuristics are specified in terms of building blocks that control their search for information in the environment (or in memory), stop that search, and use the information found to reach a decision. Because these heuristics are precisely defined, their ecological rationality can also be precisely defined: That is, we can say just what information structures in the environment will enable a given heuristic to make good decisions. For instance, one of the heuristics that we propose, called *Take The Best*, decides between two available options by searching through cues in order of their validity, stopping when the first cue is found that distinguishes the options, and selecting the option indicated by the higher cue value. This heuristic will be ecologically rational in those environments in which cues are noncompensatory, declining rapidly in validity (see Martignon & Hoffrage, 1999, for details).

Correspondence should be sent to Peter M. Todd, Center for Adaptive Behavior and Cognition, Max Planck Institute for Human Development, Lentzeallee 94, 14195 Berlin, Germany. Email: ptodd@mpib-berlin.mpg.de

http://www.tandf.co.uk/journals/pp/13546783.html

Thus the meaning of ecological rationality can be clearly specified for individual heuristics (though not necessarily easily specified—determining the conditions of ecological rationality makes up much of the work of studying decision mechanisms, often involving both mathematical analysis and computer simulation). Both these precise specifications and the broader implications of an ecological rationality perspective for simple heuristics in general are discussed in *Simple Heuristics*, and we review some of the implications in this article. Over (2000), however, focused on some topics not intended to be covered by the book, including massive modularity, evolutionary psychology, accurate representations of the world, frequency versus probability formats for information, and higher-level reasoning (e.g., deontic and causal inferences), so we address some of these additional points here as well.

ECOLOGICAL RATIONALITY

The traditional notion of classical rationality is framed largely in terms of theoretical reasoning, centring on the task of constructing true and general representations of the world.[1] Hence, it follows that classical rationality should adopt evaluative criteria that promote that outcome: logical consistency, coherence, and an effort to incorporate all available pieces of information. The kind of everyday practical reasoning that people spend much of their time involved in, however, is concerned with making useful decisions in the real world, for which the standards of classical rationality are often inappropriate. Instead, practical reasoning mechanisms should primarily be judged by whether or not they solve the problems that confront them (regardless of whether or not they build accurate representations of the world in the process). Furthermore, they must operate effectively within the constraints facing the decision maker, such as limited time and information, as captured by the concept of bounded rationality. Finally, the structure of information in the decision environment both constrains and enables the operation of many practical reasoning mechanisms

[1]The distinction between theoretical and practical reasoning is a long-standing one, dating back at least to Aristotle. Theoretical reasoning is concerned with elucidating truths about the world, whereas practical reasoning is concerned with making sound decisions and taking appropriate actions. Although we do not dwell on this distinction in our book, we present it here to show the connections to our research programme. There are several relevant differences between these two forms of reasoning (Jonsen & Toulmin, 1988): theoretical reasoning is idealised to cover abstractions or generalisations, whereas practical reasoning deals with concrete real phenomena; theoretical reasoning is atemporal because the conclusions it delivers are supposed to be universally true, whereas practical reasoning produces responses appropriate to particular situations; and theoretical reasoning is necessary in that its inferences are the product of logically consistent deductions from accepted axioms, whereas practical reasoning is presumptive in that its conclusions are open to doubt and may fail in abnormal circumstances. We believe that human decision making is more a matter of practice than theory—performing accurately, not representing the world accurately.

(including simple heuristics), leading to the concept of ecological rationality. Environmental influences can impact decision mechanisms through both the calibration of individuals to local environments and the adaptation of populations to ancestral environments. We consider these ideas in this section.

Ecological rationality and bounded rationality

A wide variety of decision mechanisms are possible for solving everyday practical reasoning problems. But there are good reasons for expecting that the mechanisms people and other animals actually use are often simple, economical, and robust, which has led us to focus particularly on what we call fast and frugal heuristics. These simple heuristics embody bounded rationality, as originally advocated by Herbert Simon—making reasonable decisions given the constraints that they face (*not* acting irrationally because of restrictions, as it is sometimes mischaracterised). The usual assumption is that the constraints that bound human rationality are internal ones, such as limited memory and computational power. But this view leaves out an important part of the picture: namely the external world and the constraints that it imposes on decision makers (Todd, 2000).

