BEHAVIORAL AND BRAIN SCIENCES (2000) **23,** 727–780
Printed in the United States of America

Précis of *Simple heuristics that make us smart*

Peter M. Todd
Center for Adaptive Behavior and Cognition, Max Planck Institute for Human Development, 14195 Berlin, Germany
ptodd@mpib-berlin.mpg.de
www.mpib-berlin.mpg.de/abc

Gerd Gigerenzer
Center for Adaptive Behavior and Cognition, Max Planck Institute for Human Development, 14195 Berlin, Germany
gigerenzer@mpib-berlin.mpg.de
www.mpib-berlin.mpg.de/abc

Abstract: How can anyone be rational in a world where knowledge is limited, time is pressing, and deep thought is often an unattainable luxury? Traditional models of unbounded rationality and optimization in cognitive science, economics, and animal behavior have tended to view decision-makers as possessing supernatural powers of reason, limitless knowledge, and endless time. But understanding decisions in the real world requires a more psychologically plausible notion of *bounded* rationality. In *Simple heuristics that make us smart* (Gigerenzer et al. 1999), we explore *fast and frugal heuristics* – simple rules in the mind's adaptive toolbox for making decisions with realistic mental resources. These heuristics can enable both living organisms and artificial systems to make smart choices quickly and with a minimum of information by exploiting the way that information is structured in particular environments. In this précis, we show how simple building blocks that control information search, stop search, and make decisions can be put together to form classes of heuristics, including: ignorance-based and one-reason decision making for choice, elimination models for categorization, and satisficing heuristics for sequential search. These simple heuristics perform comparably to more complex algorithms, particularly when generalizing to new data – that is, simplicity leads to robustness. We present evidence regarding when people use simple heuristics and describe the challenges to be addressed by this research program.

Keywords: adaptive toolbox; bounded rationality; decision making; elimination models; environment structure; heuristics; ignorance-based reasoning; limited information search; robustness; satisficing, simplicity

1. Introduction

A man is rushed to a hospital in the throes of a heart attack. The doctor needs to decide whether the victim should be treated as a low risk or a high risk patient. He is at high risk if his life is truly threatened, and should receive the most expensive and detailed care. Although this decision can save or cost a life, the doctor must decide using only the available cues, each of which is, at best, merely an uncertain predictor of the patient's risk level. Common sense dictates that the best way to make the decision is to look at the results of each of the many measurements that are taken when a heart attack patient is admitted, rank them according to their importance, and combine them somehow into a final conclusion, preferably using some fancy statistical software package.

Consider in contrast the simple decision tree in Figure 1, which was designed by Breiman et al. (1993) to classify heart attack patients according to risk using only a maximum of three variables. If a patient has had a systolic blood pressure of less than 91, he is immediately classified as high risk – no further information is needed. If not, then the decision is left to the second cue, age. If the patient is under 62.5 years old, he is classified as low risk; if he is older, then one more cue (sinus tachycardia) is needed to classify him as high or low risk. Thus, the tree requires the doctor to answer a maximum of three yes/no questions to reach a decision rather than to measure and consider all of the several usual predictors, letting her proceed to life-saving treatment all the sooner.

This decision strategy is simple in several respects. First, it ignores the great majority of possible measured predic-

PETER M. TODD is a senior research scientist at the Max Planck Institute for Human Development in Berlin, Germany, and co-founder of the Center for Adaptive Behavior and Cognition (ABC). He has published numerous papers and books on modeling behavior, music, and evolution, and is associate editor of the journals *Adaptive Behavior* and *Animal Cognition.*

GERD GIGERENZER is Director of the Center for Adaptive Behavior and Cognition (ABC) at the Max Planck Institute for Human Development in Berlin, Germany, and a former Professor of Psychology at the University of Chicago and other institutions. He has won numerous prizes, including the AAAS Prize for Behavioral Science Research in 1991; his latest book is *Adaptive Thinking: Rationality in the Real World* (Oxford University Press, forthcoming).

© 2000 Cambridge University Press 0140-525X/00 $12.50

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

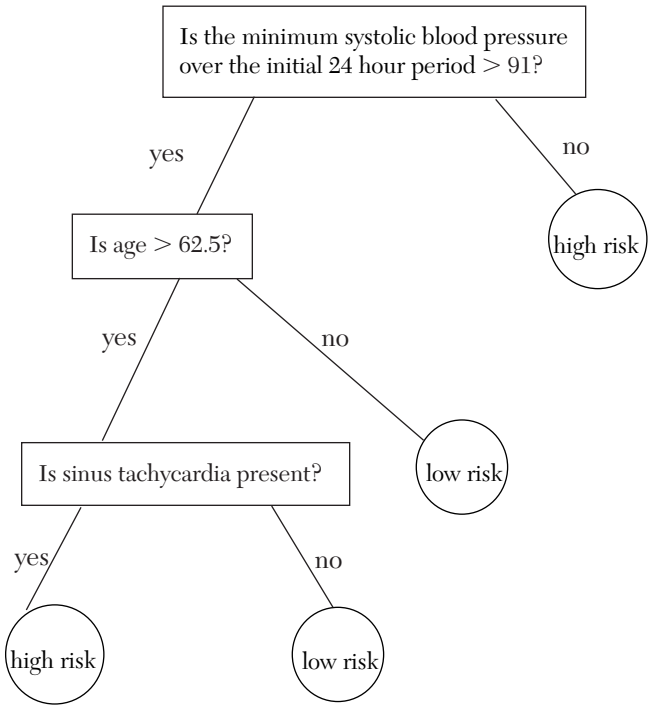


Figure 1. A simple decision tree for classifying incoming heart attack patients into high risk and low risk patients (adapted from Breiman et al. 1993).

tors. Second, it ignores quantitative information by using only yes/no answers to the three questions. For instance, it does not care how much older or younger the patient is than the 62.5 year cut-off. Third, the strategy is a step-by-step process; it may end after the first question and does not combine (e.g., weight and add) the values on the three predictors. Asking at most three yes/no questions is a fast and frugal strategy for making a decision. It is fast because it does not involve much computation, and it is frugal because it only searches for some of the available information. Its simplicity raises the suspicion that it might be highly inaccurate, compared to standard statistical classification methods that process and combine all available predictors. Yet it is actually more accurate in classifying heart attack patients according to risk status than are some rather complex statistical classification methods (Breiman et al. 1993). The more general form of this counterintuitive finding – that fast and frugal decision making can be as accurate as strategies that use all available information and expensive computation – forms one of the bases of our research program.

Our book, *Simple heuristics that make us smart* (hereafter *Simple heuristics*), is about fast and frugal heuristics for making decisions – how they work, and when and why they succeed or fail. These heuristics can be seen as models of the behavior of both living organisms and artificial systems. From a descriptive standpoint, they are intended to capture how real minds make decisions under constraints of limited time and knowledge. From an engineering standpoint, these heuristics suggest ways to build artificially intelligent systems – artificial decision-makers that are not paralyzed by the need for vast amounts of knowledge or for extensive computational power. These two applications of fast and frugal heuristics do not exclude one another – indeed, the decision tree in Figure 1 could be used to describe the behavior of an unaided human mind or could be built into an emergency-room machine. (Note that while decision trees are generally easy to use, their construction in the first place can be computationally expensive. The simple heuristics presented in the book can also avoid this costly construction phase.)

In this précis we describe the framework of our exploration of fast and frugal heuristics and summarize some of the results that have been obtained so far by the ABC Research Group. We begin by placing the study of simple heuristics within the context of bounded rationality, distinct from traditional views of unbounded rationality or optimization under constraints. We then describe the building blocks that go together to make up simple heuristics, and in section 4 we show how they can be combined into a variety of decision mechanisms for choice, categorization, estimation, and other tasks. Next we introduce the concept of ecological rationality, and explain how fast and frugal heuristics can achieve reasonable performance by fitting particular information structures in the environment and being robust to environmental change. In section 6 we cover the ways that the performance of these heuristics can be measured, and some of the evidence to date that people use such simple reasoning in particular decision tasks. We next relate our research to other recent notions of heuristics in section 7, and describe in section 8 the metaphor of the adaptive toolbox which organizes the mind's collection of simple heuristics. We conclude with a set of questions remaining to be explored, and a summary of the view of bounded rationality presented in the book.

2. Visions of rationality: From demons to bounded rationality

Humans and animals make inferences about their world with limited time, knowledge, and computational power. In contrast, many models of rational inference view the mind as if it were a supernatural being possessing demonic powers of reason, boundless knowledge, and all of eternity with which to make decisions. Such visions of rationality often conflict with reality. But we can use them as points of comparison to help clarify our own vision of *ecological rationality* – adaptive behavior resulting from the fit between the mind's mechanisms and the structure of the environment in which it operates.

We start by considering two conceptual revolutions. The first is the demise of the dream of certainty and the rise of a calculus of uncertainty – probability theory – during what is known as the *probabilistic revolution* (Gigerenzer et al. 1989; Krüger et al. 1987). The probabilistic revolution has shaped our picture of the mind in fields ranging from cognitive science to economics to animal behavior. Mental functions are assumed to be computations performed on probabilities and utilities (Gigerenzer & Murray 1987). In this view, the laws of probability describe or prescribe sound reasoning, judgment, and decision making. Probabilistic conceptions of the mind have led to many elegant theories, but also to thorny problems. The moment one moves beyond simple constrained settings such as the ones that psychologists and computer scientists typically study to real-world situations that people actually live through, the time, knowledge, and computation that probabilistic models demand grow unfeasibly large.

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

In this book, we push for a second revolution, one which provides a different vision of how minds deal with the uncertain world. We propose replacing the image of an omniscient mind computing intricate probabilities and utilities with that of a bounded mind reaching into an adaptive toolbox filled with fast and frugal heuristics. Our premise is that much of human reasoning and decision making can be modeled by such heuristics making inferences with limited time and knowledge. These heuristics do not involve much computation, and do not compute quantitative probabilities and utilities. They are models of bounded rationality. This world view embraces the earlier probabilistic revolution's emphasis on uncertainty without sharing its focus on probability theory, either as a description or as an attainable norm of human behavior. But this second revolution is only just beginning – four major visions of rationality still continue to struggle with each other today, as shown in Figure 2.

Rationality comes in many forms. The first split in Figure 2 separates models that assume the human mind has essentially unlimited demonic or supernatural reasoning power from those that assume we operate with only bounded rationality. There are two species of demons: those that exhibit *unbounded rationality,* and those that *optimize under constraints.* There are also two main forms of bounded rationality: *satisfying heuristics* for searching through a sequence of available alternatives, and *fast and frugal heuristics* that use little information and computation to make a variety of kinds of decisions. We now explore each vision of rationality in turn.

2.1. Unbounded rationality

In 1814, the astronomer-philosopher Pierre Simon Laplace contemplated the ultimate genius, an omniscient superintelligence he characterized as follows:

> Given . . . an intelligence which could comprehend all the forces of which nature is animated and the respective situation of the beings who compose it – an intelligence sufficiently vast to submit these data to analysis . . . nothing would be uncertain and the future, the past, would be present to its eyes. (Laplace 1814/1951, p. 1325)

Earlier, John Locke (1690/1959) had contrasted the omniscient God with us humble humans living in the "twilight of probability"; Laplace secularized this opposition with his fictitious superintelligence. From the perspective of God and Laplace's superintelligence alike, Nature is deterministic and certain; but for humans, Nature is fickle and uncertain. Mortals cannot precisely know the world, but must rely on uncertain inferences, on bets rather than on demonstrative proof. Although omniscience and certainty are not attainable for any real system, the spirit of Laplace's superintelligence has survived nevertheless in the vision of unbounded rationality exemplified in various modern-day incarnations built around probability theory, such as the maximization of expected utility and Bayesian models.

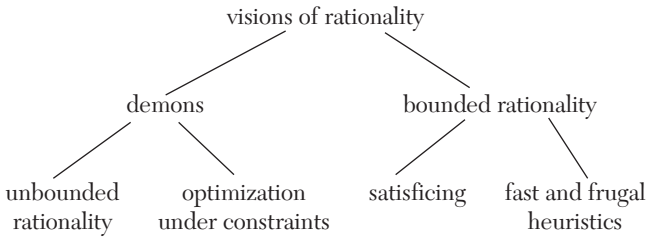


Figure 2. Visions of rationality.

Proponents of this vision paint humans in a divine light. God and Laplace's superintelligence do not worry about limited time, knowledge, or computational capacities. The fictional, unboundedly rational human mind does not either – its only challenge is the lack of heavenly certainty. In Figure 2, unbounded rationality appears in the class of models labeled "demons." We use the term in its original Greek sense of a divine (rather than evil) supernatural being, as embodied in Laplace's superintelligence.

The greatest weakness of unbounded rationality is that it does not describe the way real people think. Not even philosophers, as the following story illustrates. One philosopher was struggling to decide whether to stay at Columbia University or to accept a job offer from a rival university. The other advised him: "Just maximize your expected utility – you always write about doing this." Exasperated, the first philosopher responded: "Come on, this is serious."

Because of its unnaturalness, unbounded rationality has come under attack in the second half of the twentieth century. But when one (unboundedly rational) head has been chopped off, another very similar one has usually sprouted again in its place: its close demonic relative, optimization under constraints.

2.2. Optimization under constraints

To think is to take a risk, a step into the unknown. Our inferences, inevitably grounded in uncertainty, force us to "go beyond the information given," in Jerome Bruner's famous phrase. But the situation is usually even more challenging than this, because rarely is information given. Instead we must *search* for information – cues to classify heart attack patients as high risk, reasons to marry, indicators of stock market fluctuation, and so on. Information search is usually thought of as being internal, performed on the contents of one's memory, and hence often in parallel via our biological neural networks. But it is important to recognize that much of information search is external and sequential (and thus more time consuming), looking through the knowledge embodied in the surrounding environment. This external search includes seeking information in the socially distributed memory spanning friends and experts and in human artifacts such as libraries and the internet.

The key difference between unbounded rationality and the three other visions in Figure 2 is that the latter all involve *limited information search,* whereas models of unbounded rationality assume that search can go on indefinitely. In reasonable models, search must be limited because real decision makers have only a finite amount of time, knowledge, attention, or money to spend on a particular decision. Limited search requires a way to decide when to stop looking for information, that is, a *stopping rule.* The models in the class we call "optimization under constraints" assume that the stopping rule optimizes search with respect to the time, computation, money, and other resources being spent. More specifically, this vision of rationality holds that the mind should calculate the benefits and costs of searching for each further piece of information and stop search as soon as the costs outweigh the benefits (e.g., Anderson & Milson 1989; Sargent 1993; Stigler 1961). The rule "stop search when costs outweigh benefits" sounds

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

plausible at first glance. But a closer look reveals that optimization under constraints can require even more knowledge and computation than unbounded rationality (Vriend 1996; Winter 1975).

The motivation for replacing unbounded rationality with optimization under constraints was originally to build empirically more realistic models that respect the limitations of human minds. The paradoxical approach is to model "limited" search by assuming that the mind has essentially unlimited time and knowledge with which to evaluate the costs and benefits of further information search. The dream of optimization, threatened in its instantiation in unbounded rationality, is thus salvaged by being incorporated into an apparent competitor – constrained optimization invites unbounded rationality to sneak in through the back door.

Of course, few would argue that real humans have the time and knowledge necessary to perform the massive computations required for constrained optimization. Instead, this vision of rationality is usually presented as a lofty ideal that human reasoning *should* aspire to. But such aspirations make real human reasoning look flawed and irrational in comparison. In our view, it is these aspirations that are flawed – we will argue that reasoning can be powerful and accurate without requiring unlimited time and knowledge.

What certain forms of optimization under constraints can offer – in contrast to unbounded rationality – is an analysis of the structure of environments. For instance, in Anderson's rational analysis framework (Anderson 1990; Oaksford & Chater 1994) constraints from the environment, rather than on the decision maker, are used to modify one's understanding of what is optimal behavior in a particular context. Such an analysis does not directly address the question of what mental mechanisms could possibly yield behavior approaching the optimal norm, but at least it allows us to create a more realistic standard for assessing proposed mechanisms.

Instead of these demonic visions of reason, we turn to the idea of bounded rationality. But many, if not most, researchers in cognitive science, economics, and animal behavior interpret the term "bounded rationality" as synonymous with optimization under constraints, a (mis)use we strongly reject. This interpretation may be responsible for the frequent dismissal of bounded rationality in favor of good old-fashioned demonic visions. The economist Thomas Sargent (1993), for instance, in interpreting bounded rationality as constrained optimization, argues that when one models people as "bounded" in their rationality, one's models use a greater number of parameters and become more demanding mathematically. He believes that the reason why researchers (particularly economists) stick with models incorporating unbounded rationality is that their desire for models with fewer parameters is not met by the bounded approach: "a reduction is not what bounded rationality promises" (p. 4). But this is a misleading interpretation of bounded rationality – rationality need not be optimization, and bounds need not be constraints.

2.3. Bounded rationality: Satisficing

The "father" of bounded rationality, Herbert Simon, has vehemently rejected its reduction to optimization under constraints: "bounded rationality is not the study of optimization in relation to task environments" (Simon 1991). Instead, Simon's vision of bounded rationality has two interlocking components: the limitations of the human mind, and the structure of the environments in which the mind operates. The first component of his vision means that models of human judgment and decision making should be built on what we actually know about the mind's capacities rather than on fictitious competencies. In many real-world situations, optimal strategies are unknown or unknowable (Simon 1987). Even in a game such as chess, where an optimal (best) move does in fact exist at every point, no strategy can calculate that move in a reasonable amount of time (either by human minds or computers), despite the well-defined nature of the possibilities to be searched. In less well-defined natural situations, our hope of identifying a useable optimal strategy is even further diminished. Because of the mind's limitations, 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. We explore these classes of methods at length in our book.

The second component of Simon's view of bounded rationality, environmental structure, is of crucial importance because it can explain when and why simple heuristics perform well: if the structure of the heuristic is adapted to that of the environment. Simon's (1956a) classic example of this component concerns imaginary organisms foraging according to simple heuristics whose behavior can only be understood by looking at the structure of the information in the environment. Simon was not the only one to make this important point; it was made both before his work (e.g., Brunswik 1943) and at various times since (e.g., Anderson 1990; Shepard 1990), including more extreme emphasis on studying the environment rather than the mechanisms of the mind (e.g., Gibson 1979). But in general the second part of Simon's (1956a) paper title, "Rational choice and the structure of environments," has been neglected in mainstream cognitive sciences (even by Simon himself – see Simon 1987).

We use the term *ecological rationality* to bring environmental structure back into bounded rationality. A heuristic is ecologically rational to the degree that it is adapted to the structure of an environment (see below). Thus, simple heuristics and environmental structure can both work hand in hand to provide a realistic alternative to the ideal of optimization, whether unbounded or constrained.

One form of bounded rationality is Simon's concept of *satisficing* – a method for making a choice from a set of alternatives encountered sequentially when one does not know much about the possibilities in advance. In such situations, there may be no optimal method for stopping searching for further alternatives – for instance, there would be no optimal way of deciding when to stop looking for prospective marriage partners and settle down with a particular one (see Ch. 13 for more on satisficing in mate search). Satisficing takes the shortcut of setting an aspiration level and ending the search for alternatives as soon as one is found that exceeds the aspiration level (Simon 1956b; 1990), for instance leading an individual with Jack-Sprat-like preferences to marry the first potential mate encountered who is over a desired width.

2.4. Bounded rationality: Fast and frugal heuristics

Satisficing is a way of making a decision about a set of alternatives that respects the limitations of human time and

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

knowledge: it does not require finding out or guessing about all the options and consequences the future may hold, as constrained optimization does. However, some forms of satisficing can still require a large amount of deliberation on the part of the decision maker, for instance to set an appropriate aspiration level in the first place, or to calculate how a current option compares to the aspiration level (Simon 1956b). Rather than let overzealous mental computation slip back into our picture of human rationality, we narrow our focus still more to concentrate on fast and frugal heuristics for decision making.

Fast and frugal heuristics employ a minimum of time, knowledge, and computation to make adaptive choices in real environments. They can be used to solve problems of sequential search through objects or options, as in satisficing. They can also be used to make choices between simultaneously available objects, where the search for information (in the form of cues, features, consequences, etc.) about the possible options must be limited, rather than the search for the options themselves. Fast and frugal heuristics limit their search of objects or information using easily computable stopping rules, and they make their choices with easily computable decision rules. We thus see satisficing and fast and frugal heuristics as two overlapping but different categories of bounded rationality: there are some forms of satisficing that are fast and frugal, and others that are computationally unreasonable; and there are some fast and frugal heuristics that make satisficing sequential option decisions, and some that make simultaneous option choices (see sect. 4). We consider fast and frugal heuristics to represent bounded rationality in its purest form.

A prime example of the classes of fast and frugal heuristics that we explore in our book is *one-reason decision making,* in which only a single piece of information is used to make a choice (we describe particular instances of this class in more detail below). There is a sound rationale for basing a decision on only one reason rather than on a combination of several: Combining information from different cues requires converting them into a common currency, a conversion that may be expensive if not actually impossible. Standard models of optimization, whether constrained or unbounded, assume that there is a common currency for all beliefs and desires, namely, quantitative probabilities and utilities. Although this is a mathematically convenient assumption, the way humans look at the world does not always conform to it. Some things do not have a price tag, and cannot be reduced to and exchanged for any common currency (Elster 1979). Love, true friendship, military honors, and Ph.D.s, for example, are supposed to be priceless, and therefore incommensurable with items for sale in a shopping mall. When reasons cannot be converted to a single currency, the mind may do best by employing a fast and frugal strategy that bases its decision on just one good reason. As we demonstrate (in Chs. 4–6), however, incommensurability is not the only reason for one-reason decision making.

Before we take a closer look at fast and frugal heuristics, let us sum up our discussion so far. Bounded rationality has become a fashionable term in many quarters, and a plethora of proposed examples have been thrown together under this term, including optimization under constraints. Figure 2 helps to make clear the distinctions between bounded rationality and the demonic visions of rationality. Unbounded rationality is not concerned with the costs of search, while bounded rationality explicitly limits search through stopping rules. Optimization under constraints also limits search, but does so by computing the optimal stopping point, that is, when the costs of further search exceed the benefits. In contrast, bounded rationality "bets" on the effectiveness of simple ways of guiding and stopping information search (described in the next section) that do not attempt to optimize. Finally, the purest form of bounded rationality is to be found in fast and frugal heuristics, which employ limited search through objects (in satisficing) or cues and exploit environmental structure to yield adaptive decisions.

3. The ABCs of fast and frugal heuristics

In *Simple heuristics* we explore the view that people operate with bounded rationality to make the majority of their inferences and decisions – a framework that is also useful for studying other animals and for developing decision-making heuristics for artificial agents. This exploration of boundedly rational heuristics involves (1) designing computational models of candidate simple heuristics, (2) analyzing the environmental structures in which they perform well, (3) testing their performance in real-world environments, and (4) determining whether and when people (and other animals) really use these heuristics. The results of the investigatory stages (2), (3), and (4) can be used to revise the next round of theorizing in stage (1). The different stages of this research program rest on multiple methods, including theoretical modeling of heuristics, computer simulation of their performance, mathematical analysis of the fit between heuristics and specific environments, and laboratory experimentation. Across the next four sections we consider each of these stages in turn.

A computational model of a heuristic specifies the precise steps of information gathering and processing that are involved in generating a decision, such that the heuristic can be instantiated as a computer program. For a fast and frugal heuristic, this means the computational model must include principles for guiding search for alternatives or information (or both), stopping that search, and making a decision, as we now describe.

3.1. Heuristic principles for guiding search

Decisions must be made between alternatives, and based on information about those alternatives. In different situations, those alternatives and pieces of information may need to be found through active search. The heuristic principles for guiding search, whether across alternatives or information, are what give search its direction (if it has one). For instance, cues can be searched for in a random manner, or in order of some pre-computed criterion related to their usefulness (Ch. 6), or based on a recollection about which cues worked previously when making the same decision (Ch. 4). Search for alternatives can similarly be random or ordered.

Fast and frugal search-guiding principles do not use extensive computations or knowledge to determine where to look next. But such simplicity need not lead to a disadvantage in decision accuracy, because simple search strategies can help heuristics to be more robust than those that attempt to optimize their information search. For instance, the choice heuristics we focus on (Ch. 4) use cue orderings that are easy to compute, ignoring dependencies between cues just as people have been reported to do (e.g., Armelius & Armelius 1974). If instead the heuristics computed con-

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

ditional probabilities between cues to determine search order, or tried all of the enormous number of cue orders to find the optimal one for a given data set, they might be slightly more accurate – but only when fitting the data set they already know. When making predictions about new data, simple information search methods that ignore dependencies between cues can actually yield more accurate choices (Ch. 6).

3.2. Heuristic principles for stopping search

In our conception of bounded rationality, the temporal limitations of the human mind (or that of any realistic decision-making agent) must be respected as much as any other constraint. This implies in particular that search for alternatives or information must be terminated at some (preferably early) point. Moreover, to fit the computational capacities of the human mind, the method for determining when to stop search should not be overly complicated. For example, one simple stopping rule is to cease searching for information and make a decision as soon as the first cue or reason that favors one alternative is found (Ch. 4). This and other cue-based stopping rules do not need to compute an optimal cost-benefit trade-off as in optimization under constraints; in fact, they need not compute any utilities or costs and benefits at all. For searching through alternatives (rather than cues), simple aspiration-level stopping rules can be used, as in Simon's original satisficing notion (Simon 1956b; 1990; see also Ch. 13).

3.3. Heuristic principles for decision making

Once search has been guided to find the appropriate alternatives or information and then been stopped, a final set of heuristic principles can be called upon to make the decision or inference based on the results of the search. These principles can also be very simple and computationally bounded. For instance, a decision or inference can be based on only one cue or reason, whatever the total number of cues found during search (see Chs. 2–6). Such one-reason decision making does not need to weight or combine cues, and so no common currency between cues need be determined. Decisions can also be made through a simple elimination process, in which alternatives are thrown out by successive cues until only one final choice remains (see Chs. 10–12).

3.4. Putting heuristic building blocks together

These heuristic principles are the building blocks, or the ABCs, of fast and frugal heuristics. Given that the mind is a biological rather than a logical entity, formed through a process of successive accrual, borrowing, and refinement of components, it seems reasonable to assume that new heuristics are built from the parts of the old ones, rather than from scratch (Pinker 1998; Wimsatt 2000a). In this light, we have used two main methods to construct computational models of fast and frugal heuristics: combining building blocks and nesting existing heuristics. Heuristic principles can be combined in multiple ways, such as the several guises in which we find one-reason decision making throughout our book, though of course not arbitrarily: For instance, a fast and frugal heuristic for two-alternative choice that stops information search at the first cue on which the alternatives differ must also use a decision principle based on one-reason decision making. Additionally, entire fast and frugal heuristics can themselves be combined by nesting one inside another. As an example, the recognition heuristic (Chs. 2 and 3) works on the basis of an elementary cognitive capacity, recognition memory, but it can also serve as the first step of heuristics that draw on other capacities, such as recall memory (Chs. 4 and 5; see also sect. 8 below on combining tools in the adaptive toolbox). Recognition memory develops earlier than recall memory both ontogenetically and evolutionarily, and the nesting of heuristics can similarly be seen as analogous to the addition of a new adaptation on top of an existing one.

4. Classes of heuristics

All of the heuristics that the ABC Research Group has been exploring can be thought of as enabling a choice of one or more objects or options from a (larger) set of possibilities. How many options there are in a particular decision situation, and how many are to be chosen, will partly determine the heuristics that can be employed. The amount and kind of cues available to make this choice can further constrain the set of appropriate mental tools. Together, these features divide the heuristics we have developed into the four main classes presented in this section.

4.1. Ignorance-based decision making

The simplest kind of choice – numerically, at least – is to select one option from two possibilities, according to some criterion on which the two can be compared. Many of the heuristics described in our book fall into this category, and they can be further arranged in terms of the kinds and amount of information they use to make a choice. In the most limited case, 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* (Ch. 2): When choosing between two objects (according to some criterion), if one is recognized and the other is not, then select the former. For instance, if deciding at mealtime between Dr. Seuss's famous menu choices of green eggs and ham (using the criterion of being good to eat), this heuristic would lead one to choose the recognized ham over the unrecognized odd-colored eggs.

Following the recognition heuristic will be adaptive, yielding good choices more often than would random choice, in those decision environments in which exposure to different possibilities is positively correlated with their ranking along the decision criterion being used. To continue with our breakfast example, those things that we do not recognize in our environment are more often than not inedible, because humans have done a reasonable job of discovering and incorporating edible substances into our diet. Norway rats follow a similar rule, 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). We have used a different kind of example to amass experimental evidence that people also use the recognition heuristic: Because we hear about large cities more often than small cities, using recog-

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

nition 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). In our experiments, over 90% of the participants act in accordance with the recognition heuristic, even after they have been taught further information about the recognized cities that should lead them to stop following this decision rule. Employing the recognition heuristic can lead to the surprising *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 – knowing (i.e., recognizing) more than this will actually *decrease* performance (Ch. 2).

The recognition heuristic can be generalized to cases in which several options are to be chosen from a larger set of possibilities, for instance when several social partners are to be chosen for some collaborative activity such as resource exchange or hunting. We have investigated a modern-day equivalent of this type of choice: selecting companies for investment. When deciding which companies to invest in from among those trading in a particular stock market, the recognition heuristic would lead investors to choose just those that they have heard of before. Such a choice can be profitable assuming that more-often-recognized companies will typically have some of the better-performing stocks on the market – a testable, but not obvious, assumption.

We tested precisely this assumption, and this approach to fast and frugal investing, by asking several sets of people what companies they recognized and forming investment portfolios based on the most familiar firms (Ch. 3). In this (admittedly short-term) trial of a simple heuristic in an unforgiving and often chaotic real social environment, we found that ignorance-driven recognition alone could match and often beat the highly trained wisdom of professional stock pickers. This does not, of course, prove that people use the recognition heuristic when making such choices (though common investment advice suggests this is so) – at this point we only have evidence that using this heuristic can be a surprisingly adaptive strategy in complex environments. Experimental examination of whether or not people employ the recognition heuristic (and others) in these types of social domains remains an important upcoming challenge.

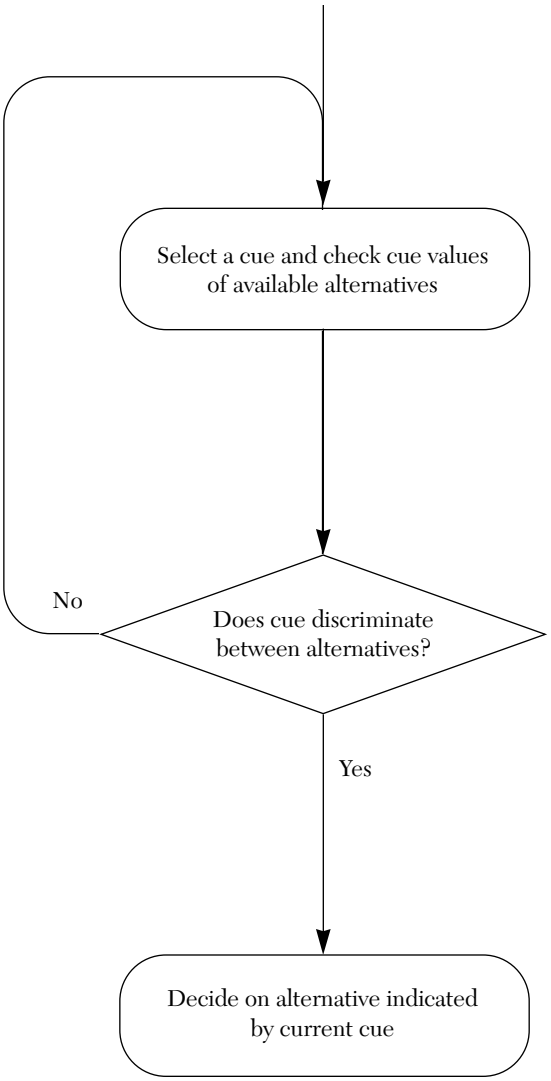


Figure 3. A flowchart of one-reason decision making. First, search for a cue and the corresponding cue values of each alternative; next, check whether the values for that cue discriminate between the alternatives; if so, then choose the indicated alternative; if not, select another cue and repeat this process. (Random choice can be used if no more cues are available.)

4.2. One-reason decision making

Returning to choices of one of two options, most of the time we have more information than just a vague memory of recognition to go on, so that other heuristics can be employed. When multiple cues are available for guiding decisions, how can a fast and frugal reasoner proceed? The most frugal approach is to use a stopping rule that terminates the search for information as soon as enough has been gathered to make a decision. In particular, as mentioned earlier, one can rely on one-reason decision making (Ch. 4): Stop looking for cues as soon as one is found that differentiates between the two options being considered. This allows the decision maker to follow a simple loop, as shown in Figure 3: (1) select a cue dimension and look for the corresponding cue values of each option; (2) compare the two options on their values for that cue dimension; (3) if they differ (e.g., if one value is larger or if there is positive information for one option but not for the other), then stop and choose the option with the cue value indicating the greater value on the choice criterion; (4) if they do not differ, then return to the beginning of this loop (step 1) to look for another cue dimension.

This little four-step loop incorporates two of the important building blocks of simple heuristics: a stopping rule (here, stopping after a single cue is found that enables a choice between the two options) and a decision rule (here, deciding on the option to which the one cue points). One other building block remains to be specified, however, before we can build a particular heuristic. We must determine just how cue dimensions are "looked for" in step 1 – that is, we must pick a specific information search rule. We have developed three fast and frugal one-reason decision heuristics that differ only in their search rule (Ch. 4; see also Gigerenzer & Goldstein 1996a). Take the Best searches for cues in the order of their validity – that is, how often the cue has indicated the correct versus incorrect options. Take the Last looks for cues in the order determined by their past success in stopping search, so that the cue that was used for the most recent previous decision (whether or not it was correct) is

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

checked first when making the next decision. Finally, the *Minimalist heuristic* selects cues in a random order.

What we found when we tested the performance of these one-reason decision-making heuristics was again surprising: Despite (or often, as we found later, because of) their simplicity and disregard for most of the available information, they still made very accurate choices. We compared these heuristics against a set of more traditional information-combining methods such as multiple regression, which weights and sums all cues in an optimal linear fashion, and a simple linear strategy (dubbed Dawes's Rule) that counts up all of the cues for and against a choice and looks at the difference. We found that the simple heuristics always came close to, and often exceeded, the proportion of correct inferences achieved by multiple regression and Dawes's Rule. This unexpected performance was found first with the data set that we have used as our simple "drosophila" example in both human and simulation experiments: choosing the larger of two German cities (Ch. 4). We then confirmed the inference accuracy of these simple heuristics in a further 19 data sets selected for their variety in both number of objects and number of cues available (Ch. 5).

The overall average performance across all 20 data sets for two simple heuristics and two traditional decision methods is shown in Table 1 (under "Fitting"). The high accuracy of Take the Best and Minimalist was achieved even though they looked through only a third of the cues on average (and decided to use only one of them), while multiple regression and Dawes's Rule used them all (see Table 1, "Frugality"). The advantages of simplicity grew in the more important test of generalization performance, where the decision mechanisms were applied to a portion of each data set that they had not seen during training. Here, Take the Best outperformed all three other algorithms by at least two percentage points (see Table 1, "Generalization"). The finding that a simple heuristic can outstrip its less frugal brethren particularly when generalizing to new decisions demonstrates the potential robustness of fast and frugal reasoning. These heuristics even performed nearly as well as much more sophisticated Bayesian methods that employ complex calculations to approach optimal behavior (Ch. 8). (These results also show the well-known "flat maximum" result that a linear model with equal-sized weights, e.g., Dawes's Rule, can predict about as well as multiple regression; see Dawes 1979.) Thus, making good decisions need not rely on the standard rational approach of collecting all available information and combining it according to the relative importance of each cue – simply betting on one good reason, even a reason selected at random, can do the trick.

Table 1. *Performance of different decision strategies across 20 data sets*

		Accuracy (% correct)	
Strategy	Frugality	Fitting	Generalization
Minimalist	2.2	69	65
Take the Best	2.4	75	71
Dawes's Rule	7.7	73	69
Multiple regression	7.7	77	68

Performance of two fast and frugal heuristics (Minimalist, Take the Best) and two linear strategies (Dawes's rule, multiple regression) across 20 data sets. The mean number of predictors available in the 20 data sets was 7.7. "Frugality" indicates the mean number of cues actually used by each strategy. "Fitting accuracy" indicates the percentage of correct answers achieved by the strategy when fitting data (test set = training set). "Generalization accuracy" indicates the percentage of correct answers achieved by the strategy when generalizing to new data (cross validation, i.e., test set ≠ training set). (Data from *Simple heuristics,* Ch. 5.)

But how? We turned to mathematical analysis (Ch. 6) to uncover the secrets of success of one-reason decision making. These simple heuristics are noncompensatory, meaning that once they have used a single cue to make a decision, no further cues in any combination can undo or compensate for that one cue's effect. When the information in the decision environment is structured in a matching noncompensatory fashion (i.e., the importance or validity of cues falls off rapidly in a particular pattern), the Take the Best heuristic can exploit that structure to make correct decisions as often as compensatory rules. Take the Best also performs comparatively well when information is scarce, that is, when there are many more objects than cues to distinguish them. Finally, the particular ordering of cues used by Take the Best, based on their ecological validity rather than other possible measures of validity, seems to give this heuristic great robustness when generalizing to new choices. We discuss this issue of exploiting environment structure further in section 5.1.

One-reason decision making may be at work in more than just consciously deliberated choices. We hypothesize that simple heuristics such as Take the Best can also play a role in memory reconstruction, updating and amending our recollection of the past in a rapid manner when further information is encountered (Ch. 9). But this adaptive updating in memory can cause as a side effect the curious phenomenon of hindsight bias – the erroneous belief that one's past judgments were closer to one's present state of knowledge than they actually were ("I knew it all along"). A memory model incorporating Take the Best can make precise predictions about when individuals will show hindsight bias, something that previous models have not allowed.

Single reasons can also suffice in situations where there are more than two options – particularly, when individual cues are fine-grained enough (or at least have enough possible values) to differentiate all the options. We have looked at the implications of this sort of single-cue decision making in the domain of parental investment (Ch. 14), specifically asking: How can a parent decide which of several offspring it should give resources to first?

Parent birds, for instance, returning to their nest with a juicy bug, typically face a number of gaping mouths that they must decide between. The parent can use the weight, hunger, age, or fixed position of each chick in the nest when picking which one to feed. As in other tasks described earlier, decision-making approaches based on traditional notions of rationality (e.g., in Gary Becker's economic analysis of the family; see Becker 1991) would dictate that the parent should assess and combine all of these cues to come up with the best choice (where "best" in this case means the choice that will lead to the greatest growth of the nestlings). But because each of these cues provides a full ordering of all the chicks (e.g., one is heaviest, one is next heaviest, and so on), only one cue is necessary for an unambiguous decision. We found that one-cue feeding rules are not only pos-

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

sible, but can also be advantageous – they perform significantly better (again in terms of total chick growth) than rules that combine all the available information in an attempt to look forward in time and predict the optimal course of action (Ch. 14). This is another way that the simplicity of fast and frugal rules can become an advantage: In situations in which repeated decisions must be made (as in feeding and raising offspring), a simple cue-based heuristic that sticks to present knowledge can outperform rules that attempt to predict an uncertain future, because it avoids the compounded noise that accumulates the further forward one strains to look.

