

Bounded Rationality

The Adaptive Toolbox

Edited by

G. Gigerenzer and R. Selten

Program Advisory Committee:

G. Gigerenzer and R. Selten, Chairpersons
R. Boyd, P. Hammerstein

The MIT Press

Cambridge, Massachusetts
London, England
2001

10

Group Report: Why and When Do Simple Heuristics Work?

Daniel G. Goldstein, Rapporteur
Gerd Gigerenzer, Robin M. Hogarth, Alex Kacelnik,
Yaakov Kareev, Gary Klein, Laura Martignon,
John W. Payne, and Karl H. Schlag

In a remote stream in Alaska, a rainbow trout spies a colorful dimple on the undersurface of the water with an insect resting atop it. Darting over with mouth agape, the fish bites down and turns in search of its next victim. It does not get far, however, before the “insect” strikes back. The trout is yanked from the quiet stream by the whiplike pull of a fly fisherman’s rod. In a world without fishermen, striking at all that glitters is adaptive; it increases chances for survival. In a world with predators, however, this once-adaptive strategy can turn a feeding fish into a fisherman’s food.

Herbert Simon provided the metaphor of a pair of scissors for thinking about rational behavior: one blade has to do with the psychology of the organism and the other with the structure of the environment. In this chapter, we think of the two blades as the simple heuristics used by organisms and the structures of environments that govern their performance. As the case of the poor trout illustrates, a strategy cannot be evaluated without taking into account the environment in which it operates. Attempts to model just one blade of the scissors simply will not cut it. We provide several examples of simple heuristics at work in the world and discuss how the environment determines their success or failure. To do this, however, we are led to ask what counts as success or failure for a strategy. Focusing on the capabilities of the organism and the structure of the environment, we explore the question of how organisms may choose among the heuristics available to them. In conclusion, we generalize about the properties of

environments that allow heuristics to work and speculate about the psychological building blocks of decision strategies.

SIMPLE HEURISTICS AT WORK IN THE WORLD

Strategies in the adaptive toolbox are fast and frugal. Fast refers to the relative ease of computation the strategies entail, which has been measured as order of complexity or with elementary information-processing steps (Payne et al. 1993; Czerlinski et al. 1999). Frugal refers to the very limited amount of information these strategies need. Each strategy is described along with a discussion of how the structure of the environment can cause it to succeed or fail.

Imitation

Imitation is a fast and frugal strategy that saves an organism from having to extract information from the environment anew, or from calculating from scratch. For instance, New Caledonian crows use imitation to make tools from Pandanus tree leaves to extract grubs from trees (Hunt 1996). Evidence of imitation is seen in the local similarity of the tools — birds of a particular region tend to make tools of a similar shape (such as facing left or right, for instance). Humans embrace imitation as well and have been observed to imitate those who are compatible with themselves. Hogarth et al. (1980) studied industrial plant relocation decisions of small firms in a town in southern France. The owners of these firms had great difficulty in making these decisions and did not want to commit themselves even though they could see large foreign firms taking advantage of conditions offered by local government. However, when they observed the move made by a local mid-sized firm with a high reputation and whose “values” were compatible with theirs, they quickly adopted the same strategy and moved. A couple of years after these decisions, the small business owners were glad they had relocated. Imitation certainly has its advantages; however, it may not lead to desirable outcomes in all situations, as we shall soon see.

What makes it work? To answer this question, it is perhaps most illuminating to think about when imitation will *not* work. If the environment is changing rapidly, imitation will fail because a successful strategy at the time of observation may no longer be effective, or even executable, at a later time. If the environment masks or obscures what individuals are doing, imitation will be impossible or unwise. The environment also includes other organisms, and the available choice of whom to imitate matters as well. Sometimes, this is a Hobson’s choice, and the results can be costly if the sole alternative is not the ideal one. The Mapuche Indians of the high country of Chile have grown wheat for several generations, but before the emergence of a crop disease long ago, they grew barley. Today, barley would be a more profitable crop for them, and modern-day pesticides have solved the disease problem (Henrich 1999). However, attempts by

developers to reintroduce the crop have been met with the Mapuche’s response that barley farming “is not the present custom” and that they would rather stay with their “traditional” crop. In a world without pesticides, there is no need to pass down knowledge of how to grow crops that are prone to diseases. In a world with access to new ideas and technologies from other cultures, lack of willingness to try the new can hinder progress.