There are two particularly important classes of constraints that stem from the nature of the environment around us. First, because the external world is uncertain—we never face exactly the same situation twice—our mental mechanisms must be robust, able to generalise well from old instances to new ones. One of the best ways to be robust is to be simple, for instance by employing a mechanism containing few parameters. As a consequence, external uncertainty can impose a bound of simplicity on our mental mechanisms. Second, because the world is competitive and time is money, or at least energy, our decision mechanisms must generally be fast. The more time we spend on a given decision, the less time we have available for other important activities. To be fast, we must minimise the information or alternatives we search for in making our decisions (assuming that time needed to search for information outstrips that needed for mental computation, as is often the case). That is, the external world also constrains us to be frugal in what we search for.

But the external world does not just impose the bounds of simplicity, speed, and frugality on us—it also provides the means for staying within these bounds. A decision mechanism can stay simple and robust by relying on some of its work being done by the external environment—that is, by making use of patterns of structured information in the environment rather than trying to use a structured model of the world in one's head. A mechanism can also be frugal by taking environment structure into account in guiding what pieces of information to search for and in what order, rather than seeking all available cues in any order as if no structure were present. Thus simplicity, frugality, speed, and robustness go hand in hand with exploiting the structure of information in the environment-and these characteristics are the foundations of our conception of ecological rationality.

Herbert Simon has described human rationality as being shaped by a scissors whose two blades are the computational capabilities of individuals and the structure of the environments they face—if these two blades are not closely matched, then decision making will be ineffectual. What Simon's metaphor suggests—and what ecological rationality promotes—is that we should expand the bounds of bounded rationality outside the head and into the environment. This does not mean *replacing* the mind with the environment, which the research programmes of Egon Brunswik and J.J. Gibson might lead to, but rather exploring the adaptive mesh between the two, as has been advocated by Roger Shepard. Narrowly emphasising one of the two blades, either mind or environment, obscures their adaptive, interlocking fit. Amos Tversky and Daniel Kahneman began their study of heuristics by talking about this ecological fit to the environment, before their research programme focused on studying the biases that constrained rationality by itself can lead to. The ecological rationality perspective is intended to reinvigorate the programme of studying the mind in its environmental context.

Ecological rationality, environments, and domain-specificity

Ecological rationality implies a two-way relationship between simple heuristics and their environments. First, the success of simple heuristics is defined with respect to pragmatic goals in a particular environmental context. Thus external pragmatically relevant criteria, such as making adaptive behavioural choices in an efficient manner, are used to judge the performance of mental mechanisms, rather than internal coherence criteria, such as making logically consistent choices. This does not mean that the behavioural outputs of simple heuristics cannot ever be described by some rules of logic—often they can be—but it does imply that these decision mechanisms were not specifically built (by natural selection or learning) to *instantiate* such rules. Observing behaviour that follows coherence (or other) criteria does not by itself specify the form of the mechanism producing that behaviour.

Second, the success of simple heuristics is enabled by their fit to environmental structure. Exploiting the information structure of environments, and thereby letting the environment do some of the work of decision making, is what allows effective heuristics to be simple. Different environment structures can be exploited by—and hence call for—different heuristics, just as different tasks (e.g. two-option choice versus categorisation) call for different heuristics.[2]

[2]Although the nature of practical reasoning is such that universal problem-solving procedures are unlikely (Jonsen & Toulmin, 1988), there are at present so many opinions about how domain-specific mechanisms are instantiated in the mind (e.g. Fodor, 1983; Jackendoff, 1992; Karmiloff-Smith, 1992; Mithen, 1996; Sperber, 1994; Tooby & Cosmides, 1992), coupled with a scarcity of solid evidence in favour of any one specific view, that we can currently say little more than this about the domain-specificity of simple heuristics.

Over (2000) feels that such observations imply that the mind must be massively modular. But matching heuristics to environment *structure* does not mean that every new *environment* or problem demands a new heuristic—the simplicity of these mechanisms implies that they can often be used in multiple, similarly structured domains with just a change in the information they employ (Czerlinski, Gigerenzer, & Goldstein, 1999). As a consequence, while we have explored several heuristics found in the mind's adaptive toolbox, this does not wed ecological rationality to a view of massive modularity any more than the investigation of multiple heuristics and biases did for that research programme (and in fact, the emphasis on simplicity and robustness in the ecological rationality perspective more strongly limits the expected number of heuristics in the adaptive toolbox).

Neither relying on the environment nor being judged in terms of the environment mean that ecologically rational heuristics *always* yield good decisions or correct or accurate choices, as Over (2000) seems to expect. Simple heuristics are, by their very nature and by the nature of the uncertain world, often wrong, and often in systematic ways, inevitably producing errors just as do standard normative procedures. The question is whether or not heuristics are *more* often wrong than other decision mechanisms that require more information and more computation. Our research has shown that this is often not the case—simple heuristics can match the performance of more complex mechanisms in many situations. The errors they *do* make can, as the heuristics-and-biases programme showed, be very useful in elucidating just which heuristics are being used by people and when (Todd, et al., 2000).