4.3. Elimination heuristics for multiple-option choices

As the bird-feeding example just given shows, not all choices in life are presented to us as convenient pairs of options – often we must choose between several alternatives. In situations where each available cue dimension has fewer values than the number of available alternatives, one-reason decision making will usually not suffice, because a single cue will be unable to distinguish between all of the alternatives. For instance, knowing whether or not each of 15 cities has a river is not enough information to decide which city is most habitable. But this does not doom the fast and frugal reasoner to a long process of cue search and combination in these situations. Again, a simple stopping rule can work to limit information search: Only seek cues (in an order specified by the search rule) until enough is known to make a decision. But now a different type of decision rule is needed instead of relying on one reason. One way to select a single option from among multiple alternatives is to follow the simple principle of *elimination:* Successive cues are used to eliminate more and more alternatives and thereby reduce the set of remaining options, until a single option can be decided upon.

The QuickEst heuristic (Ch. 10) is designed to estimate the values of objects along some criterion while using as little information as possible. The estimates are constrained to map onto certain round numbers (for instance, when estimating city population sizes, QuickEst can return values of 100,000, 150,000, 200,000, 300,000, and other "spontaneous" numbers, following Albers 1997), so this heuristic can be seen as choosing one value from several possibilities. QuickEst is designed to work well in environments characterized by a J-distribution, where there are many more objects at one end of a criterion range than at the other. To exploit this environmental structure, QuickEst first looks at a cue that separates the most common objects from all of the others (e.g., because most small cities in Germany do not have a professional soccer team, this cue should be one of the first checked when estimating a German city's population). QuickEst then looks at the next cue that separates the remaining most common objects from the rest, and so on until an estimate can be made. To estimate the criterion value of a particular object, the heuristic looks through the cues or features in this order until it comes to the first one that the object does not possess, at which point it stops searching for any further information (e.g., if a city possesses the first several features in order but lacks an exposition site, search will stop on that cue). QuickEst then gives the "rounded" mean criterion value associated with the absence of that cue as its final estimate (e.g., the mean size of all cities without an exposition site). Thus in effect QuickEst uses features that are present to eliminate all common criterion categories, and absent features to eliminate all less common criterion categories, so that only one criterion estimate remains. No cue combination is necessary, and no adjustment from further cues is possible.

QuickEst proves to be fast and frugal, as well as accurate, in environments in which small values are frequent and large values are rare, a distribution that characterizes a variety of naturally occurring phenomena including many formed by accretionary growth. This growth pattern applies to cities (Makse et al. 1995), and indeed big cities are much less common than small ones. As a consequence, when applied to the data set of German cities, QuickEst is able to estimate rapidly the small sizes that most of them have.

We have also used the principle of elimination to build a categorization heuristic called Categorization by Elimination (Ch. 11; see also Berretty et al. 1997). In this case, the task is to choose the one category, from several possible, that a given object falls into. The simple Categorization by Elimination heuristic makes accurate category judgments by using each successive cue to whittle away the set of possible categories to which the object in question could belong, until only a single possible category remains. Its performance comes within a few percentage points of the accuracy of traditional categorization algorithms including exemplar and neural network models, and yet in our tests it uses only about a quarter of the information that these other models employ. In situations in which categorization must be performed quickly and cues take time to search for, this fast and frugal approach has clear advantages.

Such advantages are obvious in the case of trying to ascertain and categorize the intentions of other animals (including humans) we happen to encounter. If we can decide quickly and with few cues whether an approaching person or bear is interested in fighting, playing, or courting, we will have more time to prepare and react accordingly (though in the case of the bear all three intentions may be equally unappealing). Some of the most obvious cues of intention that can be assessed at a distance (as opposed to facial expression, for instance, which requires closer scrutiny) are contained in an organism's motion: Is it coming at me or heading away, slowly or quickly, directly or indirectly? We have investigated just what motion cues (including velocity, heading, and curvature of path) people can use along with the Categorization by Elimination heuristic to judge the intention of another organism in a fast and frugal manner (Ch. 12; see also Blythe et al. 1996). We determined a set of simple motion cues that can be combined (e.g., by a neural network) to indicate intention correctly in over 80% of our laboratory trials; Categorization by Elimination uses only half of these cues and still correctly predicts two-thirds of the intentions, similar to the performance of trained human observers.

4.4. Satisficing heuristics

All the heuristics that we have discussed so far for choosing one option from more than one operate with the assumption that all the possible options are presently available to the decision maker: For instance, all the possible categories of motion are known, and all the chicks are sitting patiently in the nest. But a different strategy is called for when alternatives themselves (as opposed to cue values) take time to find, appearing sequentially over an extended period or

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

spatial region. In this type of choice task, a fast and frugal reasoner need not (only) limit information search, but (also) must have a stopping rule for ending the search for alternatives themselves. One instance of this type of problem is the challenge that faces individuals searching for a mate from a stream of potential candidates met at different points in time. Here, Simon's (1956b, 1990) notion of a satisficing heuristic can be adaptive: An aspiration level is set for the selection criterion being used, and the search for alternatives is stopped as soon as the aspiration is met.

We have begun our study of satisficing heuristics for sequential search, including mate search, by simulating their performance in different mating environments (Ch. 13), focusing on simple methods for setting the aspiration level. The goal was to find satisficing heuristics that would limit both the time needed to determine a good aspiration level and the average number of potential mates that had to be considered before one was found exceeding the aspiration level. We have identified a class of simple learning heuristics that do indeed determine such adaptive aspiration levels, while still coming close to the criterion-selection performance of more optimal (and much slower) search rules. The next step, of course, is to test these theoretically plausible heuristics against data gleaned from observations of real people engaged in the mating game.

5. Why and when do simple heuristics work? The basics of ecological rationality

Traditional definitions of rationality are concerned with maintaining internal order of beliefs and inferences (sect. 6.1). But real organisms spend most of their time dealing with the external disorder of their environment, trying to make the decisions that will allow them to survive and reproduce (Tooby & Cosmides 1998). To behave adaptively in the face of environmental challenges, organisms must be able to make inferences that are fast, frugal, and accurate. These real-world requirements lead to a new conception of what proper reasoning is: *ecological rationality.* Fast and frugal heuristics that are matched to particular environmental structures allow organisms to be ecologically rational. The study of ecological rationality thus involves analyzing the structure of environments, the structure of heuristics, and the match between them, as we demonstrate throughout our book.

How is ecological rationality possible? That is, how *can* fast and frugal heuristics work as well as they do, and escape the tradeoffs between different real-world criteria including speed and accuracy? The main reason for their success is that they make a tradeoff on another dimension: that of generality versus specificity. What works to make quick and accurate inferences in one domain may well not work in another. Thus, different environments can have different specific fast and frugal heuristics that exploit their particular information structure to make adaptive decisions. But specificity can also be a danger: if a different heuristic were required for every slightly different decision-making environment, we would need an unworkable multitude of heuristics to reason with, and we would not be able to generalize to previously unencountered environments. Fast and frugal heuristics avoid this trap by their very simplicity, which allows them to be robust when confronted by environmental change and enables them to generalize well to new situations.

5.1. Exploiting environment structure

Fast and frugal heuristics can benefit from the way information is structured in environments. The QuickEst heuristic described earlier, for instance (Ch. 10), relies on the skewed distributions of many real-world variables such as city population size – an aspect of environment structure that traditional statistical estimation techniques would either ignore or even try to erase by normalizing the data. Standard statistical models, and standard theories of rationality, aim to be as general as possible, so they make as broad and as few assumptions as possible 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, Karl 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.

Because ecological rationality is a consequence of the match between heuristic and environment, we have investigated several instances where structures of environments can make heuristics ecologically rational:

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 (Ch. 6).

Scarce information. Take The Best outperforms a class of linear models on average when few cues are known relative to the number of objects (Ch. 6).

J-shaped distributions. The QuickEst heuristic estimates quantities about as accurately as more complex information-demanding strategies when the criterion to be estimated follows a J-shaped distribution, that is, one with many small values and few high values (Ch. 10).

Decreasing populations. In situations where the set of alternatives to choose from is constantly shrinking, 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 (Ch. 13).

By matching these structures 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 – that is, they can escape the curse of overfitting, which often strikes more complex, parameter-laden models, as described next. 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.

5.2. Robustness

How can simple domain-specific heuristics ever be about as accurate as complex general strategies that work with many free parameters? One answer lies in not being *too* specific. Simple heuristics are meant to apply to specific environments, but they do not contain enough detail to match any

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

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 fiddle with can be a hindrance. This failure of generalization, a phenomenon known as *overfitting* (e.g., Geman et al. 1992; Massaro 1988), stems from assuming that every detail is of utmost relevance. As we show in various chapters, models with many free parameters, from multiple linear regression to neural networks, can suffer from trying to make sense of every piece of information they encounter.

Thus, there is an important difference between the two typical applications of a strategy, *fitting* (modeling decisions for a given set of data) and *generalization* (predicting or inferring based on 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, in contrast, 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-guzzling strategy that overfits.

Robustness goes hand in hand with speed, accuracy, and especially information frugality (Table 1). 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 to some degree – for instance, April is likely to be associated with showers in northern locations year after year. In contrast, the random fluctuations of noise and the effects of smaller systematic factors may frequently change – for instance, May flowers may depend on many variable factors like temperature, rainfall, seed dispersal, and insect pests that collectively vary more from one year to the next. Because of this pattern, fast and frugal heuristics that pay attention to systematic informative cues while overlooking more variable uninformative cues can ride out a degree of environmental change without suffering much decrement in performance. Laplace's superintelligence would never overfit because it does not have to make uncertain predictions. But models of inference that try to be like a Laplacean superintelligence are doomed to overfitting, when they swallow more data than they can digest.

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 where too much information and too much information processing can hurt. Cognition is the art of focusing on the relevant and deliberately ignoring the rest. We take the same approach to *modeling* cognition.

6. How can simple heuristics be evaluated?

6.1. Performance in real-world environments

As mentioned earlier, bounded rationality is often characterized as a view that takes into account the cognitive limitations of thinking humans – an incomplete and potentially misleading characterization. If we want to understand how real human minds work, we must look not only at how our reasoning is "limited" compared to that of supernatural beings, but also at how our minds are adapted to real-world environments. This two-sided conception of bounded rationality should inform our choice of criteria with which to evaluate the performance of heuristics.

It is not enough merely to strive to compare human behavior to some optimal standard. As mentioned in section 2.3, many real-world situations do not have implementable optimizing strategies. Many other situations have too many possible optimizing strategies, because different definitions of optimality follow from different assumptions about the situation or the decision-maker's goals. Where the assumptions must be uncertain, an optimizing approach becomes uncertain as well, potentially leading to suboptimal outcomes if the wrong guesses are made. Alternatively, a nonoptimizing fast and frugal strategy can nonetheless get lucky and yield optimal outcomes. Hence, whether or not a decision strategy attempts to optimize its performance is not a sufficient evaluation criterion.

One set of criteria that is often used to assess judgments and decisions is the laws of logic and probability theory. These are often called *coherence* criteria because they are primarily concerned with the internal logical coherence of judgments rather than with how well they help us to make useful decisions in the real world. Most experimental research programs aimed at demonstrating the rationality or (usually) irrationality of humans and animals have used abstract coherence criteria. For instance, many claims that there are systematic irrational fallacies in human reasoning are based entirely on a violation of some rule or other of logic or probability (e.g., Tversky & Kahneman 1983; Wason 1983; see sect. 7).

In *Simple heuristics* we adopt a different, adaptive view of rational behavior. We do not compare human judgment with the laws of logic or probability, but rather examine how it fares in real-world environments. The function of heuristics is not to be coherent. Rather, their function is to make reasonable, adaptive inferences about the real social and physical world given limited time and knowledge. Hence, we should evaluate the performance of heuristics by criteria that reflect this function. Measures that relate decision-making strategies to the external world rather than to internal consistency – measures such as accuracy, frugality, and speed – are called *correspondence* criteria (Hammond 1996). As Egon Brunswik (1957) observed, the mind and the environment are like a husband and wife couple who must come to terms with one another by mutual adaptation. However, owing to the focus on coherence in much research on reasoning and decision making, the couple has become estranged. Our aim is to get this couple corresponding again, even if they cannot be coherent.

Indeed, the two kinds of criteria, coherence and correspondence, can sometimes be at odds with each other. For instance, in social situations, including some competitive games and predator-prey interactions, it can be advantageous to exhibit inconsistent behavior in order to maximize adaptive unpredictability and avoid capture or loss (Driver & Humphries 1988). In Chapters 4 and 5, we introduce a similarly illogical heuristic – the Minimalist heuristic – that violates transitivity but nevertheless makes fairly robust and accurate inferences in particular environments. Thus, logic and adaptive behavior can be logically distinct.

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

To conclude: Heuristics are not simply hobbled versions of optimal strategies. There are no optimal strategies in many real-world environments in the first place. This does not mean, though, that there are no performance criteria in the real world. 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 stemming from the laws of logic and probability with multiple correspondence criteria relating to real-world decision performance. But there is a further difference between these two sets of multiple criteria: While *all* coherence criteria must be met for a decision method to be deemed rational, correspondence criteria can be considered in relation to each other. In some environments, for instance, it may be more important to make a decision quickly than completely accurately. However, one of the surprising empirical results reported in our book is that simple heuristics need not always make such tradeoffs. We show that, when compared to some standard benchmark strategies on a range of decision tasks, fast and frugal heuristics can be faster, more frugal, and more accurate at the same time. No tradeoff need be considered.

6.2. Do people use fast and frugal 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? Exploring 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 if we are also concerned with 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?

We know rather little about the heuristic principles of limited search and stopping that people and animals use. One major reason for this is that the typical experimental task eliminates search in the first place (but see, e.g., Connolly & Gilani 1982; Payne et al. 1993; Saad & Russo 1996). Researchers usually sidestep questions of search by using tasks in which all pieces of information – usually only two or three – are already conveniently laid out in front of the participant. Theories of cognition and the experimental tasks used to test those theories often conspire hand in hand to overlook limited search and stopping rules. More is thus known about the heuristic decision principles that people employ (e.g., Payne et al. 1993), and we have begun to investigate this with some of the fast and frugal heuristics described in the book as well. Additionally, we have started to make some inroads into the questions surrounding information search by using an experimental setting in which cues must be actively sought. We now give two brief examples of the kinds of empirical evidence we are gathering.

How can we distinguish whether people are using a simple versus a more complex decision strategy? One way is to compare the decision performance of humans and algorithms, using *outcome measures* that focus on the final decision behavior. Experiments designed to test whether or not people use the recognition heuristic, for instance (Ch. 2), showed that in 90% of the cases where individuals could use the recognition heuristic when comparing the sizes of two cities (i.e., when they recognized one city but not the other), their choices matched those made by the recognition heuristic. This does not prove that participants were actually using the recognition heuristic to make their decisions, however – they could have been doing something more complex, such as using the information about the recognized city to estimate its size and compare it to the average size of unrecognized cities (though this seems unlikely). Additional evidence that the recognition heuristic *was* being followed, though, was obtained by giving participants extra information about recognized cities that contradicted the choices that the recognition heuristic would make – that is, participants were taught that some recognized cities had cues indicating small size. Despite this conflicting information (which could have been used in the more complex estimation-based strategy described above to yield different choices), participants still made 92% of their inferences in agreement with the recognition heuristic. Furthermore, participants typically showed the less-is-more effect predicted by the earlier theoretical analysis of the recognition heuristic, strengthening suspicions of this heuristic's presence.

But often outcome measures are insufficient to distinguish between simple and complex heuristics, because they all lead to roughly the same level of performance (the "flat maximum" problem). Furthermore, comparisons made only on selected item sets chosen to accentuate the differences between algorithms can still lead to ambiguities or ungeneralizable findings (Ch. 7). Instead, *process measures* can reveal differences between algorithms that are reflected in human behavior. For instance, noncompensatory algorithms, particularly those that make decisions on the basis of a single cue, would direct the decision maker to search for information about one cue at a time across all of the available alternatives. In contrast, compensatory algorithms that combine all information about a particular choice would direct search for all of the cues of one alternative at a time. We have found such evidence for fast and frugal heuristics in laboratory settings where participants must actively search for cues (Ch. 7), especially in situations where time-pressure forces rapid decisions. However, there is considerable variability in the data of these studies, with many participants appearing to use more complex strategies or behaving in ways that cannot be easily categorized. Thus much work remains to be done to provide evidence for when humans and other animals use simple heuristics in their daily decisions.

7. How our research program relates to earlier notions of heuristics

The term "heuristic" is of Greek origin, meaning "serving to find out or discover." From its introduction into English in the early 1800s up until about 1970, "heuristics" referred to useful, even indispensable cognitive processes for solving problems that cannot be handled by logic and probability theory alone (e.g., Groner et al. 1983; Polya 1954). After 1970, a second meaning of "heuristics" emerged in the

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

fields of psychology and decision-making research: limited decision-making methods that people often misapply to situations where logic and probability theory should be applied instead (e.g., Tversky & Kahneman 1974). We use the term in the same positive sense as the earlier theorists, emphasizing their beneficial role in guiding search, and following Simon and Newell's emphasis on creating precise computational models. However, we break with the past tradition of using well-defined artificial settings for the study of heuristics, such as mathematical problems (Polya 1954) or the games of chess and cryptarithmetic that Newell and Simon (1972) investigated. Instead, our research addresses how fast and frugal heuristics can make inferences about unknown aspects of real-world environments.

The research most closely related to the ABC program on fast and frugal heuristics is that on adaptive decision making and on simple classification rules in machine learning. In their study of the "adaptive decision maker," Payne et al. (1993) studied the trade-off between accuracy and effort for various choice strategies, including lexicographic rules and Elimination by Aspects (Tversky 1972). Payne et al. emphasized that a decision maker has a multitude of strategies available and chooses between them depending on their costs and accuracy given constraints such as time pressure. One important distinction from the ABC program is that Payne et al. focused on preferences, such as between hypothetical job candidates or randomly selected gambles, rather than on inferences whose correct answer can be assessed, such as which soccer team will win or which of two cities is larger. As a consequence, they measured a strategy's accuracy by how closely it matched the predictions of a weighted additive rule, the traditional gold standard for rational preferences. Thus, in Payne et al.'s research a heuristic can never be better than a weighted additive rule in accuracy (though it may require less computational effort). In contrast, by measuring the performance of all competing strategies against external real-world criteria, we find that fast and frugal heuristics can be more accurate than a weighted additive rule both in theory (Ch. 4) and in practice (Ch. 5). Research in machine learning does typically focus on inferences about real-world environments, allowing accuracy to be measured objectively. Work on simple classification rules that use only one or a few cues (e.g., Holte 1993; Rivest 1987) has demonstrated that fast and frugal methods can be accurate, as well as being robust generalizers owing to their limited parameter use.

A very different notion emerged in psychology in the early 1970s, emphasizing how the use of heuristics can lead to systematic errors and lapses of reasoning that indicate human irrationality. This "heuristics-and-biases" program launched by Tversky and Kahneman (1974) tainted the idea of simple mental mechanisms by attaching them to the value-laden "bias" term in a single inseparable phrase. Within this program, heuristics were often invoked as the explanation when errors – mainly deviations from the laws of probability – were found in human reasoning. Although Tversky and Kahneman (1974) repeatedly asserted that heuristics sometimes succeed and sometimes fail, their experimental results were typically interpreted as indicating some kind of fallacy, which was usually attributed to one of three main heuristics: representativeness (judgments influenced by what is typical), availability (judgments based on what come easily to mind), or anchoring and adjustment (judgments relying on what comes first). The reasoning fallacies described by the heuristics-and-biases program have not only been deemed irrational, but they have also been interpreted as signs of the bounded rationality of humans (e.g., Thaler 1991, p. 4). Equating bounded rationality with irrationality in this way is as serious a confusion as equating it with constrained optimization. Bounded rationality is neither limited optimality nor irrationality.

Our research program of studying fast and frugal heuristics shares some basic features with the heuristics-and-biases program. Both emphasize the important role that simple psychological heuristics play in human thought, and both are concerned with finding the situations in which these heuristics are employed. But these similarities mask a profound basic difference of opinion on the underlying nature of rationality, leading to very divergent research agendas: In our program, we see heuristics as the way the human mind can take advantage of the structure of information in the environment to arrive at reasonable decisions, and so we focus on the ways and settings in which simple heuristics lead to accurate and useful inferences. Furthermore, we emphasize the need for specific computational models of heuristics rather than vague one-word labels like "availability." In contrast, the heuristics-and-biases approach does not analyze the fit between cognitive mechanisms and their environments, in part owing to the absence of precise definitions of heuristics in this program. Because these loosely defined heuristics are viewed as only partly reliable devices commonly called on, despite their inferior decision-making performance, by the limited human mind, researchers in this tradition often seek out cases where heuristics can be blamed for poor reasoning. For arguments in favor of each of these views of heuristics, see the debate between Kahneman and Tversky (1996) and Gigerenzer (1996).

To summarize the place of our research in its historical context, the ABC program takes up the traditional notion of heuristics as an essential cognitive tool for making reasonable decisions. We specify the function and role of fast and frugal heuristics more precisely than has been done in the past, by building computational models with specific principles of information search, stopping, and decision making. We replace the narrow, content-blind norms of coherence criteria with the analysis of heuristic accuracy, speed, and frugality in real-world environments as part of our study of ecological rationality. Finally, whereas the heuristics-and-biases program portrays heuristics as a possible hindrance to sound reasoning, we see fast and frugal heuristics as enabling us to make reasonable decisions and behave adaptively in our environment.

8. The adaptive toolbox

Gottfried Wilhelm Leibniz (1677/1951) dreamed of a universal logical language, the Universal Characteristic, that would replace all reasoning. The multitude of simple concepts constituting Leibniz's alphabet of human thought were all to be operated on by a single general-purpose tool such as probability theory. But no such universal tool of inference can be found. Just as a mechanic will pull out specific wrenches, pliers, and spark-plug gap gauges for each task in maintaining a car's engine rather than merely hitting everything with a large hammer, different domains

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

of thought require different specialized tools. This is the basic idea of the *adaptive toolbox:* the collection of specialized cognitive mechanisms that evolution has built into the human mind for specific domains of inference and reasoning, including fast and frugal heuristics (see also Bettman 1979; Cosmides & Tooby 1992; Payne et al. 1993). The notion of a toolbox jumbled full of unique one-function devices lacks the beauty of Leibniz's dream of a single all-purpose inferential power tool. Instead, it invokes the more modest but surprising abilities of a "backwoods mechanic and used parts dealer" (as Wimsatt 2000a describes Nature) who can provide serviceable solutions to most any problem with just the things at hand.

The adaptive toolbox contains psychological (as opposed to morphological or physiological) adaptations (Tooby & Cosmides 1992). These include so-called "lower-order" perceptual and memory processes which can be fairly automatic, such as depth perception, auditory scene analysis, and face recognition, as well as "higher-order" processes that are based on the "lower" processes and can be at least partly accessible to consciousness. Higher-order mental processes include the examples we have discussed earlier of inferring whether a heart attack victim should be treated as a high- or low-risk patient and deciding whom to marry. The focus of *Simple heuristics* is on fast and frugal heuristics for higher-order cognitive processes that call upon lower-order processes of cue perception and memory. We also apply this constructive view to the mental tools themselves, creating heuristics from combinations of building blocks and other heuristics, as described in section 3. This feature distinguishes the adaptive toolbox image from the similar metaphor of the mind as a Swiss Army knife (Cosmides & Tooby 1992). Both analogies emphasize that the mind uses a collection of many specifically designed adaptive strategies rather than a few general-purpose power tools, but the toolbox metaphor puts more emphasis on the possibility of recombining tools and building blocks and the nesting of heuristics.

Lower-order perceptual and memory processes such as face and voice recognition are complex and difficult to unravel, in part because they make use of massively parallel computations. No one has yet managed to build a machine that recognizes faces as well as a two-year-old child. Now consider a higher-order decision mechanism that makes inferences based on these processes, the recognition heuristic introduced in Chapter 2. This fast and frugal heuristic uses recognition to make rapid inferences about unknown aspects of the world. Although the mechanisms of recognition memory may be intricate and complex, the recognition heuristic can be described as an algorithm just a few steps long. There is thus no opposition between simple and complex processes operating in the mind – both have their place, and can be studied somewhat independently. We do not need to know precisely how recognition memory works to describe a heuristic that relies on recognition. This example illustrates an apparently paradoxical thesis: *Higher-order cognitive mechanisms can often be modeled by simpler algorithms than can lower-order mechanisms.*

This thesis is not new. It has been proposed in various forms over the past century, as for example by proponents of the Würzburg school of psychology in the early twentieth century (Kusch 1999) and more recently by Shepard (1967). The thesis has limits as well, of course: Some higher-order processes, such as the creative processes involved in the development of scientific theories or the design of sophisticated artifacts, are most likely beyond the purview of fast and frugal heuristics. But we believe that simple heuristics can be used singly and in combination to account for a great variety of higher-order mental processes that may at first glance seem to require more complex explanation, as we demonstrate throughout our book.

9. Remaining challenges

Simple heuristics presents our efforts to date at advancing a vision of ecological rationality arising from fast and frugal decision mechanisms matched to their task environments. Our successes have been modest in the face of the challenges that remain. Here we indicate the directions that this research program must explore for us to gain a fuller understanding of how minds can make use of simple heuristics.

Cognitive tasks. The first challenge is to explore fast and frugal heuristics for solving tasks beyond those we considered so far. What other classes of decisions can be made by simple mechanisms? How can fast and frugal cognition help in tasks that extend over time such as planning or problem solving? Can simple heuristics be applied to perceptual mechanisms as well? We expect so – a few researchers have called perception a "bag of tricks" (e.g., Ramachandran 1990), full of quick and sometimes dirty mechanisms that evolved not because of their consistency but because they worked.

Adaptive problems. The next challenge is to study how fast and frugal heuristics are applied to important adaptive problems – how domain-specific should we expect simple heuristics to be? The discovery of domain-specific heuristics for important adaptive problems may help clarify how the mind is organized – for instance, if heuristics used for sequential mate search differ from heuristics for sequential habitat search, this may indicate that mate choice and habitat choice are distinct domains with specialized mechanisms. What heuristics apply to adaptive problems such as food choice (including modern forms of dieting), health preservation (including visiting doctors and taking drugs), and navigation (including getting from one end of a city to another)?

Social norms and emotions. Simple heuristics can also be advantageous for navigating the complexities of social domains, and can be learned in a social manner, through imitation, word of mouth, or cultural heritage. We suspect that social norms, cultural strictures, historical proverbs, and the like can enable fast and frugal social reasoning by obviating cost-benefit calculations and extensive information search. We also speculate that emotions may 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. Where can we find further evidence for the decision-making functions of these cultural and emotional processes, and how can they serve as building blocks in precise models of fast and frugal heuristics?

Ecological rationality. We do not yet have a well-developed language 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

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

perspectives. For instance, the statistical measures of two-dimensional patterns developed in spatial data analysis (e.g., Upton & Fingleton 1985) can be used when assessing heuristics for spatial search in foraging or habitat selection.

Performance criteria. How should the performance and usefulness of heuristics be measured? Ultimately, ecological rationality depends on decision making that furthers an organism's adaptive goals in the physical or social environment. How can measures of decision speed, frugality, and accuracy be augmented by and combined with measures of adaptive utility? We have tested the generalization ability of heuristics so far only in cross-validation tests. How can we measure predictive accuracy and robustness in environments that are in a state of continual flux, with new objects and cues appearing over time? Finally, we have focused on adaptive goals in terms of correspondence criteria (e.g., accuracy, speed, and frugality) as opposed to the coherence criteria (e.g., consistency, transitivity, and additivity of probabilities) traditionally used to define rationality. Is there any role for coherence criteria left? Should one follow Sen (1993) in arguing that consistency is an ill-defined concept unless the social objectives and goals of people are specified?

Selecting heuristics. How does the mind know which heuristic to use? Following our bounded rationality perspective, a fast and frugal mind need not employ a meta-level demon who makes optimal cost-benefit computations when selecting 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 (Ch. 1). But we have not yet addressed how individual heuristics are selected from the adaptive toolbox for application to specific problems.

Multiple methodologies. The combination of conceptual analysis, simulation, and experimentation has deepened our understanding of fast and frugal heuristics. However, more evidence must be amassed for the prevalence of simple heuristics in human and animal reasoning. This need not be done solely through laboratory experiments, where we often find that alternative mechanisms can equally account for the observed behavior (as discussed in Ch. 7). Collecting data from the field – whether that field is a jungle habitat or an airplane cockpit – is also vital for discovering new heuristics and teasing competing mechanisms apart.

10. Summary of the ABC view of rationality

The research program described in *Simple heuristics* is designed to elucidate three distinct but interconnected aspects of rationality (see also Chase et al. 1998):

1. *Bounded rationality.* Decision-making agents in the real world must arrive at their inferences using realistic amounts of time, information, and computational resources. We look for inference mechanisms exhibiting bounded rationality by designing and testing computational models of fast and frugal heuristics and their psychological building blocks. The building blocks include heuristic principles for guiding search for information or alternatives, stopping the search, and making decisions.

2. *Ecological rationality.* Decision-making mechanisms can exploit the structure of information in the environment to arrive at more adaptively useful outcomes. To understand how different heuristics can be ecologically rational, we characterize the ways that information can be structured in different decision environments and how heuristics can tap that structure to be fast, frugal, accurate, and adaptive at the same time.

3. *Social rationality.* The most important aspects of an agent's environment are often created by the other agents it interacts with. Thus, predators must make crucial inferences about the behavior of their prey (Ch. 12), males and females must make decisions about others they are interested in mating with (Ch. 13), and parents must figure out how to help their children (Ch. 14). Social rationality is a special form of ecological rationality, and to study it we design and test computational models of fast and frugal heuristics that exploit the information structure of the social environment to enable adaptive interactions with other agents. These heuristics can include socially adaptive building blocks, such as social norms and emotions of anger and parental love, which can act as further heuristic principles for search, stopping, and decision.

These three aspects of rationality look toward the same central goal: to understand human (and animal) behavior and cognition as it is adapted to specific environments, both ecological and social, and to discover the heuristics that guide adaptive behavior. In some ways, this view leaves behind a certain sense of beauty and morality associated with the dream of optimal thought. Leibniz's universal calculus exhibits the aesthetics and the moral virtue of this lofty ideal, as does Laplace's omniscient superintelligence. Cognitive scientists, economists, and biologists have often chased after the same beautiful dreams by building elaborate models endowing organisms with unlimited abilities to know, memorize, and compute. These heavenly dreams, however, tend to evaporate when they encounter the physical and psychological realities of the waking world: Mere mortal humans cannot hope to live up to these standards, and instead appear nightmarishly irrational and dysfunctional in comparison.

In the face of this dilemma, many researchers have still preferred to keep dreaming that humans can approximate the exacting standards of optimality, rather than surrendering to an ungodly picture of human irrationality and stupidity. The choice, however, is not between an unrealistic dreaming rationality and a realistic nightmare irrationality. There is a third vision that dispenses with this opposition: rationality through simplicity, and accuracy through frugality. In *Simple heuristics*, we strive to paint in a few more of the details of this hopeful vision.

NOTE

1. Gerd Gigerenzer, Peter M. Todd, and the ABC Research Group are the authors of *Simple heuristics that make us smart* (1999) Oxford University Press.

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

Open Peer Commentary

Commentary submitted by the qualified professional readership of this journal will be considered for publication in a later issue as Continuing Commentary on this article. Integrative overviews and syntheses are especially encouraged.

The evolution of rational demons

Colin Allen

Department of Philosophy, Texas A&M University, College Station, TX 77843-4237. **colin-allen@tamu.edu www-phil.tamu.edu/~colin/**

Abstract: If fast and frugal heuristics are as good as they seem to be, who needs logic and probability theory? Fast and frugal heuristics depend for their success on reliable structure in the environment. In passive environments, there is relatively little change in structure as a consequence of individual choices. But in social interactions with competing agents, the environment may be structured by agents capable of exploiting logical and probabilistic weaknesses in competitors' heuristics. Aspirations toward the ideal of a demon reasoner may consequently be adaptive for direct competition with such agents.

Gigerenzer, Todd, and the ABC Research Group (Gigerenzer et al. 1999) provide a compelling account of real-world decision making. I very much like what they have accomplished and I do not expect to say anything here that they would strongly disagree with. But as someone whose livelihood depends to a considerable extent on attempting to inculcate "the baggage of the laws of logic" (p. 365) into minds of varying impressionability, I wonder how much to fear that my attempts to present first-order logic as a standard of good reasoning will be consigned to lectures to the Flat Earth Society.

I take it that people (and not just philosophers) *are* sometimes concerned that their reasoning satisfies what the authors call "coherence criteria" – criteria encoded in classical logics and probability theory. I take it also that dreams of "demon" rationality were not implanted in us by a Cartesian deity. Humans, under their own impetus, aspire to consistency in their beliefs and likelihood assessments, and they reason accordingly. From whence comes this aspiration? And is it an adaptation?

One of Gigerenzer et al.'s important principles is that fast and frugal heuristics are successful because environments are structured in specific ways (Gigerenzer et al. 1999). The recognition heuristic works because of a usually dependable relationship between the saliency of names in the environment and target properties such as population size or market performance. Other fast and frugal heuristics take into account additional cues that are correlated with target properties. Simple heuristics that use only a subset of the available cues to discriminate options can lead to remarkably accurate assessments of the target properties. But fast and frugal heuristics can also deliver inconsistent judgments. As the authors point out, some heuristics can, in principle, return judgments such as that city A is larger in population than city B, B is larger than C, and C is larger than A.

Gigerenzer et al. (1999) provide no empirical evidence that people actually make such inconsistent judgments. But if they do, I would also like to know whether, in a given session, subjects hesitate or show other signs of uncertainty when they produce a judgment that is inconsistent with others just produced. Except under the most severe time pressures, I would be astounded if people blithely produced the inconsistent series of judgments shown above without revealing some sign of discomfort. By assigning a penalty for wrong answers – a trick sometimes employed by my colleagues who use multiple-choice examinations – one might perhaps expect to induce consistency checking. So, in considering when people might not use simple heuristics, I would like to suggest investigating those situations where the cost of being wrong exceeds the benefit of being right. Perhaps it is no accident that Darwin attempted to consider all his reasons in deciding whether to marry!

Where else might attempting to live up to the ideal of a demon rationalizer be an adaptive thing to do? Here I think it may be interesting to distinguish between the "passive" environment and the "agentive" environment. According to this distinction the structure of the passive environment is relatively independent of an individual's behavior or dispositions. The passive environment contains what would be thought of as the nonsocial environment as well as large-scale aspects of the social environment (such as the stock market) which are not responsive to the isolated acts of individuals. For example, the environmentally-provided relationships between fruit color and nutritional value or between company name recognition and company stock performance are not going to change (in the short run) if an individual tends to rely heavily on that cue to make feeding or investing decisions. (In the long run, of course, there may be coevolution between plants and feeders, and there will be an effect on the market if many individuals pick stocks on name recognition alone.) In contrast, in dealing directly with other agents, an individual's selection tendencies can be exploited by a savvy opponent. Thus, for example, if you know that I rely on presence of a university as the best cue for estimating city size, you could lower my overall success rate by presenting me with disproportionately many pairs of cities that violate the norm. Here, the structure of the encountered environment is not independent of the subject's dispositions. Failures of logical or probabilistic consistency provide additional vulnerabilities that a logically or statistically proficient opponent might exploit.

We might see our imperfect implementation of demon rationality as an evolutionary artifact, the byproduct of other adapted systems, or as itself a cognitive adaptation. I have never really liked the often-produced hypothesis that intelligence is a social adaptation but here, perhaps, is a place where it finds a home. Individuals who are smart because of fast and frugal heuristics alone can be exploited in social interactions by those who are more proficient in logic and probability. There may therefore be selection pressure to develop more demon-like forms of rationality. Fast and frugal heuristics certainly are part of what makes us smart. But the "baggage" of logic and probability theory could be well worth carrying when journeying among competing agents. I shall continue to load up my students with as much as they can bear.

Where does fast and frugal cognition stop? The boundary between complex cognition and simple heuristics

Thom Baguley,[1] and S. Ian Robertson[2]

[1]*Department of Human Sciences, Loughborough University, Loughborough, LE11 3TU, England;* [2]*Department of Psychology, University of Luton, Luton, LU1 3JU, England.* **t.s.baguley@lboro.ac.uk ian.robertson@luton.ac.uk**

Abstract: *Simple heuristics that make us smart* presents a valuable and valid interpretation of how we make fast decisions particularly in situations of ignorance and uncertainty. What is missing is how this intersects with thinking under even greater uncertainty or ignorance, such as novice problem solving, and with the development of expert cognition.

Reading *Simple heuristics* led to a curious feeling of familiarity. Many of the ideas were familiar, yet at the same time there was a sense of discovery. The notion that evolution and ecological constraints influence cognition is topical in recent research (Lansdale 1998), student texts (Robertson 1999), and popular science (Pinker 1997). The idea that decisions draw on evidence provided by our memories is also well established. Although several impor-

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

tant contributions are made by the book we focus on what we consider the most exciting one: the proposal of a flexible, modular, ecologically rational toolbox of fast and frugal heuristics.

First, the toolbox approach offers the potential to balance the costs of domain-general and domain-specific solutions to problems. Second, it attempts detailed process models of cognition. Third, it may offer a way of integrating a range of work within and without cognitive science. None of these potential benefits has been fully realized, yet *Simple heuristics* is a promising sketch of how such a toolbox might work (e.g., of how skewed environmental distributions can constrain our selection of search or decision rules). The close relationship between environmental structure and heuristics suggests that a drawing together of research on perception, memory, and cognition is required. An earlier response to such a proposal led to the development of large, far-from-frugal, unified theories of cognition (Newell 1990). Unified theories such as SOAR or ACT-R try to model everything at once, whereas the adaptive toolbox approach tries to find niches where simple strategies can be applied in a modular way.

ACT-R and SOAR emerged from a tradition of research on high-level cognition such as problem solving (Newell & Simon 1972). Where does fast and frugal cognition stop and where does slow and lavish cognition such as reasoning or problem solving begin? One extreme view is that all cognition is fast and frugal: thinking always involves some computational shortcut to solve a syllogism, prove a theorem, or discover a solution. For example, skill and expertise are often readily modeled by recognition of familiar problems and application and modification of routine methods (see Payne 1988). This position does not seem that far from the sympathy expressed with Brooks's program of research which eschews "building elaborate mental models" of the world (Ch. 15, Goodie, Ortmann, Davis, Bullock & Werner, p. 334). At the other extreme, it may be argued that because fast and frugal heuristics require a certain "beneficial" level of ignorance this limits them to a few narrow niches. There is no reason to take on either extreme (at least where human cognition is concerned). However, it is necessary to explicate the relationship between fast and frugal heuristics and the existing body of research on complex cognition such as problem solving.

When fast and frugal approaches fail are we then unable to achieve our desired goal directly and therefore resort to problem solving? We think that the fast and frugal approach can be more intimately linked to traditional problem solving approaches than this. Take one of several issues left open in *Simple heuristics:* that of selecting heuristics (Ch. 16, Gigerenzer et al. 1999, p. 364). This occurs in two guises: how heuristics become available for selection, and, once available, how one is selected. Humans may learn new heuristics, new principles for selecting heuristics, or develop expertise with their application through problem solving. Selecting and applying a heuristic seems to require a more-or-less elaborate representation of a domain. Adults are able to use heuristics in a sophisticated way because they have a relatively rich conceptual representation of the world compared to infants. Many tasks that appear simple for adults are predicated on rich expertise (Wood 1988). Problem solving and reasoning are important for exploring unfamiliar situations and building up representations of them (Baguley & Payne 2000). This takes time and effort (even if relatively slow and frugal weak methods such as imitative problem solving are used; Robertson, in press), but the dividends may be very large. Once those representations are established, experts are able to apply domain-specific solutions to problems just as mature humans can apply heuristics such as "Take the Best" in familiar environments with the appropriate structure. Reading *Simple heuristics* sometimes gave us the impression that evolution is the only candidate for the origins of the toolbox components. This needs qualifying. Humans are innately equipped to develop and learn certain classes of heuristics.

