Simple Heuristics That Make Us Smart

Gerd Gigerenzer
Peter M. Todd
and the ABC Research Group

New York Oxford
Oxford University Press
1999

16

What We Have Learned (So Far)

Peter M. Todd
Gerd Gigerenzer

> At first it was thought that the surest way would be to take as a foundation for the psychological analysis of the thought-processes the laws of logical thinking, as they had been laid down from the time of Aristotle by the science of logic. These norms only apply to a small part of the thought-processes. Any attempt to explain, out of these norms, thought in the psychological sense of the word can only lead to an entanglement of the real facts in a net of logical reflections. We can in fact say of such attempts, that measured by the results they have been absolutely fruitless. They have disregarded the psychological processes themselves.
>
> *Wilhelm Wundt*

We have reached the end of our initial exploratory foray across the landscape of fast and frugal heuristics. Along the way, we have found unexpected vistas and surprising terrain. In chapter 1 we presented a rough map of the journey to come; now it is time to turn that map around and look at where we have been from different angles.

The heuristics in our mind's adaptive toolbox can be organized and accessed in a number of ways. They can be classified by type of decision task being faced, or according to the adaptive problem that needs to be solved—that is, in terms of a problem's domain-independent form or its domain-specific content. The first scheme divides heuristics into those for estimation, classification, two-alternative choice, and so on. The second organizes heuristics into those for food choice, mate choice, parental investment, inferring intentions from motions, and so on. In this book, we have taken both perspectives, and so can decision makers when seeking the appropriate tool for the problem at hand.

There is a third point of view: Heuristics can be categorized in terms of their building blocks—the particular heuristic principles they employ. In chapter 1 we described three classes of building blocks, namely principles for directing information search, for stopping that search, and for making a decision based on the search results. These can be used to define classes of heuristics that share one or more building blocks. These classes cut across decision tasks and adaptive problems. We now review the major classes of heuristics and the associated visions of rationality before ending this book with open questions still awaiting answers.

Classes of Heuristics

Ignorance-Based Decision Making

Good decisions need not require amassing large amounts of information; they can also be made on the basis of *lack* of knowledge. A basic cognitive adaptation is the ability to recognize faces, voices, smells, names, and other environmental features. There is a class of very simple heuristics based on this adaptation that share one building block: only search for recognition information. This may not sound like much for a decision maker to go on, but there is often information implicit in the *failure* to recognize something, and it can be exploited by these *ignorance-based heuristics*.

The simplest exemplar in this class is the *recognition heuristic*, for which we have found strong experimental evidence (chapter 2). Outside the laboratory, the recognition heuristic influences the behavior of organisms as widely varied as wild rats searching for food and humans deciding on a restaurant for lunch, and underlies the proliferation of identical fast-food chain outlets in much of the world (Schlosser, 1998). It can also be generalized to the task of choosing a subset of objects from a larger set, as in selecting a set of stocks based on recognition alone (chapter 3).

By analyzing and simulating the performance of the recognition heuristic, we arrived at a surprising prediction: Using this heuristic, a person who knows *less* than another can make systematically *more* accurate inferences. In chapter 2 we showed that this *less-is-more effect* is borne out by empirical data. Counterintuitive consequences such as this, which are not predicted by other theories or by common sense, are important indications of the empirical validity and theoretical significance of fast and frugal heuristics.

One-Reason Decision Making

Heuristics in the class of *one-reason decision making* search for reasons or cues beyond mere recognition, either in recall memory or in external stores of information. They use only a single piece of information for making a decision—this is their common building block. Therefore, they can also stop search as soon as the first reason is found that allows a decision to be made. We proposed and studied a variety of heuristics in this class, including the Minimalist, Take The Last, Take The Best (all in chapters 4, 5, and 6), and parental feeding heuristics (in chapter 14).

We were surprised by how accurate one-reason decision heuristics can be, for example outperforming multiple regression across 20 decision environments (chapter 5) and coming within a few percentage points of the accuracy of computationally complex Bayesian models (chapter 8). Thus there seem to be many situations without a trade-off between making a decision fast and frugally and making it accurately. Simplicity can come without a heavy price. The environmental conditions that explain this bargain are the topic of the study of ecological rationality.

There is growing empirical evidence that people actually use lexicographic heuristics such as Take The Best, particularly when time is limited (chapter 7). We have also investigated Take The Best as part of a memory-updating mechanism that underlies hindsight bias, successfully providing the first process model of this phenomenon (chapter 9). The quest for empirical evidence, however, is still burdened with a methodological problem. Policy-capturing methods for tracing thought processes still lack the power to distinguish exactly which heuristic an individual may be using—developing more sensitive methods is a challenge for further research.