Ecological rationality and evolutionary psychology

The perspective of ecological rationality shares many similarities with evolutionary psychology, but the two are not synonymous. Evolutionary psychology is grounded in ecological rationality (rather than the other way around, as Over, 2000, suggests): It assumes that our minds were designed by natural selection to solve practical problems in an efficient and effective manner. However, evolutionary psychology focuses specifically on ancestral environments and practical problems with fitness consequences, while ecological rationality additionally encompasses decision making in present environments without privileging problems with fitness consequences. This gives ecological rationality a broader purview. For instance, we have explored how the Take The Best heuristic mentioned earlier can be ecologically rational in a modern-day decision task with little importance for the fitness of most individuals: the *German cities problem*, in which people must decide which of two German cities is larger. Beyond such modem applications, we believe that Take The Best and other simple heuristics we have explored are also good candidates for evolved mental mechanisms.

But it is important to be clear about the level of description at which we can make evolutionary claims about these heuristics. For instance, Take The Best is not a *German cities heuristic* nor a *German cities adaptation*. It is a decision heuristic suitable for making a choice among two or more options in a variety of content domains sharing the same environmental structure, namely a noncompensatory distribution of cue validities (Martignon & Hoffrage, 1999). Clearly, German cities were not a part of the ancestral environment, but the noncompensatory structure of information that Take The Best exploits most likely *has* been an enduring feature of our world. Hence, our evolutionary arguments in this case centre on Take The Best's adaptive fit to specific environmental structures of information and not to German cities, *per se*. Thus, the concrete content that a heuristic takes as its input and the abstract computational problem that it solves—including the information structure that it relies on—should not be conflated. A heuristic that evolved because of its ability to solve a set of problems characterised by a common (information) structure can embody a mechanism that is independent of, but also nonetheless applicable to, domain-specific information in some particular situation.

ENVIRONMENTAL COMPLEXITY VERSUS MECHANISM COMPLEXITY

The error in questioning whether German city comparison is an adaptive problem lies in drawing too strong a link between the concrete content of a task and the mechanism that provides its solution. A related error is to assume a linear relationship between (perceived) problem complexity and the complexity of the mechanism that solves the problem: One cannot assume that the mechanism solving the problem is as complex as the concrete form of the problem. Some organisms (e.g., sedentary and nonsocial species) inhabit simple environments where there are few choices to be made and few cues to consider in making them, while other organisms (e.g., migratory and social species) experience more complex environments rich with decision points and information. Clearly the former can get by with simple heuristics, and we would only expect more complex forms of cognition to evolve in the latter set of species. But it does not then follow that environmental complexity *requires* mental complexity, as Over (2000) implies. The question addressed in *Simple Heuristics* is not whether there are simple *problems*, but whether, given even a complex problem, there might be simple *solutions*. The point of exploring ecological rationality is to see to what extent simple mechanisms can be useful and adaptive even in complex environments.

In fact, even if complex mechanisms are greatly superior in some context, the opportunism of the evolutionary process may still serve to give preference to simpler solutions: Once evolution has discovered simple mechanisms for simple problems, it may often co-opt them for more complex problems if they still yield reasonable behaviour (a similar argument applies to heuristics that people learn). Furthermore, even when complex mechanisms underlie some parts of an organism's information processing (e.g., visual processing involved in face recognition), the psychological decision mechanisms built atop those processes may best be simply described (e.g., with the recognition heuristic, which makes choices based on assessed recognition of alternatives; see Todd, 1999, for a discussion of simplicity built on complexity and the relationship to Morgan's Canon). Thus, complexity of environment does not imply complexity of decision mechanism (nor does complex *behaviour* imply complex mechanisms—see Godfrey-Smith, 1996, for an overview of the philosophical debate surrounding this issue).

Another way of thinking about environmental complexity (besides many decisions and much information) is in terms of the heterogeneity of the environment: the more different types of structures encountered and decisions called for in an environment, the more complex it can appear. Here ecological rationality would propose that there may be different simple heuristics for the different environmental challenges (i.e., some degree of domain-specificity, provided that the environments are different enough, pulling against the trends towards simplicity and robustness). The question then arises, how is the appropriate heuristic to be selected and invoked in different circumstances? Some have proposed that this opportunity for higher-level choice will require complex cognition of the sort that simple heuristics cannot achieve (e.g., Cooper, 2000; Feeney, 2000). However, here too there is no logical necessity to let complex mechanisms in the back door—heuristic selection could be accomplished through simple "zero intelligence" reinforcement-learning techniques (Erev & Roth, 2000) as well as through higher-level selection heuristics themselves (Morton, 2000).

