

Précis of Simple Heuristics That Make Us Smart

Peter M. Todd and **Gerd Gigerenzer**

Center for Adaptive Behavior and Cognition,
Max Planck Institute for Human Development, Berlin, Germany

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.

Key words: 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.

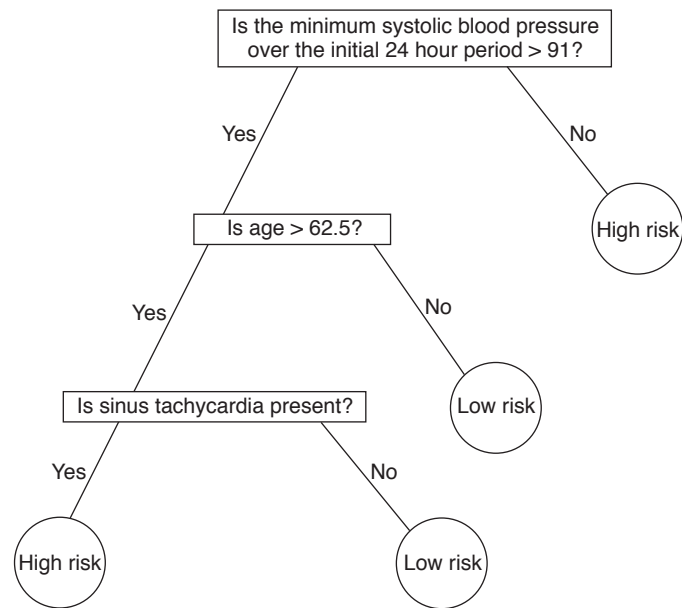


Figure 1. A simple decision tree for classifying incoming heart attack patients into high risk and low risk patients (adapted from Breiman et al., 1993).

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

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

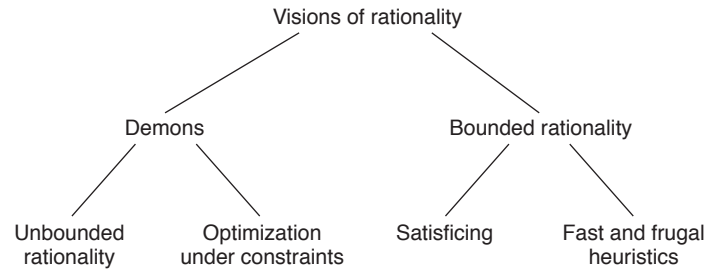


Figure 2. Visions of rationality.

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: *satisficing 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.

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 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 Chap. 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-Spratlike 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 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 Sec. 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 Chaps. 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 (Chap. 6), or based on a recollection about which cues worked previously when making the same decision (Chap. 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 (Chap. 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 conditional 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 (Chap. 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 (Chap. 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 Chap. 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 Chaps. 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 Chaps. 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 (Chaps. 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 (Chaps. 4 and 5; see also Sec. 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* (Chap. 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’ 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 recognition to decide which of two cities is larger will often yield the correct answer (in those cases where one city is recognized and the other is not). 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 (Chap. 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 (Chap. 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.

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 (Chap. 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

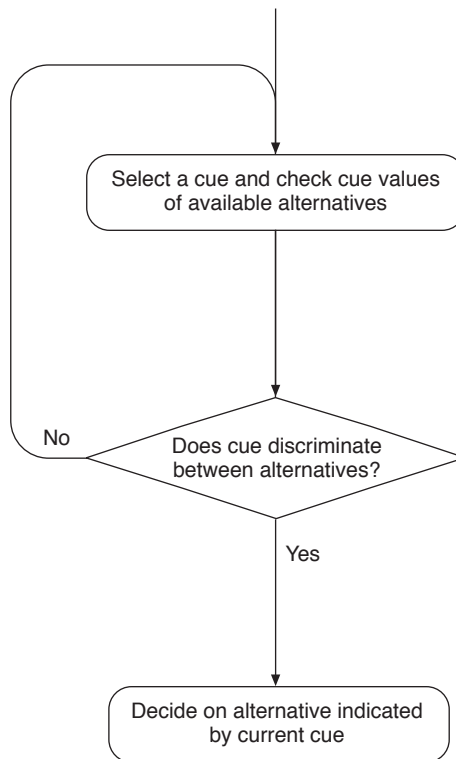


Figure 3. A flowchart of one-reason decision making. First, search for a cue and the corresponding cue values of each alternative; next, check whether the values for that cue discriminate between the alternatives; if so, then choose the indicated alternative; if not, select another cue and repeat this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

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 (Chap. 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 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' 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' 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 (Chap. 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 (Chap. 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' 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 (Chap. 8). (These results also show the well-known "flat maximum" result that a linear model with equal-sized weights, e.g., Dawes' 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

Table 1
Performance of Different Decision Strategies Across 20 Data Sets

Strategy	Frugality	Accuracy (% correct)	
		Fitting	Generalization
Minimalist	2.2	69	65
Take the Best	2.4	75	71
Dawes' 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' 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 \neq training set). (Data from *Simple heuristics*, Chap. 5.)

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.

But how? We turned to mathematical analysis (Chap. 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 (Chap. 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 (Chap. 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 possible, 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 (Chap. 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 (Chap. 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 (i.e., 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 (Chap. 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, i.e., 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 (Chap. 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 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 (Chap. 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 (Sec. 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 trade-offs between different real-world criteria including speed and accuracy? The main reason for their success is that they make a trade-off 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 (Chap. 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 data sets 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 (Chap. 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 (Chap. 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 (Chap. 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 (Chap. 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 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 Sec. 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.

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 trade-offs. 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 trade-off 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 topical 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 (Chap. 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 (Chap. 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 (Chap. 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 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 cryptarithmic 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 (Chap. 4) and in practice (Chap. 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 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 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 (Chap. 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 Chap. 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 (Chap. 12), males and females must make decisions about others they are interested in mating with (Chap. 13), and parents must figure out how to help their children (Chap. 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.