Elimination Heuristics

Ignorance-based and one-reason decision heuristics (especially where the one cue has only two values) are most appropriate for tasks where one of two options must be selected. Other tasks call for a different class of heuristics. In categorization, for instance, one category must be chosen from several possibilities. For this type of task we can make fast and frugal, but still accurate, decisions by using an elimination rule rather than one-reason decision making. The class of elimination heuristics uses cues one by one to whittle down the set of remaining possible choices, stopping as soon as only a single category remains. QuickEst (chapter 10) can be seen as taking an elimination approach to estimation. The Categorization by Elimination heuristic (chapter 11) came within a few percentage points of the accuracy of traditional categorization algorithms including exemplar and neural network models, despite using only about a quarter of the available information. In chapter 12, we explored how Categorization by Elimination can be used to make rapid decisions about the intentions of other organisms from their motion cues alone, helping individuals avoid costly conflict or even predatory ingestion. In situations in which categorization must be performed quickly and cues take time to search for, this fast and frugal approach has clear advantages.

Satisficing

The previous three classes of heuristics are designed for situations in which all of the possible options are immediately available to the decision maker: For instance, the categories of possible intentions are all known, and the chicks to be fed are all sitting patiently in the nest (chapters 12 and 14). But a different class of heuristics is needed when alternatives (as opposed to cue values) take time to find, appearing sequentially over an extended period. In this type of choice task, a fast and frugal reasoner should limit not only the search for information (cues) about each alternative, but also the search for alternatives themselves. Herbert Simon (1955a, 1990) has used the term "satisficing" for heuristics that solve this problem by relying on an aspiration level to stop search.

We investigated satisficing heuristics for sequential mate search in chapter 13. Our focus was on simple mechanisms that limit both the time needed to determine a useful aspiration level, and the average number of potential mates considered before finding one who exceeded the aspiration. Simple learning heuristics can indeed find such adaptive aspiration levels, while still coming close to the mate choice performance of more "optimal" (and much slower) rules. The design of heuristics that search through *both* objects and cues, sequentially or simultaneously, is one of the unresolved issues for future research.

Visions of Rationality

We began this book with a triadic vision of bounded, ecological, and social rationality. The three are intimately linked: The success of boundedly rational heuristics depends on their ability to exploit the information structures in the ecological and social environment. Thus, the interaction of these three perspectives is essential for our notion of rationality.

Bounded Rationality

Our research program contributes to the study of bounded rationality on two levels. First, we have laid out a general framework for the construction of fast and frugal heuristics from a small set of building blocks. New heuristics can be formed through the combination of simple principles for guiding information search, stopping search, and reaching a decision. Second, we have explored a variety of specific heuristics that make accurate inferences despite being bounded by limited time, knowledge, and computation. These examples provide clear evidence that a demonic level of power and resources is not necessary for rationality.

Ecological Rationality

There are two reasons for the surprising performance of fast and frugal heuristics: their exploitation of environment structure and their robustness (generalizing appropriately to new situations as opposed to overfitting—see chapter 1). Ecological rationality is not a feature of a heuristic, but a consequence of the match between heuristic and environment. For instance, we have investigated the following structures of environments that 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 (as defined in chapter 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 (as defined in chapter 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 (as described in chapter 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 (as described in chapter 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. 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.

Social Rationality

Some of the most challenging decisions faced by social species are those arising from an environment comprising the decisions of conspecifics. Social environments are characterized by the speed with which they can change and by the need to consider the decisions being made by others. These two features make social rationality an important and distinct form of ecological rationality. We have shown in this book that fast and frugal heuristics can guide behavior in these challenging domains, when the environment is changing rapidly as a result of others' behavior (e.g., in stock market investment—chapter 3), when the environment requires many decisions to be made in a successively dependent fashion (e.g., in parental investment—chapter 14), or when decisions must be made in coordination with other individuals (e.g., in mutual mate choice—chapter 13). These particular features of social environments can be exploited by heuristics that make rapid decisions rather than gathering and processing information over a long period during which a fleeter-minded competitor could leap forward and gain an edge.

These three perspectives on rationality are all defined in terms of an organism's adaptive goals: making decisions that are fast, frugal, accurate, and beneficial, in social and nonsocial situations. Thus, we see rationality as defined by decisions and actions that lead to success in the external world, rather than by internal coherence of knowledge and inferences. Theories of mind that focus on internal coherence have led, in artificial intelligence, economics, and elsewhere (see chapter 15), to models that assume that an individual must create elaborate representations of knowledge and solve impressive equations when making up its mind. The challenge ahead is not to construct models of omniscient minds, but rather of adaptive minds that can act quickly and reliably in their environments.