One cannot get specific about how the environment affects the success of imitation without getting specific about the exact kind of imitation at hand, as Schlag (1998) has proven. One rule considered, called “Imitate If Better,” assumes that individuals imitate all others who are more successful than themselves, and stick with their current strategies otherwise. Surprisingly, in risky environments this can lead the entire population to choose the alternative with the lowest expected payoff. On the other hand, another quite simple rule called “Proportional Imitation” — which dictates imitating those who are more successful than oneself with a probability that is proportional to the difference between the observed and current degrees of success — always leads the population to the expected payoff-maximizing action. Imitation is promising as one of the simplest and most effective heuristics an organism can employ. As Schlag and others have shown, the success of imitation depends upon the environment, its inhabitants, and the exact kind of imitation being applied. For a discussion of how imitation compares with individual learning, see also Boyd and Richerson (this volume).

Equal Weighting

Multiple regression and Bayesian networks are exemplary of complex models that capture desired properties of classical rationality: they consider all information, and they integrate this information in a compensatory way. Additionally, they apply complicated routines to find the optimal set of weights (the beta weights in regression, or the internode weights in Bayesian networks) that the model will use. Some thinkers (Dawes 1979; Einhorn and Hogarth 1975) have considered a simpler class of model that replaces real-valued weights with simple unit weights (such as +1 and -1). Though one might expect performance to drop appreciably, in some situations these simple models can outperform multiple regression (see also Gigerenzer and Goldstein 1996).

What makes it work? As an example, imagine ranking applicants to a graduate program where candidates’ profiles are measured on a number of variables such as test scores, quality of previous education, letters of recommendation, and so on. If you believe that these variables are compensatory or trade-off, then simply give all variables equal weight in the evaluation. The advantage of this procedure is that you cannot make the error of reversing the size of weights attached to two variables. In other words, if — unknown to you — one variable is more important than another but, because of sampling fluctuations, estimation

by regression reverses the “true” relative sizes of the weights, regression will make inferior predictions. Equal weighting insures you against making this kind of error. More generally, Einhorn and Hogarth (1975) showed that equal weighting makes better predictions than regression as (1) the number of predictor (or independent) variables increase, (2) average inter-correlation between predictors increases, (3) the ratio of predictors to data points (on which regression weights are estimated) increases, and (4) the R^2 of the regression model decreases. To see how equal weighting can help you pick the “all star” basketball team or decide how many and which forecasters you should consult in making a “consensus” forecast, see Einhorn and McCoach (1977) and Hogarth (1978).

Take The Best

Take The Best is a heuristic from the adaptive toolbox (Gigerenzer and Goldstein 1996) that neither looks up nor integrates all available information. It is a lexicographic procedure (similar to the LEX model tested by Payne and colleagues) that uses a rank ordering of cues to make inferences and predictions (Martignon and Hoffrage 1999). Cues are searched through one at a time, until a cue that satisfies a stopping rule is found. The decision is made on the basis of the cue that stopped search, and all other cues are ignored. In empirical tests, Take The Best used less than a third of all information available to it. Remarkably, despite its simplicity, Take The Best can make predictions that are more accurate than those made by multiple regression and approximates the accuracy of Bayesian networks (Martignon and Laskey 1999; Czerlinski et al. 1999).

What makes it work? In Take The Best, the decision made by a higher-ranked cue cannot be overruled by the integration of lower-ranked cues. Its predictions are equivalent to those of linear models with noncompensatory families of weights (Martignon and Hoffrage 1999), e.g., consider the linear model:

$$y = 8x_1 + 4x_2 + 2x_3 + 1x_4 \quad (10.1)$$

where x_i is a binary (1 or 0) cue for $i = 1, 2, 3, 4$. Each term on the right-hand side cannot be equaled or exceeded by the sum of all the terms with lesser weights. If cues are not binary but have positive real values, which become neither infinitely small nor infinitely large (i.e., bounded from below and from above by strictly positive real numbers), it is always possible to find weights that determine such a noncompensatory linear model equivalent to Take The Best in performance. If the “true” weights of the cues (i.e., those of an optimal model like regression) are noncompensatory, then Take The Best cannot be beaten by any other linear model when fitting data.