ACCURATE VERSUS USEFUL REPRESENTATIONS

Our minds can usefully exploit the structure of information in the environment, even if that information does not lead us to build completely accurate representations of the environment. Consider the case of probabilistic reasoning, which Over (2000) concentrates on, but which our group has written about extensively as a domain for ecological rationality only outside *Simple Heuristics*. Gigerenzer and Hoffrage (1995) proposed that people make probability judgements on the basis of concrete, naturally sampled frequencies—the form in which such information is naturally experienced—instead of a theoretical abstraction, single-event probabilities. Over (2000) warns that without forming higher-level hypotheses about causation or independence, we would be stuck with what can be misleading information from natural sampling. But we must first ask whether this theoretical knowledge is really required to solve the practical task at hand. For Bayesian inferences with one cue (Gigerenzer &

Hoffrage, 1995) or more (Krauss, Martignon, & Hoffrage, 1999), natural frequency format information does not include exact assessments of causality or independence, and yet can be much less misleading than probabilities (e.g., as in Eddy, 1982, where doctors misjudge the probability of having breast cancer after a positive mammography as 80% instead of the correct 8%). In other tasks, such as judging whether a given cue should be used in addition to another one, determining independence can be important. Thus the usefulness of higher-level reasoning mechanisms (e.g., about independence or causality) must not be assumed, and instead the ability to solve a decision task with simple mechanisms matched to the available information structure should first be assessed.[3]

Over (2000, p. 190) also warns against the simplistic use of naturally sampled frequencies in the following example, "We do not necessarily conclude from 4 heads out of 5 spins of a coin that the probability of heads is 4 out of 5 … our observation of the symmetry of the coin, and our model of the causal process of the spinning, may give us the hypothesis that the coin is fair. This may lead us to override any inclination to think that the probability of heads is 4 out of 5." While true, this same example also illustrates the power of natural frequencies compared with probabilities. Jacob Bernoulli once remarked that the law of large numbers is a rule that "even the stupidest man knows by some instinct of nature *per se* and by no previous instruction" (quoted in Gigerenzer et al., 1989, p. 29). Given the intuition that sample size is small, we would be careful in giving a probability of 4 heads out of 5. However, if we were to sample 400 heads out of 500 spins, we would certainly be more confident that the probability of heads is indeed 4 out of 5. The advantage of frequency formats for preserving information about sample size (Brase, Cosmides, & Tooby, 1998) disappears with percentages, which abstract away this useful information to leave a probability of 80% for both samples.

MAKING CONNECTIONS

Ecological rationality is a distinct perspective connecting to, and expanding on, several other schools of thought. It is based on the ideas of decision making within realistic bounds of limited time and information as captured by bounded rationality, extending these bounds beyond the mind and into the world. It

[3]Over's concern may be that higher-order reasoning is a more important aspect of decision making than those aspects we have so far dealt with (see, e.g., Sternberg, 2000; Todd, et al., 2000). We think that very many of the decisions that people (and other animals) commonly make are guided by simple fast and frugal heuristics, rather than by more complex causal, or Bayesian, or logical reasoning, and the vast literature in probabilistic reasoning also supports this belief. But it would be very useful to examine the differing beliefs about what are the most common and what are the most important decisions that people face during their lives in an empirical manner, for instance through content analyses of daily decision making. Furthermore, we would like to see more research done on the possibility of simple heuristics providing useful shortcuts for more complex forms of inference such as causal reasoning.

encompasses evolutionary psychology's pragmatic view of cognition, exploring domain-specific mental mechanisms that are well adapted to their particular environment, while also considering adaptation via learning to modern environments. It reflects Gibson's ecological psychology in expecting the solutions to many problems to already lie in the world, but matches this with intelligent heuristics in the mind that can best capitalise on the world's resources. And it shares common foundations with the heuristics-and-biases research programme in looking for the simple shortcuts that people actually use to make decisions, but differs by building precise models of environment-exploiting heuristics that work well according to ecological, and not just logical, norms. We hope that the perspective of ecological rationality will allow further connections to be drawn between these and related fields and provide a unified understanding of decision making in its environmental context.