In *Walden* (1854/1960), Henry Thoreau thought deeply about the relationship that people have with their environment, albeit from a different perspective. His advice is equally appropriate for modeling minds in their environments: "Our life is frittered away by detail. . . . I say, let your affairs be as two or three, and not a hundred or a thousand. . . . Simplify, simplify" (p. 66). Such simplicity in models has a certain aesthetic appeal. The mechanisms are readily understood and communicated, and are amenable to step-by-step scrutiny. Furthermore, Popper (1959) has argued that simpler models are more falsifiable, and Sober (1975) deems them more informative. But the transparency, falsifiability, or informativeness of *models* are not the only grounds to argue for the simplicity of actual mental *mechanisms*. We have provided evidence that simple heuristics are also adaptive for those who actually use them (see also Forster & Sober, 1994). Simplicity can have both aesthetic appeal and adaptive value.

Looking Ahead

In this book we have proposed a variety of fast and frugal heuristics for making adaptive inferences and decisions. For each new heuristic we have endeavored to ask three main questions: How good is it—how well does it perform in comparison with decision mechanisms adhering to traditional notions of rationality? How is it ecologically rational—when and why does it work in real environments? And finally, do people or other animals actually use this heuristic? We certainly do not have all the answers to these queries. In fact, to date, most of our attention has been focused on the first (easiest) question, and while we are starting to gain some understanding about the second, our efforts to answer the third (very difficult and in some ways most significant) question are just beginning. This imbalance needs to be redressed. Looking ahead, there are many open challenges that follow from these three questions. Facing these challenges will not entail lone discovery of wholly new lands: Building on results already found by others, often in other fields and expressed in different languages, will accelerate the process of finding new answers.

Cognitive Tasks

The first challenge is to explore fast and frugal heuristics for solving tasks beyond those we have considered here. 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? 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—the second organizing scheme for the adaptive toolbox mentioned at the beginning of this chapter. The program of carving up an organism's life and behavior into separate adaptive domains, each containing several adaptive problems, has proved to be a great challenge (see Hirschfeld & Gelman, 1994, for the current plethora of approaches). But the discovery of domain-specific heuristics for important adaptive problems may help clarify some of the divisions—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)? Why do people often prefer to solve adaptive problems using socially transmitted information, for instance, deciding what medical risks to take on the basis of hearsay rather than statistical evidence (while at the same time often obsessing about baseball statistics)?

Social Norms and Emotions

Simple heuristics can 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 have also speculated occasionally in this book 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? This is one of the most important areas still to be mapped out.

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 perspectives. For instance, the statistical measures of two-dimensional patterns developed in spatial data analysis (see, e.g., Upton & Fingleton, 1985) can be used when assessing heuristics for spatial search in foraging or habitat selection. Evolutionary psychology reminds us to reflect on possible differences between present and past environments, by considering the important adaptive problems our ancestors faced, the information available in their environment to solve those problems, and how these inputs have changed in the modern world (Cosmides & Tooby, 1987, p. 302).

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 mostly 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 focussed on adaptive goals in terms of correspondence criteria (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? 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 perspective of bounded rationality, a fast and frugal mind need not employ a metalevel demon who makes optimal cost-benefit computations to select a heuristic. The fact that heuristics are designed for particular tasks rather than being general-purpose strategies solves part of the selection problem by reducing the choice set (see chapter 1). But we have not 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 chapter 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.

The Rational Meets the Psychological

Some years ago, sequestered in the hills overlooking Stanford, a gathering of economists and psychologists engaged in an animated conversation on the nature of reasoning. We argued over the latest stories about this or that paradox or stubborn irrationality until finally one of the economists concluded the discussion by throwing down the gauntlet. "Look," he said with the conviction of his field, "either reasoning is rational, or it's psychological." To him, this inviolable dichotomy implied an intellectual division of labor: Rational judgment is defined by the laws of logic and probability, and thus should be the domain of rigorous economists and mathematicians; what we know about the human mind is irrelevant for defining sound reasoning. Only when things go wrong should psychologists be called in to explain why people can be irrational.

We hope that the simple heuristics analyzed in this book exemplify a way to break down this unfortunate but widespread belief in an opposition between the rational and the psychological. This misleading idea has cursed the cognitive sciences since the antipsychologism of nineteenth-century philosophy, and it continues to obscure a realistic view of cognition to this day. A bit of trust in the abilities of the mind and the rich structure of the environment may help us to see how thought processes that forgo the baggage of the laws of logic and probability can solve real-world adaptive problems quickly and well.

Models of reasoning need not forsake rationality for psychological plausibility, nor accuracy for simplicity. The mind can have it both ways.