When making predictions on new data, the frugality and simplicity of Take The Best are responsible for its robustness. Here the predictive accuracy of Take The Best is comparable to that of subtle Bayesian models, often surpassing optimal linear models (which tend to overfit). A variant of Akaike’s Theorem

(Forster and Sober 1994; Martignon, this volume) explains when and why Take The Best makes more accurate predictions than models that base their decisions on larger sets of adjusted parameters.

Take The First

Experts, such as pilots, firefighters, and chess players, have a simple strategy at their disposal: when faced with a problem to solve, often the best course of action to take is the first (or only) one that comes to mind. The strategy of evaluating solutions as they come to mind, and stopping with the first one that satisfies an aspiration level, is called Take The First. (Actually, this strategy has been described as the recognition-primed decision [RPD] model [Klein 1998], but for purposes of this volume, we use a language that better conforms with the heuristics described by the ABC Research Group [Gigerenzer et al. 1999].) Experiments have been carried out to test the quality of solutions that come to the minds of experienced chess players (Klein et al. 1995). When experts were asked to rate all possible moves from given board situations, they rated only one in six as worthy of consideration. However, when looking at the set of first moves that came to the minds of seasoned players, the experts evaluated four out of six of these moves worthwhile. Furthermore, even if the first move considered was not playable, the flaws were usually discovered quickly, so the cost of rejection was low in terms of time and effort.

What makes it work? Take The First is argued to be effective because, for an expert, part of recognizing or categorizing a situation as typical is to recall what to do in that situation. Options generated are not random but may come to mind in order of quality. Take The First is less successful in domains where the decision maker is not an expert or in completely novel situations within a domain of expertise. (Prescriptive decision analysis methods also struggle under these conditions.) It has limited effectiveness in domains where learning is difficult, such as domains with noisy feedback, or where there are low costs for making errors.

Small-sample Inferences

We all detect meaningful covariation fairly rapidly, seemingly despite the fact that we have only small samples to draw upon. Kareev (1995) has shown that, for correlations, the limitation in sample size imposed upon us by working-memory capacity actually acts as an amplifier of correlations. This allows organisms to detect relationships in the environment because of, not despite, their limited working-memory capacity. With the degree of amplification negatively related to sample size, this effect is more pronounced for children, who both have a smaller capacity and are more in need of detecting correlations.

What makes it work? This model of correlation detection gets its power from a mathematical truth, namely, that small samples tend to overestimate Pearson correlations. It benefits an organism in tasks and domains where the costs of missing a relationship in the world is high.

The Recognition Heuristic

When German and American subjects were asked which of San Diego or San Antonio had a greater population (Goldstein and Gigerenzer 1999), all of the Germans correctly answered that San Diego was larger. Only two-thirds of the Americans got the right answer. Many of the Germans had only heard of San Diego, and not of San Antonio, and chose on the basis of what is called the recognition heuristic. In one formulation, the recognition heuristic dictates that when choosing between two objects, if only one of them is recognized, then choose the recognized object. As it did with the Germans and Americans, this heuristic can lead to a counterintuitive state of affairs in which those who know more perform worse than those who know less, the so-called *less-is-more effect*. The adaptive advantage of this strategy is that it exploits a resource most organisms have in abundance: missing knowledge.

What makes it work? The recognition heuristic depends foremost on missing knowledge—only when some objects are not recognized can the heuristic come into play. However, not any kind of missing knowledge will do when being correct is crucial. Only when ignorance is systematically, rather than randomly, distributed can it be used to make accurate inferences. If the cities people did *not* recognize were not particularly small ones, this missing information could not be used to increase accuracy. It could even hurt it. Luckily, in a wide array of domains (such as the deadliness of diseases, the length of rivers, or the success of sporting teams), the objects people recognize stand out on dimensions they find important. Indeed, it is because people find these objects important that they talk about them, and in so doing assure them a place in the recognition memory of others.