We suggest, therefore, that fast and frugal heuristics supplement traditional problem solving research. They may be brought to bear on selecting and applying operators under means-ends analysis (e.g., where a recognition heuristic is used to avoid returning to previous states). In conclusion, the adaptive toolbox has much to offer. Important aspects remain to be fleshed out, and we have reservations about some elements of the approach. The way the toolbox draws on supposed end-products of memory (given that the "binary quality of recognition" is itself the output of an incompletely understood decision process) may be an over-simplification. It may also be unwise to reject normative approaches out of hand, although we agree that norms should not be applied unquestioningly.

Keeping it simple, socially

Louise Barrett[1,2] and Peter Henzi[2]

[1]*Evolutionary Psychology and Behavioural Ecology Research Group, School of Biological Sciences, University of Liverpool, L69 3BX, United Kingdom;* [2]*Behavioural Ecology Research Group, School of Social Anthropology and Psychology, University of Natal, Durban 4041, South Africa.*
louiseb@liv.ac.uk **henzi@mtb.und.ac.za**

Abstract: Fast and frugal heuristics function accurately and swiftly over a wide range of decision making processes. The performance of these algorithms in the social domain would be an object for research. The use of simple algorithms to investigate social decision-making could prove fruitful in studies of nonhuman primates as well as humans.

Three and a half years ago, when we began a study of baboons in the Western Cape, South Africa, we used a very simple one-reason decision rule to pick our study troops: we selected the ones who slept on the cliffs directly opposite our research house. We have never regretted this decision, although a more thorough investigation of the baboon population might have yielded troops with a more intriguing demographic composition or more unusual ranging patterns. Our relatively uninformed search and decision-making process yielded baboons that were good enough. As advocates of one-reason decision-making, we were therefore delighted to discover that Gerd Gigerenzer, Peter Todd, and the ABC Research Group had conducted a much more systematic investigation of simple heuristics, and were very impressed with the results (Gigerenzer et al. 1999). One of the most satisfying elements of their analyses was the way in which the various algorithms were tested in real-world situations, rather than against a perfectly rational or optimal ideal. In the field of animal behaviour, we are all too often confronted with optimality models that make unrealistic assumptions ("imagine your baboon is a sphere moving in a vacuum") and treat animals as perfect rational agents with complete knowledge. If something as simple as the recognition heuristic can beat the stock market, that most fiendishly complicated of human inventions, then, as Gigerenzer et al. (1999) themselves suggest, it seems likely that this simplicity may also pay dividends within the animal kingdom.

The "fast and frugal" approach is particularly interesting to us as primatologists because of the essentially anthropocentric approach taken with studies of primate cognition. If humans, with their relatively enormous brains, can achieve more with less knowledge as Gigerenzer et al. (1999) have demonstrated, then it seems reasonable to suppose that the same might be true for our less well-endowed cousins. This may seem obvious, but primates are often imputed to have extremely advanced cognitive skills and reasoning abilities largely on the basis of their close relationship with humans, rather than any firm demonstration that this is the case. Monkeys and apes are often assumed to use complex mentalist analyses of their social situation and be able to adopt manipulative courses of action, as befits their increased neocortex size relative to other animals (see, e.g., Byrne & Whiten 1988). Perhaps, however, this increased cortex is actually needed to help animals think more simply, rather than in more complicated ways. [See also Whiten & Byrne: "Tactical Deception in Primates" BBS 11(2) 1988.]

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

Blurton-Jones et al. (1999) recently suggested that the need to learn increasing amounts of information over the course of human evolution would have selected for new ways of learning and processing information. Fast and frugal heuristics could have been one such mechanism to help speed decision-making in an ever more complicated social world. The same argument could, of course, be applied to the increase in social complexity over the course of primate evolution as a whole, and could explain the increase in primate brain size over that of the other mammals. It is therefore an interesting empirical issue as to whether fast and frugal heuristics are the result of our primate heritage, or whether they truly do represent a new way of using information that is not available to other primate species. Is a large neocortex needed to make the initial inferences about ecological (social) validities, correlations between cues and the structure of the social environment that allow fast and frugal algorithms to work so speedily and accurately? Rather ironically, is it the case that large brains are needed in order to reason in the simplest way possible? The success of the fast and frugal algorithm of parental investment suggests that this need not be the case, but the social world does present animals with an inherently more complex environment consisting as it does of other individuals with their own goals and interests at heart. Decisions made in the social domain may therefore require more complex assessments, but once the initial correlations and cue values are learned, the generalisation to a fast and frugal heuristic would seem likely to pay dividends. Testing the performance of fast and frugal heuristics in the social arena would seem to be the logical next step for Gigerenzer, Todd, and their colleagues to take. It would be interesting to know the limits of these algorithms under these more dynamic and interactive conditions. Would a point be reached where simple heuristics failed to make us smart, and more complex abilities such as "mind-reading" and the attribution of mental states were needed to perform well? Or do these abilities constitute fast and frugal heuristics in themselves?

The use of fast and frugal heuristics as an explanation of the social decision-making of other primate species also warrants further attention. Our own data on baboons show that females make short-term tactical decisions with regard to grooming partners and social interactions and, therefore do not support the idea that long-term strategic alliance formation is the key to primate sociality (Barrett et al. 1999). It would be interesting to investigate whether these short-term decisions can be modelled by a fast and frugal algorithm such as Take the Best and whether females exploit their ignorance to good advantage. The fact that other social decisions requiring mutual choice, such as finding a mate, can be modelled using fast and frugal algorithms suggests to us that social decision-making of the kind we observe would be an excellent candidate for investigating the power of simple heuristics from a comparative perspective.

Rationality, logic, and fast and frugal heuristics

José Luis Bermúdez

Department of Philosophy, University of Stirling, Stirling FK9 4LA, Scotland
jb10@stir.ac.uk www.stir.ac.uk/philosophy/cnw/webpage1.htm

Abstract: Gigerenzer and his co-workers make some bold and striking claims about the relation between the fast and frugal heuristics discussed in their book and the traditional norms of rationality provided by deductive logic and probability theory. We are told, for example, that fast and frugal heuristics such as "Take the Best" replace "the multiple coherence criteria stemming from the laws of logic and probability with multiple correspondence criteria relating to real-world decision performance." This commentary explores just how we should interpret this proposed replacement of logic and probability theory by fast and frugal heuristics.

The concept of rationality is Janus-faced. It is customary to distinguish the psychological laws governing the actual processes of reasoning from the normative theories according to which such reasoning is to be evaluated (Nozick 1993).[1] Authors who make this distinction often enjoin us to ignore one of the faces to concentrate on the other – either instructing us, as Frege (1918–1919) famously did, to ignore the messy details of psychology to focus on the objective relations between thoughts that are the domain of logic, or, like Willhelm Wundt (1973, cited on p. 357), imploring us to ignore the subtleties of the logician in order to explain the psychology of reasoning. It is clear that Gigerenzer et al. (1999) favour the second of these positions. In fact, they seem to be making an even stronger claim. They seem to be suggesting that, where the dictates of the logician and the heuristics of the quick and dirty reasoner come into conflict, it is logic and probability theory that must be sacrificed. But what exactly does this mean in detail? Why should the prevalence of fast and frugal heuristics have any implications at all for the normative theories of logic and probability theory?

The notion of a fast and frugal heuristic has been around for a long time, although not perhaps under that name (Kahneman & Tversky 1973; Simon 1982). There is a standard way of interpreting such heuristics so that they are not really in conflict at all with the normative theories of logic and probability theory. One might say, for example, that, although the normative theories provide the standards by which practical reasoning ought to be judged (that is to say, they tell us what *ought* to be done, or what it is *rational* to do, in a particular situation when one possesses such-and-such information), there are nonetheless computational reasons why it is often not actually possible to use such theories in the actual process of decision-making. Because it is not always possible in practice to assign numerical probabilities to possible outcomes, or to give them numerical desirability ratings (let alone to use these figures to work out the course of action that best maximises expected utility), we are impelled to use shortcuts. But what justifies these shortcuts (what makes it *rational* to adopt them) is that they lead us to do more or less what we would have done had we actually put the normative theories into practice by working out the figures and crunching the numbers. Proponents of optimal foraging theory think that something like this holds of the "rules of thumb" converging on optimality that natural selection has thrown up (Dawkins 1995, Ch. 2). And, as the authors point out (Gigerenzer & Todd, p. 26), the accuracy of fast and frugal heuristics is often assessed in the literature by measuring how closely they match the predictions of a weighted additive rule, like the rule of maximising expected utility.

But this is emphatically not how Gigerenzer et al. view the operation of heuristics like "Take the Best." It can, by their lights, be rational to use such heuristics even when they result in courses of action that contravene the dictates of the normative theories. The heuristics can trump the normative theories. Again, there is a relatively innocuous way of understanding such a claim. One might think, for example, that it might be rational to use a heuristic even in a situation where it does not match the predictions of the normative rule simply because, when one takes a sufficiently long-term view, the overall benefits of using the heuristic outweigh the occasional benefits to be had by crunching the numbers each time. But all this really amounts to is the claim that we should evaluate a strategy rather than a particular application of that strategy. The normative standards according to which we do the judging are not themselves changed. Gigerenzer et al. certainly want to say something more controversial than this. But what?

Gigerenzer et al. say that they want to impose a new set of criteria for judging the rationality of decision-rules – what they term "multiple correspondence criteria relating to real-world decision performance" (Gigerenzer et al. 1999, p. 22). For example, it becomes rational to select the companies in one's investment portfolio by a simple version of the recognition heuristic (viz. only put companies with a high recognition factor into one's portfolio) because, as it happens, such portfolios seem to outperform both the relevant indices and random portfolios of stocks (see Ch. 3). That is, success becomes the sole determinant of rational decision-making.

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

There are several problems with this; they point to a difficulty with the research programme as a whole. First, there is a very important equivocation in how the stockmarket experiment is being described. The authors write as if their strategy was a pure application of the recognition heuristic. But this seems wrong. They did not invest in companies that they recognised. Rather, they invested in companies that had a high national and/or international recognition factor, where this is calculated statistically by comparing the recognition judgments of several different populations. These are two very different things. The first would have been a pure fast and frugal heuristic. The second, in contrast, seems much closer to a calculated investment strategy. What makes this equivocation important is that the notion of rationality applies very differently in the two cases. It is hard to see how anything other than a pure success-based criterion of rationality could be applied to the fast and frugal version of the recognition heuristic. Or, to put this another way, it is hard to see what reasons there might be for holding that it is rational to make one's investment decisions solely according to whether one has heard of the companies in question other than that the strategy more or less works over time – and it is equally hard to see how, if the strategy doesn't work, it could then possibly be described as rational. But the same does not hold of the sophisticated investment strategy of investing only in companies with a statistically attested high recognition. There are all sorts of reasons why this is a rational strategy to adopt – quite apart from the well-documented "big company effect" in bull markets (to which the authors themselves draw attention) and the simple thought that a company with a high recognition factor will correspondingly have a high market share. That is to say, even if it did turn out (as it probably would in a bear market) that the strategy did not beat the index it might well still count as a rational strategy to have adopted.

What this points us to is an important discussion in the concept of rationality. A workable concept of rationality must allow us to evaluate the rationality of an action without knowing its outcome. Without this the concept of rationality cannot be a useful tool in the control, regulation and evaluation of decision-making as and when it happens. And it is precisely such a way of evaluating the rationality of an action that we are offered by the orthodox normative theories of expected utility maximisation and so forth. But is far from clear that Gigerenzer and his co-workers have offered a genuine alternative to this. They claim to have replaced criteria of rationality based upon logic and probability theory with a heuristic-based criteria of real-world performance. but it doesn't look as if they've offered us criteria of *rationality*.

NOTE

1. For discussion of ways to strike the balance between these two facets of rationality see Bermúdez 1998; 1999a; 1999b; 1999c; 2000 and the essays in Bermúdez and Millar (in preparation).

How smart can simple heuristics be?

Nick Chater

Department of Psychology, University of Warwick, Coventry, CV4 7AL, United Kingdom **nick.chater@warwick.ac.uk**

Abstract: This commentary focuses on three issues raised by Gigerenzer, Todd, and the ABC Research Group (1999). First, I stress the need for further experimental evidence to determine which heuristics people use in cognitive judgment tasks. Second, I question the scope of cognitive models based on simple heuristics, arguing that many aspects of cognition are too sophisticated to be modeled in this way. Third, I note the complementary role that rational explanation can play to Gigenerenzer et al.'s "ecological" analysis of why heuristics succeed.

Gigerenzer, Todd, and the ABC Research Group have provided a series of impressive demonstrations of how simple "fast and frugal" cognitive heuristics can attain surprisingly impressive levels of performance, comparable to human performance in a range of tasks. They show, for example, that decision making based on a single piece of evidence, rather than integrating across all available evidence, can lead to close optimal performance in a wide range of estimation tasks (Gigerenzer et al. 1999, Ch. 4, p. 75, Gigerenzer & Goldstein). Gigerenzer et al. interpret these results as having radical implications for cognition in general – in particular, as undercutting the view that cognition must involve well-optimized cognitive machinery which behaves in accordance with classical rational norms of probability theory, logic, and decision theory. This line of thought raises the attractive possibility that the complexity of the mind may have been dramatically overestimated. Perhaps the mind is really just a collection of smart heuristics, rather than a fantastically powerful computing machine. This is an exciting and important thesis. This commentary focuses on three challenges to this approach, which may open up avenues for future research.

1. Empirical evidence. Gigerenzer et al. focus on providing a feasibility proof for the viability of a particular kind of simple reasoning heuristic. This task primarily involves providing computer simulations showing that simple heuristics give good results on specific decision problems, in comparison to conventional methods such as linear regression, and to other heuristic approaches, such as unit-weighted regression. But there is little by way of experimental evidence that people actually do reason in this way, aside from important but preliminary evidence reported in Chapter 7. This is particularly important precisely because the simulations in this book show that a wide range of algorithms give very similar levels of performance. Hence, prima facie, all these algorithms are equally plausible candidates as models of how people might perform on these problems.

In the absence of a broader set of experimental tests there is some reason to doubt that people make decisions by relying on one cue only. As Gigerenzer et al. note, in perception and language processing there is ample evidence that multiple cues are integrated in recognition and classification, in extremely complex ways (e.g., Massaro 1987). Gigerenzer et al. propose that these cases are in sharp contrast to the operation of conscious decision-making processes – determining whether this divide is a real one is an important area for empirical research.

2. Scope. One of the most startling findings in psychology is that, across a very wide range of judgment tasks, including medical diagnosis, expert performance does not exceed, and is frequently poorer than, results obtained by linear regression over sets of features of the cases under consideration (Meehl 1954; Sawyer 1966).

An equally startling finding, this time from artificial intelligence and cognitive science, has been that in everyday reasoning, people vastly outperform any existing computational model (Oaksford & Chater 1998a). Even the inferences involved in understanding a simple story draw on arbitrarily large amounts of world knowledge, and people must integrate and apply that knowledge highly effectively and rapidly. Attempts to model such processes computationally have become mired in the nest of difficulties known as the "frame problem" (Pylyshyn 1987).

So cognition is, in some regards, remarkably weak; and in other regards it is remarkably powerful. In the present context, the crucial point is that the simple heuristics discussed in this book are aimed at modeling areas where cognition is weak – indeed, where cognitive performance is already known to be frequently outperformed by linear regression. But it is by no means clear that the picture of the mind as a set of simple heuristics will generalize to everyday reasoning, where cognitive performance appears to be remarkably strong. Indeed, it may be that it is not that simple heuristics make us smart (as Gigerenzer et al.'s title suggests); rather it may be that we resort to simple heuristics to do the very thing we are *not* smart at.

3.* Why *do heuristics work? Gigerenzer et al. downplay the importance of traditional conceptions of rationality in their discussion of reasoning methods. Indeed, they note that a heuristic such

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

as Take the Best has not been derived from "rational" principles of probability or statistics. Instead, they focus on an ecological notion of rationality – does the heuristic work in practice on real world data?

The viewpoint may appear to be an alternative to more traditional notions of rationality as used in psychology (Anderson 1990; Chater et al. 1999; Oaksford & Chater 1998b), economics (Kreps 1990) and behavioral ecology (McFarland & Houston 1981), in which behavior is assumed to approximate, to some degree, the dictates of rational theories, such as probability and decision theory. But it may be more appropriate to see the two viewpoints as complementary. Gigerenzer et al. (1999) are concerned to demonstrate rigorously *which* particular heuristics are successful, by computer simulation on realistic data sets. Traditional rational theories aim to explain *why* heuristics work. They characterize the optimization problem that the cognitive process, economic actor or animal faces; using rational theories (probability, decision theory, operations research) to determine the "rational" course of action; and conjecture that the heuristics used in actual performance approximate this rational standard to some degree. From this point of view, rational methods can be viewed as compatible with the "ecological" view of rationality outlined in Gigerenzer et al. (1999). Focusing on simple cognitive heuristics does not make the application of rational standards derived from formal calculi unnecessary. Instead, it gives a defined role for rational explanation – to explain why and under what conditions those heuristics succeed in the environment. This perspective is, indeed, exemplified in Gigerenzer et al.'s formal analysis of the conditions under which the Take the Best heuristic is effective (Ch. 6) and consistent with Gigerenzer et al.'s valuable comparisons between Take the Best and Bayesian algorithms (Ch. 8).

This book shows an important direction for research on human reasoning. It should act as a stimulus for empirical, computational, and theoretical developments in this area.

Simple heuristics could make us smart; but which heuristics do we apply when?

Richard Cooper

School of Psychology, Birkbeck College, University of London, London, WC1E 7HX United Kingdom **r.cooper@psychology.bbk.ac.uk**
www.psyc.bbk.ac.uk/staff/rc.html

Abstract: Simple heuristics are clearly powerful tools for making near optimal decisions, but evidence for their use in specific situations is weak. Gigerenzer et al. (1999) suggest a range of heuristics, but fail to address the question of which environmental or task cues might prompt the use of any specific heuristic. This failure compromises the falsifiability of the fast and frugal approach.

Gigerenzer, Todd, and the ABC Research Group (1999) are right to criticise much contemporary psychological decision-making research for its focus on mathematically optimal approaches whose application requires unbounded time and knowledge. They have clearly demonstrated that an agent can make effective decisions in a range of ecologically valid decision-making situations without recourse to omniscient or omnipotent demons. They have also cogently argued that biological decision-making agents cannot have recourse to such demons. The question is therefore not "Do such agents use heuristics?", but "Which heuristics do such agents use (and when do they use them)?" Gigerenzer et al. acknowledge that this question is important, but address it only in passing.

Gigerenzer et al.'s failure to specify conditions that might lead to the use of specific fast and frugal heuristics compromises the falsifiability of the fast and frugal approach. Difficult empirical results may be dismissed as resulting from the application of an as-yet-unidentified fast and frugal heuristic or the combination of items from the "adaptive toolbox" in a previously unidentified way. Gigerenzer et al. criticise the heuristics-and-biases approach of Tversky and Kahneman (1974) on much the same grounds. They note, for example, that both base-rate neglect and conservatism (two apparently opposing phenomena) can be "explained" by appealing to the appropriate heuristic or bias (because Tversky & Kahneman provide insufficient detail on the conditions that are held to evoke particular heuristics or biases). Gigerenzer et al. contend, quite reasonably, that such a post-hoc appeal does not amount to an adequate explanation of the behaviour. It is suggested that the use by the ABC Research Group of precise computational simulation techniques avoids this criticism. The computational simulations are very welcome and add a dimension often lacking in decision making research, but they do not, in their disembodied form, address the question of which heuristic might be applied when.

The issue of heuristic selection is not entirely ignored by Gigerenzer et al. The suggestion is that heuristics are either selected (or built from components within the adaptive toolbox) on a task-by-task basis. The challenge therefore lies in specifying: (1) the conditions under which different established heuristics are employed; (2) the conditions that provoke the construction of novel, task-specific heuristics; (3) the basic components available in the adaptive toolbox; and (4) the mechanisms by which appropriate heuristics may be constructed from these components. Of these, only the third is discussed at any length – the building blocks are held to involve elements that direct search, stop search, and decide based on the results of search – but that discussion of these elements is insufficient to allow the fourth concern to be addressed.

At several points in the book heuristic selection is conceived of as a meta-level decision-making task, suggesting that one might use a (presumably fast and frugal) heuristic to decide which fast and frugal heuristic to apply in a given situation. Two issues of minor concern are the possibility of an infinite regress (How do we select the fast and frugal heuristic to decide which fast and frugal heuristic to apply in the first place?) and the possibility that (assuming the infinite regress is avoided) the fast and frugal heuristic doing the selection might not yield the optimal fast and frugal heuristic for the original decision-making task.

Speculation as to the environmental and task cues that might lead to selection of the fast and frugal heuristics discussed in the book does not yield an obvious solution. For example, it is implied that one-reason decision making is appropriate (and presumably employed) for binary choice and that Categorisation By Elimination (CBE) is appropriate (and presumably employed) for multiple choice. However, both heuristics might be applied in both situations. The question of which heuristic to apply when remains.

The importance of heuristic selection is compounded by the lack of evidence presented in favour of the use by human decision makers of many of the heuristics discussed. For example, both QuickEst and CBE are presented solely as heuristics that can be shown to be fast and frugal. No comment is made on the psychological reality of either of these heuristics. This seems particularly odd when robust psychological findings that would appear to be of relevance (such as those addressed by Tversky and Kahneman's [1974] heuristics-and-biases approach) are ignored. How, for example, might the fast and frugal approach address the phenomena that Gigerenzer et al. use to demonstrate the difficulties present in the heuristics-and-biases approach (base-rate neglect and conservatism), or the confirmation bias often seen in diagnosis versions of categorisation? The latter, in particular, appears to be in direct conflict with the only categorisation heuristic proposed (CBE).

Fast and frugal heuristics have great promise. Human decision making cannot result from the application of algorithms with unbounded costs. Gigerenzer et al. have shown that fast and frugal heuristics can yield good decisions. They have not shown that humans use such heuristics, and by not addressing the question of which heuristics might be applied when, they have, like Tversky and Kahneman, given us a theory of human decision making that is unfalsifiable.

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

Psychological research on heuristics meets the law

Christoph Engel

Max-Planck-Projektgruppe, Recht der Gemeinschaftsguter, Poppelsdorfer Allee 45, D 53115 Bonn, Germany. **engel@mpp-rdg.mpg.de**
www.mpp-rdg.mpg.de

Abstract: Heuristics make decisions not only fast and frugally, but often nearly as well as "full" rationality or even better. Using such heuristics should therefore meet health care standards under liability law. But an independent court often has little chance to verify the necessary information. And judgments based on heuristics might appear to have little legitimacy, given the widespread belief in formal rationality.

The mind is one of the last unknown territories of this world. Gigerenzer, Todd, and the ABC Research Group (1999) make an expedition into this territory and find it inhabited by strange and unexpected beings. They have baptized them fast and frugal heuristics, with given names like Take the Best. (1) What can a lawyer learn from these insights into our internal geography? And (2) is there anything he could give psychologists in return?

(1) The book starts with what reads like a legal case history. A man is rushed to a hospital in the throes of a heart attack. The doctor needs to decide quickly whether the victim should be treated as a low-risk or a high-risk patient. Only in the latter case does he receive expensive care. Assume the doctor has read the book and applies the heuristic reported therein. The doctor tests systolic blood pressure first. If it is below 91, the patient is immediately treated as high risk. If not, the doctor checks the patient's age. A patient under 62.5 years is classified as low risk. For older patients the doctor searches for a third cue, *sinus tachycardia.* If it is present, the patient gets full and immediate treatment. If not, the patient is once and forever treated as low risk. Let this be our patient, and assume that he dies from heart disease on the next day. His relatives sue the hospital for malpractice. They point to a host of other diagnostic means that are state of the medical art. Does the hospital nonetheless have a chance to win in court?

The answer depends on the standard of health care. It would be of no use for the hospital to recall what the book reports on the limitations of the individual mind. At least not insofar as these limitations can be overcome by technology, other specialists, anticipatory training or better organisation. Within this framework, time pressure might count. And if, in that legal order, cost-benefit analysis may be applied to questions of life and death, decision costs might be an argument. The hospital would be in a stronger position if the doctor had to choose among several patients who looked similarly ill at the outset. But the attorney of the hospital would focus on another argument: other hospitals that tested all the scientifically available cues made on average less accurate choices.

Is the law hindered from using this knowledge? There are two major obstacles. Going to court is a form of third party settlement. This third party needs reliable information on the facts of the case. In cases of alleged malpractice, this information is not easy to collect, and even more difficult to prove. Information economics speaks about nonverifiable information. Standard court practice overcomes the problem by formal rules of consent, and by relying on state of the art treatment. If the doctor proved both, he would normally not be liable. This strategy is difficult to transfer to heuristics, for they perform well precisely because they are not general, but issue-specific or even situation specific. They exploit local knowledge. Local knowledge is often unverifiable. And the general reliability of a search and decision tool is hard to prove, if its very essence is to be "ecologically rational," not general. One way out might be an ex ante contract that explicitly empowers the doctor to use those heuristics that are approved by a formalized doctors' community.

The second obstacle stems from the fact that the courts are part of government. What they do must appear legitimate. More traditional forms of rationality easily produce what political scientists call "output legitimacy." The term refers to a form of legitimacy that does not rely on acts of empowerment by the electorate, but on the fact that government ostensibly improves welfare, which is why court decisions must have reasons. That heuristics may do better than formal rationality is a highly counter-intuitive finding, and it is provocative in that it destroys the illusion of full rationality. If the courts accept decisions based on heuristics, they have to confront the parties with the fact that they live in an insecure world. Put more precisely: heuristics will only have a fair chance in court if people learn that they can trust heuristics at least as much as they now believe they can trust formal rationality.

(2) Heuristics are not the sole form of reasoning to overcome the limitations of formal rationality. Heuristics should rather be tested against the other forms of applied rationality, not against utopian scientific models. Some of the legal experiences with applied rationality may be worth considering for psychologists. Applied rationality is not for Robinson Crusoe, but for a world previously shaped by institutions, be they formal legal rules, social norms, personal ties, cultural and historical embeddedness, or common beliefs. The stronger institutions narrow the situation down, the longer the illusion of formal rationality can be maintained at relatively low cost.

The way a lawyer is trained to find his decisions is not formal rationality, but also not heuristics. One might call it holistic decision making. It relies strongly on common sense and judgment (*Urteilsvermoegen*). The lawyer starts by relatively formal and general rules. But if his reading of the facts before him makes him doubt the wisdom of the rule, a rich toolbox allows him to take tailor-made steps aside. The further he leaves the general rule behind, the stronger his reasons must be. And the single courtroom is part of the legal system. The single unusual case can therefore set a learning process for the legal system into motion.

The picture becomes even richer if we look at the social function of legal rules. In the great majority of cases, legal rules are simply obeyed. Often the parties do not even know that there is a formal rule. The rule is mirrored in social norms or simple routines. But the formal legal system is always present in the background. Each party has at any time the right to trigger its application. One can interpret this as different levels of rationalisation, depending on the controversial character of the case, or the resources a party wants to invest into its settlement.

How good are fast and frugal inference heuristics in case of limited knowledge?

Edgar Erdfelder and Martin Brandt

Department of Psychology, University of Bonn, D-53117 Bonn, Germany. **{erdfelder; brandt}@uni-bonn.de www.psychologie.uni-bonn.de/~erdfel_e**

Abstract: Gigerenzer and his collaborators have shown that the Take the Best heuristic (TTB) approximates optimal decision behavior for many inference problems. We studied the effect of incomplete cue knowledge on the quality of this approximation. Bayesian algorithms clearly outperformed TTB in case of partial cue knowledge, especially when the validity of the recognition cue is assumed to be low.

Gigerenzer et al.'s *Simple heuristics that make us smart* provides a powerful demonstration of Simon's (1956a) satisficing principle: Human inference mechanisms can be both simple *and* accurate. A class of reasoning heuristics characterized by speed and frugality is shown to perform very efficiently in a variety of areas and across a wide range of inference problems. The most prominent example is the Take the Best (TTB) rule of decision making under uncertainty, discussed in more than half of the chapters of the book. TTB uses the principle of one-reason decision making in an apparently simplistic step-by-step manner. In fact, had we not read the pioneering article by Gigerenzer and Goldstein (1996a) prior to thinking about the effectiveness of TTB, we would most

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

likely have sided with those believing that TTB is easily outperformed by alternative inference rules.

What struck us even more than the good performance of TTB compared to more costly decision heuristics are the results of Martignon and Laskey's (Ch. 8) comparisons to Bayesian algorithms. The computationally intensive "profile memorization method" (PMM) defines the optimum that can be achieved for cue-based inferences under uncertainty. Let a and b denote any two of n objects. For each object, m binary cue values C_1, C_2, . . . , C_m are given, and the task is to predict which of the two objects scores higher on a criterion variable X. Then PMM predicts $X_a > X_b$ if

$$L_{\mathrm{PMM}}(a,b) := \frac{P\left(X_a > X_b \mid \left(C_1(a),C_2(a),\ldots,C_m(a)\right),\left(C_1(b),C_2(b),\ldots,C_m(b)\right)\right)}{P\left(X_a \le X_b \mid \left(C_1(a),C_2(a),\ldots,C_m(a)\right),\left(C_1(b),C_2(b),\ldots,C_m(b)\right)\right)} > 1. \quad (1)$$

Thus, if there are $d \le n$ different cue profiles, then $d\cdot(d-1)/2$ conditional probabilities need to be computed (or estimated from a sample). In contrast, TTB does not even need the m cue validities to determine the sequence of steps in the decision tree; their rank order is sufficient.

A Bayesian algorithm more similar to TTB in terms of frugality is the "Naive Bayes" (NB) rule. NB assumes conditional independence of cue values given the criterion values. If this assumption holds, then the PMM decision rule reduces to the NB rule: Predict $X_a > X_b$ if the product

$$L_{\mathrm{NB}}(a,b) := \prod_{j=1}^{m} \frac{P\left(X_a > X_b \mid C_j(a),C_j(b)\right)}{P\left(X_a \le X_b \mid C_j(a),C_j(b)\right)} > 1. \quad (2)$$

If the independence assumption is violated then, in general, $L_{\mathrm{NB}}(a, b)$ differs from $L_{\mathrm{PMM}}(a, b)$ and NB is just an approximation to PMM that has the advantage of requiring less parameters for the prediction task.

Despite its frugality, TTB comes close to the PMM benchmarks in almost all inference tasks summarized in Table 8-1 (p. 182), and it performs as well as and sometimes even better than the NB rule. For example, for the city population task previously analyzed by Gigerenzer and Goldstein (1996a), the percentages of correct decisions for TTB, NB, and PMM are 74.2%, 74.0%, and 80.1%, respectively.

We were interested in the boundary conditions that must be met to guarantee that TTB approximates optimal decision behavior so closely. This is a topic in some chapters of the book, and the message of the authors is clear: It is certainly possible to conceive inference problems in which TTB compares unfavorably with PMM, NB, and alternative decision heuristics. For example, TTB runs into difficulties if a large number of cues (m) is used to rank a small number of objects (n) and if the cue validities are approximately equal (see Chs. 5 and 6). However, inference problems characterized by these attributes appear to be "artificial" in the sense that they do not conform to the structure of inference problems that are typical of real-world settings.

Is it permissible to argue that the more an inference problem resembles real-world problems, the more TTB approximates optimal decision behavior? To answer this question, we ran another competition between TTB and Bayesian algorithms, this time considering the more realistic case of limited knowledge. A simple way of incorporating incomplete knowledge into Bayesian algorithms is to compute expected probabilities whenever some of the cue values are missing. For example, if we need to compute $P(X_a > X_b|(C_1(a) = ?, C_2(a) = 0), (C_1(b) = 0, C_2(b) = 1)$, where "1" means a cue is present, "0" it is absent, and "?" cue value is unknown, then we use the conditional expected value

Table 1 (Erdfelder & Brandt). *Percent correct inferences of the Take-the-Best heuristic and two Bayesian decision algorithms (PMM Bayes, Naive Bayes) for the city population inference task given different degrees of cue knowledge*

% cues known	Take the Best	Naive Bayes	PMM Bayes
0%	50.0%	50.0%	50.0%
10%	55.4%	58.7%	58.9%
20%	59.4%	63.1%	63.7%
50%	66.8%	69.5%	71.9%
75%	70.9%	72.2%	76.6%
100%	74.2%	74.0%	80.1%

$$\begin{aligned} &P(X_a > X_b|(C_1(a) = 1, C_2(a) = 0), (C_1(b) = 0, C_2(b) = 1))\\ &\cdot P(C_1(a) = 1|(C_2(a) = 0), (C_1(b) = 0, C_2(b) = 1))\\ &+ P(X_a > X_b|(C_1(a) = 0, C_2(a) = 0), (C_1(b) = 0, C_2(b) = 1))\\ &\cdot P(C_1(a) = 0|(C_2(a) = 0), (C_1(b) = 0, C_2(b) = 1)). \end{aligned} \quad (3)$$

We employed this principle in a computer simulation of the city population inference task. In each simulation run, 90, 80, 50, and 25% of the cue values were randomly deleted from the profiles. Table 1 summarizes the results for the case that all $n = 83$ German cities are recognized. As expected, the TTB scores match the results reported by Gigerenzer and Goldstein (1996a, p. 656). Somewhat unexpectedly, however, TTB performs best in the unrealistic case of perfect cue knowledge. If cue knowledge is not perfect, NB outperforms TTB quite clearly. Similarly, the relative performance of TTB with respect to PMM is slightly worse if 75% rather than 100% of the cue values are known.

Figure 1 illustrates how these results generalize to more realistic scenarios where not all of the cities are recognized and only 50% of the cue values are known. Obviously, the disadvantage of TTB with respect to the Bayesian algorithms depends heavily on the validity of the recognition cue which is always evaluated first in TTB. If the recognition cue validity is .80, then TTB is almost as efficient as Bayesian procedures (Fig. 1a). However, if the recognition cue is not correlated with the criterion – a realistic assumption for several inference problems discussed in the book – then TTB performs worse the smaller the number of cities recognized (Fig. 1b).

We conclude that TTB has no built-in mechanism that makes it suitable for all types of real-world inference problems. Rather, TTB and other fast and frugal heuristics need to be complemented by meta-mechanisms that determine the choice among the heuristics and their variants. Gigerenzer et al. (1999) hold a similar position (pp. 364–65). We are thus looking forward to future research that helps surmount the remaining problems.

ACKNOWLEDGMENT

The preparation of this work was supported by grants ER 224/1-1 and ER 224/1-2 of the Deutsche Forschungsgemeinschaft. Correspondence should be addressed to Edgar Erdfelder, Psychologisches Institut, Universität Bonn, Römerstraße 164, D-53117 Bonn, Germany. erdfelder@uni-bonn.de.

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

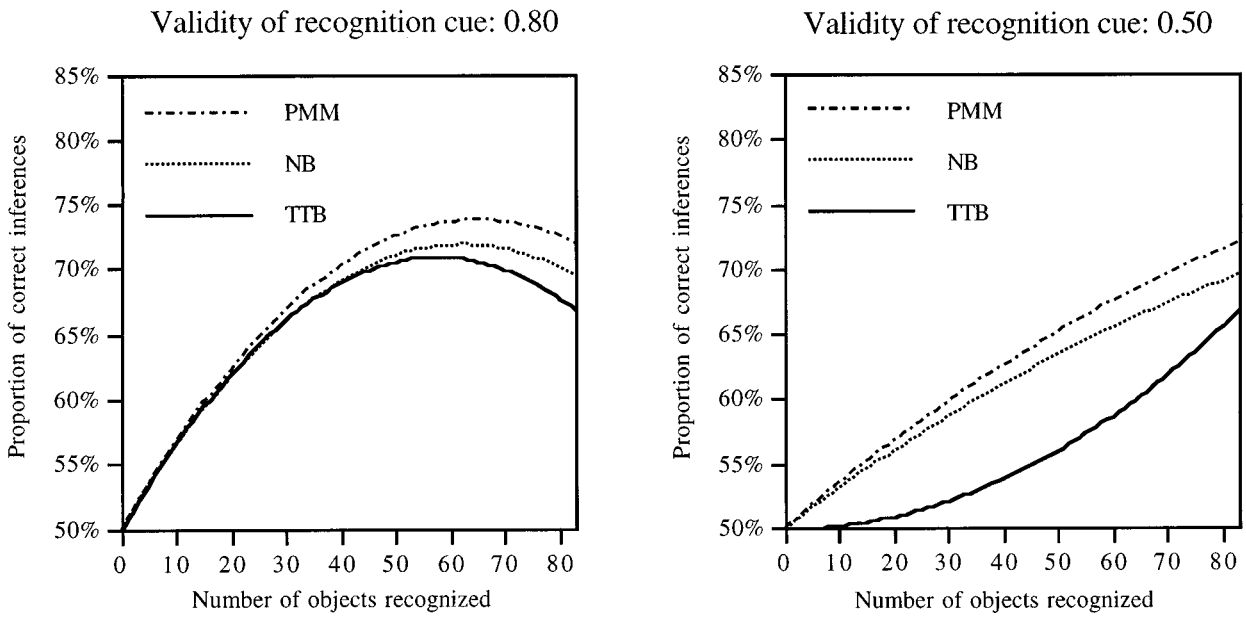


Figure 1 (Erdfelder & Brandt). Percent correct inferences of the Take the Best heuristic and two Bayesian decision algorithms (PMM Bayes, Naive Bayes) for the city population inference task with 50% cue knowledge. Performance is plotted as a function of the number of cities recognized. Figure 1a (left) refers to a recognition cue validity of .80, Figure 1b (right) to a recognition cue validity of .50.

Simple heuristics: From one infinite regress to another?

Aidan Feeney

Department of Psychology, University of Durham, Science Laboratories, Durham DH1 3LE, England. **aidan.feeney@durham.ac.uk**
http://psynts.dur.ac.uk/staff/feeney.html

Abstract: Gigerenzer, Todd, and the ABC Research Group argue that optimisation under constraints leads to an infinite regress due to decisions about how much information to consider when deciding. In certain cases, however, their fast and frugal heuristics lead instead to an endless series of decisions about how best to decide.

Gigerenzer, Todd, and the ABC Research Group's alternative approach to decision making is based upon a dissatisfaction with the notions of unbounded rationality and rationality as optimisation under constraints. They claim that whilst the notion of unbounded rationality is not informed by psychological considerations, optimisation under constraints seeks to respect the limitations of information processors, but in a manner which leads to an infinite regress. Specifically, when an information processor is making a decision, the optimisation under constraints approach requires him/her to further decide whether the benefits of considering more information relevant to that decision are outweighed by the costs of so doing. As such a cost/benefit analysis must be performed every time a decision is made, the decision making process becomes endless.