Manuscript received 10 August 2000

REFERENCES

Brase, G.L., Cosmides, L., & Tooby J. (1998). Individuation, counting and statistical inference: The role of frequency and whole object representation in judgment under uncertainty. *Journal of Experimental Psychology: General*, *127*, 3–21.

Cooper, R. (2000). Simple heuristics could make us smart, but which heuristics do we apply when? *Behavioral and Brain Sciences*, *23*(5), 746.

Czerlinski, J., Gigerenzer, G., & Goldstein, D.G. (1999). How good are simple heuristics? In G. Gigerenzer, P.M. Todd, & the ABC Research Group, *Simple heuristics that make us smart* (pp. 97–118). New York: Oxford University Press.

Eddy, D.M. (1982). Probabilistic reasoning in clinical medicine: Problems and opportunities. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty . Heuristics and biases* (pp. 249–267). Cambridge, UK: Cambridge University Press.

Erev, I., & Roth, A.E. (2000). Learning, reciprocation, and the value of bounded rationality. In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality: The adaptive toolbox* (Dahlem Workshop Report). Cambridge, MA: MIT Press.

Feeney, A. (2000). Simple heuristics: From one infinite regress to another? *Behavioral and Brain Sciences, 23*(5), 749–750.

Fodor, J. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.

Gigerenzer, G., & Hoffrage, U. (1995). How to improve Bayesian reasoning without instruction: Frequency formats. *Psychological Review, 102*, 684–704.

Gigerenzer, G., Swijtink, Z., Porter, T., Daston, L., Beatty, J., & Krüger, L. (1989). *The empire of chance: How probability changed science and everyday life*. Cambridge, UK: Cambridge University Press.

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

Godfrey-Smith, P. (1996). Precis of *Complexity and the function of mind in nature. Adaptive Behavior*, *4*, 453–465.

Jackendoff, R. (1992). *The languages of the mind*. Cambridge, MA: MIT Press.

Jonsen, A., & Toulmin, S. (1988). *The abuse of casuistry: A history of moral reasoning*. Berkeley, CA: University of California Press.

Karmiloff-Smith, A. (1992). *Beyond modularity: A developmental perspective on cognitive science*. Cambridge, MA: MIT Press.

Krauss, S., Martignon, L., & Hoffrage, U. (1999). Simplifying Bayesian inference: The general case. In L. Magnani, N. Nersessian, & P. Thagard (Eds.), *Model-based reasoning in scientific discovery* (pp. 165–179). New York: Plenum Press.

Martignon, L., & Hoffrage, U. (1999). Why does one-reason decision making work? A case study in ecological rationality. In G. Gigerenzer, P.M. Todd, & the ABC Research Group, *Simple heuristics that make us smart* (pp. 119–140). New York: Oxford University Press.

Mithen, S. (1996). *The prehistory of the mind*. New York: Thames & Hudson.

Morton, A. (2000). Heuristics all the way up? *Behavioural and Brain Sciences*, *23*(5), 758–759.

Mysterud, I. (2000). Take the best (review of *Simple heuristics that make us smart*). *Trends in Ecology and Evolution*, *15*, 343–344.

Over, D.E. (2000). Ecological rationality and its heuristics. *Thinking and Reasoning*, *6*(2), 182–192.

Sperber, D. (1994). The modularity of thought and the epidemiology of representations. In L. Hirschfeld & S. Gelman (Eds.), *Mapping the mind: Domain-specificity in cognition and culture* (pp. 39–67). New York: Cambridge University Press.

Sternberg, R.J. (2000). Damn it, I still don't know what to do. *Behavioural and Brain Sciences*, *22*(5), 764.

Todd, P.M. (1999). Simple inference heuristics versus complex decision machines. *Minds and Machines*, *9*(4), 461–477.

Todd, P.M. (2000). Fast and frugal heuristics for environmentally bounded minds. In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality: The Adaptive toolbox* (Dahlem Workshop Report). Cambridge, MA: MIT Press.

Todd, P.M., Gigerenzer, G., & the ABC Research Group (2000). How can we open up the adaptive toolbox? (Reply to 30 commentaries on *Simple heuristics that make us smart*). *Behavioral and Brain Sciences*, *23*(5), 767–780.

Tooby, J., & Cosmides, L. (1992). The psychological foundations of culture. In J. Barkow, L. Cosmides, & J. Tooby (Eds.), *The adapted mind: Evolutionary psychology and the generation of culture* (pp. 19–136). New York: Oxford University Press.