MEASURING SUCCESS

We have looked at a short list of simple strategies thought to be used by human decision makers and discussed structures of the environment that affect their chances of success or failure. A large question remains, however: What does it mean for a strategy to succeed, to fail, or to be optimal?

Optimality

Optimality can refer to outcomes or processes. Optimal outcomes are considered to be the best behavioral consequences an organism can achieve given the

information it has available. Any means may be used by scientists to compute optimal outcomes, from analytic proof, to simulation methods, to mechanized optimization routines. Similarly, many processes can lead an organism to achieve these outcomes. For instance, a seedling may achieve optimal orientation towards a light source by differential growth of the two sides of its stem. Optimizing processes, on the other hand, are considered to be those that satisfy various criteria of rationality, for instance, that all available information is considered, or that information is weighted in an optimal (for instance, in a least-squares minimizing) way. Optimality models of animal behavior utilize the assumption of optimality at the level of outcomes, but not at the level of processes. Indeed, in many cases, identification of the process used by the organism can explain why an optimal outcome is *not* achieved under specific circumstances, such as those created in the laboratory. If the process used by the organism was selected in an environment different from that in which it is being tested, this result is not surprising.

The more accurate the specification of the optimal outcome, the more helpful it will be in guiding research into the processes controlling behavior. In the following example, the contrast between a predicted optimum and observed behavior serves to guide research into behavioral mechanisms.

Consider an animal choosing between two foraging sites that differ in prey distribution. In habitat *A*, prey are found regularly at intervals of *F* seconds. In habitat *B*, prey are found in pairs, so that they take intervals of either *S* (for short interval) or *L* (for long interval) seconds between prey.

We may ask how long *F* should be so that the two environments will be equally desirable to an organism. The scientist first interprets the problem faced by the organism as that of maximizing the overall rate of gain. Equalizing rates (Rate *A* = Rate *B*), an animal should be indifferent between the sites when:

$$\frac{1}{F} = \frac{2}{S+L}, \text{ or } F = \frac{1}{2}(S+L). \tag{10.2}$$

With this normative solution, or “optimum,” worked out, the matter can be tested empirically. Bateson and Kacelnik (1996, 1997) conducted a series of experiments with starlings, where the birds had to choose between two colored keys that delivered food according to schedules as shown in Figure 10.1, with *S* and *L* equal to 3 seconds and 18 seconds, respectively. In a typical experiment, they used a titration procedure: when the subject chose “*A*,” *F* grew by 1 second, but when it chose “*B*,” *F* became 1 second shorter. They found that birds were indifferent when *F* ≅ 5.14 seconds. This result differs from the expected *F* = 10.5




Figure 10.1 The occurrence of food at different time intervals in various environments. Each dot represents one second of search time and each x represents a prey capture.

seconds predicted by Equation 10.2. This proves puzzling in light of the fact that the birds would be getting more food per time unit by acting otherwise. For instance, if they faced a habitat A where $F = 8$ seconds, they would choose B , even though they would be getting one prey every 8 seconds in A and only one prey every 10.5 seconds in B .

Since there was a gulf between the computed optimum and what was observed in the data, it was appropriate to go back to the blackboard to make another attempt at specifying an algorithm that may be a better predictor. The empirical data turned out to match the equation:

$$\frac{1}{F} = \frac{1}{2} \left(\frac{1}{S} + \frac{1}{L} \right), \tag{10.3}$$

which describes F as the harmonic mean of S and L , rather than the arithmetic mean described by Equation 10.2. Similar results have been reported in other species and with other experimental protocols, and the cost of employing such policies is well known (Gilliam et al. 1982; Mazur 1984; Gibbon et al. 1988). Why would birds use such a policy? The answer may lie in the process of choice.

The birds chose between colored keys which had no intrinsic value, other than that acquired by association with food. The process by which these associations are acquired can be modeled as shown in Figure 10.2.

To use this model in the context of choice between the two sites, it is assumed that choices are controlled by the subjective value of the stimulus (colored key) signaling each place. V_i indicates the value of option i .

