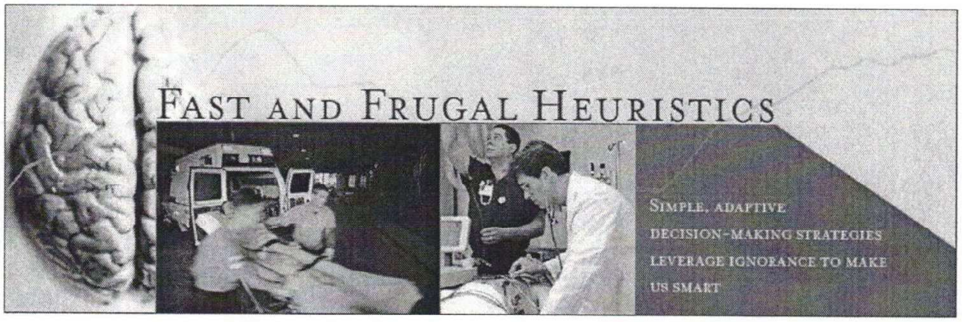

Fast and Frugal Heuristics

by Gerd Gigerenzer, Max Planck Institute, Berlin, 2001

Contributed by the author

Based on *Simple Heuristics That Make Us Smart* by Gerd Gigerenzer, Peter M. Todd, and the ABC Research Group, Oxford University Press, 1999

Summary

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.

Fast and frugal heuristics – simple rules in the mind's adaptive toolbox for making decisions with realistic mental resources – 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. These simple heuristics perform comparably to more complex algorithms, particularly when generalizing to new data – that is, simplicity leads to robustness.

A heuristic could save your life

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 and colleagues 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

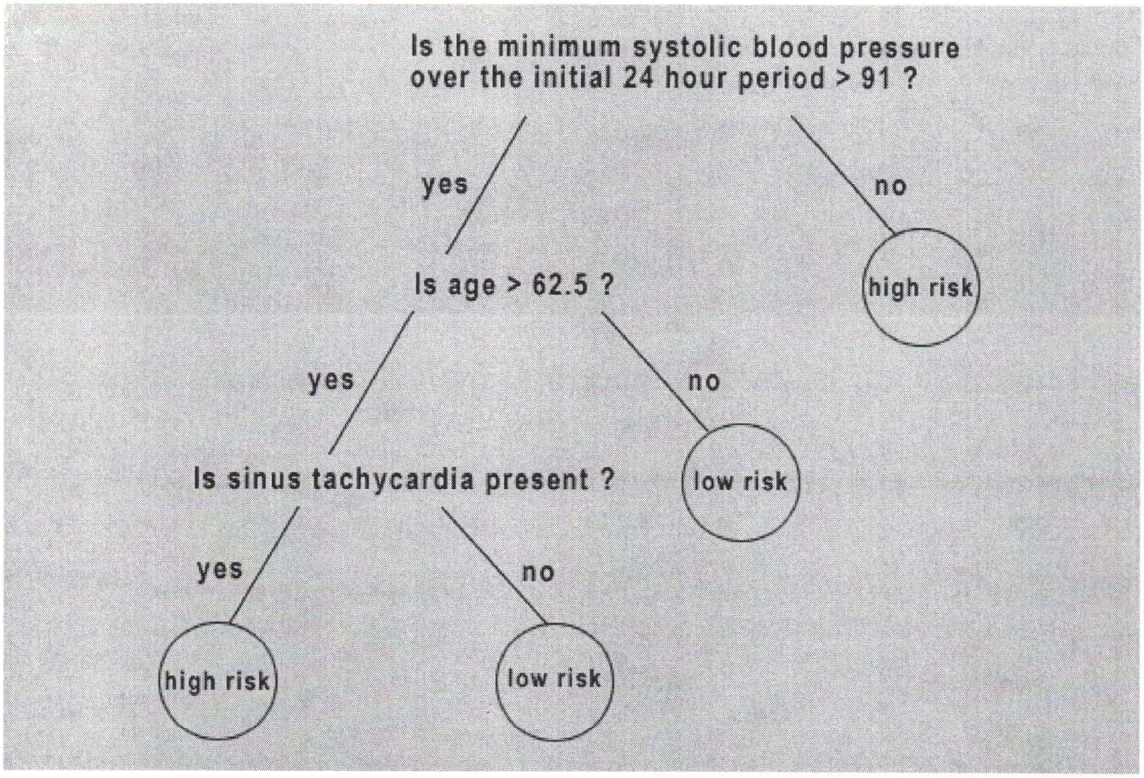

This decision strategy is simple in several respects. First, it ignores the great majority of possible measured predictors. Second, it ignores quantitative information by using only yes/no answers to the three questions. For instance, it does not care how much older or younger the patient is than the 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.

Bounded rationality

Since the Enlightenment, the main model of rational judgment in an uncertain world has been probability theory. The laws of probability, however, do not deal with the constraints in time, information, memory, and other resources that are characteristic of the decision making of actual humans (and machines). As a consequence, the underlying vision of rationality has been termed "unbounded rationality"—an omniscient and omnipotent fiction with little or no regard for the limitations in time, knowledge, and computational capacities that humans face. To make rationality more human than God-like, the concept of "bounded rationality" has been proposed. The key difference between unbounded and bounded rationality is the concept of *limited search*, to be defined by a *stopping rule*. The vision of bounded rationality, however, is not of one kind.