Gigerenzer et al.'s alternative is to specify a range of fast and frugal heuristics which, they claim, underlie decision making. Their stated goals (Gigerenzer et al. 1999, p. 362) are as follows:

1. To see how good fast and frugal heuristics are when compared to decision mechanisms adhering to traditional notions of rationality.
2. To investigate the conditions under which fast and frugal heuristics work in the real environment.
3. To demonstrate that people and animals use these fast and frugal heuristics.

As the authors admit (p. 362), they are undoubtedly much closer to achieving their first two aims than they are to achieving the third. Before their third objective can be reached however, Gigerenzer et al. must attempt to answer a further question: How do decision makers decide how to decide?

Applying the fast and frugal analysis seems least objectionable in cases where the algorithm used to make a decision is predetermined (by evolution or by a programmer, etc.). In principle, fast and frugal heuristics seem nonproblematic for phenomena such as a parent bird's decisions about provisioning or some of the relatively automatic processes involved in, for example, memory retrieval (all that remains is a demonstration that animals and humans do employ them). Their application is nonproblematic in these cases because, given certain environmental conditions, there is likely to be only one way in which, for example, an individual bird will feed its young. The relevant heuristic is pre-programmed so, as there exists only one decision method for that decision under those conditions, there is no confusion about which method to apply.

Now consider the example used by Rieskamp and Hoffrage (Ch. 7, p. 141). They ask the reader to consider Mr. K., who is trying to figure out how his friends (Mr. Speed and Mr. Slow) make investment decisions. Here decision methods requiring more processing resources than do the fastest and most frugal of the fast and frugal heuristics are treated as alternative strategies that people are more likely to use in the absence of time pressure. This sounds suspiciously like the kind of situation that could, potentially, lead to an infinite regress. The decision maker is free to decide between a range of decision-making strategies, leaving him/her to make decisions about how to make decisions ad infinitum.

Rieskamp and Hoffrage's only attempt to explain why the theory does not predict an infinite regress is the following: "Based on an individual's prior experience of decision making, a particular situation could prompt her or him to use a particular decision strategy" (p. 147). This explanation seems inadequate. For one thing, part of our advantage over other species is our ability to make decisions in situations for which prior experience has not prepared us. For example, Wason's abstract indicative selection has recently been characterised as a decision-making task rather than as a problem best thought of as requiring deductive inference (Evans & Over 1996; Oaksford & Chater 1994). It is certainly a task requiring decisions (which card or cards should be selected in order to test the experimental rule) for which the naïve participant's prior experience is unlikely to have prepared him. Given

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

such a lack of preparedness one might expect participants, when selecting cards, to rely wholly on the inappropriate application of heuristics successful in superficially similar domains. Indeed, there is substantial evidence that this may be the case (see Evans et al. 1993 for a review). However, more recent evidence (Feeney & Handley, in press) suggests that, at least under certain experimental conditions, participants will reason deductively on the abstract selection task. It is not clear how prior experience might explain such differential use of heuristic and deductive strategies on Wason's task.

Rieskamp and Hoffrage contrast their "prior experience" solution with the view that limited time and limited knowledge constraints are supplementary criteria evaluated in the course of making a decision. The implication seems to be that evaluating these supplementary criteria whilst making a decision leads to an infinite regress, whereas an appeal to prior experience does not. As we have seen, prior experience cannot be relied upon to prepare people for certain situations, such as attempting Wason's selection task, where decisions are called for. Yet it has been shown that under different experimental conditions, people may use different strategies to make these decisions. Of course, this is not to claim that Rieskamp and Hoffrage are incorrect in claiming that very often, prior experience will guide the strategy we select to make a decision.

Unfortunately, Gigerenzer et al.'s analysis may lead to an infinite regress when a decision maker faces an unfamiliar decision situation and is free to adopt one of a number of decision strategies or decision heuristics. Given that Gigerenzer et al.'s central complaint about the optimisation under constraints view of rationality is that it leads to an infinite regress, it is ironic that unless their own account can be supplemented with metacognitive principles it must be adjudged merely to have replaced one infinite regress with another.

ACKNOWLEDGMENT

The author would like to thank David Over for helpful comments on this commentary.

Against an uncritical sense of adaptiveness

Steve Fuller

Department of Sociology, University of Warwick, Coventry, CV4 7AL, United Kingdom. **s.w.fuller@warwick.ac.uk**

Abstract: The "adaptive toolbox" model of the mind is much too uncritical, even as a model of bounded rationality. There is no place for a "meta-rationality" that questions the shape of the decision-making environments themselves. Thus, using the ABC Group's "fast and frugal heuristics," one could justify all sorts of conformist behavior as rational. Telling in this regard is their appeal to the philosophical distinction between coherence and correspondence theories of truth.

Despite its popular evolutionary resonances, the word "adaptive" normally makes rationality theorists reach for their wallets, since adaptive conceptions of rationality usually deliver much less than they promise. Here I do not mean that they lack the rigor, precision, and systematicity of the standard conceptions. This is already granted by Gigerenzer, Todd, and the ABC Research Group (1999) as the price the theorist must pay for a model of rationality that explains how subjects achieve their goals under the time and resource constraints of realistic decision-making environments. I am also happy to grant them this point (cf. Fuller 1985). My concern is that the "adaptive toolbox" lacks any provision for meta-rationality, or metacognition more generally, other than as an iterated version of the first-order mental processes that the ABC Group have identified.

Symptomatic of the problem, as well as a sense of the stakes, is the discussion of the trade-off between what Gigerenzer et al. call "generality" and "specificity" of a particular heuristic's adaptiveness (1999, p. 18). Always writing with an eye toward the history of philosophy, they associate these two dimensions with, respectively, the coherence and correspondence theories of truth. Although such connections give their model a richness often lacking in contemporary psychology, they also invite unintended queries.

It is clear what the ABC Group mean by the distinction. To a first approximation, their "fast and frugal" heuristics "correspond" to the particular environments that a subject regularly encounters. But these environments cannot be so numerous and diverse that they create computational problems for the subjects; otherwise their adaptiveness would be undermined. In this context, "coherence" refers to the meta-level ability to economize over environments, so that some heuristics are applied in several environments whose first-order differences do not matter at a more abstract level of analysis.

Yet philosophers have come to recognize that correspondence and coherence are not theories of truth in the same sense. Whereas correspondence is meant to provide a definition of truth, coherence offers a criterion of truth. The distinction is not trivial in the present context. Part of what philosophers have tried to capture by this division of labor is that a match between word (or thought) and deed (or fact) is not sufficient as a mark of truth, since people may respond in a manner that is appropriate to their experience, but their experience may provide only limited access to a larger reality. A vivid case in point is Eichmann, the Nazi dispatcher who attempted to absolve himself of guilt for sending people to concentration camps by claiming he was simply doing the best he could to get the trains to their destinations. No doubt he operated with fast and frugal heuristics, but presumably one would say something was missing at the meta-level.

The point of regarding coherence as primarily a criterion, and not a definition, of truth is that it forces people to consider whether they have adopted the right standpoint from which to make a decision. This Eichmann did not do. In terms the ABC Group might understand: Have subjects been allowed to alter their decision-making environments in ways that would give them a more comprehensive sense of the issues over which they must pronounce? After all, a model of bounded rationality worth its salt must account for the fact that any decision taken has consequences, not only for the task at hand but also for a variety of other environments. One's rationality, then, should be judged, at least in part, by the ability to anticipate these environments in the decision taken.

At a more conceptual level, it is clear that the ABC Group regard validation rather differently from the standard philosophical models to which they make rhetorical appeal. For philosophers, correspondence to reality is the ultimate goal of any cognitive activity, but coherence with a wide range of experience provides intermittent short-term checks on the pursuit of this goal. By regarding coherence-correspondence in such means-ends terms, philosophers aim to delay the kind of locally adaptive responses that allowed Ptolemaic astronomy to flourish without serious questioning for 1,500 years. Philosophers assume that if a sufficiently broad range of decision-making environments are considered together, coherence will not be immediately forthcoming; rather, a reorientation to reality will be needed to accord the divergent experiences associated with these environments the epistemic value they deserve.

All of this appears alien to the ABC Group's approach. If anything, they seem to regard coherence as simply facilitating correspondence to environments that subjects treat as given. Where is the space for deliberation over alternative goals against which one must trade off in the decision-making environments in which subjects find themselves? Where is the space for subjects to resist the stereotyped decision-making environments that have often led to the victimization of minority groups and, more generally, to an illusory sense that repeated media exposure about a kind of situation places one in an informed state about it? It is a credit to his honesty (if nothing else), that Alfred Schutz (1964) bit the bullet and argued that people should discount media exposure in favor

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

of personal cumulative experience. But this also led him to distrust public opinion polls and mass democratic votes as reliable indicators of normatively acceptable social policies.

For all their talk of "social" and "ecological" rationality, the ABC Group are conspicuously silent on the normative implications of their research. It would be easy to epitomize their findings as implying that people with a conventional awareness of the social environments in which they act should trust their gut feelings when making decisions. The concept of "adaptive preference formation," popularized by Leon Festinger (1957), seems to ring no alarms for them. Consequently, they do not consider how prior successes in one's own (or others') decision-making environments might reinforce certain responses that, in a more reflective setting, might come to be doubted and perhaps even reversed. Certainly, their chapter on the latently positive virtues of "hindsight bias" do not inspire confidence (Gigerenzer et al. 1999, Ch. 9).

Is less knowledge better than more?

Alvin I. Goldman

Department of Philosophy, University of Arizona, Tucson, AZ 85721-0027.
goldman@u.arizona.edu w3.arizona.edu/~phil/faculty/goldman.html

Abstract: When a distinction is drawn between "total" knowledge and "problem-specific" knowledge, it is seen that successful users of the recognition heuristic have more problem-specific knowledge than people unable to exploit this heuristic. So it is not ignorance that makes them smart, but knowledge.

Gigerenzer et al. (1999) present a fascinating case for the accuracy of fast and frugal heuristics. In Chapter 2, "The recognition heuristic: How ignorance makes us smart," Goldstein and Gigerenzer contend that (in certain cases) the recognition heuristic (RH) enables comparatively ignorant people to get correct answers to questions more often than knowledgeable people. This finding, they claim, is paradoxical: "The recognition heuristic can thus lead to a paradoxical situation where those who know more exhibit lower inferential accuracy than those who know less" (Gigerenzer et al. 1999: pp. 45–46). Does RH really imply that comparatively ignorant people fare better than comparatively knowledgeable ones in certain cases? Yes and no; it depends on how one conceptualizes comparative knowledgeability.

In a Goldstein-Gigerenzer thought experiment (later substantiated by experimental work), the Scottish MacAlister brothers take a quiz composed of two-alternative questions about the population sizes of the 50 largest German cities. The youngest brother recognizes none of these cities, the middle brother recognizes 25 of them, and the eldest brother recognizes all 50. Because the middle brother recognizes some cities but not others, he can sometimes apply RH, where RH consists in the following rule: "If one of two objects is recognized and the other is not, then infer that the recognized object has the higher value" (p. 41). In the present case the "higher value" is the larger population size. So when the middle brother is presented with a city that he recognizes and a city that he does not, he can apply RH and infer that the recognized city has a larger population than the unrecognized one. By assumption, he will be right on 80% of these choices. The older MacAlister brother has no opportunity to use RH, because he recognizes all 50 of the cities. Goldstein and Gigerenzer assume that when a brother recognizes both cities in a pair, there is a 60% chance of making the correct choice. Thus, it turns out that the middle brother scores better on the quiz than the older brother, despite the fact that, as Goldstein and Gigerenzer put it, the older brother knows more.

Is this a good way to describe the comparative knowledge states of the older and middle brothers? Let us distinguish two bodies of knowledge: *total* knowledge and *problem-specific* knowledge. Plausibly, the older brother has more total knowledge about German cities than the middle brother; but does he have more problem-specific knowledge? On the contrary, I suggest, the middle brother has more problem-specific knowledge; so there is nothing "paradoxical" about the fact that the middle brother performs better on the quiz. It is not unambiguously true of the middle brother, as Goldstein and Gigerenzer claim, that ignorance "makes him smart." Yes, he is comparatively ignorant in total knowledge of German cities, but he is not comparatively ignorant on the crucial dimension of problem-specific knowledge.

Consider an analogous case. Mr. Savvy knows a lot about politics. Concerning the current Senatorial campaign, he knows the voting records of each candidate, he knows who has contributed to their campaigns, and so forth. Mr. Kinny has much less knowledge on such matters. So Savvy certainly has more total political knowledge than Kinny. It does not follow, however, that Savvy has more problem-specific knowledge on every question concerning the Senatorial race. Even such a crucial question as "Which candidate would do a better job vis-à-vis your interests?" might be one on which Kinny has superior problem-specific knowledge than Savvy. Suppose, as Kinny's name suggests, that he has a well-informed cousin who always knows which of several candidates is best suited to his own interests, and the cousin also has the same interests as Kinny himself. So as long as Kinny knows which candidate is favored by his cousin, he will select the correct answer to the question, "Which candidate would do a better job vis-à-vis my (Kinny's) interests?" (In Goldman 1999, Ch. 10, I call this sort of question the "core voter question.") In recent years political scientists have pointed out that there are shortcuts that voters can use to make crucial political choices, even when their general political knowledge is thin. Kinny's holding the (true) opinion that his cousin knows who is best is such a shortcut. Indeed, it is an invaluable piece of knowledge for the task at hand, one that Savvy may well lack.

Return now to the German cities problem. When the middle brother gets a test item in which he recognizes exactly one of the two named cities, he acquires the information that this city is more recognizable to him than the other. As it happens, this is an extremely reliable or diagnostic cue of the relative population sizes of the cities. If, as Goldstein and Gigerenzer assume, he also believes that a recognized city has a larger population than an unrecognized one, then these two items of belief constitute very substantial problem-specific knowledge. The older brother has no comparable amount of problem-specific knowledge. From the point of view of problem-specific knowledge, then, the brothers are misdescribed as a case in which ignorance makes the middle brother smart. What he knows makes him smart, not what he is ignorant of. More precisely, his ignorance of certain things – the names of certain German cities – creates a highly diagnostic fact that he knows, namely, that one city is more recognizable (for him) than the other. No comparably useful fact is known to his older brother. So although the older brother has more total knowledge, he has less problem-specific knowledge.

I do not take issue with any substantive claims or findings of Goldstein and Gigerenzer, only with the spin they put on their findings. RH can be exploited only when a person knows the true direction of correlation between recognition and the criterion. In the cities case, he must know (truly believe) that recognition correlates with greater population. When this kind of knowledge is lacking, RH cannot be exploited. Suppose he is asked which of two German cities has greater average annual rainfall. Even if there is a (positive or negative) correlation between recognition and rainfall, he won't be able to exploit this correlation if he lacks an accurate belief about it. Goldstein and Gigerenzer do not neglect this issue. They cast interesting light on the ways someone can estimate the correlation between recognition and the criterion (1999: pp. 41–43). But they fail to emphasize that only with a (tolerably) accurate estimate of this sort will applications of RH succeed. For this important reason, knowledge rather than ignorance is responsible for inferential accuracy. There is no need to rush away from the Goldstein-Gigerenzer chapter to advise our chil-

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

dren or our students that the new way to get smart is to stop reading books or drop out of school. Knowledge of relevant premises, not ignorance, is still the best way to acquire further knowledge through inference.

Heuristics in technoscientific thinking

Michael E. Gorman

School of Engineering and Applied Science, University of Virginia, Charlottesville, VA 22903-2442 **meg3c@virginia.edu**
repo-nt.tcc.virginia.edu

Abstract: This review of Gigerenzer, Todd, and the ABC Research Group's *Simple heuristics that make us smart* focuses on the role of heuristics in discovery, invention, and hypothesis-testing and concludes with a comment on the role of heuristics in population growth.

In this commentary, I will explore the role of heuristics in technoscientific thinking, suggesting future research.

Scientific discovery. Herbert Simon and a group of colleagues developed computational simulations of scientific discovery that relied on a hierarchy of heuristics, from very general or weak ones to domain-specific ones. The simplest of the discovery programs, BACON.1, simulated the discovery of Kepler's third law by looking for relationships between two columns of data, one corresponding to a planet's distance from the sun, the other its orbital period (Bradshaw 1983). The program's three heuristics:

1. If two terms increase together, compute ratio.
2. If one term increases as another decreases, compute product.
3. If one term is a constant, stop.

Given the data nicely arranged in columns, these heuristics will derive Kepler's third law. Simon and his colleagues termed this an instance of data-driven discovery. Gigerenzer et al. (1999) argue that heuristics save human beings from having to have complex representations of problem domains. All of Kepler's struggles to find a way to represent the planetary orbits might be unnecessary – all he had to do was apply a few simple heuristics.

But the importance of problem representation is cueing the right heuristics. No one but Kepler would have looked for this kind of relationship. He had earlier developed a mental model of the solar system in which the five Pythagorean perfect solids fit into the orbital spaces between the six planets. This pattern did not fit the data Kepler obtained from Tycho Brahe, but Kepler knew there had to be a fundamental geometrical harmony that linked the period and revolution of the planets around the sun. A heuristics-based approach can help us understand how Kepler might have discovered the pattern, but we also need to understand how he represented the problem in a way that made the heuristics useful.

Kulkarni and Simon (1998) did a computational simulation of the discovery of the Ornithine cycle, involving multiple levels and types of heuristics. To establish ecological validity, they used the fine-grained analysis of the historian Larry Holmes. Unlike Kepler, Krebs did not have to come up with a unique problem representation; the problem of urea synthesis was a key reverse-salient in chemistry at the time.

Gigerenzer et al. need to pay more attention to the way in which problems and solutions come from other people. An alternative to recognizing a solution is recognizing who is most likely to know the solution. Could "Take the Best" apply to a heuristic search for expertise?

For Krebs, the key was a special tissue-slicing technique he learned from his mentor, Otto Warburg. This hands-on skill was Krebs's "secret weapon." These hands-on skills are at least partly tacit, because they have to be learned through careful apprenticeship and practice, not simply through verbal explanation. To what extent are heuristics tacit as well?

In his experimental notebook on February 21, 1876, Alexander Graham Bell announced that he was going to follow the analogy of nature and build a telephone based on the human ear. "Follow the analogy of nature" is a common heuristic used by inventors (Gorman 1998), but using it depends on having a good mental model of the source of the analogy. Bell had built devices that gave him hands-on tacit knowledge of the workings of the human ear, so his use of the nature heuristic was his "secret weapon."

To what extent are heuristics tacit? Is there value in making them explicit, as Bell tried to do?

Confirmation bias. In Chapter 11, p. 235, Berretty, Todd, and Martignon discuss the categorization by elimination (CBE) heuristic. In 1960, Peter Wason discussed how participants failed to eliminate hypotheses on the 2–4–6 task. Wason's finding sparked a literature on confirmation bias (Gorman 1992), the gist of which is that people do not ordinarily eliminate hypotheses when confronted with negative evidence. [See Stanovich & West: "Individual Differences in Reasoning" *BBS* 23(5) 2000.]

Klayman and Ha (1987) tried to separate positive and negative test heuristics from confirmation. In the case of Bell, for example, most of his experiments were directed towards obtaining a positive signal, but failure to obtain the hoped-for result usually meant a problem with the experimental apparatus, not a disconfirmation (Gorman 1995). Falsification is useful in a late stage of problem-solving when the possibility of error is low. Would the same be true for CBE on a complex, multi-dimensional classification where evidence had to be gathered over time? Could CBE be used to direct the search for additional evidence?

Population growth. In Chapter 14, p. 309, Davis and Todd discuss heuristics for parental investment. In humans, the heuristic with the greatest global consequences may be the one that determines how many children to have. The greater the risk of losing a child, the more children humans want. Instead of investing heavily in one child, this is a "spread the risk" heuristic; in large parts of the world, people feel they have little control over disease, war, famine, and other factors that increase the risk of child mortality. In more affluent societies with low child mortality and educated, empowered women, parents invest in fewer offspring (Sen 1994). The role of heuristics in population growth is a promising area for future research. [See also Vining: "Social versus Reproductive Success" *BBS* 9(1) 1986."]

Evolution, learning, games, and simple heuristics

Peter Hammerstein

Innovationskolleg Theoretische Biologie, Humboldt-Universität zu Berlin, 12247 Berlin, Germany **p.hammerstein@itb.biologie.hu-berlin.de**

Abstract: Humans are incapable of acting as utility maximisers. In contrast, the evolutionary process had considerable time and computational power to select optimal heuristics from a set of alternatives. To view evolution as the optimising agent has revolutionised game theory. Gigerenzer and his co-workers can help us understand the circumstances under which evolution and learning achieve optimisation and Nash equilibria.

Gigerenzer and his co-workers present an impressive study (Gigerenzer et al. 1999) of how human decision making is guided by fast and simple procedures. They also theorise about the adaptive nature of these procedures and demonstrate that simple heuristics can even be superior to statistical methods, such as multiple regression analysis, if only one makes appropriate assumptions about the environments in which decisions are made. This seems to indicate that in the environments under consideration the "heuristic that makes us smart" is better than many alternatives. The authors thus implicitly invoke optimality considerations, at least for a constrained set of alternative procedures. At first sight

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

this may appear inconsistent with some of their statements in which they suggest that we should dismiss optimality considerations in the study of decision making. My comment aims at showing how this inconsistency can be resolved, and why optimality still has a place in decision theory even if one assumes like the authors that humans do not possess sufficient mental tools to actually solve optimisation problems in their daily lives.

I agree with them that considering humans as utility maximisers is like entering the world of science fiction. We know too well that most humans cannot deal properly with probabilities and that the task of computing conditional expectations is too difficult even for professional mathematicians when they have to decide quickly. Businessmen are perhaps making more money than mathematicians but again this is not due to their talents in optimisation. For example, in many auctions it can be observed that the bidders who actually get the object that is put up for auction are those who overestimate the value of this object and pay more for it than it is worth to them. This phenomenon is known as the "winner's curse." Rational decision makers would anticipate the winner's curse but experiments dealing with this phenomenon show that humans easily step into the pitfall set up by the auctioneer.

In contrast, suppose that the process of biological evolution would select bidding strategies for auctions that take place generation after generation in a population in which strategies are genetically inherited. The maladaptive winner's curse phenomenon would then disappear! In other words, evolution and not the human mind has an optimising tendency. Our fast and frugal heuristics are probably better than many alternative procedures because they have been subject to Darwinian evolution.

The history of game theory can be used to further illustrate this point. John Nash invented the central game-theoretic solution concept which is now named after him. In a Nash equilibrium, every player uses a strategy that maximises this player's payoff, given that all the others use their equilibrium strategies. A fast but not frugal way of justifying this concept is to make the assumption that "every player knows that every player knows that everybody in the game is rational." Of course, this is an absurd assumption. Why then has it been used for decades? A cynical answer to this question would be that this assumption created an ideal "playground" and many positions for mathematicians, enabling them to do work of apparently great importance without ever having to bother much about empirical findings. Another answer would be that people in the field did not really trust rationality reasoning and had vague ideas about alternative foundations of the Nash equilibrium in mind.

Modern evolutionary game theory demonstrates nicely that such alternative foundations are possible as long as one invokes evolution or learning as those processes that actually make the choice of strategy. The first major step in this direction was made by Maynard Smith and Price and in the biological literature following their work. More recently, economists, such as Weibull, have also adopted evolutionary game theory. Even Nash himself had interpretations along the lines of evolutionary game theory in mind. In his Ph.D. thesis he already introduced the "mass action interpretation" of the Nash equilibrium. This passage of his thesis did not show up in any of Nash's publications. Ironically, theorems about mathematical existence of Nash equilibria seemed to be more acceptable for publication than conceptual arguments about why these equilibria are of any importance. Had Nash not received the Nobel award, forcing some people to think about laudatory speeches, we would still not know about his mass action interpretation.

Seeing this partially unfortunate history of conceptual thinking in game theory it appears to me that the book by Gigerenzer et al. is an important step in exactly the direction of research that has been ignored for too long. However, one should not throw out the baby with the bath water. Evolution certainly has an optimising tendency if one makes reasonable assumptions about the space of alternative mental mechanisms from which natural selection "picks a winner." This is where evolutionary biology and psychology mutually depend on one another. Biology provides the background about the optimising process and psychology can help to understand what the space of strategies is that evolution acts upon. Not only evolution but also learning can optimise and produce Nash equilibria but it does so less reliably. For example, the literature on learning direction theory, which is partially cited in the book, shows that unlike evolution, learning does not cause the winner's curse phenomenon to disappear in experiments where people have a chance to get experience at auctions.

Why do people not learn the optimal behaviour in this case? Probably because fast and too simple mental procedures set the conditions for how to change bidding tendencies according to experience from previous rounds of the same experiment. Once these tendencies are inappropriately specified, the resulting learning procedure fails to act as an optimising agent. This shows that even in the context of learning one has to think about the effects of simple heuristics. I suggest placing more emphasis on this problem in future work of the ABC group. Needless to say that it seems like a very promising task for both biologists and psychologists to elaborate on the ideas that Gigerenzer and his co-workers have put together in their seminal contribution to decision theory with empirical content.

On the descriptive validity and prescriptive utility of fast and frugal models

Clare Harries[1] and Mandeep K. Dhami[2]

[1]*Department of Psychology, University College London, London WC1E6BT, United Kingdom;* [2]*Department of Psychology, City University, London EC1V 0HB, United Kingdom.* **clare.harries@ucl.ac.uk**
www.psychol.ucl.ac.uk/clare/self.html **m.k.dhami@city.ac.uk**

Abstract: Simple heuristics and regression models make different assumptions about behaviour. Both the environment and judgment can be described as fast and frugal. We do not know whether humans are successful when being fast and frugal. We must assess both global accuracy and the costs of Type I and II errors. These may be "smart heuristics that make researchers look simple."

Are humans really fast and frugal? Should humans be fast and frugal? Human judgment may be described on a number of dimensions such as the amount of information used and how it is integrated. The choice of a model dictates how these dimensions are characterised, irrespective of the data. For example, regression models in judgment and decision-making research are linear and compensatory, and researchers have assumed that humans are, too (Cooksey 1996). The fast and frugal models proposed by Gigerenzer, Todd, and the ABC Research Group (1999) use few cues and are often noncompensatory, assuming humans are, too.

Are humans really using less information in the fast and frugal model than the regression model? People can chunk information (e.g., Simon 1979). Regression models are characterised as complex in terms of use of multiple cues, but they often contain few significant cues (on average three; Brehmer 1994). This challenges the argument that fast and frugal models are more frugal than regression models, at least in terms of the number of cues searched. Unlike standard practice (Tabachnik & Fidell 1996), in their regression analyses, Gigerenzer et al. retain nonsignificant cue weights. A fairer test would compare fast and frugal models against parsimonious regression models.

Regression models have been used to describe the relationship between judgments and the cues (the judgment system), and the relationship between outcomes and the cues (the environment system; Cooksey 1996). In both cases the underlying structure of the cues is similar (e.g., they are correlated). Fifteen years of machine learning research demonstrates that fast and frugal models can describe environments (Dutton & Conroy 1996). It is not surprising that these models should also be good at describing human

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

judgment. Both types of models are useful for fitting data per se. Chapter 7 demonstrates that fast and frugal models are good at describing human judgment, under conditions of time pressure with participants neither experienced nor familiar with the choice task. At other points in the book, results of simulations seem to be generalised with little justification to humans. However, research has shown that fast and frugal models are valid descriptions of professionals' judgment behaviour in the legal and medical domains (Dhami & Ayton 1998; Harries & Dhami 1998). So far we have learned that the environment and humans can be described as fast and frugal. We do not yet know whether humans are successful when being fast and frugal because studies either don't collect judgments or don't include outcomes.

We have doubts about the prescriptive utility of fast and frugal models. The book argues (at least implicitly) that if fast and frugal models are good at predicting the environment, then humans should (and do) use these strategies to make accurate decisions. However, in many situations global accuracy is not the first concern: the two types of errors (Type I and Type II) are differentially weighted. In medicine, for example, all tests involve a trade off between Type I and Type II errors. Researchers developing machine-learning models in this domain incorporate the costs of the two types of error (Kukar et al., 1999). In criminal justice, the opposing concepts of due process and of crime control attempt to reduce crime whilst minimising different types of error (Type I–the number of people falsely convicted and Type II–the number of guilty people acquitted, respectively). It so happens that due process expects a regression-like behaviour: all available information is searched, weighted, and integrated. The fast and frugal strategy of crime control emphasises information associated with guilt, encouraging "conveyor belt justice." Justice is synonymous with due process, so the judge should behave like a regression model. Of course, unlike medicine, the socially constructed nature of criminal justice implies that we could change our notion of justice to one that reflects fast and frugal behaviour. Before Gigerenzer et al. encourage this we recommend that they evaluate models on other criteria in addition to global accuracy. In sum, the descriptive success of simple heuristics does not, of itself, imply their prescriptive utility.

On a differing point, if we do find that fast and frugal models have prescriptive utility, their potential as cognitive aids or in cognitive feedback vastly outweighs that of regression models (which have done pretty well; Cooksey 1996). Fast and frugal models are easy to understand and to apply while regression models are difficult to use without an aid and without knowing the range of cases on which they were formed.

In short, we welcome fast and frugal models because they make us re-think the dimensions of human judgment. The danger is that they will be automatically adopted as tools to describe human judgment, like regression models have been for 50 years. All models are paramorphic not isomorphic. So, we should be wary of "smart heuristics that make researchers look simple."

Decision rules in behavioural ecology

Alasdair I. Houston

Centre for Behavioural Biology, School of Biological Sciences, University of Bristol, Woodland Road, Bristol BS8 1UG, United Kingdom.
a.i.houston@bristol.ac.uk
www.bio.bris.ac.uk/research/behavior/behavior.htm

Abstract: Gigerenzer, Todd, and the ABC Research Group give an interesting account of simple decision rules in a variety of contexts. I agree with their basic idea that animals use simple rules. In my commentary I concentrate on some aspects of their treatment of decision rules in behavioural ecology.

Like Gigerenzer et al. (1999) I believe that animals are likely to base their decisions on simple rules that are appropriate for a particular environment (see for example Houston 1987; Houston & McNamara 1984; 1999; Houston et al. 1982; McNamara & Houston 1980). I feel, however, that Gigerenzer et al. may have underemphasised the use of rules within the context of optimisation. Some of my concerns can be described in the context of patch-use.

Gigerenzer et al. (1999, Ch. 15, p. 327, Goodie et al.) discuss the work of Green (1984) concerning when to leave a patch of food. The optimal solution is given by the marginal value theorem (Charnov 1976): leave when the rate of gain on the current patch falls to the maximum possible rate of energetic gain γ^*. Gigerenzer et al. point out that this result does not tell us how an animal could know that it was time to leave a patch. The impression they give is that the animal has to know too much. McNamara and Houston (1985) address this issue. They consider a simple learning rule that starts with an estimate γ_0 for γ^*. The rule states that the animal leaves the first patch when the rate on the patch falls to γ_0. On arrival at the next patch a new estimate γ_1 of the rate γ^* is calculated on the basis of food obtained and time taken. This rate is used to decide when to leave the second patch that is encountered, and on arrival at the next patch a new estimate γ_2 is calculated. McNamara and Houston show that such a rule produces a sequence of estimates γ_0, γ_1, γ_2 . . . that tend to γ^*. In other words, the rule learns the correct optimal rate. McNamara and Houston also discuss the problems that arise when the environment changes. In the context of heuristics, it can be pointed out that the marginal value theorem can be viewed as a form of satisficing, with γ^* acting as an aspiration level. The rule proposed by McNamara and Houston is a form of satisficing in which the aspiration level is modified by experience (cf. Simon 1959).

Gigerenzer et al. conclude (1999; p. 341) that Green's work "could be, and often has been, viewed as optimizing under constraints of information gathering ability, but . . . Green's analysis explicitly does not try to optimize. . . . There are no claims that this is the best an animal could possibly do." I disagree with this suggestion that Green's (1984) work does not involve optimization. Green considers three sorts of rule: (i) fixed-time rule (ii) giving-up time rule, and (iii) assessment rule. In each case, Green finds the best rule (1984; see pp. 33–35). In case (i) Green says "the best thing to do is to stay until each patch has been exhausted." (p. 33). In case (iii) he says (p. 33) "The best rule is found by dynamic programming." In fact Green's use of dynamic programming in this context involves a recursive estimation of γ^* that provides a robust and efficient way of finding optimal behaviour (see Houston & McNamara 1999 for further discussion).

The use of optimality to investigate rules is not unusual in behavioural ecology. This can be thought of as optimization under constraints, but it does not correspond to the restrictive way in which this term is defined in Chapter 1 (p. 3; Gigerenzer & Todd). It is often desirable to gain an idea of how well a rule performs (see Ch. 13, p. 287, Todd & Miller and Ch. 14, p. 309, Davis & Todd). This may involve finding an optimal rule. It may also involve finding the optimal behaviour. Finding the optimal behaviour enables us to find the optimal performance, which provides a way in which the performance of rules can be evaluated (Houston 1987; Houston & McNamara 1984; Houston et al. 1982).

Gigerenzer et al. make a good job of the difficult task of both explaining the debate on matching versus maximising and putting operant schedules in a naturalistic setting. I agree with their general conclusion that matching is likely to be the consequence of a rule that evolved through natural selection (see Houston & McNamara 1988). I would have liked to see a clear statement that matching per se cannot be the fundamental principle. This follows from the fact that matching does not uniquely specify behaviour (e.g., Houston 1983; Houston & McNamara 1981). Gigerenzer et al. do mention melioration as being able to remove some ambiguities in matching. It does not, however, remove the fundamental non-uniqueness concerning behaviour on concurrent variable-interval schedules. Furthermore, it is not clear from Gigerenzer et al. that biased matching rather than simple matching describes

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

behaviour when an animal can choose between a variable-interval and a variable-ratio schedule (Williams 1988).

What's in the adaptive toolbox: Global heuristics or more elementary components?

Oswald Huber

Department of Psychology, University of Fribourg, CH-1700 Fribourg, Switzerland. **oswald.huber@unifr.ch**
www.unifr.ch/psycho/Allgemeine/Oswald_Huber.htm

Abstract: From the standpoint of decision research, investigating global heuristics like LEX is not fruitful, because we know already that people use partial heuristics instead. It is necessary (1) to identify partial heuristics in different tasks, and (2) to investigate rules governing their application and especially their combination. Furthermore, research is necessary into the adequate level of resolution of the elements in the toolbox.

Fast and frugal heuristics in the Gigerenzer, Todd, and the ABC Research Group (1999) book are assumed to be specialized higher order cognitive processes that are in the mind's adaptive toolbox. If we postulate elements in the toolbox, we have to make assumptions about the adequate grain of these elements. As I interpret the authors, they regard global heuristics like the Lexicographic heuristic (LEX), Weighted Pros or Elimination-by-Aspects (EBA) as elements of this toolbox, in the decision making context.

Results of decision theory contradict such an assumption. In many experiments where decision makers have to choose one of several alternatives (presented simultaneously), the following very stable result is observed (cf. e.g., Ford et al. 1989): The decision process can be divided into two main phases. In the first, people use, for example, some parts of EBA, not to make the final choice, but in order to reduce the set of alternatives quickly to a short list. When they have arrived at a short list they use, in the second phase, another partial heuristic in order to select the best alternative. For these alternatives, usually much more information is processed than for those not in the short list, and in a more alternativewise manner (transition to another attribute of the same alternative). From these results as well as from those obtained with verbal protocols, we conclude that decision makers do not apply a complete heuristic like LEX, and so on, but use a part of a heuristic or combine partial components of heuristics.

One such partial heuristic is, for example, Partial EBA, which is characterized by a sequence of the following two steps:

Step 1: Select a criterion (attribute, dimension . . .) from set *C* of criteria.
Step 2: Eliminate all alternatives from the set *A* of alternatives which do not surpass an acceptance level on the selected criterion.

Global EBA consists of repeated applications of Partial-EBA, but most often Partial-EBA is combined flexibly with other partial heuristics. If the decision maker faces several alternatives she may apply Partial-EBA once or twice to reduce the set of alternatives to a more manageable short list, and then use another partial heuristic to make the final choice.

Thus, from the point of view of decision research, the toolbox has to contain smaller units than global strategies. The experiments reported by Gigerenzer et al. (1999) do not contradict these well established results. They were not designed to study the question of global heuristics versus components, and are restricted to the special case of only two alternatives. However, even in cases with two alternatives we cannot exclude the possibility that people combine parts of heuristics (e.g., LEX and the Conjunctive heuristic).

For research on the elements of the toolbox, this has two consequences:

1. In the context of decision making and judgment, the investigation of global heuristics is not a fruitful research strategy, because we already know that people do not use them. It is rather necessary to identify partial heuristics in different tasks. A variety of partial heuristics has already been investigated in multi-attribute decision making (see, e.g., Montgomery & Svenson 1989; Payne et al. 1993), but we do not know much, for example, about nonlottery risky decision situations.

2. If we have identified partial heuristics as elements, we need a theory that explains the rules of the use and the combination of partial heuristics. Even in multi-attribute decision making we do not have such a theory, able to model the decision process as a sequence of partial heuristics in such a way that the detailed decision behavior can be predicted. The problem of a theory that explains the combination of partial components is not treated in the book.

If we decompose global heuristics into smaller components as the results of decision research make necessary, the question arises at which level of decomposition to stop, because the partial heuristics again can be decomposed, and so on, until elementary operators are reached. We have to determine which level of resolution is adequate for the elements of the toolbox.

Consider, as an example, Partial EBA as described above as a sequence of two steps. At one level of resolution we could define each of the steps 1 and 2 as a basic component of the toolbox. An advantage of this resolution is frugality, because step 1 could serve as a component in many partial or even global heuristics (e.g., Partial LEX or Weighted Pros). At another less fine-grained level, we consider the two steps together as *one* unit. In my opinion, even if Partial-EBA is on a coarser level than the level where its two components are treated as separate units, it is a more adequate basic element in the toolbox. There are two reasons: (1) In people's task representations of a decision situation, probably alternatives and sets of alternatives play a central role. Therefore, people can easily form a subgoal to reduce the set of alternatives and search for an operator to reach that subgoal. In this process, details of how this operator functions are not relevant to the decision maker. Thus, partial heuristics are, on the one hand, big enough to ease the workload on the thinking process. (2) On the other hand, partial heuristics are small enough to be combined in a flexible manner.

The question of the adequate level of resolution of the elements of the toolbox is not treated in the book. It is especially important if we assume that the elements are genetically fixed (because they have been put into the toolbox by evolution).

ACKNOWLEDGMENT
I am grateful for discussions with Siegfried Macho.

The role of mathematics in heuristic performance

Paul C. Kainen

Department of Mathematics, Georgetown University, Washington, DC 20057
kainen@math.georgetown.edu
www.math.georgetown.edu/faculty/kainen/

Abstract: A mathematical approach to heuristics is proposed, in contrast to Gigerenzer et al.'s assertion that laws of logic and probability are of little importance. Examples are given of effective heuristics in abstract settings. Other short-comings of the text are discussed, including omissions in psychophysics and cognitive science. However, the authors' ecological view is endorsed.

The book by Gigerenzer, Todd, and the ABC Research Group (1999) proposes that heuristics enjoy some sort of ecological rationality. However, the arguments do not seem entirely consistent; on pp. 18–19, heuristics are taken to be so simple that they apply to many cases, but complex enough to exploit environmental dependencies.

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

Gigerenzer et al. (1999) seem to believe that no formal structure is needed to understand the remarkable performance of heuristics. On p. vii, they assert: "The laws of logic and probability play little if any role." I believe that this is false and that, on the contrary, heuristics derive their power precisely by their utilization of mathematical theory. Indeed, heuristics can be clearly seen in an entirely abstract context. I'll give two examples that show the idea (Kainen 1997).