Figure 10.2 shows that when

$$F = \frac{1}{2}(S + L), \quad V_F < \frac{1}{2}(V_S + V_L). \tag{10.4}$$

Actually, if

$$m(t) \equiv \frac{1}{t}, \quad V_F = \frac{1}{2}(V_S + V_L) \text{ when } \frac{1}{F} = \frac{1}{2} \left(\frac{1}{S} + \frac{1}{L} \right), \tag{10.5}$$

as found experimentally. In this approach, the researchers “explain” why the subject fails to optimize (namely, maximize rate of gain over time) by the mechanisms that it uses to choose. This finding leads to another target for research, namely the adaptive significance of the process by which stimuli acquire value. Biologically, what needs to be explained is why natural selection has not eliminated this mechanism. One answer is that associative learning is a tool that has a much wider domain than the problem posed by the foraging task. Associative learning is a rather general tool to predict events by temporal contiguity (this is a deliberate simplification of the laws of learning). In terms of the blades of Simon’s scissors, one could say that the organism has evolved in environments where the advantages of employing such an associative learning process for predicting relations between events outweigh the losses due to the occasional foraging costs the mechanism may entail. In Figure 10.3, we summarize the research path.




Figure 10.2 A simple hypothetical model for the acquisition of value by arbitrary stimuli (colored keys) while foraging in an environment as depicted in Figure 10.1. The open circle marks the time in which the key lights up with a color and the solid symbols marks the times at which food is delivered. The solid circle indicates the typical time of delivery in site A and the solid squares the two possible times in site A . $m(t)$ shows the “memory trace” of the onset of the key light. We assume that the key gains associative strength in proportion to the value of $m(t)$ when food occurs.

In this and many other foraging examples, “optimality” is assumed and defended at the evolutionary level, but is not proposed as the mechanism of choice used by the subject. When faced with a gulf between hypothetical optimum and outcome, the model was revised to take account of the psychological properties of the organism. In the end, it was posited that the deviation from optimality was

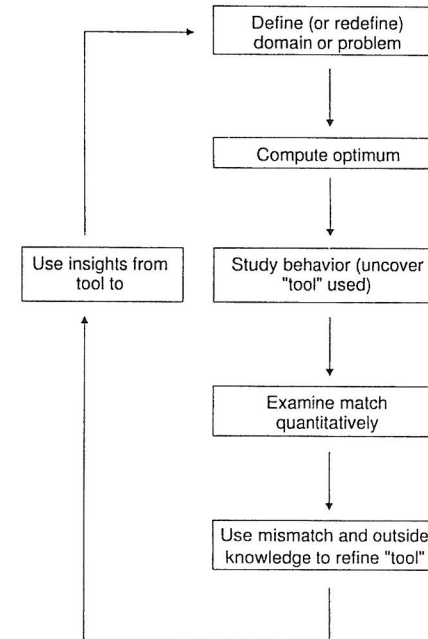


Figure 10.3 A research path.

due to the fact that a more general associative learning mechanism, not evolved to handle this particular task, was being used, and thus the individual is found performing at a less than optimal level. This explanation is reminiscent of Campbell's (1959) hypothesis, i.e., when one observes biases in behavior, it is important to check whether this is not due to a competing response function of the same mechanism. Organisms are less complex than their environments and thus certain responses must be able to handle multiple tasks. Depending on environmental payoffs, trade-offs should be expected.

Other ways for a researcher to explain deviations from predicted optima have to do with changing the level of analysis. Many decisions that are considered nonoptimal for the individual could make sense when one considers that the individual exists within a group. Most small businesses fail, so it arguably does not make sense for an individual to open one. A collection of overoptimistic entrepreneurs, however, each perhaps making the "wrong" move by attempting to open a small business, could, under certain assumptions, lead to an economy in which they are on average wealthier than if they had all made the "right" decision, not been enterprising, and stagnated the economy. Clearly some apparent "irrational" decisions appear more rational upon further considerations of a decision maker's goals for a task, or for instance, when looking at equilibria instead of simple optima. In evolutionary biology, the computation of optima is by no means the only or preferred research tool. The analysis of evolutionary stability is often more useful, as the complexity of frequency-dependent problems tends to make analytical identification of optima impossible. The vast literature on the evolution of cooperation is dominated by the identification (mostly through simulation) of putative evolutionarily stable interaction rules, such as Tit-for-Tat, Pavlov, or image scoring (Nowak and Sigmund 1998a, b). The example of how self-motivated businessmen may be led to equilibria where all players overestimate their individual chances may require this form of treatment, as might the coexistence of sellers and buyers for each share in the market. We should add, however, that all attempts to "rationalize" behavior are disputed by some who believe that ample evidence of nonoptimal decisions and nonoptimal equilibria can be seen in both the laboratory and in the "real world" (e.g., Thaler 1991; Tversky 1996).