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. Optimization under constraints means optimization given various constraints, that is, limited resources such as attention, time, money, or information. The vision of constrained optimization is that minds would calculate the optimal trade-off between the benefits and costs of further search at regular time intervals, and stop search when the costs would outweigh the benefits. The rule "stop search when costs > benefits" sounds plausible at first glance, but a closer look reveals that this… can demand even more knowledge and computation than unbounded rationality.

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.

Figure 2: Visions of rationality

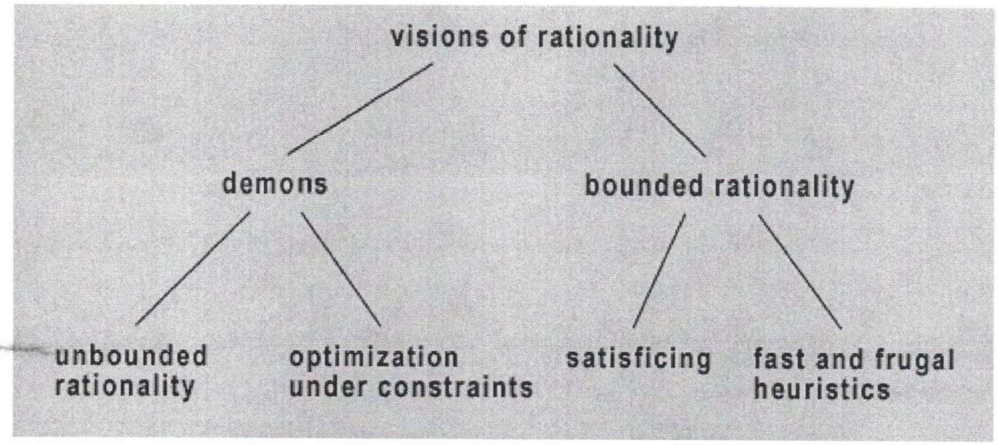

What are fast and frugal heuristics?

1. **Fast and frugal heuristics are psychological adaptations – not incomplete or distorted versions of "optimal" strategies**. There are two opposite starting points for thinking about satisficing. One is to start with an optimal model (of foraging or risk taking, for instance) and prune it down with simplifying constraints, such as deleting information the organism cannot know or does not pay attention to. In this tradition, "descriptive" accounts of human judgment and decision-making… often start with a "normative" rule, and then explain actual behavior as an incomplete or distorted version of the same rule that neglects, overweighs, or underweighs some parameters. The heuristics that people and animals use will be interpreted by this approach as "suboptimal" or even "irrational."

In contrast, we start with the psychology of an organism rather than with some statistical norm. In our view, heuristics are simple rules that build on highly complex adaptations, such as recognition memory and recall memory. Heuristics are not simply handicapped versions of optimal strategies; there are no optimal strategies in many or most real-world environments in the first place. This should not be confused with the absence of performance criteria: heuristics can be measured in terms of their accuracy, speed, information use, and so on.

2. **Fast and frugal heuristics embody limited information search by using simple stopping rules**. Simple stopping rules do not involve cost-benefit computations to try to find an "optimal" point at which to stop searching for information on which to base a decision. For instance, here is a simple stopping rule that can be used when searching for information for choosing between two alternatives: "terminate search when the first good reason is found that speaks for one alternative over the other". No other cues are looked up after this point, and no cost-benefit computations are made. Limited search with simple stopping rules allows an organism to predict its world with limited time, limited knowledge, and limited computational capacities.

An example: ignorance-based decision making

The *recognition heuristic* exploits the vast and efficient capacity of recognition to make inferences about unknown aspects of the world. The processes underlying recognition, such as face, voice, or name recognition, are anything but simple. These are complex mechanisms, the products of millennia of evolution, and are still far from being understood. What they give us, however, is a sense of recognition that can be exploited by a very simple heuristic. The knowledge demands of this simple heuristic are so low that it actually requires a beneficial lack of knowledge to work.

When the task is to infer which of two objects has a higher value on some criterion (e.g., which is larger, safer, stronger), the recognition heuristic is simply stated: *If one of two objects is recognized and the other is not, then infer that the recognized object has the higher value.*

Let me illustrate the way this heuristic works with one example: *Which US City Has More Inhabitants: San Diego or San Antonio?* We posed this question to students at the University of Munich and the University of Chicago. The latter, who have a reputation for being among the most knowledgeable in the US, were correct 62% of the time. Yet 100% of the Germans got the correct answer 100% of the time. How did the Germans infer that San Diego was larger? All of the Germans had heard of San Diego, but many of them did not recognize San Antonio. They were thus able to apply the recognition heuristic and make a correct inference. The American students were not *ignorant* enough to be able to apply the recognition heuristic.

How good are fast and frugal heuristics?

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: Stop looking for cues as soon as one is found that differentiates between the two options being considered...

We have developed three fast and frugal one-reason decision heuristics that differ only in their search rule: *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's Rule) that counts up all of the cues for and against a choice and looks at the difference. We found that the simple heuristics always came close to, and often exceeded, the proportion of correct inferences achieved by multiple regression and Dawes's Rule...

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's Rule	7.7	73	69
Multiple regression	7.7	77	68

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

1,900 words

- Heuristics aid decision-making by providing stopping rules - at some point more information gathering can impede rather than support decisions
- Human beings have adaptive mechanisms for exploiting limited information - pattern recognition and decision rules allow for robust

decision-making under uncertainty

- Traditional models of rationality are unrealistic - complete exploration of probabilities and cost-benefit analysis is not feasible given real world constraints

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