Consider the search for an endpoint in a large tree. This is a well-defined task and can be implemented via software. If the tree has n vertices, then the usual algorithm, which merely follows a path until it finally terminates at an endpoint, will need order-log-n steps. However, if the tree is selected "at random" from all possible trees (on the same set of labeled vertices) and if a vertex is then selected at random, then there is a chance of approximately 1/e (i.e., 0.37 . . .) that the vertex will *be* an endpoint. Guessing 40 times, one has a probability of less than $(2/3)^{40}$ of being wrong, independent of n. Using random choice as a heuristic beats the algorithm by as much as you like once n gets sufficiently large. Since many situations, for example, the solutions of games, can be cast in terms of tree search for endpoints, this mathematical result can be applied to many concrete examples.

Notice that the guessing heuristic allows parallel processing and makes no demands on sharing information. Moreover, if the data contains an extra non-tree edge, then the algorithm could be trapped in a loop but the heuristic will only be influenced slightly because each extra edge can kill off at most only two of the roughly n/e vertices which are endpoints in the tree.

Another heuristic, called First Fit Decreasing, has been analyzed for a resource-allocation problem. It turns out that using two reasonable strategies together (do the hardest first and just take the first available resource – i.e., be greedy) will get one to the essentially best possible solution.

While the book includes many examples showing how well some heuristics do in special test conditions, for example, how we infer intent from mouse movements on a computer, it does not go on to consider detailed computations from the perspective of psychophysics. For example, movement attracts human vision, guiding the direction of focal attention, as do other salient features. Heuristics may be involved because the speed of visual response and minimum requirements for information pretty much eliminate any conceivable algorithm. Another example may be the process of linearization so well known in mathematics, which has turned up in neurophysiological data for speech (Sussman et al. 1998) and also for reaching behaviors (Desmurget et al. 1996). I wish the book had gone further in these directions, which certainly seem congruent with the authors' basic slant.

Heuristic ideas can also be extended to cover issues relevant to cognitive science – G. Lakoff (1980) and H. Plotkin (1993) come to mind. These connections ought to be pointed out in future editions of this work. That physical metaphors fit so well with linguistic and psychological data may be evidence for the "unreasonable effectiveness" of mathematics even in the biological domain (Kainen 1998).

It would be nice to have an encyclopedia of basic heuristics. However, this might be a Sisyphean task. As the statistician B. Efron said, "*Good* simple ideas are our most precious intellectual commodity," but there may be no end to such good ways to organize knowledge.

The worst aspect of the book is its verbosity. Surely, the story could be told more succinctly. The authors must have used the phrase "fast and frugal heuristic" a thousand times. Homer got away with this trick, but he had a better ear. One might like to hear, rather, several reports, each of a few pages, on how using a simple heuristic has led to improvement. For example, here is one I stumbled across while working at Bell Laboratories: Distrust software (or systems) whose name is a third-order acronym.

This useful heuristic revealed a logistical horror story. A program had been written to determine the deployment of a new device which allowed multiple conversations to share the same transmission facilities. But this powerful multiplexing device was reduced to only a few percent of its actual capability by an elaborate algorithm – fossilized from birth because of the vast distance from programmers to factory-floor network environment, as the acronym indicated. Once the conceptual errors in the program were recognized, the problem was readdressed (successfully) via heuristic methods, and that would make another case history.

It would improve the next publication of this group if they could recount concrete instances like these. Although heuristic concepts can have some overarching features, the richness of the context is essential to see them *in vivo*.

In their preface, Gigerenzer and Todd refer to an interdisciplinary spirit "put[ting] everyone . . . together . . . learning the language and the skills of the others." This type of collaboration is called *transdisciplinary* by Ivan Havel, who distinguishes it precisely by the degree, depth, and longevity of the interactions. The authors have certainly made a contribution by promoting awareness of heuristic methods. If there is any area in which multiple research streams are converging, it is right here. Heuristics provide a vital tool for ongoing problems of science and industry.

Two cheers for bounded rationality

Raanan Lipshitz

Department of Psychology, University of Haifa, Haifa, Israel 31905
raanan@psy.haifa.ac.il

Abstract: Replacing logical coherence by effectiveness as criteria of rationality, Gigerenzer et al. show that simple heuristics can outperform comprehensive procedures (e.g., regression analysis) that overload human limited information processing capacity. Although their work casts long overdue doubt on the normative status of the Rational Choice Paradigm, their methodology leaves open its relevance as to how decisions are actually made.

Hastie (1991, p. 137) wondered whether the field of decision making "will ever escape the oppressive yoke of normative 'Rational' models" and guessed that the Expected Utility Model will "fade away gradually as more and more psychologically valid and computationally tractable revisions of the basic formation overlay the original" (p. 138). Gigerenzer et al.'s work shows that escaping the "Rational" paradigm requires not a gradual revision of the "original," but radical reformulation of the purpose of psychology and the nature of rationality.

Following Brunswik's dictum that psychology is the study of people in relation to their environment, Gigerenzer et al. (1999) suggest that "the function of heuristics is not to be coherent. Rather, their function is to make reasonable, adaptive inferences about the real social and physical world given limited time and knowledge" (p. 22). Replacing internal coherence by external correspondence as a criterion for rationality, Gigerenzer et al. manage to cast doubt on the normative underpinnings of the rational paradigm. Despite repeated demonstrations of its poor descriptive validity, the model continues to influence Behavioral Decision Theory both descriptively (e.g., the operationalization of decision as choice and of uncertainty as probability estimates), and prescriptively (i.e., guiding evaluation and improvement of decision quality). This resilience is due, in large part, to Savage (1954), who suggested that the SEU model can be interpreted as (1) a description of an actual human decision maker, or (2) a description of an ideal decision maker who satisfies some "rational" criteria (e.g., transitivity of preferences). Because most human decision makers cannot be regarded as ideal, neither the fact that they systematically deviate from the model nor the fact that they are incapable of the comprehensive information processing required by the model is relevant to its coherence-based normative status. Uniquely in science, the rational paradigm in the study of decision making takes lack of descriptive validity as grounds for improving

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

its "suboptimal" subjects instead of its invalid models (Cohen 1993).

Using correspondence in place of coherence, Gigerenzer et al. use rational methods (regression and Bayesian analyses) as competitors (rather than yardsticks) of simple heuristics which require a fraction of the information, time, and computational power required by the former. Their result – that under certain conditions fast and frugal heuristics win the competition – should dampen enthusiasm for unrealistic, wasteful, and resource guzzling recipes epitomized by Janis and Mann's (1977) prescription:

1. Thoroughly canvass a wide range of alternative courses of action;
2. Survey a full range of objectives and values;
3. Carefully weigh all known positive and negative consequences;
4. Search for new information intensively;
5. Reexamine all alternatives and considerations before making the decision;
6. Make detailed implementation plans, prepare for potential contingencies.

The principal contribution of Gigerenzer et al.'s work leads to its principal weakness. To rigorously refute the normative basis of the rational paradigm they use simple judgment problems and computer simulations. It is certainly reasonable to begin with simple tasks and controlled methodologies before moving on to more complex and less tractable, but the authors themselves admit that this strategy leaves the really interesting (and "messy") questions unanswered: "[I]f we are also concerned with the principles that guide natural human and animal behavior, we must . . . [ask] what is the evidence that humans or animals use specific fast and frugal heuristics?" (p. 23).

An alternative research strategy is to forego traditional standards of rigor and tackle real-world decision processes in their naturalistic settings (Chi et al. 1982; Klein et al. 1993; Rogoff & Lave, 1984). Two findings and one conclusion from this research seems particularly pertinent to the question which Gigerenzer et al. leave unanswered:

1. Recognition does play a major role in decision making, but not as a mere sense of familiarity. It is a complex (still ill-understood) process of pattern matching which focuses attention on critical cues and generates goals, expectations, and action alternatives (Klein et al. 1993).

2. Experienced decision makers do use heuristics that save cognitive effort (Cohen et al. 1996; Scribner 1984). These, however, are highly domain specific and rely on considerable amount of domain specific knowledge and cannot be reduced either to "maximize expected utility" or to "choose the best."

3. Thus, effective real-wold use of "fast and frugal" heuristics requires considerable background knowledge, for example, "Which is the best indicator?" "What is the informational structure of the environment?" Similar to the Rational paradigm, which they so effectively undermine, Gigerenzer et al. pay insufficient attention to the fact that boundedly-rational human decision makers compensate the limited capacity of their short-term memories by the considerable capacities of their long-term memories and learning.

Fast, frugal, and surprisingly accurate heuristics

R. Duncan Luce

Social Science Plaza, Departments of Cognitive Sciences and Economics, University of California, Irvine, CA 92697-5100. **rdluce@uci.edu**

Abstract: A research program is announced, and initial, exciting progress described. Many inference problems, poorly modeled by some traditional approaches, are surprisingly well handled by kinds of simple-minded Bayesian approximations. Fuller Bayesian approaches are typically more accurate but rarely are they either fast or frugal. Open issues include codifying when to use which heuristic and to give detailed evolutionary explanations.

This volume (Gigerenzer et al. 1999) is a preliminary progress report – and progress it does indeed report – of the Center for Adaptive Behavior and Cognition (ABC) of the Max Planck Institute for Human Development in Berlin. Although not identified as such, it is an edited work with 16 signed chapters that are more than usually well integrated and coherent. Presumably Gigerenzer and Todd are the editors. The theme is that evolution has led people and other creatures to infer heuristics that are, in an often repeated phrase, "fast and frugal" and surprisingly effective.

The message seems directed primarily to psychologists interested in decision theory, but in reality the audience should be far broader than that. One claim is that our standard arsenal of statistical methods – those based on the familiar additive expressions of analysis of variance and linear regression – really seem to rest on a singularly bad representation of a great deal of reality; this matters greatly. Virtually all scientific psychologists should heed this message.

Examples of inference and categorization problems repeatedly arise that exhibit the following two features. First, each of a number of observables is correlated with a measure of interest whose value is not known directly. These observables may or may not be independent. For example, suppose comparative city size is a measure of interest and the observables are such things as having a professional soccer team, being the home of a university, being a state capital, and so on. Second, to decide which of two cities is larger, suppose one simply proceeds through the observables in descending order of correlation stopping at the first observable that discriminates between the cities: that dictates the choice. Surprisingly, this "take-the-best" heuristic regularly outperforms a variety of additive measures. Because no trade-offs occur, not all of the available information needs to be examined exhaustively, and very little computation is needed. So the heuristic tends to be fast as well as accurate. This can be seen as one way to implement Herbert Simon's (1956a) call for limited search and satisficing in decision making.

Many of the inference problems posed seem to fit more closely a Bayesian perspective than they do the classical additive trade offs, and indeed when fast and frugal heuristics are compared in Chapter 8 (p. 169, Martignon & Laskey) with certain Bayesian analyses, the latter indeed do somewhat better. But the expense in time and computation is considerable and the authors believe, but hardly prove, that evolution has found it effective to trade off a bit of accuracy for very fast procedures, often a condition for survival. "Fast and frugal heuristics can have their place in everyday affairs where time is limited and knowledge is scarce, and Bayesian tools can be the choice of someone who is in no hurry and has access to a computer" (p. 186).

Indeed, a major theme of the book is just that: Calculations of optimal behavior just are not practical for much real world decision making, and people should not be assumed even to approximate such calculations. Chapter 13 (p. 287, Todd & Miller) takes up the problems typified by selecting a secretary or a mate, where each party is dealing with a somewhat unknown pool of alternatives, both must agree on a match, and one cannot return to individuals previously passed over. The optimal rule is known for a given, fixed pool size, say 100, and assuming that the candidate selected will accept the offer: Examine 37 of the applicants, the best of whom establishes an aspiration level, and then select the next person exceeding that value. The chance of actually finding the optimal person is also 37% and, on average, 74 people are interviewed. Clearly, this is a costly and far from reliable procedure. Indeed, should one's goal be the optimum? Simulations show that a criterion based on the first 14 has an 83% chance of finding someone in the top 10% of the pool.

Another part of the so-called Bayesian revolution are attempts to take into account people's preferences over alternatives and uncertainties about the environment giving rise to them. This aspect has been embodied in subjective expected utility theory and a number of variants. Although they allude to this line of work, which I suspect they do not much like, they do not approach it directly: "we have focussed on adaptive goals in terms of correspondence crite-

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

ria (e.g., accuracy, speed, and frugality) as opposed to coherence criteria (e.g., consistency, transitivity, additivity of probabilities) traditionally used to define rationality. Is any role left for coherence criteria?" (p. 364, Todd & Gigerenzer). "Models of reasoning need not forsake rationality for psychological plausibility, nor accuracy for simplicity. The mind can have it both ways" (p. 365).

A large area is delineated, and it will require many scientists – not just psychologists – and much time to explore and codify it. First, how does one classify problems and decide upon which of several possible fast and frugal heuristics to employ? They speak of having an "adaptive tool box," and the question is when to use which tool. Over three quarters of a century of effort has gone into systematizing experimental design and analysis based on standard statistics, and something comparable probably will be needed to show how to proceed more effectively in modeling inference and classification. Second, how does one decide which tool each participant in an empirical study is actually using? This is an issue because the problems encountered neither automatically nor uniquely point to a particular analysis, and so it is only reasonable to expect variation in tool use. And third, how do we test in a nontautological fashion the claim that some of these procedures have been selected for by evolution? As we all are well aware, this is a very tricky issue, and evolutionary arguments, at least in psychology, typically invite a good deal of skepticism and, perhaps, have been misused in some cases. Should these developments actually occur, this book will be seen as seminal.

Simple heuristics that make us dumb

Howard Margolis

The Irving B. Harris Graduate School of Public Policy Studies, University of Chicago, Chicago, IL 60637. **hmarg@uchicago.edu**

Abstract: The simple heuristics that may indeed usually make us smart – or at least smart enough – in contexts of individual choice will sometimes make us dumb, especially in contexts of social choice. Here each individual choice (or vote) has little impact on the overall choice, although the overall choice is compounded out of the individual choices. I use an example (risk aversion) to illustrate the point.

Gigerenzer, Todd, and the ABC Research Group (1999) argue that simple rule-of-thumb responses will often be strikingly effective relative to logically better but pragmatically hard to learn and hard to use responses. Hence, their message: "simple heuristics can make us smart." And they are right. But simple heuristics can also make us dumb. In fact, the very same simple heuristics.

As Gigerenzer et al. (1999) make clear, their simple heuristics come to be in place because experience in the world has trained us to rely on them, or sometimes even because experience of our species in evolutionary time has entrenched the responses genetically. We can expect that entrenched responses will be good enough to be favored by whatever selection process is operating. The net effect (the benefits when it works, net of the costs when it doesn't) must be positive enough, relative to feasible alternatives, to make sense of how the response comes to be entrenched.

But this argument only works for contexts close to those that favored the response. Even in just those contexts, the response will be subject to false alarms. It may be prompted when it is irrelevant or ineffective, and so does no good while perhaps pre-empting a less easily prompted response that would actually have worked. But in contexts remote from those that account for why a response becomes entrenched, these ordinarily favorable responses may be *far* from good. There is no reason to assume they would continue to be even moderately reliable in contexts far from their home territory, and they might even be completely unreliable. In unfamiliar contexts, the anchor-and-adjust response that guides our intuition might trigger a simple heuristic that is not at all a good proxy for what a more considered, more logical analysis would recommend.

I stress here, in particular, the contrast between contexts of individual choice in commonplace circumstances against contexts of social choice, where the aggregate effects of individual responses have large consequences entirely beyond the scope of ordinary experience. Then, the effect of any one individual choice is socially microscopic, hence also the motivation for an individual to think hard about that choice. Indeed, in the individual context, if a particular response is what a person feels comfortable with, that is likely to be all we need to know to say that response is a good one, or at least a good enough one. But in social contexts that very often is not true at all.

Here is an illustration. Our genetically entrenched propensities with respect to risk would have evolved over many millennia of hunter-gatherer experience, where life was commonly at the margin of subsistence, and where opportunities for effective stockpiling of surplus resources were rare. In such a context, running risks to make a large gain would indeed be risky. If you succeed, you may be unable to use what you have gained, and if you fail (when you are living on the margin of subsistence), there may be no future chances at all.

In contrast, such a life would encourage risk taking with respect to averting losses. If you have little chance to survive a substantial loss, you might as well take risks that could avert it. So we have a simple account of why (Kahneman & Tversky 1981) we are ordinarily risk-averse with respect to gains but risk-prone in the domain of losses. And this yields simple heuristics: with respect to gains, when in doubt, don't try it, and the converse with respect to averting a loss.

But we do not now live on the margin of subsistence, so now these simple heuristics that make us smart may easily be dumb: especially so with respect to a social choice that is going to affect large numbers of individual cases, where (from the law of large numbers) maximizing expectation makes far more sense than either risk aversion or its opposite. But because entrenched responses are (of course!) hard to change, and especially so because their basis is likely to be scarcely (or even not at all) noticed, effectively altering such choices will be socially and politically difficult.

Although space limits preclude developing the point here, I think a strong case can be made that the consequences of such effects are not small. One example (but far from the only one), is the case of new drug approvals in the United States, where there has been a great disparity between attention paid to the risks of harmful side effects and to the risks of delaying the availability of a medicine to patients whose lives might be saved.

Nor can we assume that faulty judgments governed by simple heuristics will be easily or soon corrected. I ran across one recently that had persisted among the very best experts for 400 years (Margolis 1998).

I think a bit of old-fashioned advice is warranted here: "If we may admire the power of ideas when they lead us and inspire us, so must we learn also their sinister effects when having served their purpose they oppress us with the dead hand." (Sir Thomas Clifford 1913) Same for simple heuristics.

Heuristics all the way up?

Adam Morton

Department of Philosophy, University of Bristol, Bristol BS8 1TB, United Kingdom **adam.morton@bristol.ac.uk** **mail.bristol.ac.uk/~plam**

Abstract: I investigate whether heuristics similar to those studied by Gigerenzer and his co-authors can apply to the problem of finding a suitable heuristic for a given problem. I argue that not only can heuristics of a very similar kind apply but they have the added advantage that they need not incorporate specific trade-off parameters for balancing the different desiderata of a good decision-procedure.

Assume that most of the claims made by Gigerenzer et al. (1999) are true, and that each person possesses a toolbox of relatively

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

naive procedures which, when used in suitable environments, give results almost as good as sophisticated statistical reasoning, for a fraction of the psychological cost. These procedures will give extremely bad results when used in the wrong environments. The simplest example of a bad environment is one where the values of a quantity to be estimated by the recognition heuristic are inversely correlated with the familiarity of the cues. There is thus a very non-trivial problem of matching heuristics to problems. How do we know which tool to take from the box?

The matching could be done by some approximation to the statistical considerations of Chapters 6 and 8. Or it could itself use some simple heuristics, which in suitable environments gave good results. Or, it could work in some entirely different way. Let us explore the second possibility. (The first seems unattractive, given the general ethos of the project, and although the third might well be true we can only explore it once we see the limitations of the second.) We must thus see how, given a problem of decision or estimation, we can choose from among a set of available simple procedures one which will give an accurate-enough and efficient-enough solution. People do seem to adapt their decision-making procedures to the circumstances of the problem at hand, as remarked in Chapter 6 (Martignon & Hoffrage, p. 140), referring to Payne et al. (1988; 1993). But the core theory of frugal heuristics is itself neutral on the procedures by which the right heuristic for the problem at hand is chosen, as Chapter 16 (p. 364, Todd & Gigerenzer) explicitly accepts.

When we pose the meta-choice problem we are immediately faced with a question of incomparability. The choice procedure is asked simultaneously to optimize accuracy and frugality, without being given a trade-off function between them. This might be taken to be another dimension of difficulty, and indeed recent philosophical literature on decision-making often treats incomparability as a basic conceptual problem of decision, along with risk and cooperation (see Chang 1997; Morton 1990). But the simple heuristics point of view has a very significant card to play. Frugal heuristics often manage incomparability as a matter of course.

Suppose, for example, that we are faced with the problem of choosing a restaurant, where both quality and price are desiderata. If we apply a heuristic with a simple stopping rule, such as Take the Best we can find that we do not need to decide how to balance the desiderata. We must phrase the stopping rule neutrally: stop when you find a characteristic that correlates with either of the targets. But then we can simply stop searching when we find a candidate and a characteristic that is linked either to good price or good quality. In a series of such choices a balance of quality and price will be struck, but the form of the balance will depend on the environment to which the heuristic is applied and not on any parameter of its definition.

The same can apply with a meta-choice. Suppose we have a toolbox of heuristics and a database of cases, which would specify the categorization of a problem, whether one of the heuristics was applied, and whether it gave a result that was acceptable in terms of time or accuracy. Then, given a new problem falling into some category, a person can apply a meta-heuristic that involves a stopping rule as described above to select one that is acceptable either on the one criterion or the other. The result of applying this heuristic to this problem then expands the database to make it a more effective basis for subsequent meta-choices.

The choice of heuristics is unlikely to be made by any rule that is exactly parallel to any first order heuristic. But some features of first order heuristics are very attractive as attributes of meta-choice, in particular the use of a simple stopping rule and the determination of some basic parameters by the environment. This suggests to me that there is another break to make with a traditional conception of rationality. We must not only be wary of an ideal of unlimited cognitive capacity; we must also be suspicious of ideals of self-containment, according to which all the parameters of a choice procedure are fixed by internal features of cognition.

What is an ecologically rational heuristic?

Stephen E. Newstead

Department of Psychology, University of Plymouth, Plymouth PL4 8AA United Kingdom. **s.newstead@plymouth.ac.uk**

Abstract: The notion of ecological rationality, although plausible, does not readily lead to testable predictions. This is illustrated with respect to heuristics in syllogistic reasoning. Several possible heuristics have been proposed but ecological rationality does not appear to offer a sensible rationale for choosing between these.

There is much that I agree with in the book by Gigerenzer, Todd, and the ABC Research Group. I am sure they are right that many human judgments are based on heuristics. These heuristics are indeed often sensible ways of dealing with the situation in hand, and are sometimes more effective than logical responding. I find myself less convinced by their claim that these heuristics are "ecologically rational," in the sense that they are "adapted to the structure of an environment" (Gigerenzer et al. 1999, p. 13).

The problem is in translating this claim into operational terms and using it to make predictions. Any heuristic response will involve responding on the basis of only part of the information presented or using only limited search. But this claim has value only if it can predict which aspects of the situation will be selectively responded to. If the relevant features can be indicated only after responses are given then the theory has little or no predictive validity.

Consider, for example, syllogistic reasoning. This is precisely the kind of artificial reasoning task in which people deviate from the dictates of logic and where heuristics might come into play. There has, in fact, been no shortage of possible heuristics invoked to explain syllogistic reasoning errors. Possible contenders include:

Atmosphere: The claim that people base their responses on surface features of the quantifiers – whether these are negative or positive, and universal or particular (Woodworth & Sells 1935).

Probability: The claim that people use the least informative premise as the basis for their conclusion (Chater & Oaksford 1999).

Caution: The claim that people choose the response which commits them to the least general statement (Sells 1936).

Limited processing: The claim that people reach a decision after constructing just one representation of the premises (Newstead et al. 1999).

Believability: The claim that people give the conclusion they find most believable (Evans et al. 1983).

All of these predict reasonably well the responses that people actually give, and they are not mutually exclusive. The point is that there seems to be no sensible rationale based on ecological rationality for predicting which of these biases (if any) would be expected to operate. Intuitively, both the limited processing and the believability heuristics might make sense in a variety of other situations. Limited processing is, of course, one of the heuristics discussed by Gigerenzer et al.; and responding according to believability makes sense, as it would in many situations be irrational to accept an argument that was known to lead to an untrue conclusion. But much the same could be said of probability and caution. Oaksford and Chater claim that their probability heuristic has considerable generality, and caution could be a sensible heuristic to adopt in a number of very different situations. Atmosphere seems to be the only heuristic motivated solely by considerations of the syllogistic reasoning task, and even this could be seen as a more general tendency to match responses to surface characteristics of the information given.

And there's the rub. For if any of these heuristics can be argued to be ecologically rational, how can one predict which one will occur? And if one cannot make a prediction, how can the theory be tested? It is not enough to see which heuristic best fits the evidence and then claim that this is the most ecologically rational; this would clearly be circular.

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

A further point worth noting with respect to syllogistic reasoning is that there is overwhelming evidence that people *are* capable of responding according to logic. Studies of belief bias, for example, almost invariably find an effect of logic as well as an effect of believability (see Evans et al. 1983). It is not clear whether individuals respond partly on the basis of logic and partly on the basis of belief, or whether there are individual differences in the tendency to respond in these ways. Whichever is the case, this is clearly a complicating factor in any attempt to explain performance in terms of heuristics alone.

None of this implies that Gigerenzer et al. are wrong. On the contrary, what is needed is for them to develop their theory further so that more specific predictions can be put to the test.

Speed, frugality, and the empirical basis of Take-The-Best

Mike Oaksford

School of Psychology, Cardiff University, Cardiff, CF1 3YG, Wales, United Kingdom. **oaksford@cardiff.ac.uk** **www.cardiff.ac.uk/psych/oaksford**

Abstract: This commentary questions the claim that Take-The-Best provides a cognitively more plausible account of cue utilisation in decision making because it is faster and more frugal than alternative algorithms. It is also argued that the experimental evidence for Take-The-Best, or non-integrative algorithms, is weak and appears consistent with people normally adopting an integrative approach to cue utilisation.

Although I agree with much of Gigerenzer, Todd, and the ABC Research Group's approach – indeed, Nick Chater and I believe it can be extended to classical reasoning tasks such as syllogisms (Chater & Oaksford 1999) – I argue that their comparison with other algorithms provides less support for Take-The-Best than they suggest and that its empirical basis is weak and consistent with alternative interpretations.

Gigerenzer et al. (1999, Ch. 4, p. 94, Gigerenzer & Goldstein; Ch. 10, p. 231, Hertwig, Hoffrage & Martignon) cite the results of an unpublished study (Chater et al. 1999), comparing Take-The-Best with integrative algorithms such as neural networks and the generalised context model (GCM). Our interpretation of these results differs from that of Gigerenzer et al. (1999). They claim that Take-The-Best is particularly attractive because it is *fast* (it involves a small number of serial processing steps) and *frugal* (it draws on very limited information, because it is non-integrative). However, neither consideration straightforwardly gives Take-The-Best an advantage over the available alternatives.

There are two points to consider concerning speed. First, rapid integration of large amounts of information is believed to occur in language processing, perception, motor control, and commonsense reasoning. Hence integrative processing may be fast enough to account for the relatively slow human responses in the city size estimation task. the only empirical evidence the authors provide (Ch. 7, p. 141, Rieskamp & Hoffrage) is that people use non-integrative algorithms like Take-The-Best only under time pressure. However, most of the decision domains to which Gigerenzer et al. have directed their modeling efforts (Ch. 5, pp. 99–100, Czerlinski et al.), do not require decisions to be made under time pressure. For example, predicting which of two cities has the higher rate of homelessness is not a decision that people would normally need to make rapidly. Consequently, the emphasis on speed as a deciding factor between algorithms may be inappropriate.

Second, Gigerenzer et al.'s measure of speed depends on assumptions about the architecture of the cognitive system (Chater & Oaksford 1990; Oaksford & Chater 1991; 1993; 1995). On a serial architecture, where information is searched in memory at a constant rate, Take-The-Best would be more rapid than the accounts Chater et al. considered. But in a parallel architecture, processing speed will not generally be related to the amount of information searched in memory, because large amounts of information can be searched simultaneously. For example, both the learning and application of multiple regression can be implemented in parallel using a connectionist network with a single layer of connections. This implementation could operate very rapidly – in the time it takes to propagate activity across one layer of connections (e.g., Hinton 1989). Similarly in an instance-based architecture, where instances can be retrieved in parallel, GCM would be the quickest. Take-The-Best only has a clear advantage over other algorithms if cognitive processes are assumed to be serial. Given the extensive research programs aimed at establishing the viability of instance-based and connectionist architectures as general accounts of cognitive architecture (e.g., Kolodner 1993; Rumelhart et al. 1986), it seems that considerable caution must be attached to a measure of speed that presupposes a serial architecture.

Gigerenzer et al. (1999) also compare the computational complexity of different algorithms (pp. 183–86). Using this measure they find that the complexity of Take-The-Best compares favourably with integrative algorithms. However, Gigerenzer et al. concede that the relevance of these worst-case sequential analyses to the presumably highly parallel human implementation is not clear (although the speed up could only be by a constant factor). Moreover, as we now argue there are many cognitive functions that require the mind/brain to perform rapid, integrative processing.

Take-The-Best is undoubtedly a very frugal algorithm. Rather than integrating all the information that it is given (all the features of the cities), it draws on only enough feature values to "break the tie" between the two cities. But does the frugality of Take-The-Best make it more cognitively plausible? Comparison with other domains suggests that it may not.

In other cognitive domains, there is a considerable evidence for the integration of multiple sources of information. We consider two examples. First, in speech perception, there is evidence for rapid integration of different cues, including cues from different modalities (e.g., Massaro 1987). This integration even appears to obey law-like regularities (e.g., Morton 1969), which follow from a Bayesian approach to cue integration (Movellan & Chadderdon 1996), and can be modelled by neural network learning models (Movellan & McClelland 1995). Second, recent work on sentence processing has also shown evidence for the rapid integration of multiple "soft" constraints of many different kinds (MacDonald et al. 1994; Taraban & McClelland 1988).

Two points from these examples are relevant to the cognitive plausibility of Take-The-Best. First, the ability to integrate large amounts of information may be cognitively quite natural. Consequently, it cannot be taken for granted that the non-frugality of connectionist or exemplar-based models should count against their cognitive plausibility. Second, the processes involved require rich and rapid information integration, which cannot be handled by a non-integrative algorithm such as Take-The-Best. Thus Take-The-Best may be at a disadvantage with respect to the generality of its cognitive performance.

A possible objection is that evidence for rapid integration of large amounts of information in perceptual and linguistic domains does not necessarily carry over to the reasoning involved in a *judgment* task, such as deciding which of two cities is the larger. Perhaps here retrieval from memory is slow and sequential, and hence rapid information integration cannot occur. However, the empirical results that Gigerenzer et al. report (Ch. 7) seem to show performance consistent with Take-The-Best only when people are under time pressure. But, as we suggested above, the judgments used in Gigerenzer et al.'s experiment do not normally need to be made rapidly. Consequently, the authors' data are consistent with the view that under normal conditions, that is, no time pressure, participants wait until an integrative strategy can be used. That is, for the particular judgments that Gigerenzer et al. consider, it is more natural for participants to adopt an integrative strategy; they use a non-integrative strategy only when they do not have the time to access more cues.

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

In sum, whether Take-The-Best is fast, is architecture dependent, and whether frugality is a virtue are questioned by the many other cognitive functions that require fast, parallel, and integrative approaches. Moreover, the weak empirical basis for Take-The-Best seems consistent with people normally adopting an integrative approach (albeit with limited cues).

Sub-optimal reasons for rejecting optimality

David R. Shanks and David Lagnado

Department of Psychology, University College London, London WC1E 6BT United Kingdom. **{d.shanks; d.lagnado}@ucl.ac.uk**
www.psychol.ucl.ac.uk/david.shanks/Shanks.html.

Abstract: Although we welcome Gigerenzer, Todd, and the ABC Research Group's shift of emphasis from "coherence" to "correspondence" criteria, their rejection of optimality in human decision making is premature: In many situations, experts can achieve near-optimal performance. Moreover, this competence does not require implausible computing power. The models Gigerenzer et al. evaluate fail to account for many of the most robust properties of human decision making, including examples of optimality.

A paradox in the rationale for fast and frugal algorithms. There is a curious paradox in Gigerenzer et al.'s argument for the role of fast and frugal algorithms in human decision-making (Gigerenzer et al. 1999). They suggest that psychologists have been led astray by focusing on behavior from an optimization perspective and they imply that optimization models are implausible, intractable, and require demonic capacities. Instead, they urge us to explore fast and frugal algorithms. Yet they admit on p. 237 that categorization performance can often be optimal, but if that is the case, then surely the case for fast and frugal algorithms – which will almost certainly never achieve this level of performance – evaporates?

To evaluate Gigerenzer et al.'s case for minimal complexity in cognitive processes it is critical to determine whether decision making is truly optimal. The jury is still out on this issue, of course, but what is indisputable is that near-optimal performance can be achieved by experts in many realms including categorization (Anderson 1991; Ashby & Maddox 1992) and choice (Binmore 1991; Davis et al. 1993).

Gigerenzer et al. repeatedly ridicule what they take to be "optimal" theories (e.g., multiple linear regression, MLR) on the grounds that they require unrealistic amounts of computation (e.g., p. 76), but this is a highly misleading claim. Contrary to the impression made by Gigerenzer et al., it is possible to find a regression solution in minimal time without doing any computation at all. Imagine a set of points each represented by a peg on a two-dimensional board (the solution also works in principle in n dimensions). Then attach a long thin rod by a set of elastic bands to the pegs. By minimizing the allocation of tension across the elastic bands, the rod will align exactly according to the regression equation. As another example, consider the well-known Travelling Salesman problem in which the shortest route must be computed that visits each of a number (N) of cities exactly once. Despite its computational complexity (the computing time needed to solve this problem increases faster than any power of N), near-optimal solutions can be achieved by parallel neural networks in the blink of an eye (Hopfield & Tank 1985). What objection is there to the view that the human cognitive system approximates optimality by use of parallel constraint-satisfaction processes?

It is also troubling that Gigerenzer et al. take multiple linear regression (MLR) as one of their benchmark models throughout the book. Humans can learn highly nonlinear judgment rules in a variety of domains (Ashby & Maddox 1992; Ceci & Liker 1986) so MLR is simply not an appropriate model. If TTB (Take The Best) and CBE (categorization by elimination) approximately match the performance of MLR and if human experts significantly outperform MLR then the obvious conclusion is that TTB and CBE are inadequate models of human performance.

Implausibility of the CBE model. We believe that the candidate fast-and-frugal model for categorization which Gigerenzer et al. present, the CBE model, is wholly inadequate for human performance. First, it is unable to predict one of the benchmark phenomena of categorization, namely the ubiquitous "exemplar effect," that is, the finding that classification of a test item is affected by its similarity to specific study items, all else held constant (e.g., Whittlesea 1987). Even in the case of medical diagnosis, decision-making in situations very like the heart-attack problem Gigerenzer et al. describe is known to be strongly influenced by memory for specific prior cases (Brooks et al. 1991). The recognition heuristic is not adequate to explain this effect because it concerns the relative similarity of previous cases, not the absolute presence versus absence of a previous case. If a heavy involvement of memory in simple decision tasks seems to characterize human performance accurately, then plainly models which ignore this feature must be inadequate.

Second, there is strong evidence against deterministic response rules of the sort embodied in Gigerenzer et al.'s fast-and-frugal algorithms (Friedman & Massaro 1998): for instance, Kalish and Kruschke (1997) found that such rules were rarely used even in a one-dimensional classification problem. Thirdly, CBE is not a model of learning: it says nothing about how cue validities and response assignments are learned. When compared with other current models of categorization such as exemplar, connectionist, and decision-bound models, which suffer none of these drawbacks, the CBE model begins to look seriously inadequate.

Methodology of testing the models. By taking a tiny domain of application (and one which is artificial and highly constrained), Gigerenzer et al. find that the CBE model performs competently and conclude that much of categorization is based on the application of such algorithms. Yet they mostly do not actually fit the model to human data. The data in Tables 5-4, 11-1, and so on, are for objective classifications, not actual human behavior. It is hard to see how a model's ability to classify objects appropriately according to an objective standard provides any evidence that humans classify in the same way as the model.

Even in the cases they describe, the models often seriously underperform other models such as a neural network (Table 11-1). Compared to the more standard approach in this field, in which researchers fit large sets of data and obtain log-likelihood measures of fit, the analyses in Chapters 5 and 11 are very rudimentary. Gigerenzer et al. report percent correct data, which is known to be a very poor measure of model performance, and use very small datasets, which are certain to be highly nondiscriminating. The difference between the CBE model and a neural network (e.g., up to 9% in Table 11.1) is vast by the standards of categorization research: for instance, Nosofsky (1987) was able to distinguish to a statistically-significant degree between two models which differed by 1% in their percentages of correct choices.

Melioration as a fast-and-frugal mechanism. The algorithms explored by Gigerenzer et al. (TTB, CBE, etc.) share the common feature that when a cue is selected and that cue discriminates between the choice alternatives, a response is emitted which depends solely on the value of that cue. Gigerenzer et al. (Ch. 15, p. 327, Goodie et al.) consider the application of such models to the simplest possible choice situation in which a repeated choice is made between two alternatives in an unchanging context. The prototypical version of such a situation is an animal operant choice task in which, say, a food reinforcer is delivered according to one schedule for left-lever responses and according to an independent schedule for right-lever responses. As Gigerenzer et al. point out (p. 343), fast-and-frugal algorithms predict choice of the alternative with the highest value or reinforcement rate. Although this may seem like a sensible prediction, human choice does not conform to such a "momentary-maximization" or "melioration" process. In situations in which such a myopic process does not maximize overall reinforcement rate, people are quite capable of adopting

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

better response strategies. A well-known example is the Harvard Game (Herrnstein 1997) in which one response alternative (say, right) always pays more at any moment than the other (left), but where overall reinforcement rate is maximized by allocating all responses to left. People's behavior is often seen to approach optimality under such conditions (Herrnstein et al. 1993; R. Tunney & D. Shanks, unpublished data). Yet again we have an example of humans' ability to achieve near optimal levels of performance, exactly as the "demonic" theory of rational choice predicts.

The selection problem. Gigerenzer et al. say very little about how individual heuristics are selected for application to specific problem domains. Such meta-level decisions will typically require some prior knowledge about the structure of the environment (e.g., whether it is non-compensatory, J-shaped, etc.), which may add substantially to the overall processing costs of a fast and frugal model. This would reduce its advantage over those models that can learn about the environment and have general applicability (thus cutting out the metadecision stage). In the majority of the simulations in the book, one particular heuristic is pre-selected to operate in a specific environment. More rigorous tests would place a complete fast and frugal system (different heuristics plus metadecision heuristics) in a variety of different environments, and compare its performance against an alternative general-purpose learning model.

The precision/accuracy trade-off. On a more positive note, we welcome Gigerenzer et al.'s shift of emphasis from "coherence" to "correspondence" criteria. This is an important step towards a more complete understanding of rationality, and removes some of the obstacles placed by the heuristics-and-biases school. In addition to the examples cited in the book, the inadequacy of coherence criteria has been demonstrated in various experiments in which people trade the probability of being correct for increased precision in their judgments (Yaniv & Foster 1995), thereby sacrificing probabilistic coherence for a possible gain in informational content. It is not clear, however, how readily this fits into the fast and frugal picture. Computing a trade-off between precision and accuracy would appear to place additional processing demands on the decision maker, contrary to the spirit of speed and frugality.

We believe that this problem is resolvable by identifying the appropriate correspondence criterion, and the cognitive mechanisms attuned to this criterion. Recent work in human causal induction (López et al. 1998) suggests that predictive judgments are mediated by associative mechanisms sensitive to real-world statistical contingencies. Furthermore, it can be shown that networks sensitive to such a measure automatically compute a precision/accuracy trade-off. This suggests, contra Gigerenzer et al., that although the mechanisms underlying our mental algorithms may be simple, the computations which they embody need not be.