Coherence

Many psychological and economic research programs concern themselves not just with an organism's behavior compared to an optimum, but also with how several instances of behavior form patterns. Coherence is the degree to which an organism's patterns of behavior satisfy various criteria of rationality, such as transitivity, consistency over time, and so on.

Incoherence, like deviations from optimality, can also be explained away. Consider the example of the child who hates fish and refuses to eat it, but who

can be tricked into eating and liking it through the use of food coloring. When the parent accuses her of being inconsistent, she replies that she has been consistent in her true objective: to nonplus the parent.

Some scholars address incoherence by observing that although a strategy may in principle admit some flaws from the standpoint of classical rationality, the structure of the environment is such that these deviations will have little or no consequence in practice (Fraser et al. 1992). For instance, a heuristic called Minimalist studied by Gigerenzer and Goldstein (1996) admits intransitivities; however, when tested in a real-world environment, it was nearly as accurate, or better, than some linear models that never make intransitive choices. Contrasting such views, many researchers feel that coherence is a major concern for human decision makers, since we live in a constructed economic world, rigged with traps set by those who wish to exploit vulnerable, incoherent decision strategies. Thus, unless people are consistent in their assessment of probabilities, others may exploit their behavior so that they become, in effect, "Dutch books" (Savage 1954). In other words, whatever bets they place, they will always lose. Similar fates await those who would consistently use intransitive choice rules because they could be turned into "money pumps" (Tversky 1969).

CHOOSING AMONG STRATEGIES

A review on strategy selection by Payne and Bettman (this volume) shows how various structures of the environment (such as the number of alternatives, correlational structure, or the presence of time pressure) and the concerns of the decision maker (such as accuracy, effort, or ease of justification) affect which decision strategies human decision makers choose. We look now at the meta-decision problem and then suggest some ways in which the capabilities (perceptual and cognitive) of the actor and the structure of the environment may do much of the work in strategy selection.

Deciding How to Decide

The deciding-how-to-decide problem is inherent in the idea that there is an adaptive toolbox for the solving of decision problems; that is, individuals are postulated to have a toolbox of different heuristics, and these different heuristics perform differentially across task environments. If no single heuristic works well in every environment, this suggests that an individual must choose the appropriate heuristic for a given situation, i.e., decide how to decide. The question of how people decide how to decide is an ongoing research question subject to debate. In this section we briefly review some of the issues and evidence relating to how people select a tool or tools from the toolbox of decision strategies.

One view of strategy selection is that a decision maker, when faced with a judgment or choice task, evaluates the available tools in his or her toolbox in

terms of relative benefits and costs and selects the one that is best fitted for solving the decision problem. This “top-down” view of strategy (tool) selection is consistent with the evidence that people do plan how to solve problems (Anderson 1983) in a variety of cognitive tasks. For example, as discussed by Payne and Bettman (this volume), there is evidence of such planning in solving decision problems as well. There is also evidence that a crucial difference among students in terms of reading skill is the ability to adjust different reading processes (e.g., reading approaches) to different reading tasks (Garner 1987). More generally, Gagné (1984) argues that good problem solving requires strategic knowledge of when and how to use procedural (tool) and declarative (fact) knowledge.

The top-down view of tool (or strategy) selection has the potential for infinite regress. On the basis of benefits and costs, one decides how to decide, how to choose, ... The infinite regress problem in deciding how to decide is similar to the problem in game theory regarding one’s beliefs about the beliefs of others, who have beliefs about one’s beliefs about the beliefs of others (Lipman 1991). The potential for infinite regress in strategy selection is a serious issue. However, there are both theoretical reasons and empirical evidence to suggest that the sequence of decisions on deciding how to decide may quickly converge to some fixed point (Lipman 1991; Nagel 1995). In her work on games, for example, Nagel (1995) reports that the depth of reasoning about what others are thinking about what you are thinking does not extend much beyond one or two levels.

A different view of strategy selection sees the issue of deciding how to decide as involving a much more bottom-up process. Instead of a conscious decision on how to decide, the selection of a tool from the toolbox may reflect a learned response that has related various task factors to the effectiveness and efficiency of various strategies (for a related view, see Rieskamp and Hoffrage 1999). For instance, people may have learned that the Take The Best rule works very well for certain tasks, and that rule is then typically evoked whenever those situations are encountered. This bottom-up view of strategy selection avoids the infinite regress problem. However, the learned response approach does raise the problem of deciding how to decide when faced with a novel situation.

Of course, it is quite likely that the strategy selection problem is sometimes solved in a top-down way and at other times in a much more bottom-up fashion. One possibility is that the top-down approach to strategy selection will be seen more often when people are faced with complex problems and have the time to decide how to decide. It is also possible that people may start with an approach to a decision problem and then constructively adjust their processing during the course of making the decision as they learn more about the problem structure. Processing, in other words, can change on the spot in an “opportunistic” fashion (Hayes-Roth and Hayes-Roth 1979). In any event, asking the question of “when do simple heuristics work?” clearly raises important issues, such as when and how heuristics that work differentially in different situations are more or less

likely to be selected. Below, we explore the strategy selection question as it is affected by the interplay between the capacities of the organism (both perceptual and cognitive) as well as the structure of the environment.

Perceptual Specialization as a Strategy Selector

The particular perceptual system of an organism might cause it to choose among strategies in a way an onlooker will not suspect. Heyes and colleagues (Heyes 1993; Heyes and Dawson 1992) studied imitation by placing two rats face to face in adjoining cages. One rat, the demonstrator, learned to push a joystick to the left to be rewarded with food, while another rat, an observer, looked on. From the observer’s head-on perspective, the joystick appeared to move to the right, not left, to trigger the release of food. Later, the observer was later placed in the demonstrator’s cage. Which way would it push the joystick? It correctly pushed the joystick to the left. What perceptual strategy did the rat use to compute the direction? The scientists wondered if rats could map observed actions onto their own bodies when undertaking imitative behavior. However, in later experiments, the experimenters learned that by giving the joystick a twist, they could make the observer push the joystick in the opposite direction (Heyes, pers. comm.). The observed behavior was caused by sniffing at the side of the joystick where it detected the scent of the demonstrator. The tendency to show olfactory interest in places where there were signs of conspecific activity was capable of reaping the benefits of imitation, but by another means than the human observers originally suspected. The specialization of a perceptual system can keep classes of strategies out of the choice set (and thereby reduce the meta-decision problem without any utility computations) and favor the selection of strategies other organisms might not have at their disposal.

Cognitive and Knowledge Limitations as Strategy Selectors

In addition to perceptual capacities, cognitive capacities vary between species and individuals. Humans differ in their working memory capacity, and experiments performed by Kareev and his associates (Kareev et al. 1997), on the detection of correlations, show how these differences in cognitive capacity lead to different outcomes in the detection of correlation. The larger a person’s working-memory capacity, the larger the sample they were assumed to consider in assessing correlations. Individuals with smaller working memories were actually more effective at detecting correlations. This is a surprising result, but predicted by the simple mechanism Kareev proposes: people with small working memories consider small samples, which in turn amplify correlations to a larger degree than do the larger samples likely to be considered by people with larger capacity. Another little-studied dimension on which individuals differ is their degree of ignorance (lack of recognition) in certain domains. The recognition heuristic

depends on a lack of knowledge to be applicable, and can become less and less effective as recognition knowledge is gained. Here again, a deficit on the part of the organism (missing knowledge) enables a strategy which can lead to better performance in specific environments than can be achieved without the deficit (for an explanation of why this happens, see Goldstein and Gigerenzer 1999).