Fast and frugal heuristics: What about unfriendly environments?

James Shanteau and Rickey P. Thomas

Department of Psychology, Kansas State University, Manhattan, KS 66506-5302 **{shanteau; thomas}@ksu.edu www.ksu.edu/isbr/shanteau.htm**

Abstract: *Simple heuristics that make us smart* offers an impressive compilation of work that demonstrates fast and frugal (one-reason) heuristics can be simple, adaptive, and accurate. However, many decision environments differ from those explored in the book. We conducted a Monte Carlo simulation that shows one-reason strategies are accurate in "friendly" environments, but less accurate in "unfriendly" environments characterized by negative cue intercorrelations, that is, tradeoffs.

Gigerenzer, Todd, and the ABC Research Group's book offers an impressive compilation of theoretical and empirical work on how one-reason decision making can be both efficient and accurate in appropriate environments (Gigerenzer et al. 1999). The interaction between mind and environment is central to the book (see Gigerenzer & Todd, Ch. 1). Simon's "bounded rationality" concerns the mind side of the interaction, that is, how the human mind operates with limited computational capacity. "Ecological rationality" concerns the environment side of the interaction, namely, how decision strategies implemented by the mind are adapted to the environment. The findings in the book demonstrate convincingly that fast and frugal (F&F) heuristics can be simple, adaptive, and accurate.

The environments examined are sampled from real-world domains and share four important properties; choices are based (1) on judgments of estimation/inference (2) between two alternatives (3) with dichotomous cues (4) where cue intercorrelations are positive. Such simple environments are an obvious place to begin in a research program.

Many decision environments, however, possess different properties than those explored here (as recognized by Todd & Gigerenzer, Ch. 16). For instance, choice is often based on (1) preferences (2) between more than two alternatives (3) with continuous cues (4) in "unfriendly" environments. "Friendly" environments (of the sort analyzed in this book) have intercorrelations that are positive for the most valid cues. Rieskamp & Hoffrage (Ch. 7) are aware of findings that suggest negative cue intercorrelations may significantly influence the decision process. We argue that unfriendly environments must be considered when evaluating decision rules.

In many multiattribute problems cues are negatively correlated, which implies the need to make tradeoffs (Stillwell et al. 1981). For example, consumer decisions require a tradeoff between price and quality, that is, an increase in quality (a good thing) can lead to a price increase (a bad thing). Note that reverse scoring of cues does not eliminate the tradeoff. In an unfriendly world, no alternative exists that simultaneously maximizes all attribute dimensions (McClelland 1978). Decision-makers often make cascaded (multistage) choices, in which alternatives are successively eliminated if they fail to meet certain minimal requirements (Todd & Miller, Ch. 13). Such exclusion stages can dramatically change cue intercorrelations. McClelland (1978) proves analytically that multiattribute cues are necessarily negatively correlated for non-dominated alternatives. Johnson et al. (1989) further argue that negative cue intercorrelations arise in consumer decisions whenever the market/consumer eliminates dominated alternatives from consideration.

Does it matter whether environments are friendly or unfriendly? The following simulation study was designed to provide answers.

We compared three decision strategies in a Monte-Carlo simulation: Lexicographic (LEX) strategy (Coombs 1974), Dawes's equal-weight (EQ) rule (Dawes & Corrigan 1974), and Multiattribute utility (MAU) with Rank-Sum weights (Stillwell et al. 1981). LEX (a one-reason decision strategy) selects the alternative with the highest utility for the most valid cue – a generalization of Take The Best (Rieskamp & Hoffrage, Ch. 7). EQ serves as a nonnormative alternative and MAU (RS weights) serves as a normative alternative.

The decision strategies were evaluated for choices in three environments: friendly, neutral, and unfriendly. All decision environments consist of two cues and three alternatives. The cues are continuous bivariate normal with cue intercorrelations of .75 (friendly), 0 (neutral), and −.75 (unfriendly). Five thousand decision problems were simulated for each environment. All results were compared to the optimal weighted additive difference (WAD) model using true weights. The entries in Table 1 are mean selection accuracy, that is, the average proportion of times the decision strategy selects the same alternative as WAD (true weights).

The accuracy results for all decision strategies decrease as the environment changes from friendly to neutral to unfriendly. Just as the authors contend, a one-reason strategy (LEX) performs as well as other strategies in friendly environments (Martignon & Hoffrage, Ch. 6; Todd & Gigerenzer, Ch. 16). However, the selection accuracy of LEX is less in unfriendly environments.

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

Table 1 (Shanteau & Thomas). *Average selection accuracy of three decision strategies in friendly, neutral, and unfriendly decision environments*

	Strategy		
Environment	LEX	EQ	RS
Friendly	0.90	0.93	0.96
Neutral	0.80	0.84	0.90
Unfriendly	0.76	0.70	0.84

Note. The average selection accuracy estimates are based on comparisons with WAD across five true weight vectors ({.50, .50}, {.60, .40}, {.70, .30}, {.80, .20}, and {.90, .10}).

From other simulations, we have found the selection accuracy of LEX generally decreases in unfriendly environments as (1) the steepness of the tradeoff increases, (2) the number of important cues increases, and (3) the number of alternatives increases. Moreover in some unfriendly environments, the selection accuracy of LEX is lower than a random choice process.

We believe our findings are supported by common sense; one good reason is not adequate for making decisions when there are tradeoffs. When consumers buy a VCR, for instance, they do not automatically go with the lowest price because the selected model is likely to be of inferior quality. As Keeney and Raiffa (1993) note, LEX is "naively simple" and "will rarely pass a test of reasonableness" (pp. 77–78).

A related problem is that one-reason heuristics are not descriptive of expertise. Experts generally work in complex environments involving multiattribute decisions with multiple tradeoffs, that is, negative cue intercorrelations. They typically follow a hierarchical process, with different stages of decision making. Although a one-reason rule may describe the final stage, one cannot conclude the entire process consists of such simple processes (Phelps & Shanteau 1976).

In conclusion, one-reason decision making works well in the environments assumed in the book – two alternative estimation (or inference) tasks where the cues are dichotomous and positively intercorrelated (for examples, see Czerlinski et al., Ch. 5). Our simulations show that F&F rules also perform well when choices are based on preference and the cues are continuous, as long as the environment is friendly. One-reason decision making, however, does not perform well in unfriendly environments (characterized by tradeoffs) when there are more than two alternatives. Future research is needed to determine how well modified (multistage) F&F rules might do in such environments, for example, using elimination and satisficing heuristics (Hertwig et al., Ch. 10; Berretty et al., Ch. 11; Todd & Gigerenzer, Ch. 16). We look forward to learning of the results from Gigerenzer, Todd, and the ABC Research Group.

ACKNOWLEDGMENT

Preparation of this paper was supported in part, by Grant 98-G-026 from the Federal Aviation Administration, Department of Transportation.

Heuristics and development: Getting even smarter

Gregg E. A. Solomon

Decision Management Group and Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139.
geas@psyche.mit.edu

Abstract: There are parallels between Gigerenzer et al.'s emphasis on the rationality of adults' reasoning in terms of simple heuristics and developmental researchers' emphasis on the rationality of children's reasoning in terms of intuitive theories. Indeed, just as children become better at using their theories, so might some people, experts, become better at using simple heuristics.

Gigerenzer, Todd, and the ABC Research Group (1999) present a welcome counter to the prevailing interpretation of human decision making: Left to our own devices, we tend to be irrational. The issues raised in the debate over the extent to which people observe the rules of logic recalls the debate in cognitive development concerning the extent to which *children* observe the rules of logic and how they become better thinkers.

The traditional, staunchly Piagetian, interpretation is that young children's reasoning does not adhere to the rules of formal logic. Development, on such a view, can be described as the acquisition of increasingly more powerful, domain-general reasoning structures. Just as the adult rationality literature has focused on robust and compelling phenomena, the Piagetian research tradition has yielded a collection of striking failures of logic. Not surprisingly, the conception of the young child implicit in many science education curricula in Euro-America is that of a thinker incapable of formal logical reasoning. Consequently, for many of these curricula, the preliminary goal is to teach children how to think logically (e.g., Lawson & Thompson 1988). But the attempt to improve children's understanding of science through tuition in the principles of logic has not been terribly successful (Bruer 1983). The literature on judgment errors would similarly indicate that even adults who have been trained in logic and probability are not immune to such errors, suggesting a lack of transfer of abstract logical principles to real-world decision making.

In recent years, cognitive science has broken with the concept of the child as a fundamentally different kind of thinker (e.g., Carey 1985). Many researchers have come to describe them as possessing powerful intuitive or framework theories that guide their understanding of the world about them (Wellman & Gelman 1992). Gigerenzer and Todd share this notion that we can have dispositions to reason within particular domains in terms of specific principles, and that such privileged ways of reasoning are triggered by specific cues. An integral part of the framework theory is still an account of how children's thinking develops, of how they get better at making predictions about the world. The educational implications of this characterization of cognitive development are an emphasis on teaching facts, as well as on addressing the frameworks in which children interpret their new facts.

Gigerenzer et al.'s model, unfortunately, gives development short shrift. This is not to say that their ABC research program precludes a developmental component. Far from it; their lack of attention to educational implications of heuristics is more likely a consequence of their (delightfully perverse) emphasis on demonstrations of superior novice performance. There is again a parallel in the framework theory literature. Many developmental researchers, in their rush to demonstrate the brilliant rationality of little children, have done so at the expense of a recognition of differences between children's early theories and the target adult theories. So too do the ABC researchers pay little attention to possible expert/novice differences. If there are such things as experts (granting that there is no single kind of expertise), then there is a place in this research program for the study of development.

What, if anything, do experts do that novices do not in their use of heuristics? Borrowing again from the debate in cognitive de-

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

velopment, we might expect two broad courses of development: continuous and discontinuous. It can be argued that young children reason about particular phenomena (e.g., folk-psychology) in qualitatively the same ways as adults do, but their predictions often differ simply because they do not know as much as adults do. Do experts in a particular domain use the same heuristics as do novices, but merely use them better, more efficiently? For example, surely the Take The Best heuristic yields improved performance when employed by individuals with more experience on which to base their estimations of which cues have been more successful.

By contrast, there are other domains of thought in which children's theories would appear to undergo conceptual change. For example, young children may respond to cues in "living things" that lead them to reason in a manner more appropriate to folk-psychology than folkbiology (see Carey 1985). Children arguably come to interrelate concepts concerning living things in a qualitatively different manner; there is a change in which cues are recognized as centrally important. Similarly, it may be that experts and novices alike invoke heuristics in a particular domain, but they could differ immensely in their choices of search, stop, or decision making strategies. Indeed, novices, given their relative ignorance, might tend to rely on default dispositions to recognize particular cues in specific domains, whereas experts might be driven by explicit theories to attend to particular cues and in terms of particular heuristics (e.g., in the domain of wine tasting, see Solomon 1997). The expert strategies might be more ecologically rational, more suited to the particular environmental structures, but novices might not know enough to realize this (not unreasonably, in the case of wine tasting, novices' decision strategies would appear more suited to making predictions about beverages in general rather than wines in particular). Finally, it remains to be seen to what extent novices (and children) can be trained in more efficient use of heuristics. At the very least, the nature of such interventions ought to reflect an attention to the nature of the specific differences between experts and novices, rather than a general attention to logical ability.

Damn it, I still don't know what to do!

Robert J. Sternberg

Department of Psychology, Yale University, New Haven, CT 06520-8205.
robert.sternberg@yale.edu www.yale.edu/rjsternberg

Abstract: The simple heuristics described in this book are ingenious but are unlikely to be optimally helpful in real-world, consequential, high-stakes decision making, such as mate and job selection. I discuss why the heuristics may not always provide people with such decisions to make with as much enlightenment as they would wish.

Simple heuristics that make us smart (Gigerenzer et al. 1999) is a wonderful, ingenious, and even brilliant book. I found it one of the most interesting psychology books I have read in quite a while. It is a milestone that helps move human decision-making research into a new post-Kahneman-and-Tversky era.

I was delighted to receive the book when I did, because it arrived when I was (and still am) in the throes of a very difficult, high-stakes decision – whether to remain at the institution where I am currently teaching or to accept another job offer. Much as I enjoyed reading the book, when I finished it, I threw up my hands and said to myself, "Damn it, I still don't know what to decide." Why?

Stakes in decision making. Most (although certainly not all) of the decisions described in the book are for very low stakes. How much do I care whether San Diego has more inhabitants than San Antonio, how much houses cost in Erie, Pennsylvania, how fertile 395 Arctic char are, or how much oxygen is absorbed by cow manure? In making these decisions, I would happily go for any of the one-reason heuristics for decision making. But I would be much more reluctant to do so when an important life decision is at stake. Although the kinds of problems used by Gigerenzer, Todd, and their colleagues are several steps up from poker chips, most of them are still a step down from consequential life decisions.

Some of the decisions described in the book are potentially consequential, such as investing in the stock market, choosing a mate, and investing in children. Whether most people will feel comfortable investing solely on the basis of name recognition of companies remains to be seen. But I found the advice on mate-selection decisions and parental investment less than useful. Why?

What are the attributes and what are their weights? Issues with regard to mate selection are described in terms of a dowry problem, where one is trying to select the mate who will produce the highest dowry. This decision and others in the chapter differ from real mate-selection decisions, however, in that the attributes involved are clearly specified and unidimensional. But in real-life mate selection, at least for many people, there is no obvious attribute, such as a dowry, that can be readily quantified and expressed unidimensionally. Often it is not even clear what the attributes are, or if they are identified, which ones should matter at all.

The chapter repeatedly refers to one mate choice being better than another or in the top such-and-such percent, but much of the problem is figuring out the attributes to use to decide who is better than whom. And having figured out the attributes, one still needs to figure out their weights. These issues, discussed in so much detail for the less consequential decisions, are bypassed for this kind of more consequential decision.

Multiple interests. The chapter on mate selection as well as that on parental investment are especially interesting because they touch on a further issue that characterizes real-life high-stakes decisions but that does not characterize most of the decisions considered in this book. In everyday life, many interests are involved. In truth, mate selection does not just involve the self and the partner. It usually also involves interests of parents, friends, members of reference groups, and so forth. A career decision, too, often involves many different parties, as do union-management negotiations, international negotiations, and the like.

Gigerenzer et al. deserve credit for looking at many decision-making contexts. But many of the contexts are rather different from the decisions that truly matter in one's life.

What is missing. Other issues such as those mentioned above could be raised, but my goal is not to be exhaustive but rather to make a more general point. The kind of decision making described here, in many cases, is reminiscent of the kinds of problem solving found on conventional tests of mental abilities. The items clearly measure important skills, but perhaps not those that will matter most to who succeeds in real-life decision making or problem solving so that colleagues, spouses, and friends unaware of test or task scores will label the person as "smart."

Consequential real-world decision making and problem solving are often different in kind from the decision making and problem solving represented by low-stakes problems that to participants may seem contrived. Our own data suggest that there is little or no correlation between performance on the two kinds of tasks (Sternberg et al. 2000). Even the real-world tasks studied, such as mate selection and parental investment, are not studied in a way that is likely to be maximally helpful to people facing serious everyday constraints. For example, it is not at all clear that parents can or will "rely on simple rules to guide their investment in their children" (p. 309), or that we would label as "smart" parents who do so.

One book later, I am still stuck as to my employment decision. I could try just counting the top two considerations, if I could figure out what they are or even should be. But that is much of the problem. Wonderful though the book may be, I suspect that many others seeking decision rules for the high-stakes decisions they encounter in their lives will not find that the rules in this book will make their decisions all that easy. Damn.

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

ACKNOWLEDGMENT

Preparation of this report was supported by Grant R206R950001 from the U.S. Department of Education, Office of Educational Research and Improvement. Such financial support does not constitute agreement with or support for any of the statements made in this article.

Smart people who make simple heuristics work

Annika Wallin and Peter Gärdenfors

Lund University Cognitive Science, Kungshuset, Lundagård, 222 22 Sweden.
{annika.wallin; peter.gardenfors}@lucs.lu.se
www.lucs.lu.se/People/{Annika.Wallin; Peter.Gardenfors@lucs.lu.se

Abstract: To evaluate the success of simple heuristics we need to know more about how a relevant heuristic is chosen and how we learn which cues are relevant. These meta-abilities are at the core of ecological rationality, rather than the individual heuristics.

Gigerenzer, Todd, and the ABC Research Group (1999) focus on simple heuristics for decisions instead of optimization procedures that presume unbounded rationality. We agree that this is an important step toward an understanding of the cognitive processes underlying human (and animal) decision making. However, Gigerenzer et al. mainly explain the success of simple heuristics as an exploitation of the structure of our natural environment. We wish to add that it is not the simple heuristics in themselves that make us smart. Knowing how to choose the right heuristic in the right context and how to select relevant cues is just as important in the decision process (regardless of whether theses choices are conscious or unconscious). In brief, we are smart enough to make simple heuristics work and before we can evaluate the role of simple heuristics, we must know more about how people choose to apply a particular heuristic in a given decision situation.

A heuristic must be applied in a context where it can reliably utilize the world's natural structure. For instance, the recognition heuristic is most sensibly used when there is a (causal) connection between the fact that we recognize something, and whatever factor it is we are trying to determine. In the examples presented, the environmental criteria presumed by the heuristics *are fulfilled by the selection of examples.* However, there are plenty of real world situations where this is not the case. If these heuristics are applied in such situations, they may not be as successful as Gigerenzer et al. claim. To repair this, one must add to the description of the heuristics how they take advantage of the environmental structure *through* our ability to find and understand certain regularities.

In order to apply most heuristics successfully, it is also necessary to know the value of the cues that are utilized. Another feature of Gigerenzer et al.'s examples is that knowledge concerning the relevant cues is accessible to the decision maker. The selection and ecological ordering of cues had already been made in the context the examples came from (mostly statistical textbooks). The ecological rationality of a heuristic such as Take The Best cannot be evaluated until we know more about how the cues are selected.

The value of a cue is judged by its ecological validity, which Gigerenzer et al. define as the proportion of correct predictions generated by the cue. Knowledge about the ecological validity of different cues is necessary for successful application of several of the heuristics studied by Gigerenzer et al. However, in a practical decision situation, agents must select the cues themselves and have no guarantee that the most relevant ones have been found. In such a situation, there is often *no way of knowing whether the best decision was made.* Hence there is a double difficulty in determining the validity of the cues.

We believe that ecological validity should be seen as only a secondary effect of the fact that a decision maker aims at forming hypothesis about *causal connections* between the cues and the decision variable. The causal reasoning involved in this process may better explain how the decision makers act than the statistical correlations that are used in Take The Best and the other heuristics. Unfortunately, Gigerenzer et al. do not discuss this kind of causal reasoning (Glymour 1998; Gopnik 1998).

Even if we stick to the ecological validity studied by Gigerenzer et al., it will be important to know how humans learn the correlations. One reassuring finding is that humans are very good at detecting covariations between multiple variables (Billman & Heit 1988; Holland et al. 1986). (But we don't know *how* we do it.) This capacity is helpful in finding the relevant cues to be used by a heuristic. The ability can be seen as a more general version of "ecological validity" and it may thus be used to support Gigerenzer et al.'s arguments.

Another aspect of the role of the experience of the agent is that the agent has some meta-knowledge about the decision situation and its context which influences the attitude of uncertainty to the decision. If the type of situation is well-known, the agent may be confident in applying a particular heuristic (since it has worked well before). But the agent may also be aware of her own lack of relevant knowledge and thereby choose a different (less risk-prone) heuristic. The uncertainty pertaining to a particular decision situation will also lead the agent to greater attentiveness concerning which cues are relevant in that kind of situation.

We have focused on two problems that have been neglected by Gigerenzer et al.: How the decision maker chooses the relevant heuristics and how the decision maker learns which cues are most relevant. We believe that these meta-abilities constitute the core of ecological rationality, rather than the specific heuristics that are used (whether simple or not). In other words, the important question concerning the role of heuristics is not whether the simple heuristics do their work, but rather *whether* we as humans possess the right expertise to use a heuristic principle successfully, and *how* we acquire that expertise.

From Simon's scissors for rationality to ABC's adaptive toolbox

X. T. Wang

Psychology Department, University of South Dakota, Vermillion, SD 57069.
xtwang@usd.edu

Abstract: The smartness of simple heuristics depends upon their fit to the structure of task environments. Being fast and frugal becomes psychologically demanding when a decision goal is bounded by the risk distribution in a task environment. The lack of clear goals and prioritized cues in a decision problem may lead to the use of simple but irrational heuristics. Future research should focus more on how people use and integrate simple heuristics in the face of goal conflict under risk.

1. A scissors missing one blade. Bounded rationality, according to Herbert Simon, is shaped by a scissors whose two blades are "the structure of task environments and the computational capacities of the actor" (1990, p. 7). However, an overview of the studies of human reasoning and decision making shows an unbalanced achievement. We have gained a great deal of knowledge about human computational capacities over the last several decades, but have learned little about the roles of the structure of task environments played in human rationality.

Although persistent judgmental errors and decision biases have been demonstrated in cognitive studies, biologists, anthropologists, and ecologists have shown that even young monkeys are adept at inferring causality, transitivity, and reciprocity in social relations (e.g., Cheney & Seyfarth 1985) and foraging birds and bees are rational in making risky choices between a low variance food source and a high variance one based on their bodily energy budget (e.g., Real & Caraco 1986; Stephens & Krebs 1986). This picture of "rational bees and irrational humans" challenges the Laplacian notion of unbounded rationality and calls for attention to the

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

mapping between our mental structure and the structure of evolutionary, ecological, and social environments.

In answering such a call, Gigerenzer et al. (1999) provide us with a toolbox of fast and frugal heuristics designed to work in different task environments. The book offers a ground-breaking synthesis of Simon's concept of satisficing (e.g., 1956a; 1990) and Brunswik's concept of vicarious functioning (e.g., 1940) with appealing theoretical ideas, thorough computer simulations and initial empirical testing. This team work of the ABC (Adaptive Behavior and Cognition) group represents a major advance in understanding human rationality under real-world constraints.

2. Satisficing goal setting and cue ranking. The smartness of simple heuristics depends on their fit to the structure of task environments. A lack of these constraints in task environments would devastate the search process for cues. An important issue raised by the book concerns the priority structure of decision cues. The proposed simple heuristics are primarily procedural heuristics dealing with the selection and integration of cues. However, these fast and frugal heuristics would not work effectively if one did not know the priority ranking of relevant cues. Thus, how to sort out valid cue ranking in a task environment, particularly when risks are involved, becomes essential for the success of simple heuristics. Although cue ranking determines the use of simple heuristics, goal setting in a task affects the ranking of relevant cues.

The functional values of fast and frugal heuristics can be viewed in a framework of goal setting and problem solving in which a decision maker is expected to maximize the opportunity of reaching a task-specific goal (see also, Lopes 1987). Being fast and frugal becomes psychologically demanding when a decision goal is bounded by the risk distribution in a task environment. The simple heuristics become even more useful when the risk factors are taken into account. Under risk, one needs not only consider a task-specific goal but also a task specific minimum requirement (e.g., to get x amount of y before the deadline z). In such a situation, to search for a satisficing solution under task constraints is to maximize the likelihood of reaching a goal and minimize the likelihood of falling below a minimum requirement at the same time. In contrast to the normative concept of maximizing expected utility values, a choice alternative yielding the highest expected value may not have a risk-variance distribution that satisfies the task constraints as measured by both the goal level and minimum requirement.

Goals not only define the relevance of decision cues but also define psychological accounting in making decisions. For instance, as demonstrated by Tversky and Kahneman (1981), one is more likely to buy a theater ticket for $10 after losing a $10 bill than to buy a ticket again after losing a $10 ticket. Presumably, the loss of a $10 bill is a goal-irrelevant event and thus a small loss in an individual's overall wealth. In contrast, the loss of a $10 ticket is related to the goal of seeing a play and thus a larger loss in the account that is set up along with the goal.

On the other hand, the lack of clear goals and prioritized cues in a decision problem may lead to the use of simple but irrational heuristics. To illustrate the point consider the framing effect demonstrated in the Asian disease problem by Tversky and Kahneman (1981). The observed reversal in risk preference as a result of ways in which a choice problem is presented violates the descriptive invariance principle of normatively defined rationality. As a matter of fact, this irrational reversal in risk preference can be attributed to the use of a simple heuristic which directs the decision maker to be risk-averse when choice outcomes are framed positively, but risk-seeking when the same outcomes are framed negatively.

Our studies (Wang 1996a; 1996b) suggest that the use of such simple but dumb heuristics is a result of the lack of clear task goals. This irrational framing effect disappears when the problem is presented in a kinship context or when the group size of endangered lives is reduced to a single- to two-digit number which approximates a typical group size throughout most of human evolution. The kinship cue and group size cue have made a sure outcome of saving one-third of endangered lives unacceptable and forced the decision makers to choose the gamble option irrespective of the hedonic cues of verbal framing. Consistent with the Take-The-Best heuristic, subjects' risky choice is determined by the most dominant decision cue whose presence overwhelms the secondary cue of outcome framing. Kinship, group size, and other social and biological concepts may help us understand goal setting and cue ranking in social decisions. Relying on valid cue ranking, fast and frugal heuristics are no longer quick and dirty expediencies but adaptive mental tools for solving problems of information search and goal conflict at risk.

Heuristics refound

William C. Wimsatt

Department of Philosophy and Committee on Evolutionary Biology, The University of Chicago, Chicago, IL 60637. **w-wimsatt@uchicago.edu**

Abstract: Gigerenzer et al.'s is an extremely important book. The ecological validity of the key heuristics is strengthened by their relation to ubiquitous Poisson processes. The recognition heuristic is also used in conspecific cueing processes in ecology. Three additional classes of problem-solving heuristics are proposed for further study: families based on near-decomposability analysis, exaptive construction of functional structures, and robustness.

This is an extremely important book. It could precipitate a second wave of interest in bounded rationality – delivering the revolution begun by Herbert Simon in the late 1940s with his "satisficing" model of decision making. Gigerenzer, Todd, and the ABC Research Group (1999) argue the natural affinity of their "fast and frugal" heuristics to new developments in social cognition; Simon's model came out of a study of corporate decision making in administrative structures – decisions which are both satisficed and distributed. Ecological rationality also figures in Simon's original work, and even more centrally in the present work. With its new tools, results, perspectives, and diverse applications, this book should give significant direction to that new wave, while confirming the original deep insights of Simon's first wave.

This is an *amicus curae* brief for this new program from a long-time devotee and user of Simon's work. Most central are the striking experimental results and simulations with the "recognition" and "take the best" heuristics, and their "ecological" analyses to determine the conditions under which they should be expected to work well, and when they should break down. These chapters are extremely well done. The robustness, economy, and simplicity of the heuristics under these conditions make them important, but the crisp analysis of their limits is also a methodological paradigm for future workers. The varied applications in Chapters 9 through 14 only begin to suggest the fruitfulness of this new paradigm. One of my favorites is the chapter on hindsight bias (Ch. 9). It is a powerful argument for the "fast and frugal" approach in that it yields such a revealing account of this ubiquitous and seemingly unrelated phenomenon – making it a signature of such methods.

The ecological rationality of these procedures noted in Chapter 10 deserves particular emphasis: the J-shaped distributions required for the robustness of these procedures should themselves be extremely robust: they are produced by power laws (emphasized by physicist Per Bak) and more generally by Poisson processes, which are ubiquitous in nature. Poisson distributions for objects and processes in nature do not guarantee Poisson distributions for cue validities (what else is required?), but are very suggestive, and worth more attention. If this connection can be fleshed out, it would be hard to imagine a more robust way of anchoring these procedures in a stochastic world.

The recognition heuristic is both inevitable and advantageous in another ecological situation: so called "conspecific cueing." Colonizing species or organisms with wide ranges may cover large amounts of territory in search of food, nest sites, or other resources. Because they would leave unsuitable areas quickly and

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

spend more time in suitable ones, the presence of a conspecific, when rare, is a good indicator of the suitability of an area. And most species already have and need means for recognizing conspecifics for mating, territory defense, migration, and other ends. Many use specific marking behaviors to indicate their presence, and can detect the age, sex, and other characteristics of other markers (Kiester 1979; Stamps 1987).

But there are other problem areas which could be fruitfully searched for "fast and frugal" strategies. The three classes of heuristics below are effective, simplifying, and extremely widely used. They would benefit from a similar analysis.

The first, "near-decomposability" heuristics, also received pioneering elaborations by Simon. They are at the core of reductionistic and "analytic" problem-solving methods in all areas (Simon 1996; Simon & Ando 1961; Wimsatt 2000a). One approaches a problem by trying to break it into sub-problems which can be solved independently and whose solutions can then be strung together as a solution, or the first try at one, for a more complex problem. It is crucial in the analysis of complex systems and the synthesis of composite systems, but also in planning for and executing any complex task – either an extended task by an individual or a coordinated plan (e.g., for hunting) by a group of individuals. Not all complex problems are solvable in this way (there may be complex interactions between subsystems which cannot be treated in this quasi-modular fashion), but the architecture of our own artifacts is likely skewed in this direction as a result of our own cognitive methods.

The second are "exaptive" heuristics. Gould and Vrba (1982) note that many things classified as adaptations were not originally created for the ends they now serve. Evolution, human engineering, science, and culture all systematically reuse constructs in new contexts that drive their elaboration in new directions. It is simply easier to take something which you already have and can make do (perhaps with simple modifications) for the new task. And it is exaptations all the way down. As a result, one cannot understand the functional organization of any complex system without considering its history. And some deep problems – the Y2K fiasco – can be traced to early trivial decisions whose consequences propagated so widely that their correction is immensely costly (Wimsatt 2000b).

A third class of "robustness" strategies use multiple presumptively independent means to locate something, triangulate on its properties, cross-check and calibrate the means of access, and more generally decide what is real and trustworthy and what is fleeting and artifactual. These strategies are widely used in our perceptual systems, and in scientific inference (Wimsatt 2000a).

Authors' Response

How can we open up the adaptive toolbox?

Peter M. Todd, Gerd Gigerenzer, and the ABC Research Group

Center for Adaptive Behavior and Cognition, Max Planck Institute for Human Development, 14195 Berlin, Germany
{ptodd; gigerenzer}@mpib-berlin.mpg.de
www.mpib-berlin.mpg.de/abc

Abstract: The adaptive toolbox is an evolutionarily inspired vision of the mechanisms of cognition, including simple decision making heuristics for specific problem domains. In *Simple heuristics* we showed how different heuristics in the adaptive toolbox could be constructed for different tasks, and how they could achieve ecological rationality (being accurate and robust) by exploiting the structure of information in the environment. Our commentators have raised a number of important challenges for further extending the study of ecological rationality. Here we summarize those challenges and discuss how they are being met along three theoretical and three empirical fronts: Where do heuristics come from? How are heuristics selected from the adaptive toolbox? How can environment structure be characterized? How can we study which heuristics people use? What is the evidence for fast and frugal heuristics? And what criteria should be used to evaluate the performance of heuristics?

In *Simple heuristics that make us smart* (Gigerenzer et al. 1999), we introduced the main concepts for studying the cognitive mechanisms that make up the *adaptive toolbox:* the idea of simple heuristics whose building blocks are precisely specified (simple rules for search, stopping, and decision), the way heuristics achieve ecological rationality by exploiting the structure of information in the environment, and how they can be accurate and robust through being fast and frugal. The adaptive toolbox is inspired by a Darwinian vision of decision making in humans, animals, and artificial agents. First, just as evolution does not follow a grand plan but results in a patchwork of solutions for specific problems, so the toolbox is structured as a collection of mechanisms that each do a particular job. Second, just as evolution produces adaptations that are bound to their particular context, the heuristics in the adaptive toolbox are not good or bad, rational or irrational, per se, but only relative to a particular environment. From these two features springs the potential power of simple heuristics: They can perform astonishingly well when used in a suitable environment.

The vision of the adaptive toolbox conflicts with several beautiful ideals, which some of the commentators share explicitly or implicitly. These ideals have their origins in a long philosophical tradition in which humans are recreated in the image of an omniscient God, or in a secularized version thereof, Laplace's superintelligence. The principles underlying the adaptive toolbox also have a long pedigree, from Darwin to Herbert Simon. We therefore begin by stating a few of the idealistic assumptions taken at face value by some of our commentators, and the opposing principles underlying the adaptive toolbox:

1. *More information is always better.* One-reason decision making is granted by **Sternberg** for everyday affairs but not for "consequential" decisions, conveying the implicit assumption that you are always better off using as much information as possible when time allows. The ideal that more information is always better is, however, misleading. Rather, *in order to make sound decisions in an uncertain world, one must ignore some of the available information.* The reason is that not all of the information available about one situation is useful for making judgments about another – a strategy that used *all* the available information would fall prey to overfitting and be unable to make robust generalizations. The trick – the job of frugal heuristics – is to ignore the *proper* pieces of information. **Engel** points out that frugality also conflicts with legal systems, which (like bureaucracies) often run on the defensive vision that more is always better.

2. *Optimization is always better.* **Shanks & Lagnado** imply that because human behavior can be optimal, fast and frugal heuristics, which do not optimize, cannot account for it. But the distinction must not be blurred between optimizing processes (e.g., computing the maximum of some function such as expected utility), which heuristics do not employ, and optimal outcomes, which heuristics *can*

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

nonetheless reach. Moreover, using optimization does not guarantee an optimal outcome for a given situation, because the choice of optimizing process must often be made based on uncertain simplifying assumptions. Similarly, the optimal strategy in a particular domain will usually not generalize to being optimal in a different domain, because of the particular assumptions on which optimization must be based. The considerations of speed, simplicity, robustness, and psychological plausibility can add to make heuristics a better choice in particular situations.

3. *Complex environments demand complex reasoning strategies.* **Allen** suggests that social environments, being "responsive" rather than "passive," are so complex and quixotic that they require demonic reasoning abilities (cf. the Machiavellian intelligence hypothesis on the social evolution of intelligence; Whiten & Byrne 1997), or at least the application of logic and probability theory. In particular, the assumption is that if you do not reason logically, you can be exploited by others. But simple social exchange rules (e.g., cheater detection; see Cosmides & Tooby 1992) can coordinate social interactions without logical reasoning. Conversely, following logic can make one predictable and hence open to exploitation (Frank 1988), making "illogical" protean behavior more adaptive in many situations (Driver & Humphries 1988).

4. *Search can be ignored.* **Oaksford** argues that information is usually integrated in decision making by pointing to examples (such as speech perception and sentence processing) where the necessary information is available simultaneously, obviating search. But integration seems less universal when one recognizes that many decision situations require cues to be searched for, whether internally in memory or externally in the world. In such cases, there is evidence for fast and frugal mechanisms that stop information search as soon as possible, coupled with decision rules that do not integrate information (see sect. 5). (This may also apply in some of the language processing examples that Oaksford draws upon, according to optimality theory's lexicographic search for rules; see Prince & Smolensky 1997.) By ignoring the need for search, we can easily overlook those situations in which information is not integrated – one of the territories of simple heuristics.

Simple heuristics is a progress report of our first three years of studying the adaptive toolbox. As many of the commentators agreed, there is a need for a new alternative to rational choice theory and other demonic visions of decision making in fields ranging from primatology, to cognitive psychology, to philosophy, and we are grateful to the commentators for pointing out the important open questions in this program and suggesting some possible answers. We have organized this reply around six open questions raised by the commentators. The three major theoretical questions are: Where do heuristics come from? How are heuristics selected from the adaptive toolbox? How can environment structure be characterized? The three methodological and empirical questions are: How can we study which heuristics people use? What is the evidence for fast and frugal heuristics? What criteria should be used to evaluate the performance of heuristics?

Before we address these pressing questions, we want to clarify one issue. We called the book "Simple heuristics *that* make us smart" rather than "Simple heuristics *make* us smart" for a good reason. We do not believe, and do not want to be misconstrued as saying, that the set of simple heuristics and the set of things that make us smart are identical. Not every reasoning task is tackled using simple heuristics – some may indeed call for lengthy deliberation or help from involved calculations. Conversely, not every simple strategy is smart. Thus, the key question organizing the overall study of the adaptive toolbox is that of ecological rationality: Which structures of environments can a heuristic exploit to be smart?

R1. Where do heuristics come from?

Several commentators have called for clarification about the origins of heuristics. There are three ways to answer this question, which are not mutually exclusive: heuristics can arise through evolution, individual and social learning, and recombination of building blocks in the adaptive toolbox.

R1.1. Evolution of heuristics

Ecologically and evolutionarily informed theories of cognition, are approved by **Baguley & Robertson** but they express concern over how the adaptive toolbox comes to be stocked with heuristics. The book, they say, sometimes leaves the reader with the impression that natural selection is the only process capable of adding a new tool to the toolbox, whereas humans are innately equipped to learn certain classes of heuristics. We certainly did not want to give the impression that evolution and learning would be mutually exclusive. Whatever the evolved genetic code of a species is, it usually enables its members to learn, and to learn some things faster than others.

For some important adaptive tasks, for instance where trial-and-error learning would have overly high costs, there would be strong selective advantages in coming into the world with at least some heuristics already wired into the nervous system. In *Simple heuristics,* we pointed out a few examples for heuristics that seem to be wired into animals. For instance, the recognition heuristic is used by wild Norway rats when they choose between foods (Ch. 2), and female guppies follow a lexicographic strategy like Take The Best when choosing between males as mates (Ch. 4). A growing literature deals with heuristics used by animals that are appropriate to their particular environment, as **Houston** points out. This includes distributed intelligence, such as the simple rules that honey bees use when selecting a location for a new hive (e.g., Seeley 2000).

Hammerstein makes the distinction between an *optimizing process* and an *optimal outcome* very clear. He discusses the view generally held in biology of how decision mechanisms including simple heuristics can arise through the optimizing process of evolution (often in conjunction with evolved learning mechanisms). This does *not* imply that the heuristics themselves (nor specific learning mechanisms) are optimizing, that is, that they are calculating the maximum or minimum of a function. But they will tend to be the best alternative in the set of possible strategies from which evolution could select (that is, usually the one closest to producing the optimal *outcome*). As Hammerstein says, biology meets psychology at this point, because investigations of psychologically plausible – and ecologically rational – decision mechanisms are necessary to delineate the set of alternatives that evolutionary selection could have operated on in particular domains.

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

Hammerstein also reminds us that the absence of optimality in observed behavior does not imply that evolution had no role in building the heuristics involved; instead, a heuristic that was evolved for adaptive use in one environment may be misapplied in another. An example is decision heuristics that appropriately neglect base rate information in some circumstances (such as in changing environments; see Goodie & Todd, submitted), but which are tripped up by this simplicity in other situations (for instance ignoring base rates when they are stable and useful). Here again psychologists can work together with biologists to understand behavior that might otherwise be relegated to the category of irrational biases.

R1.2. Learning and development of heuristics

R1.2.1. Individual learning. Certain classes of heuristics can be learned, state **Baguley & Robertson.** Certainly people can learn new problem-solving techniques, but there is a distinction to be drawn between adding a new component to the adaptive toolbox and learning to use the tools already there in new ways (parallel to **Wimsatt**'s evolutionary example of exaptations as borrowed adaptations). **Solomon**'s commentary is useful in this regard: he points out that since Piaget, the search for domain-general principles of reasoning and decision-making in children – and the cataloging of their "failures of logic" – has been eclipsed by interest in domain-specific reasoning principles that children use in particular contexts or environments. Two general conclusions that can be drawn from the developmental literature to which Solomon refers are (1) that young children, and even infants, apply different reasoning principles in different problem domains (the infant's adaptive toolbox already contains multiple tools), and (2) that children know when and how to apply these tools from a very early age (see, e.g., Baillargeon 1986; Csibra et al. 1999; Spelke 1990).

Solomon distinguishes two kinds of developmental change that map onto the distinction mentioned above between adding new tools and using old tools in new ways: change in the core "theories," concepts, or reasoning principles deployed in a particular domain – what might be called the tools in the adaptive toolbox (see, e.g., Carey 1985; Gopnik & Meltzoff 1997) – and change in "expertise" with age, that is, changes in the stores of data upon which particular tools can be brought to bear, or in how and when the tools are used. An important area to be explored is thus the developmental trajectories of different domain-specific heuristics. For example, Klayman (1985) showed that 12-year-olds tend to use more frugal strategies in complex search tasks than adults, though many of the strategies used by adults are present by age 12.

Simple heuristics did not much address individual learning and development. But this does not mean that the notion of the adaptive toolbox should be seen as irrelevant for developmental researchers. On the contrary, as **Solomon** points out, the adaptive toolbox perspective raises a set specific developmental questions: What heuristics come preloaded into the adaptive toolbox? Which ones drop out over time, which are added, and which are modified? Can the change in heuristic use over time be matched to changes in the environment facing growing children? What heuristic principles for search do children use (e.g., random search, ordered search), when do they stop search, and how do they arrive at a decision? And how do children acquire new heuristics?

R1.2.2. Social learning. Primatologists **Barrett & Henzi** wonder whether fast and frugal heuristics are the result of our primate heritage, or even a specifically human form of reasoning not available to other primate species. Although such an extreme position is almost certainly untenable (as **Houston** indicates, biologists have long been satisfied with the idea that animals of all phyla use simple decision-making rules), humans probably *are* extreme in our reliance on social learning of heuristics. This enables much more rapid increase in the contents of the adaptive toolbox and provides a foundation for culture.

Gorman points out that there are many domains of human endeavor, including scientific discovery, where heuristics can be obtained from other individuals, but wonders whether heuristics might be applicable in the search for whom to borrow from as well. In this line, **Goldman** also describes how the decisions themselves (not just the strategies) can be obtained from others: For instance, it can make sense to copy someone with similar political views but more information about the candidates, and just vote for whomever she chooses. There is a growing body of work, primarily covering other species, that explores the simple mechanisms that can be used for these sorts of social information exchange (Noble & Todd, in press). This research has focused on the ways that imitation-like behavior can be achieved without the necessity of the imitator modeling the intentions and actions of the imitated individual; indeed, what looks like imitation can often be the result of one agent merely following another around and being prompted by the same environmental cues. Explicit language-based instruction and information exchange in humans is obviously more involved than this, but the processes underlying *whom* to listen to – for instance, how to find experts, as Gorman says, or core voters, in Goldman's example – may well be guided by fast and frugal mechanisms. Important links to existing research in social psychology and sociology can be made here.

As **Gorman** discusses, scientific discovery and theory-building is a special case of social information exchange, where the environment is the cultural one of other scientists and existing theories. He addresses the role of heuristics in discovery, in particular heuristics for guessing laws from data as in the tradition of the computer program BACON. While we have not included this topic in *Simple heuristics,* some of us have analyzed heuristic processes in scientific discovery in earlier work, focussing on heuristics for finding analogies. For instance, the tools-to-theories heuristic (Gigerenzer 1991; 1994) applies to cases where new theories are inspired not by new data (as is usually thought to be the case), but rather by new research tools, such as statistical methods (Gigerenzer & Murray 1987) or computers (Gigerenzer & Goldstein 1996b). This heuristic differs from those implemented in the BACON tradition in that it does not model discovery as an inductive process, but as a projection of laboratory practice into theories of mind (e.g., Neyman-Pearson's hypothesis testing methods inspiring Tanner and Swets' 1954 theory of signal detectability).

R1.3. Construction of heuristics from building blocks

In *Simple heuristics* we discuss how heuristics can be built from simpler components, whether these are building

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

blocks for guiding information search, stopping search, and making a decision, or whole heuristics combined into composite decision mechanisms. An example of the latter combination is the use of the recognition heuristic as the first step in Take The Best. We expect such nested heuristics themselves to be used in environments to which they are matched, so that for instance the recognition heuristic/Take The Best combination would not be employed in an environment where recognition validity is too low (see the discussion of **Erdfelder & Brandt**'s simulations in sect. R2.3).

Indeed, **Huber** feels that the contents of the adaptive toolbox are more likely to be what he calls "partial" heuristics than the whole or "global" heuristics usually proposed – that is, he thinks most decision making is based on using (parts of) one heuristic first, and then switching to another. He argues that "the investigation of global heuristics is not a fruitful research strategy, because we know already that people do not use them." But in arguing this way, Huber overgeneralizes from the artificial lottery problems of risky choice that many decision experiments rely on. The fact that more than 90% of the participants in the studies reported in Chapter 2 made inferences following the recognition heuristic, and that 65% of the participants in another study made inferences following Take The Best (Bröder, in press; Experiment 4) are just two examples contradicting Huber's bold assertion. Furthermore, Huber's assertion that people use partial rather than global heuristics sets up a false opposition: both can sit comfortably in the adaptive toolbox. As mentioned above, heuristics are composed of building blocks, and heuristics themselves can be combined to form new heuristics. The possibility of combining parts into new strategies is central to the toolbox metaphor, but not to the related image of the mind as a Swiss army knife (Cosmides & Tooby 1992). Both views emphasize the existence of many domain-specific strategies (rather than a single general-purpose calculus); but the toolbox highlights the possibility of recombining tools and nesting heuristics.

For instance, in an environment where a person has an intuition about the ranking of cues based on validity, she can use Take The Best and search for cues in that order; but in another environment where cue validities are unavailable, the search rule can be changed accordingly to build a different heuristic that works with less knowledge, such as Take the Last. Similarly, if a decision is highly consequential, different building blocks may be combined. Consider **Sternberg**'s dilemma whether to accept an offer from a rival university or stay at Yale. Here, search for information is likely to be prolonged – that is, a different stopping rule may be used, such as finding all the possible relevant cues within a two-month decision period. This extended search, however, may lead to the question of how to make a trade-off between the reasons found, such as the intellectual quality of a department and the tears of one's children who do not want to move. Not all things have a common denominator (see **Morton**'s discussion of incomparability). As a consequence, an extended search may still be combined with a quick and simple decision rule, relying on the most important cue instead of exchanging all cues into a common currency and weighting and combining them. Maybe Sternberg can let us know how he made up his mind.

R1.4. Learning of Cues

It is not enough to ask where heuristics come from; we must also address the question of how to find the relevant cues, as **Wallin & Gärdenfors** point out. In general, the same sorts of possibilities exist for an organism to determine what cues to use as what heuristic to use. First, cues can be genetically encoded, as in the human reliance on bitterness as a cue to unpalatability. Second, appropriate cues can be learned through individual experience. Third, cues can be picked up socially, for instance by copying the decision strategies of others.

Knowing the appropriate cues to use in an environment and the direction of their association with a decision variable may be all that a heuristic has to go on, but this can be enough. The sufficiency of this minimum knowledge stands in contrast to **Wang**'s comment that "these fast and frugal heuristics would not work effectively if one did not know the priority ranking of relevant cues." Here is a misunderstanding that it is important to clear up: Not all heuristics need a ranking of cues according to their validity. Some do, such as Take The Best and Categorization by Elimination, while others do not, such as Take The Last and Minimalist. To repeat, the minimal knowledge needed for cue-based inference is only which cues to use and the direction in which each cue points, for instance, whether having a soccer team in the major league indicates that the associated city has a large or small population (Ch. 4). This direction can be estimated by simply counting within a small learning sample.

The Minimalist heuristic in fact knows nothing more than this, and thus has no idea about which cues are better predictors than others. Consequently, the only building block for search that Minimalist can use is to look up cues in random order. In contrast to **Wang**'s intuition, we found that Minimalist can compete well with other algorithms in terms of accuracy, particularly when knowledge is scarce (Chs. 4 and 5). A truly minimal amount of learning (or inherited knowledge) can thus be adequate to get a simple heuristic off the ground in a new domain. With a little more learning, frugal methods for ranking cues according to estimates of their validity can be used (Chs. 4 and 6). Note that these methods do not necessarily result in the "optimal" ranking of cues, but yield an ordering that is good enough and robust enough to generalize to new situations.

R2. How are heuristics selected from the adaptive toolbox?

This question has been asked by many commentators (e.g., **Cooper, Luce, Wallin & Gärdenfors**), and **Morton** has proposed an answer. In *Simple heuristics,* we spent little more than one page (pp. 32–33) on the issue of heuristic selection; here, with the help of our commentators, we can deal with this question a bit more systematically.

R2.1. How pressing is the problem of heuristic selection?

First, we should point out that heuristic selection may not always be a problem. As **Feeney** indicates, there are situations in which the need for selecting a heuristic does not even arise, for example when the use of a particular heuristic has been hardwired in a domain-specific (but possibly

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

environmentally contingent) way. When there *is* more than one available heuristic, the choice set of possibilities may still be small. One reason for this is that the heuristics in the adaptive toolbox are designed for specific tasks – like screwdrivers and wrenches, they are not universal tools. This specificity goes a long way to reduce the selection problem. For instance, when a task calls for estimating a quantitative variable, the QuickEst heuristic is a candidate, but others such as Take The Last and Categorization by Elimination, designed for choice and categorization tasks, will not be in the choice set. A second factor that reduces the set of possible heuristics to choose between is the presence or absence of particular knowledge on the part of the decision maker. For instance, Take The Last and Take The Best are both designed for the same type of choice task, but if a person has no intuition about which of several cues should be checked first (i.e., no validity-based ranking of the cues; see sect. R1.4), then Take The Best cannot be used.

R2.2. Fast and frugal selection of heuristics

Even after these task- and knowledge-specific reductions of the choice set, there may still remain a number of heuristics that are applicable for a given situation. How then can we choose between them? **Morton** suggests an answer consistent with the notion of an adaptive toolbox: a meta-heuristic which chooses between heuristics using the same principles as the fast and frugal heuristics themselves. For instance, just as Take The Best looks up cues in a particular (validity-based) order, a meta-heuristic can try heuristics in a particular (e.g., success-based) order. Furthermore, just as the cue order that Take The Best follows is not arrived at by optimizing computations – and is good enough rather than optimal – the ordering of heuristics is not achieved by optimizing calculations either, but by using simple and robust criteria, such as past success. The process of ordering itself can be modeled by a simple reinforcement learning mechanism like that described by Erev and Roth (2000). Morton's solution avoids what **Feeney** fears: an infinite regress of meta- and meta-meta-strategies that would be needed to compute the best of all heuristics and all meta-heuristics for each situation given all constraints (as in optimization under constraints). As soon as one dispenses with the ideal of finding the very best heuristics for each situation, the infinite regress that would burden an optimization approach does not arise.

Thus, **Cooper**'s and **Feeney**'s further worry that the meta-heuristics will not pick the best heuristic to use could certainly be true. But the whole point of the adaptive toolbox approach is *not aiming* at optimization. We must not let optimal requirements sneak in at the meta-level either. Moreover, because there is often more than one heuristic that can perform well in a particular situation – the flat maximum phenomenon – the choice between them may not always be critical, certainly not worth pondering an eternity over.

R2.3. How environments can select heuristics

According to **Cooper,** we have not specified the conditions that select particular heuristics. It is fair to say that most heuristic-relevant conditions must still be discovered, but we and others (e.g., Payne et al. 1993) have already filled in some examples. For instance, there are two conditions that are necessary and sufficient to make the recognition heuristic useful: recognition must be correlated with a criterion (e.g., recognition of city names is correlated with their size), and recognition must be partial (e.g., a person has heard of some objects in the test set, but not all). More formally, this means that the recognition validity α is larger than .5 and the number n of recognized objects is $0 < n < N$, where N is the total number of objects. Note that these variables, α, n, and N, can be empirically measured – they are not free parameters for data fitting.

When these conditions do not hold, using the recognition heuristic is not appropriate. **Erdfelder & Brandt,** for instance, overlooked the first condition and thereby misapplied the recognition heuristic. They tested Take The Best with the recognition heuristic as the initial step in a situation with $\alpha = .5$ (i.e., recognition validity at chance level). In this case, the recognition heuristic is not a useful tool, and Take The Best would have to proceed without it. Their procedure is like testing a car with winter tires when there is no snow. Nevertheless, their approach puts a finger on one unresolved question. Just as a driver in Maine faces the question when to put on winter tires, an agent who uses the Take The Best, Take The Last, or Minimalist heuristic needs to face the question of when to use recognition as the initial step. The recognition validity needs to be larger than chance, but how much? Is this decided on an absolute threshold, such as .6 or .7, or relative to what a person knows about the domain, that is, the knowledge validity β? There seems to be no empirical data to answer this question. Note that this threshold problem only arises in situations where objects about which an agent has not heard (e.g., new products) are compared with objects about which she has heard *and* knows some further facts. If no other facts are known about the recognized object, then the threshold question does not matter – when there is nothing else to go on, just rely on recognition, whatever its validity is.

In contrast, one condition that might seem important for the applicability of the recognition heuristic is actually irrelevant. **Goldman** says that "only with a (tolerably) accurate estimate" of the recognition validity will applications of the recognition heuristic succeed. This is not necessary – the recognition heuristic can work very well without any quantitative estimate of the recognition validity. The decision-maker just has to get the direction right. Goldman also suggests that the less-is-more effect in the example of the three brothers (Ch. 2, pp. 45–46) would be due to the fact that the middle brother knows about the validity of recognition while the eldest does not. As Figure 2-3 shows, however, the less-is-more effect occurs not only when one compares the middle with the eldest brother, but also when one compares the middle with other potential brothers (points on the curve) to the left of the eldest. Thus, the less-is-more effect occurs even when the agent who knows more about the objects also has the "substantial problem-specific knowledge," that is, the intuition that recognition predicts the criterion. Less is more.

Margolis raises the fear that environments may prompt the use of a heuristic that proves maladaptive, especially with genetically entrenched – evolved – heuristics that are resistant to modification via learning. His argument is that when environments change in a significant way (e.g., from scarce to abundant resources), heuristics adapted to past

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

conditions may prove harmful in the new state of affairs. This discrepancy can become amplified when one looks beyond the individual to aggregate social choice. Margolis's reasoning here is akin to that of Robert Ornstein and Paul Ehrlich's *New world, new mind* (1989). In this book, the argument is made that the human mind fails to comprehend the modern world, because we are more impressed by unique events and dramatic changes than by the more threatening slow-motion disasters such as the greenhouse effect, AIDS, and dwindling natural resources. The evidence for this view, however, is contested, as a widely publicized wager between Paul Ehrlich and the economist Julian L. Simon illustrates (Tierney 1990). In 1980, Ehrlich bet that the prices of five metals (which he selected) would go up within ten years because of the exploitation and exhausting of natural resources; Simon bet that they would go down because of human innovation. Simon won. Similarly, the maladaptively entrenched heuristics Margolis points to may be replaced by new more appropriate heuristics acquired through individual or social innovation and learning.

R3. How can environment structure be characterized?

The structures of environments are essential for understanding cognition and behavior in terms of adaptation, because adaptations are shaped by (past) environments. Fast and frugal heuristics that are matched to particular environmental structures allow organisms to be ecologically rational. To understand their performance, one needs conceptual languages both for the design of heuristics (e.g., rules for search, stopping, and decision) and for the structure of environments. The two classic behaviorist-inspired approaches to studying the structure of environments, Egon Brunswik's search for environmental or ecological texture and J. J. Gibson's search for invariants in the visual environment, paid little attention to heuristics or anything else going on in the mind, and therefore did not analyze the match between heuristics and environments (see Gigerenzer 2000). More recently, rational analysis (Anderson 1990; see also **Chater**), Hutchins's (1995) study of artifacts (e.g., maps, navigation instruments), and Norman's (1993) work on things that make us smart, among other research directions, epitomize a growing emphasis on the importance of environmental structure.

In *Simple heuristics*, we investigated the following types of environment structures and analyzed the degree to which specific heuristics can exploit them: environments in which lack of recognition is informative (Ch. 2); noncompensatory information (Ch. 6); scarce information (Ch. 6); J-shaped distributions (Ch. 10); and decreasing choice sets (Ch. 13). (Note that we are concerned here with the structure of the environment as it is known by the decision maker, which of course is strongly tied to the "objective" environment structure.) Clearly, this is only a beginning. We are happy that several commentators have added to the catalog of environment structures under investigation.

R3.1. Cost-benefit structure of the environment

Commentator **Allen** suggests that heuristics may not be used in situations in which the cost of being wrong exceeds the benefit of being right. In such cases, he expects people will aim to maximize their chance of producing correct responses, and therefore abandon simple heuristics for more normative or deliberative approaches. **Sternberg** suggests that heuristics can be used to good advantage in a range of everyday decisions but are inappropriate for consequential (i.e., important) real-world decisions, and **Harries & Dhami** voice similar but more prescriptive concerns for legal and medical contexts. All this sounds plausible, but what is the evidence? As a test of this plausible view, Bullock and Todd (1999) manipulated what they called the *significance structure* of an environment – the costs and benefits of hits, misses, false alarms, and correct rejections in a choice task. Against expectations, they found that lexicographic strategies can work well even in high-cost environments. They demonstrated this for artificial agents foraging in a simulated environment composed of edible and poisonous mushrooms. When ingesting a poisonous mushroom is not fatal, the best lexicographic strategies for this environment use as few cues as possible, in an order that allows a small number of bad mushrooms to be eaten in the interest of not rejecting too many good mushrooms. When poisonous mushrooms are lethal, lexicographic strategies can still be used to make appropriate decisions, but they process the cues in a different order to ensure that *all* poisonous mushrooms are passed by.

One feature of consequential decisions is that they frequently have to be justified. Often a consequential decision is reached very quickly, but much further time in the decision process is spent trying to justify and support the choice. Such cases illustrate that consequential environments may call for long and unfrugal search that is nonetheless combined with fast and frugal decision making, as discussed in section 1.3 (see also Rieskamp & Hoffrage 2000; Ch. 7).

R3.2. Friendly versus unfriendly environments

The cues in an environment (as the decision maker knows it) can all be positively correlated, or a subset of them can be negatively correlated. **Shanteau & Thomas** call the first type of environment "friendly" and the second type "unfriendly." We welcome Shanteau & Thomas's work to extend the understanding of the ecological rationality of heuristics to this type of environment structure. They have found that unfriendly environments contain tradeoffs that present a challenge to fast and frugal heuristics, particularly those that employ one-reason decision making. One way of meeting this challenge is to combine partial heuristics (as **Huber** discusses), starting with an elimination strategy to remove options with unacceptable tradeoffs, and proceeding to a lexicographic strategy to process the remaining options.

More specifically, **Shanteau & Thomas** state that the performance of the LEX heuristic decreases in unfriendly environments as the number of cues increases. But it is important to realize that there is a fundamental asymmetry between friendly and unfriendly environments. An environment with many cues can be extremely friendly, but not extremely unfriendly. That is, all correlations between cues can be positive, but not all can be negative. To see the point, imagine two cues which are correlated at $r = -1$; a third cue which is negatively correlated with one of the first two cues must be positively correlated with the other. In other words, unfriendly environments will have pockets of friendly cues, and the more cues the decision maker as-

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

sesses, the bigger those friendly groups will be. How heuristics can exploit these pockets is a question that should be explored.

We also want to point out the importance of performance measures used to compare strategies in these and other environments (see also sect. R6). **Shanteau & Thomas** measure the performance of their strategies against a weighted additive difference (WAD) model as the gold standard, as in earlier work on preferences by Payne et al. (1993). In *Simple heuristics,* in contrast, we have used real-world criteria – from attractiveness judgments to school drop-out rates – as the gold standard. This corresponds to the difference between preferences (where there is no external criterion) and inferences (where there is one). This choice has consequences. When WAD is employed as the performance standard (rather than an external criterion), a simple heuristic such as a lexicographic strategy can never be more accurate than WAD, by definition. But when real-world criteria are introduced, heuristics can outperform WAD, as we have shown in Chapters 4 and 5. Specifically, the gold standard is likely to be outperformed when the task of an agent is prediction (e.g., cross-validation) rather than mere data fitting. Thus, we are curious to know how accurate lexicographic heuristics are in unfriendly environments when they have to predict features of the environment (rather than match the outputs of a WAD model). Preliminary results obtained by the ABC Research Group indicate that in unfriendly environments, Take The Best can be more accurate than Dawes's rule (consistent with Shanteau & Thomas), but can also match or slightly outperform what we call Franklin's rule, a weighted additive model.

R3.3. Conspecifics are the environment

One important field for the study of the adaptive toolbox is social rationality, the investigation of cognitive mechanisms for dealing with an environment consisting of conspecifics. We agree with **Barrett & Henzi** that extending the simple heuristics research program further into the social domain is an important next step. This would help bring psychology and game theory together (as **Hammerstein** calls for) and emphasize the importance for social psychology of interactive strategies – the heuristics individuals actually employ to deal with others. We have begun to explore the ways in which people and animals can use simple rules to process information they gain by observing the behavior of conspecifics (Noble & Todd, in press; see also **Wimsatt**'s example linking the recognition heuristic with "conspecific cueing" in species where the rare sight of another conspecific is a good cue for the presence of resources). **Allen** and **Barrett & Henzi** raise the possibility that in the social realm, with selection for increasingly sophisticated Machiavellian behavior, simple heuristics may eventually fall prey to exploitation by more complex strategies. As mentioned earlier, however, the general assumption that increasingly complex environments can only be tamed by increasingly complex strategies is doubtful. Counterexamples are simple social heuristics such as Tit-For-Tat and its relatives, the "automaticity" of everyday life (Bargh & Chartrand 1999) and the advantage of simple and transparent rules in businesses and legal systems as opposed to cancerous growing bureaucracies (Engel 1994; Epstein 1995).

R3.4. Poisson processes

In Chapter 10, Hertwig et al. showed how the QuickEst heuristic can exploit J-shaped distributions in environments. But which environments have such structure? **Wimsatt** points out that J-distributions are likely to be very commonplace because they are produced by power laws, and more generally by Poisson processes, which are ubiquitous in the natural world. For instance, he directs us to the work of Bak (1996) who explains seemingly disparate phenomena including the formation of landscapes, the regularity of catastrophic events such as earthquakes, and the behavior of economic systems as manifestations of one principle, *self-organized criticality,* which produces power-law distributions. Whether or not one agrees with this view, Bak's work points to the questions of what kind of mechanisms and processes produce environmental structures. Linking an understanding of the processes that generate environment structure to the mechanisms that people and other animals use to understand them is an exciting challenge.

R3.5. Causal cues versus surface cues

What kind of environmental cues do heuristics use? **Wallin & Gärdenfors** propose that "ecological validity should be seen as only a secondary effect of the fact that a decision maker aims at forming hypotheses about *causal connections* between the cues and the decision variable." The implication is that causally related cues should be sought rather than merely ecologically valid cues. We disagree with relegating ecological validity to this secondary role. Causal variables may well be used as cues, if available. However, we suspect that it may be easier to assess and monitor ecologically valid cues that are covarying at the surface level with the decision variable (e.g., because both are caused by the same underlying process) rather than cues that are causally linked to the decision variable. Moreover, several causal cues may need to be taken into account to reach a decision, when only one surface-level concomitant cue might suffice. For instance, consider a baseball or cricket player who tries to catch a ball. If the player tried to compute the spot where the ball would land by relying on causal variables alone, he would have to estimate the initial velocity, initial speed, spin, air resistance, direction and intensity of wind, and a myriad of other causal factors, and then apply the proper mathematics to integrate all this information. Real players, in contrast, seem to rely on one-reason decision making, and the one reason is not a causal cue. It is the angle of gaze, that is, the angle that their eyes have to keep to stay directed at the ball. The simple heuristic is to start running and to adjust running speed so that the angle of gaze is constant (or within a certain range, see McLeod & Dienes 1996). This gaze heuristic illustrates that a smart heuristic can ignore all causal factors, and gain accuracy through its frugality.

R4. How can we study which heuristics people use?

Obtaining empirical evidence for the use of particular heuristics demands careful methodology, because of challenges such as the flat maximum phenomenon (Ch. 7). **Cooper** doubts whether this is even possible, feeling that we have presented an unfalsifiable theory of decision mak-

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

ing, akin to the heuristics-and-biases program of Kahneman and Tversky. This is incorrect – here is an example. One of the strongest tests of a heuristic is obtained when it makes a bold and new prediction. Consider the recognition heuristic, which predicts the counterintuitive less-is-more effect – that is, once a person has a critical degree of knowledge, more knowledge will produce less accurate inferences. The necessary and sufficient condition for the less-is-more effect is that a person's recognition validity is larger than her knowledge validity (i.e., $\alpha > \beta$, where both are empirically measurable variables, and no free parameters are involved; see Ch. 2). Thus, if one experimentally induces (or empirically observes) this condition, then people who use the recognition heuristic will exhibit the less-is-more effect. If the effect is not seen, then the hypothesis that participants use the recognition heuristic in this situation is falsified.

This less-is-more effect has not been repeatedly experimentally demonstrated (Goldstein & Gigerenzer, in press). There are also simpler ways to test whether people use the recognition heuristic that do not involve such bold predictions (e.g., setting up conditions as in sect. R2.3). In general, testing hypotheses about heuristics calls for a methodology that differs from the dominant approach in experimental psychology: the comparison of group-means by null hypothesis testing.

Hypothesis testing in much of experimental psychology follows a standard paradigm: vary an independent variable, or a few, and test whether the resulting group means in a dependent variable differ significantly from the null hypothesis ("chance"). This paradigm, applied to the study of reasoning strategies or heuristics, has two dramatic flaws. First, no hypotheses about mental strategies, nor their predictions, are specified, only a null hypothesis. Second, the means analyzed are aggregated across participants. Doing this, one assumes a priori that people will not use different heuristics, and researchers will hardly detect them if they do. The flaws of this methodology are well documented (e.g., Gigerenzer 1993).

The metaphor of the adaptive toolbox, in contrast, encourages a methodology that is sensitive to the existence of multiple heuristics and individual differences in their use. Such a methodology can consist of (1) specifying multiple candidate heuristics, (2) deriving the predictions for each heuristic on the experimental tasks, and (3) testing each participant's judgments against the predictions of each heuristic (e.g., Ch. 7; Gigerenzer & Hoffrage 1995). The result may be that 60% of the participants use heuristic A, 20% use heuristic B, 10% use C, and the rest use idiosyncratic strategies or cannot be categorized. Note that this methodology avoids averaging across individuals, which is unjustified when multiple strategies are in use in a population – a situation that we have found to be the rule rather than the exception. This methodology is particularly important when one analyzes ontogenetic development in heuristic use, as **Solomon** proposes. In earlier work, some of us have identified half a dozen heuristics children use for estimating size, for instance of chocolate bars (Gigerenzer & Richter 1990). These heuristics changed systematically with age, with the most frequent ones in young children dropping out later while others were added. At each age, from pre-schooler to adult, there was not one strategy used across individuals, but several.

This methodology is all too rarely used. For instance, in the *BBS* target article on individual differences in reasoning by Stanovich and West (2000), no models of reasoning heuristics were formulated, no competing predictions derived, and the analysis of individual differences was performed exclusively on whether participants' judgments deviated from some norm. This procedure can obscure evidence for the use of particular heuristics, for instance when one leads to behavior close to the norm in some situations but far from the norm in others and thus appears inconsistent (Hoffrage 2000). Unless a set of hypotheses for participants' strategies is formulated in the first place, there is little hope of identifying individual differences in heuristic use.

Thus, the methodology of the adaptive toolbox encourages precise predictions, including predictions of individual differences in performance based on individual differences in knowledge, and, in some lucky cases, bold and counterintuitive predictions. For instance, the recognition heuristic makes precise predictions of choices for each pair of objects, predicts systematic individual differences in choice on the basis of individual differences in recognition, and even indicates when the less-is-more effect might be observed. Each of these predictions can be, and has been, experimentally tested.

Chater discusses how rational analysis (Anderson 1990) can complement the study of fast and frugal heuristics: It can provide a framework to evaluate the performance of simple heuristics. (Rational analysis is thus related to the use of optimality models in biology as **Houston** discusses.) We agree. The two programs approach the first step in the above methodology – specifying multiple candidate heuristics – from two different ends. Rational analysis starts by trying to find the optimal solution to some particular problem, and then begins whittling that down through successive simplifications – all mathematically derivable and testable – until finally a psychologically plausible mechanism is found. In contrast, the program of fast and frugal heuristics starts with building blocks that people may actually use – like recognition, or forming aspiration levels – and looks at how these can be combined into good enough solutions.

R5. What is the evidence for fast and frugal heuristics?

It is widely accepted that animals use simple heuristics (more commonly called "rules of thumb") that are successful and well-adapted to their particular environments, as **Houston** points out and as we discuss in Chapters 13, 14, and 15. Here we focus on the more contentious human case. **Cooper** and **Harries & Dhami** claim that there is little or no evidence that humans use simple heuristics. However, Chapters 2, 7, and 9 of *Simple heuristics* do provide just such empirical evidence. Nevertheless, we agree that more evidence is needed, and we are happy to report here some of the efforts in this direction.

We are not the first to find evidence that people employ limited search, fast stopping rules, elimination heuristics, or one-reason decision making. For example, it has long been known that people often look up only one or two relevant cues, avoid searching for conflicting evidence, ignore dependencies between cues, and use non-compensatory strategies (e.g., Einhorn 1970; Einhorn & Hogarth 1981, p. 71; Fishburn 1988; Hogarth 1987; Payne et al. 1993; Shep-

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

ard 1967). But these "limitations" have been tossed too quickly into the box of human irrationality. For instance, ignoring cue dependencies is usually suspected as a "failure" of the human mind. But we have found that in suitable environments, one-reason decision heuristics that ignore all dependencies can be more accurate than linear strategies that carefully estimate the dependencies (Ch. 5). It may thus be profitable to return to and reinterpret some of the important results of earlier decision making research in the light of the ecological rationality perspective.

Since the publication of *Simple heuristics,* further tests of fast and frugal decision mechanisms have been performed. Bröder (in press, Experiment 4) used an ingenious experimental design to test whether people use Take The Best. He showed that when participants have to search for costly information, 65% of all participants were classified as using Take The Best. In contrast, less than 10% could be classified as using Dawes's rule, a simple linear model with unit weights. Further evidence has also been found for the earlier hypothesis of Payne et al. (1993) that lexicographic strategies are employed when time is short (Rieskamp & Hoffrage 2000). Finally, **Harries & Dhami** conclude from their work with English magistrates and other experts that fast and frugal heuristics could underlie some decisions in legal and medical domains, even when they are prescribed against by the stated goals of those professions.

Some heuristics can produce intransitive judgments. **Allen** notes that we did not provide evidence that people in fact make such inconsistent judgments, and the lack of such evidence can be taken as a strike against the use of particular heuristics. However, Lages et al. (2000) have now gathered experimental evidence for intransitivities: When people make comparisons of population size between pairs of cities (the example to which Allen refers), the total set of judgments contains about 10% of the maximum possible number of intransitive triples, that is, cases where people inferred that city A is larger than B, B is larger than C, but C is larger than A.

The reason this is important is that specific rules for search and stopping predict characteristic patterns of intransitivities, as illustrated in Gigerenzer and Goldstein (1996a, p. 664). Thus, at an individual level we can predict particular patterns of intransitivities associated with the use of particular heuristics, and then we can compare these predictions with the patterns actually observed. Lages et al. (2000) have derived the theoretically predicted patterns in detail, and have found that people's patterns of intransitivities are consistent with the use of the Take The Last heuristic. **Allen** further asks whether people show signs of discomfort when they produce intransitivities. Indeed, some participants said they were afraid of producing intransitivities. But they did not show discomfort when they actually made intransitive judgments, because they made several hundred pair comparisons and did not notice when intransitivities crept in. As Chapter 9 shows, memory of earlier judgments is limited and often even updated in hindsight, which could further erode an individual's recognition of his own intransitive judgments.

Furthermore, the common assumption that consistency per se is always better has been challenged by these findings. Consistency does not guarantee correspondence. For instance, the Take The Last heuristic, which occasionally produces systematic intransitivities, can generate more accurate judgments than linear strategies (such as Dawes's rule, see Ch. 4) which do not. Lages et al. (2000) report that in about one third of all pairs of participants, the one who generated more intransitivities was also more accurate.

Finally, **Kainen** echoes our call for examples of heuristics in other fields, such as perception and language processing, and provides examples of heuristics from engineering and other applied fields that he would like to see collected and tested – but one must be clear about how examples of these artificial mechanisms can be used to elucidate the workings of human or animal minds. **Lipshitz** sees more progress to be made in tackling the heuristics underlying naturalistic real-world decision making, of the sort that Klein (1998) investigates.

R6. What criteria should be used to evaluate the performance of heuristics?

The choice of criteria for evaluating the performance of heuristics lies at the heart of the rationality debate. Our focus on multiple correspondence criteria – such as making decisions that are fast, frugal, and accurate – rather than on internal coherence has drawn the attention of many commentators. We are happy to spark a discussion of suitable criteria for evaluating cognitive strategies, because this topic has been overlooked for too long. As an example, decades' worth of textbooks in social and cognitive psychology have overflowed with demonstrations that people misunderstand syllogisms, violate first-order logic, and ignore laws of probability, with little or no consideration of whether these actually are the proper criteria for evaluating human reasoning strategies. This lack of discussion has an important consequence for the study of cognition: The more one focuses on internal coherence as a criterion for sound reasoning, the less one can see the correspondence between heuristics and their environment, that is, their ecological rationality.

R6.1. Correspondence versus coherence

A person who states that there is a 95% chance that Elvis is still alive and a 5% chance that he has gone to his great reward makes a coherent statement (the probabilities add up to 100%). But as far as we know, the belief does not correspond to the state of the real world. Thus, there can be a tension between the two ideals of coherence and correspondence, as **Allen** and **Fuller** note. The two *need* not conflict; a person's judgments can be coherent, that is, internally consistent, and at the same time correspond well to the outside world. The ideal of complete coherence among all statements or beliefs of a person, however, can be a demonic fiction because it would require elaborate consistency checking and maintenance mechanisms and may even be empirically meaningless (Sen 1993). Research on preferences (as opposed to inferences) invites coherence as a criterion for rationality, because preferences are matters of taste for which external criteria rarely exist – similarly, there are only internal, not external, criteria in research on deductive reasoning and mental logic to which **Newstead** refers. The decision mechanisms described in *Simple heuristics,* however, are intended mainly for inferences, and here the conflict between coherence and correspondence can arise. Taking sides in this conflict, **Barrett & Henzi** welcome the move to design and test heuristics for real-

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

world situations where external correspondence criteria can reveal their success, as do **Shanks & Lagnado,** but **Allen** expects that it is still worthwhile carrying around the "baggage" of the laws of logic as a standard of good reasoning.

Bermúdez has questioned the justifiability of using fast and frugal heuristics that are evaluated on the basis of success alone. He says that "a workable concept of rationality must allow us to evaluate the rationality of an action without knowing its outcome." But basing an evaluation of success on the outcome of a decision is exactly what the processes of evolution and learning do; and as these are the processes that led to the decision making mechanisms we use, we must adopt the same success-oriented perspective if we want to uncover and understand what is going on in our heads when we make choices.

R6.2. Choosing a benchmark

To evaluate the performance of heuristics in the real world, benchmarks are useful. In *Simple heuristics,* we have employed a range of benchmarks – linear models, neural networks, and Bayesian networks, among others (not just multiple linear regression, as **Shanks & Lagnado** imply). The choice of benchmarks (like the choice of criteria) is not always straightforward, and for some, particularly benchmark models with numerous free parameters such as neural networks, it typically takes considerable search (e.g., training and testing) to find a good model for a given situation. Because this choice can be complicated, we welcome the discussion that commentators have raised of the benchmarks we used.

Harries & Dhami, for instance, argue that we should have used a modified version of multiple regression that ignores all non-significant cue weights, because this pruned regression would result in a better benchmark. Hertwig, Hoffrage, and Martignon tested this argument in Chapter 10 (p. 231), with respect to the frugality and accuracy of estimation. First, using only the significant weights did not make multiple regression much more frugal – it only decreased the number of cues used from 8 to 7.3 on average. The QuickEst heuristic, in contrast, used only 2.3 cues on average. Second, pruning improved the accuracy of regression only in one (the smallest) of the 10 training sets substantially, but the fast and frugal QuickEst still remained more accurate. Similar results were obtained when Dawes's rule, a linear strategy with equal weights, was pruned (Ch. 5, p. 112). Thus, Harries & Dhami are right that ignoring cues can improve the accuracy of linear models and make them better benchmarks (just as ignoring information can make heuristics more accurate), but we have found that this still does not change the overall pattern of comparisons between the linear models and simple heuristics.

Shanks & Lagnado also stress the problem of choosing proper benchmarks to evaluate heuristic performance, particularly challenging our assessment of the Categorization by Elimination (CBE) heuristic. They want to see more evaluation of the model in comparison to human data, and so do we – the problem is that very few experiments have ever been performed with categories comprising a large number of cues, the realistic situation for which CBE is intended. Consequently, as we explained in Chapter 11, we resorted to standard categorization data sets from the machine learning literature, which contain hundreds to thousands of instances and dozens of cues and are thus neither "very small" nor "highly non-discriminating" as Shanks & Lagnado state. But we agree that these tests only allow us to indicate the general viability of this approach to categorization, and we must extend the model to account for phenomena such as learning and then pit it against further human data (much of which will have to be obtained from new experiments; see Berretty 1998). However, the choice of human data to use as benchmarks is also not without controversy: For instance, Shanks & Lagnado say that CBE must be able to predict the so-called "exemplar effect," but the existence of this effect is still an open question. In experiments showing the exemplar effect, participants were trained on only a small number of exemplars (e.g., 10 exemplars in Whittlesea 1987, and 30 exemplars in Brooks et al. 1991). Furthermore, Thomas (1997) has shown that not all participants used exemplar information in categorization judgments after learning more realistically sized categories comprising hundreds of exemplars. Shanks & Lagnado also argue that CBE's deterministic response rule is implausible, but Maddox and Ashby (1993) have shown that humans respond with less variability than predicted by **Luce**'s choice axiom, that is, in a more deterministic fashion than expected. Thus again we confront the problem of choice of benchmarks.

R6.3. When fast and frugal heuristics conflict with established values

Engel, Fuller, and **Harries & Dhami** have extended the discussion of performance criteria for heuristics to include societal and legal values, such as accountability, legitimacy, and due process. This is an important extension that raises many fundamental questions. Harries & Dhami ask whether the use of simple heuristics should be prescribed in situations where they are successful, and point out the possible society-level legal implications of fast and frugal judgments. Engel has previously considered the normative question of bounded rationality in the law (Engel 1994). In his comment, he makes the question very clear using our introductory example, the classification of heart-attack patients by at most three yes-no questions. It is easy to imagine a lawyer who sues the doctor because he ignored information, following the assumption that more information is always better – but which party would be in the right?

We agree with **Engel** that the philosophy of fast and frugal decision making may conflict with certain interpretations of the law. The legal system faces a conflict between the ideals of optimization and bounded rationality similar to that in cognitive science. There are two ways to react to a world that is becoming more complex: to strive for perfection by designing ever more complex legal rules that govern every aspect of human behavior, or to stop this growth and strive for a few simple and robust legal rules, as Epstein (1995) proposes. Epstein argues that the half dozen rules he designed would cover 95% of all legal cases. To expect more complete coverage than this from a legal system, he says, would be an illusion.

R7. Conclusion

As the range of expertise covered by our commentators confirms, the study of bounded rationality is a multidisciplinary effort, and its results have relevance for all sciences

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

that try to understand the behavior of living organisms. Opening up the adaptive toolbox and figuring out what lies inside is a challenge that must be addressed from many directions. To the extent that fast and frugal heuristics fill the toolbox, many fields will have to rethink some of our underlying assumptions of the appropriate representation of the world, as **Luce** points out – additive linearity may be, as he says, "a singularly bad representation of a great deal of reality; this matters greatly." To find out, we must extend our understanding of ecological rationality – how environment structures and heuristic mechanisms fit together. We are grateful to the commentators for leading this exploration into new directions.

ACKNOWLEDGMENT

Members of the ABC Research Group who contributed to this response were H. Clark Barrett, Patricia M. Berretty, Seth Bullock, Valerie M. Chase, Jennifer Nerissa Davis, Thomas Dudey, Laurence Fiddick, Daniel G. Goldstein, Adam S. Goodie, Ralph Hertwig, Ulrich Hoffrage, Monika Keller, Stefan Krauss, Martin Lages, Barnaby Marsh, Laura Martignon, Jason Noble, Andreas Ortmann, Jörg Rieskamp, and Christoph Wassner.

References

Letters "a" and "r" appearing before authors' initials refer to target article and response, respectively

Albers, W. (1997) *Foundations of a theory of prominence in the decimal system.* Working papers (No. 265–271). Institute of Mathematical Economics. University of Bielefeld, Germany. [aPMT]

Anderson, J. R. (1990) *The adaptive character of thought.* Erlbaum. [NC, arPMT]

(1991) The adaptive nature of human categorization. *Psychological Review* 98:409–29. [DRS]

Anderson, J. R. & Milson, R. (1989) Human memory: An adaptive perspective. *Psychological Review* 96:703–19. [aPMT]

Armelius, B. & Armelius, K. (1974) The use of redundancy in multiple-cue judgments: Data from a suppressor-variable task. *American Journal of Psychology* 87:385–92. [aPMT]

Ashby, F. G. & Maddox, W. T. (1992) Complex decision rules in categorization: Contrasting novice and experienced performance. *Journal of Experimental Psychology: Human Perception and Performance* 18:50–71. [DRS]

Baguley, T. S. & Payne, S. J. (2000) Long-term memory for spatial and temporal mental models includes construction processes and model structure. *Quarterly Journal of Experimental Psychology* 53:479–512. [TB]

Baillargeon, R. (1986) Representing the existence and the location of hidden objects: Object permanence in 6- and 8-month old infants. *Cognition* 23:21–41. [rPMT]

Bak, P. (1996) *How nature works: The science of self-organized criticality.* Copernicus. [rPMT]

Bargh, J. A. & Chartrand, T. L. (1999) The unbearable automaticity of being. *American Psychologist* 54:462–79. [rPMT]

Barrett, L., Henzi, S. P., Weingrill, T., Lycett, J. E. & Hill, R. A. (1999) Market forces predict grooming reciprocity in female baboons. *Proceedings of the Royal Society, London, Series B* 266:665–70. [LB]

Becker, G. S. (1991) *A treatise on the family.* Harvard University Press. [aPMT]

Bermúdez, J. L. (1998) Philosophical psychopathology. *Mind and Language* 13:287–307. [JLB]

(1999a) Naturalism and conceptual norms. *Philosophical Quarterly* 49:77–85. [JLB]

(1999b) Rationality and the backwards induction argument. *Analysis* 59:243–48. [JLB]

(1999c) Psychologism and psychology. *Inquiry* 42:487–504. [JLB]

(2000) Personal and sub-personal: A difference without a distinction. *Philosophical Explorations* 3:63–82. [JLB]

Bermúdez, J. L. & Millar, A., eds. (in preparation) *Naturalism and rationality.* [JLB]

Berretty, P. M. (1998) Category choice and cue learning with multiple dimensions. Unpublished Ph. D. dissertation. University of California at Santa Barbara, Department of Psychology. [rPMT]

Berretty, P. M., Todd, P. M. & Blythe, P. W. (1997) Categorization by elimination: A fast and frugal approach to categorization. In: *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society,* ed. M. G. Shafto & P. Langley. Erlbaum. [aPMT]

Bettman, J. R. (1979) *An information processing theory of consumer choice.* Addison-Wesley. [aPMT]

Billman, D. & Heit, E. (1988) Observational learning from internal feedback: A simulation of an adaptive learning method. *Cognitive Science* 12:587–625. [AW]

Binmore, K. (1991) Rational choice theory: Sufficient but not necessary. *American Psychologist* 46:797–99. [DRS]

Blurton-Jones, N., Hawkes, K. & O'Connell, J. F. (1999) Some current ideas about the evolution of human life history. In: *Comparative primate socioecology,* ed. P. C. Lee. Cambridge Studies in Biological Anthropology. [LB]

Blythe, P. W., Miller, G. F. & Todd, P. M. (1996) Human simulation of adaptive behavior: Interactive studies of pursuit, evasion, courtship, fighting, and play. In: *From animals to animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior,* ed. P. Maes, M. J. Mataric, J.-A. Meyer, J. Pollack & S. W. Wilson. MIT Press/Bradford Books. [aPMT]

Bradshaw, G. L., Langley, P. & Simon, H. A. (1983) Studying scientific discovery by computer simulation. *Science* 222:971–75. [MEG]

Brehmer, B. (1994) The psychology of linear judgment models. *Acta Psychologica* 87:137–54. [CH]

Breiman, L., Friedman, J. H., Olshen, R. A. & Stone, C. J. (1993) *Classification and regression trees.* Chapman & Hall. [aPMT]

Bröder, A. (in press) 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.* [rPMT]

Brooks, L. R., Norman, G. R. & Allen, S. W. (1991) The role of specific similarity in a medical diagnostic task. *Journal of Experimental Psychology: General* 120(3):278–87. [DRS, rPMT]

Bruer, J. T. (1993) *Schools for thought: A science of learning in the classroom.* Bradford Books. [GEAS]

Brunswik, E. (1940) Thing constancy as measured by correlation coefficients. *Psychological Review* 47:69–78. [XTW]

(1943) Organismic achievement and environmental probability. *Psychological Review* 50:255–72. [aPMT]

(1957) Scope and aspects of the cognitive problem. In: *Contemporary approaches to cognition,* ed. J. S. Bruner, E. Brunswik, L. Festinger, F. Heider, K. F. Muenzinger, C. E. Osgood & D. Rapaport. Harvard University Press. [aPMT]

Bullock, S. & Todd, P. M. (1999) Made to measure: Ecological rationality in structured environments. *Minds and Machines* 9(4):497–541. [rPMT]

Byrne, R. W. & Whiten, A. (1988) *Machiavellian intelligence: Social expertise and the evolution of intellect in monkeys, apes and humans.* Oxford Scientific Publications. [LB]

Carey, S. (1985) *Conceptual change in childhood.* MIT Press/Bradford Books. [GEAS, rPMT]

Ceci, S. J. & Liker, J. K. (1986) A day at the races: A study of IQ, expertise, and cognitive complexity. *Journal of Experimental Psychology: General* 115:255–66. [DRS]

Chang, R., ed. (1997) *Incommensurability, incomparability, and practical reason.* Harvard University Press. [AM]

Charnov, E. L. (1976) Optimal foraging: The marginal value theorem. *Theoretical Population Biology* 9:129–36. [AIH]

Chase, V. M., Hertwig, R. & Gigerenzer, G. (1998) Visions of rationality. *Trends in Cognitive Science* 2(6):206–14. [aPMT]

Chater, N. & Oaksford, M. (1990) Autonomy, implementation, and cognitive architecture: A reply to Fodor and Pylyshyn. *Cognition* 34:93–107. [MO]

(1999) The probability heuristics model of syllogistic reasoning. *Cognitive Psychology* 38:191–258. [SEN, MO]

Chater, N., Oaksford, M., Nakisa, R. & Redington, M. (1999) Fast, frugal *and* rational: How rational norms explain behavior. Unpublished manuscript, Department of Psychology, University of Warwick, Coventry, UK. [NC, MO]

Cheney, D. L. & Seyfarth, R. M. (1985) Vervet monkey alarm calls: Manipulation through shared information? *Behaviour* 93:150–66. [XTW]

Chi, M. T. H., Glaser, R. & Farr, M., eds. (1982) *The nature of expertise.* Erlbaum. [RL]

Cohen, M. S. (1993) Three paradigms for viewing decision biases. In: *Decision making in action: Models and methods,* ed. G. A. Klein, J. Orasanu, R. Calderwood & C. Zsambok. Ablex. [RL]

Cohen, M. S., Freeman, J. T. & Wolf, S. (1996) Meta-recognition in time stressed decision making: Recognizing, critiquing and correcting. *Human Factors* 38:206–19. [RL]

Connolly, T. & Gilani, N. (1982) Information search in judgment tasks: A regression model and some preliminary findings. *Organizational Behavior and Human Performance* 30:330–50. [aPMT]

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

Cooksey, R. W. (1996) *Judgment analysis: Theory, methods and applications.* Academic Press. [CH]
Coombs, C. H. (1964) *A theory of data.* Wiley. [JS]
Cosmides, L. & Tooby, J. (1992) Cognitive adaptations for social exchange. In: *The adapted mind: Evolutionary psychology and the generation of culture,* ed. J. Barkow, L. Cosmides & J. Tooby. Oxford University Press. [arPMT]
Csibra, G., Gergely, G., Biro, S., Koos, O. & Brockbank, M. (1999) Goal attribution without agency cues: The perception of "pure reason" in infancy. *Cognition* 72:237–67. [rPMT]
Darwin, C. (1969) *The autobiography of Charles Darwin, 1809–1882,* ed. N. Barlow. Norton. (Original work published in 1887.) [aPMT]
Davis, D. G. S., Staddon, J. E. R. Machado, A. & Palmer, R. G. (1993) The process of recurrent choice. *Psychological Review* 100:320–41. [DRS]
Dawes, R. M. (1979) The robust beauty of improper linear models in decision making. *American Psychologist* 34:571–82. [aPMT]
Dawes, R. M. & Corrigan, B. (1974) Linear models in decision making. *Psychological Bulletin* 81:97–106. [JS]
Dawkins, M. S. (1995) *Unravelling animal behaviour,* 2nd edition. Longmans. [JLB]
Desmurget, I., Prablanc, C., Arzi, M., Rossetti, Y., Paulignan, Y. & Urquizar, C. (1996) Integrated control of hand transport and orientation during prehension. *Experimental Research* 110:265–78. [PCK]
Dhami, M. K. & Ayton, P. (1998) Legal decision making the fast and frugal way. Paper presented at the Meeting of the Society for Judgement and Decision Making, Dallas, TX, 1998. [CH]
Driver, P. M. & Humphries, D. A. (1988) *Protean behavior: The biology of unpredictability.* Oxford University Press. [arPMT]
Dutton, D. M. & Conroy, G. V. (1996) A review of machine learning. *The Knowledge Engineering Review* 12:341–67. [CH]
Einhorn, H. J. (1970) The use of nonlinear, noncompensatory models in decision making. *Psychological Bulletin* 73:221–30. [rPMT]
Einhorn, H. J. & Hogarth, R. M. (1981) Behavioral decision theory: Processes of judgment and choice. *Annual Review of Psychology* 32:53–88. [arPMT]
Elster, J. (1979) *Ulysses and the sirens: Studies in rationality and irrationality.* Cambridge University Press. [aPMT]
Engel, C. (1994) Legal responses to bounded rationality in German administration. *Journal of Institutional and Theoretical Economics* 150:145–62. [rPMT]
Epstein, R. A. (1995) *Simple rules for a complex world.* Harvard University Press. [rPMT]
Erev, I. & Roth, A. (2000) Learning, reciprocation and the value of bounded rationality. In: *Bounded rationality: The adaptive toolbox,* ed. G. Gigerenzer & R. Selten. MIT Press. [rPMT]
Evans, J. St. B. T., Barston, J. L. & Pollard, P. (1983) On the conflict between logic and belief in syllogistic reasoning. *Memory and Cognition* 11:295–306. [SEN]
Evans, J. St. B. T., Newstead, S. N. & Byrne, R.M. J. (1993) *Human reasoning: The psychology of deduction.* Erlbaum. [AF]
Evans, J. St. B. T. & Over, D. E. (1996) *Rationality and reasoning.* Erlbaum. [AF]
Feeney, A. & Handley, S. J. (in press) The suppression of q-card selections: Evidence for deductive inference in Wason's selection task. *Quarterly Journal of Experimental Psychology.* [AF]
Festinger, L. (1957) *A theory of cognitive dissonance.* Stanford University Press. [SF]
Fishburn, P. C. (1988) *Nonlinear preference and utility theory.* Johns Hopkins University Press. [rPMT]
Ford, J. K., Schmitt, N., Schechtman, S. L., Hults, B. M. & Doherty, M. L. (1989) Process tracing methods: Contributions, problems and neglected research questions. *Organizational Behavior and Human Decision Processes* 43:75–117. [OH]
Frank, R. H. (1988) *Passions within reason: The strategic role of emotions.* Norton. [rPMT]
Frege, G. (1918–19/1987) Thoughts. In: *Logical investigations,* ed. P. T. Geach. Republished 1987. Basil Blackwell. [JLB]
Friedman, D. & Massaro, D. W. (1998) Understanding variability in binary and continuous choice. *Psychonomic Bulletin and Review* 5:370–89. [DRS]
Fuller, S. (1985) *Bounded rationality in law and science.* Ph. D. dissertation, University of Pittsburgh, Department of History and Philosophy of Science. [SF]
Galef, B. G., Jr. (1987) Social influences on the identification of toxic foods by Norway rats. *Animal Learning and Behavior* 18:199–205. [aPMT]
Geman, S., Bienenstock, E. & Doursat, E. (1992) Neural networks and the bias/variance dilemma. *Neural Computation* 4:1–58. [aPMT]
Gibson, J. J. (1979) *The ecological approach to visual perception.* Houghton-Mifflin. [aPMT]
Gigerenzer, G. (1991) From tools to theories: A heuristic of discovery in cognitive psychology. *Psychological Review* 98:254–67. [rPMT]
(1993) The superego, the ego, and the id in statistical reasoning. In: *A handbook for data analysis in the behavioral sciences: Methodological issues,* ed. G. Keren & C. Lewis. Erlbaum. [rPMT]
(1994) Where do new ideas come from? In: *Dimensions of creativity,* ed. M. A. Boden. MIT Press. [rPMT]
(1996) On narrow norms and vague heuristics: A reply to Kahneman and Tversky. *Psychological Review* 103:592–96. [aPMT]
Gigerenzer, G. & Goldstein, D. G. (1996a) Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review* 103:650–69. [EE, arPMT]
(1996b) Mind as computer: The birth of a metaphor. *Creativity Research Journal* 9:131–44. [rPMT]
Gigerenzer, G. & Hoffrage, U. (1995) How to improve Bayesian reasoning without instructions. Frequency formats. *Psychological Review* 102:684–704. [rPMT]
Gigerenzer, G. & Murray, D. J. (1987) *Cognition as intuitive statistics.* Erlbaum. [arPMT]
Gigerenzer, G. & Richter, H. R. (1990) Context effects and their interaction with development: Area judgments. *Cognitive Development* 5:235–64. [rPMT]
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 University Press. [aPMT]
Gigerenzer, G., Todd, P. M. & the ABC Research Group (1999) *Simple heuristics that make us smart.* Oxford University Press. [CA, TB, LB, JLB, NC, RC, CE, EE, AF, SF, AIG, MEG, PH, CH, AIH, OH, PCK, RL, RDL, HM, AM, SEN, MO, DRS, JS, GEAS, RJS, aPMT, AW, XTW, WCW]
Glymour, C. (1998) Learning causes: Psychological explanations of causal explanation. *Minds and Machines* 8:39–60. [AW]
Goldman, A. I. (1999) *Knowledge in a social world.* Oxford University Press. [AIG]
Goldstein, D. G. & Gigerenzer, G. (in press) Models of ecological rationality: The recognition heuristic. *Psychological Review.* [rPMT]
Goodie, A. S. & Todd, P. M. (submitted) The ecological rationality of base-rate neglect. [rPMT]
Gopnik, A. (1998) Explanation as orgasm. *Minds and Machines* 8:39–60. [AW]
Gopnik, A. & Meltzoff, A. N. (1997) *Words, thoughts, and theories.* MIT Press. [rPMT]
Gorman, M. E. (1992) *Simulating science: Heuristics, mental models, and technoscientific thinking.* Indiana University Press. [MEG]
(1995) Confirmation, disconfirmation and invention: The case of Alexander Graham Bell and the telephone. *Thinking and Reasoning* 1(1):31–53. [MEG]
(1998) *Transforming nature: Ethics, invention and design.* Kluwer. [MEG]
Gould, S. J. & Vrba, E. (1982) Exaptation – A missing term in the science of form. *Paleobiology* 8:4–15. [WCW]
Green, R. F. (1984) Stopping rules for optimal foragers. *American Naturalist* 123:30–43. [AIH]
Groner, M., Groner, R. & Bischof, W. F. (1983) Approaches to heuristics: A historical review. In: *Methods of heuristics,* ed. R. Groner. Erlbaum. [aPMT]
Hammond, K. R. (1996) *Human judgment and social policy: Irreducible uncertainty, inevitable error, unavoidable injustice.* Oxford University Press. [aPMT]
Harries, C. & Dhami, M. K. (1998) Fast and frugal models of human judgement: Fructiferous freethinking or futile fracas? Paper presented at the Meeting of the Brunswik Society, Dallas, TX, 1998. [CH]
Hastie, R. (1991) A review from a high place: The field of judgment and decision making as revealed in its current text books. *Psychological Science* 2:135–38. [RL]
Herrnstein, R. J. (1997) *The matching law: Papers in psychology and economics,* ed. H. Rachlin & D. I. Laibson. Harvard University Press. [DRS]
Herrnstein, R. J., Loewenstein, G. F., Prelec, D. & Vaughan, W. (1993) Utility maximization and melioration: Internalities in individual choice. *Journal of Behavioral Decision Making* 6:149–85. [DRS]
Hinton, G. E. (1989) Connectionist learning procedures. *Artificial Intelligence* 40:185–234. [MO]
Hoffrage, U. (in press) Why the analysis of cognitive processes matters. *Behavioral and Brain Sciences.* [rPMT]
Hogarth, R. M. (1987) *Judgment and choice: The psychology of decision,* 2nd edition. Wiley. [rPMT]
Holland, J. H., Holyoak, K. J., Nisbett, R. E. & Thagard, P. R. (1986) *Induction: Processes of inference, learning and discovery.* MIT Press. [AW]
Holte, R. C. (1993) Very simple classification rules perform well on most commonly used datasets. *Machine Learning* 3(11):63–91. [aPMT]
Hopfield, J. J. & Tank, D. W. (1985) "Neural" computation of decisions in optimization problems. *Biological Cybernetics* 52:141–52. [DRS]
Houston, A. I. (1983) Optimality theory and matching. *Behaviour Analysis Letters* 3:1–15. [AIH]
(1987) The control of foraging decisions. In: *Quantitative analysis of behavior, vol. 6,* ed. M. L. Commons, A. Kacelnik & S. J. Shettleworth. Erlbaum. [AIH]

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

Houston, A. I., Kacelnik, A. & McNamara, J. M. (1982) Some learning rules for acquiring information. In: *Functional ontogeny,* ed. D. J. McFarland. Pitman. [AIH]

Houston, A. I. & McNamara, J. M. (1981) How to maximize reward rate on two VI paradigms. *Journal of the Experimental Analysis of Behavior* 35:367–96. [AIH]

(1984) Imperfectly optimal animals. *Behavioral Ecology and Sociobiology* 15:61–64. [AIH]

(1988) A framework for the functional analysis of behavior. *Behavioral and Brain Sciences* 11:117–63. [AIH]

(1999) *Models of adaptive behaviour.* Cambridge University Press. [AIH]

Hutchins, E. (1995) *Cognition in the wild.* MIT Press. [rPMT]

Janis, I. L. & Mann, L. (1977) *Decision making: A psychological analysis of conflict, choice and commitment.* Free Press. [RL]

Johnson, E. J., Meyer, R. J. & Ghose, S. (1989) When choice models fail: Compensatory models in negatively correlated environments. *Journal of Marketing Research* 24:255–70. [JS]

Kahneman, D. & Tversky, A. (1973) On the psychology of prediction. *Psychological Review* 80:237–51. [JLB]

(1996) On the reality of cognitive illusions: A reply to Gigerenzer's critique. *Psychological Review* 103:582–91. [aPMT]

Kainen, P. C. (1997) Utilizing geometric anomalies of high dimension: When complexity makes computation easier. In: *Computer-intensive methods in control and signal processing: Curse of dimensionality,* ed. K. Warwick & M. Karny. Birkhauser. [PCK]

(1998) Mathematical cognition. *Evolution and Cognition* 4(1):81–88. [PCK]

Kalish, M. L. & Kruschke, J. K. (1997) Decision boundaries in one-dimensional categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 23:1362–77. [DRS]

Keeney, R. L. & Raiffa, H. (1993) *Decisions with multiple objectives.* Cambridge University Press. [JS]

Kiester, A. R. (1979) Conspecifics as cues: A mechanism for habitat selection in the Panamanian grass anole (*Anolis auratus*). *Behavioral Ecology and Sociobiology* 5:323–30. [WCW]

Klayman, J. (1985) Children's decision strategies and their adaptation to task characteristics. *Organizational Behavior and Human Decision Processes* 35:179–201. [rPMT]

Klayman, J. & Ha, Y.-W. (1987) Confirmation, disconfirmation and information in hypothesis testing. *Psychological Review* 94:211–28. [MEG]

Klein, G. A. (1998) *Sources of power: How people make decisions.* MIT Press. [rPMT]

Klein, G. A., Orasanu, J., Calderwood, R. & Zsambok, C., eds. (1993) *Decision making in action: Models and methods.* Ablex. [RL]

Kolodner, J. L. (1993) *Case-based reasoning.* Morgan Kaufman. [MO]

Kreps, D. M. (1990) *A course in microeconomic theory.* Harvester Wheatsheaf. [NC]

Krüger, L., Gigerenzer, G. & Morgan, M., eds. (1987) *The probabilistic revolution. Vol. II: Ideas in the sciences.* MIT Press. [aPMT]

Kukar, M., Kononenko, I., Grošelj, C., Kralj, K. & Fettich, J. (1999) Analysing and improving the diagnosis of ischaemic heart disease with machine learning. *Artificial Intelligence in Medicine* 16:25–50. [CH]

Kulkarni, D. & Simon, H. A. (1988) The processes of scientific discovery: The strategies of experimentation. *Cognitive Science* 12:139–75. [MEG]

Kusch, M. (1999) *A social history and philosophy. Psychological knowledge.* Routledge. [aPMT]

Lages, M., Hoffrage, U. & Gigerenzer, G. (2000) How heuristics produce intransitivity and how intransitivity can discriminate between heuristics. (submitted). [rPMT]

Lakoff, G. & Johnson, M. (1980) *Metaphors we live by.* University of Chicago Press. [PCK]

Lansdale, M. W. (1998) Modeling memory for absolute location. *Psychological Review* 105:351–78. [TB]

Laplace, P. S. (1951) *A philosophical essay on probabilities,* trans. F. W. Truscott & F. L. Emory. Dover. (Original work published in 1814.) [aPMT]

Lawson, A. E. & Thompson, L. D. (1988) Formal reasoning ability and misconceptions concerning genetics and natural selection. *Journal of Research in Science Teaching* 25:733–46. [GEAS]

Leibniz, G. W. (19951) Toward a universal characteristic. In: *Leibniz: Selections,* ed. P. P. Wiener. Scribner's Sons. (Original work published in 1677.) [aPMT]

Locke, J. (1959) *An essay concerning human understanding,* ed. A. C. Fraser. Dover. (Original work published in 1690.) [aPMT]

Lopes, L. L. (1987) Between hope and fear: The psychology of risk. *Advances in Experimental Social Psychology* 20:255–95. [XTW]

López, F. J., Cobos, P. L., Caño, A. & Shanks, D. R. (1998) The rational analysis of human causal and probability judgment. In: *Rational models of cognition,* ed. M. Oaksford & N. Chater. Oxford University Press. [DRS]

MacDonald, M. C., Pearlmutter, N. J. & Seidenberg, M. S. (1994) Lexical nature of syntactic ambiguity resolution. *Psychological Review* 101:676–703. [MO]

Maddox, W. T. & Ashby, F. G. (1993) Comparing decision bound and exemplar models of categorization. *Perception and Psychophysics* 53:49–70. [rPMT]

Makse, H. A., Havlin, B. & Stanley, H. E. (1995) Modelling urban growth patterns. *Nature* 377:608–12. [aPMT]

Margolis, H. (1998) Tycho's illusion. *Nature* 382:957. [HM]

Massaro, D. W. (1987) *Speech perception by ear and eye.* Erlbaum. [NC, MO]

(1988) Some criticisms of connectionist models of human performance. *Journal of Memory and Language* 27:213–34. [aPMT]

McClelland, G. (1978) Equal versus differential weighting for multiattribute decisions: There are no free lunches. Center Report No. 207, Institute of Cognitive Science, University of Colorado, Boulder, CO. [JS]

McFarland, D. & Houston, A. (1981) *Quantitative ethology: The state space approach.* Pitman. [NC]

McLeod, P. & Dienes, Z. (1996) Do fielders know where to go to catch the ball or only how to get there? *Journal of Experimental Psychology: Human Perception and Performance* 22:531–43. [rPMT]

McNamara, J. M. & Houston, A. I. (1980) The application of statistical decision theory to animal behaviour. *Journal of Theoretical Biology* 85:673–90. [AIH]

(1985) Optimal foraging and learning. *Journal of Theoretical Biology* 117: 231–49. [AIH]

Meehl, P. E. (1954) *Clinical versus statistical predictions: A theoretical analysis and revision of the literature.* University of Minnesota Press. [NC]

Montgomery, H. & Svenson, O. eds. (1989) *Process and structure in human decision making.* Wiley. [OH]

Morton, A. (1990) *Disasters and dilemmas.* Blackwell. [AM]

Morton, J. (1969) The interaction of information in word recognition. *Psychological Review* 76:165–78. [MO]

Movellan, J. R. & Chadderdon, G. (1996) Cognition and the statistics of natural signals. In: *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society,* ed. G. Cottrell. Erlbaum. [MO]

Movellan, J. R. & McClelland, J. L. (1995) Stochastic interactive processing, channel separability and optimal perceptual inference: An examination of Morton's Law. Technical Report PDP.CNS.95.4, Carnegie Mellon University, Pittsburgh, PA. [MO]

Newell, A. (1990) *Unified theories of cognition.* Harvard University Press. [TB]

Newell, A. & Simon, H. A. (1972) *Human problem solving.* Prentice Hall. [TB, aPMT]

Newstead, S. E., Handley, S. J. & Buck, E. (1999) Falsifying mental models: Testing the predictions of theories of syllogistic reasoning. *Memory and Cognition* 27:344–54. [SEN]

Noble, J. & Todd, P. M. (in press) Imitation or something simpler? Modelling simple mechanisms for social information processing. In: *Imitation in animals and artifacts,* ed. K. Dautenhahn & C. Nehaniv. MIT Press. [rPMT]

Norman, D. A. (1993) *Things that make us smart: Defending human attributes in the age of the machine.* Addison-Wesley. [rPMT]

Nosofsky, R. M. (1987) Attention and learning processes in the identification and categorization of integral stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 13:87–108. [DRS]

Nozick, R. (1993) *The nature of rationality.* Harvard University Press. [JLB]

Oaksford, M. & Chater, N. (1991) Against logicist cognitive science. *Mind and Language* 6:1–38. [MO]

(1993) Reasoning theories and bounded rationality. In: *Rationality,* ed. K. I. Manktelow & D. E. Over. Routledge. [MO]

(1994) A rational analysis of the selection task as optimal data selection. *Psychological Review* 101:608–31. [AF, aPMT]

(1995) Theories of reasoning and the computational explanation of everyday inference. *Thinking and Reasoning* 1:121–52. [MO]

(1998a) *Rationality in an uncertain world: Essays in the cognitive science of human reasoning.* Psychology Press. [NC]

(1998b) *Rational models of cognition.* Oxford University Press. [NC]

Ornstein, R. & Ehrlich, P. (1989) *New world, new mind: Moving toward conscious evolution.* Doubleday. [rPMT]

Payne, J. W., Bettman, J. R. & Johnson, E. J. (1988) Adaptive strategy selection in decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 14:534–52. [AM]

(1993) *The adaptive decision maker.* Cambridge University Press. [OH, AM, arPMT]

Payne, S. J. (1988) Methods and mental models in theories of cognitive skill. In: *Artificial intelligence and human learning: Intelligent computer aided instruction,* ed. J. Self. Chapman & Hall. [TB]

Pearson, K. (1897) On the scientific measure of variability. *Natural Science* 11:115–18. [aPMT]

Phelps, R. H. & Shanteau, J. (1978) Livestock judges: How much information can an expert use? *Organizational Behavior and Human Decision Processes* 21:209–19. [JS]

Pinker, S. (1997/1998) *How the mind works.* Penguin/Norton. [TB, aPMT]

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press

Polya, G. (1954) *Mathematics and plausible reasoning. Vol. 1: Induction and analogy in mathematics.* Princeton University Press. [aPMT]
Potlein, H. (1993) *Darwin machines and the nature of knowledge.* Harvard University Press. [PCK]
Prince, A. & Smolensky, P. (1997) Optimality: From neural networks to universal grammar. *Science* 275:1604–10. [rPMT]
Pylyshyn, Z. W., ed. (1987) *The robot's dilemma: The frame problem in artificial intelligence.* Ablex. [NC]
Ramachandran, V. S. (1990) Interactions between motion, depth, color and form: The utilitarian theory of perception. In: *Vision: Coding and efficiency,* ed. C. Blakemore. Cambridge University Press. [aPMT]
Real, L. & Caraco, T. (1986) Risk and foraging in stochastic environments. *Annual Review of Ecology and Systematics* 17:371–90. [XTW]
Rieskamp, J. & Hoffrage, U. (2000) The use of simple heuristics: Bounded rationality and time pressure. (submitted). [rPMT]
Rivest, R. J. (1987) Learning decision lists. *Machine Learning* 2:229–46. [aPMT]
Robertson, S. I. (1999) *Types of thinking.* Routledge. [TB]
(in press) Imitative problem solving: Why transfer of learning often fails to occur. *Instructional Science.* [TB]
Rogoff, B. & Lave, J., eds. (1984) *Everyday cognition: Its development in social context.* Harvard University Press. [RL]
Rumelhart, D. E., McClelland, J. L. & the PDP Research Group (1986) *Parallel distributed processing: Explorations in the microstructure of cognition,* Vols. 1 and 2. MIT Press. [MO]
Saad, G. & Russo, J. E. (1996) Stopping criteria in sequential choice. *Organizational Behavior and Human Decision Processes* 67(3):258–70. [aPMT]
Sargent, T. J. (1993) *Bounded rationality in macroeconomics.* Oxford University Press. [aPMT]
Savage, L. J. (1954) *The foundations of statistics.* Wiley. [RL]
Sawyer, J. (1966) Measurement and prediction, clinical and statistical. *Psychological Bulletin* 66:178–200. [NC]
Schutz, A. (1964) The well-informed citizen: An essay in the social distribution of knowledge. In: *Collected papers, Vol. II,* A. Schutz. Martinus Nijhoff. [SF]
Scribner, S. (1984) Studying working intelligence. In: *Everyday cognition: Its development in social context,* ed. B. Rogoff & J. Love. Harvard University Press. [RL]
Seeley, T. D. (2000) Decision making in superorganisms: How collective wisdom arises from the poorly informed masses. In: *Bounded rationality: The adaptive toolbox,* ed. G. Gigerenzer & R. Selten. MIT Press. [rPMT]
Sells, M. B. (1936) The atmosphere effect: An experimental study of reasoning. *Archives of Psychology* 29:3–72. [SEN]
Sen, A. (1993) Internal consistency of choice. *Econometrica* 61:495–521. [arPMT]
(1994) Population: Delusion and reality. *The New York Review,* September 22, pp. 62–71. [MEG]
Shepard, R. N. (1967) On subjectively optimum selections among multi-attribute alternatives. In: *Decision making,* ed. W. Edwards & A. Tversky. Penguin. [arPMT]
(1990) *Mind sights.* Freeman. [aPMT]
Simon, H. A. (1956a) Rational choice and the structure of environments. *Psychological Review* 63:129–38. [EE, RDL, aPMT, XTW]
(1956b) Dynamic programming under uncertainty with a quadratic criterion function. *Econometrica* 24:19–33. [aPMT]
(1959) Theories of decision-making in economics and behavioral sciences. *American Economic Review* 49:253–83. [AIH]
(1979) How big is a chunk? In: *Models of thought, Vol. 1,* ed. H. A. Simon. Yale University Press. [CH]
(1982) *Models of bounded rationality.* MIT Press. [JLB]
(1987) Rational decision making in business organizations. In: *Advances in behavioral economics, Vol. 1,* ed. L. Green & J. H. Kagel. Ablex. [aPMT]
(1990) Invariants of human behavior. *Annual Review of Psychology* 41:1–19. [aPMT, XTW]
(1991) Cognitive architectures and rational analysis: Comment. In: *Architectures for intelligence,* ed. K. VanLehn. Erlbaum. [aPMT]
(1996) *The sciences of the artificial,* 3rd edition. MIT Press. [WCW]
Simon, H. A. & Ando, A. (1961) Aggregation of variables in dynamic systems. *Econometrica* 29:111–38. [WCW]
Solomon, G. E. A. (1997) Conceptual change and wine expertise. *Journal of the Learning Sciences: Special Issue on Conceptual Change* 6:41–60. [GEAS]
Spelke, E. S. (1990) Principles of object perception. *Cognitive Science* 14:29–56. [rPMT]
Stamps, J. A. (1987) Conspecifics as cues to territory quality: A preference for previously used territories by juvenile lizards (*Anolis aeneus*). *American Naturalist* 129(5):529–42. [WCW]
Stanovich, K. E. & West, R. F. (2000) Individual differences in reasoning: Implications for the rationality debate? *Behavioral and Brain Sciences* 23(5):645–726. [rPMT]
Stephens, D. W. & Krebs, J. R. (1986) *Foraging theory.* Princeton University Press. [XTW]
Sternberg, R. J., Forsythe, G. B., Hedlund, J., Horvath, J., Snook, S., Williams, W. M., Wagner, R. K. & Grigorenko, E. L. (2000) *Practical intelligence in everyday life.* Cambridge University Press. [RJS]
Stigler, G. J. (1961) The economics of information. *Journal of Political Economy* 69:213–25. [aPMT]
Stillwell, W. G., Seaver, D. A. & Edwards, W. (1981) A comparison of weight approximation techniques in multiattribute utility decision making. *Organizational Behavior and Human Performance* 28:62–77. [JS]
Sussman, M., Fruchter, D., Hilbert, J. & Sirosh, J. (1998) Linear correlates in the speech signal: The orderly output constraint. *Behavioral and Brain Sciences* 21:241–99. [PCK]
Tabachnik, B. G. & Fidell, L. S. (1996) *Using multivariate statistics,* 3rd edition. Harper Collins. [CH]
Tanner, W. P., Jr. & Swets, J. A. (1954) A decision-making theory of visual detection. *Psychological Review* 61:401–409. [rPMT]
Taraban, R. & McClelland, J. L. (1988) Constituent attachment and thematic role assignment in sentence processing: Influences of content-based expectations. *Journal of Memory and Language* 27:597–632. [MO]
Thaler, R. H. (1991) *Quasi rational economics.* Sage. [aPMT]
Thomas, R. D. (1997) Learning correlations in categorization tasks using large, ill-defined categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 24(1):119–43. [rPMT]
Tierney, J. (1990) Betting the planet. *The New York Times Magazine,* December 2, pp. 52–53, 75–79. [rPMT]
Tooby, J. & Cosmides, L. (1992) The psychological foundations of culture. In: *The adapted mind: Evolutionary psychology and the generation of culture,* ed. J. Barkow, L. Cosmides & J. Tooby. Oxford University Press. [aPMT]
(1998) Ecological rationality and the multimodular mind. (submitted). [aPMT]
Tversky, A. (1972) Elimination by aspects: A theory of choice. *Psychological Review* 79(4):281–99. [aPMT]
Tversky, A. & Kahneman, D. (1974) Judgment under uncertainty: Heuristics and biases. *Science* 185:1124–31. [RC, aPMT]
(1981) The framing of decisions and the psychology of choice. *Science* 211:453–58. [HM, XTW]
(1983) Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review* 90:293–315. [aPMT]
Upton, G. & Fingleton, B. (1985) *Spatial data analysis by example. Vol. 1: Point pattern and quantitative data.* Wiley. [aPMT]
Vriend, N. J. (1996) Rational behavior and economic theory. *Journal of Economic Behavior and Organization* 29:263–85. [aPMT]
Wang, X. T. (1996a) Domain-specific rationality in human choices: Violations of utility axioms and social contexts. *Cognition* 60:31–63. [XTW]
(1996b) Framing effects: Dynamics and task domains. *Organizational Behavior and Human Decision Processes* 68:145–57. [XTW]
Wason, P. C. (1983) Realism and rationality in the selection task. In: *Thinking and reasoning: Psychological approaches,* ed. J. St. B. T. Evans. Routledge & Kegan Paul. [aPMT]
Wellman, H. M. & Gelman, S. A. (1992) Cognitive development: Foundational theories of core domains. *Annual Review of Psychology* 43:337–75. [GEAS]
Whiten, A. & Byrne, R. W., eds. (1997) *Machiavellian intelligence II.* Cambridge University Press. [rPMT]
Whittlesea, B. W. A. (1987) Preservation of specific experiences in the representation of general knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 13(1):3–17. [DRS, rPMT]
Williams, B. A. (1988) Reinforcement, choice, and response strength. In: *Stevens' handbook of experimental psychology, Vol. 2,* ed. R. C. Atkinson, R. J. Herrnstein, G. Lindzey & R. D. Luce. Wiley. [AIH]
Wimsatt, W. C. (2000a) *Re-engineering philosophy for limited beings: Piecewise approximations to reality.* Harvard University Press. [aPMT, WCW]
(2000b) Generative entrenchment and the developmental systems approach to evolutionary processes. In: *Cycles of contingency: Developmental systems and evolution,* ed. S. Oyama, R. Gray & P. Griffiths. MIT Press. [WCW]
Winter, S. G. (1975) Optimization and evolution in the theory of the firm. In: *Adaptive economic models,* ed. R. H. Day & T. Groves. Academic Press. [aPMT]
Wood, D. J. (1988) *How children think and learn: The social contexts of cognitive development.* Blackwell. [TB]
Woodworth, R. S. & Sells, S. B. (1935) An atmospheric effect in syllogistic reasoning. *Journal of Experimental Psychology* 18:451–60. [SEN]
Wundt, W. (1973) *An introduction to psychology,* trans. R. Pintner. Arno. (Original work published 1912.) [JLB]
Yaniv, I. & Foster, D. P. (1995) Graininess of judgment under uncertainty: An accuracy-informativeness trade-off. *Journal of Experimental Psychology: General* 124:424–32. [DRS]

https://doi.org/10.1017/S0140525X00513442 Published online by Cambridge University Press