Domain Specificity as a Strategy Selector

The research program that has come to be known as evolutionary psychology has emphasized the existence of modules of the mind that are domain specific, that is, are concerned with solving specific tasks (see, e.g., Cosmides and Tooby 1997). These modules are assumed to have evolved in the so-called Environment of Evolutionary Adaptation (EEA) and are hence capable of quite inappropriate performance in different circumstances. Domain-specific strategies are advantageous because they circumvent the meta-decision problem: they restrict the set of strategies. Furthermore, heuristics from the adaptive toolbox may be specialized for solving particular tasks. For instance, Take The Best is designed for choosing between two alternatives—the question would not arise whether it should be applied to an estimation or categorization problem.

When Will Strategies Be Learned?

Any discussion of strategy selection would be incomplete without mentioning how new strategies may enter an organism’s vocabulary through learning. A conception of Hogarth (unpublished) helps us think about which environments will lead to the learning of strategies. There are situations in which the consequences of errors are large or small, and there are situations in which the quality of feedback one gets from the environment is perfect or noisy, as shown in Figure 10.4.

Common sense tells us that learning should occur most rapidly on the left side as opposed to the right, and in the upper half as opposed to the bottom. One question we found quite provocative is to speculate about what types of

		Quality of Feedback	
		Perfect	Noisy
Consequences of Errors	Large	Pressure for valid strategies – good strategies can emerge	Difficult for valid strategies to emerge
	Small	Little pressure for valid strategies – good and bad strategies can coexist	Superstitious learning quite likely

Figure 10.4 Quality of feedback and the consequences of errors.

strategies would be learned in environments denoted by the various quadrants. For example, in which quadrant would one expect complex strategies to be learned, and in which would one expect fast and frugal heuristics to be learned? One view is that complex strategies would be essential in the top half, because being correct counts a great deal in this region. An alternative view would suggest that fast and frugal heuristics would be learned in the top half, because some issues are so grave that only stubborn, inflexible strategies can be trusted to address them. Consider the strategies “learned” through natural selection. A complex, compensatory strategy—in which, for instance, a prey animal would learn through trial and error the optimal escape distance from a predator, as opposed to simply always reacting in an overcautious way—might cost a creature its life. A useful model of the relative advantage of learning and nonlearning strategies according to environmental parameters has been discussed by Stephens (1991).

WHAT MAKES SIMPLE HEURISTICS WORK?

The general point that has been introduced with Simon’s metaphor of the two blades of the pair of scissors is that the success and failure of heuristics depends on their match with the structure of environments. We summarize here, without claim for completeness, some heuristics, some structures of environments, and their match. Note that the term “structure of environment” is shorthand for the information a person, animal, or institution knows about a physical or social environment.

Systematic lack of knowledge. The recognition heuristic is a mechanism that exploits lack of knowledge. In the simplest case, the task is to predict which of two alternatives scores higher on a criterion X . The heuristic will perform above chance if the lack of recognition of alternatives is not random, but systematically correlated with X . The precise proportion of correct inferences is a function of the recognition validity and the number of objects recognized among all alternatives (see Goldstein and Gigerenzer 1999).

Presence of others to imitate. In noisy but stable environments, imitation can allow one to find the best action among those currently being employed. Imitation reduces decision costs by leaving to others the task of discovering new choices. A necessary condition for the good performance of imitation is the existence of observable individuals in similar situations to the observer.

Noncompensatory information. Take The Best is a heuristic that exploits situations in which the cues (predictors) are noncompensatory (or approximately so). For instance, in an environment where binary cues have the weights 1, 1/2, 1/4, 1/8, ..., the simple Take The Best achieves the same predictions as multiple regression (Martignon, this volume; Martignon and Hoffrage 1999).

Scarce information. Information is scarce when the number of cues is small compared to the number of objects one has to make predictions about. Specifically, if the number of cues is less than $\log_2 N$, Take The Best outperforms on

average a class of linear models including unit-weighting strategies (Martignon and Hoffrage 1999).

Abundant information. Simple unit-weighting tends to outpredict linear regression as the number of cues increases (Einhorn and Hogarth 1975).

Redundant information. Unit-weighting tends to outpredict linear regression as the average intercorrelation between the predictors increases (Einhorn and Hogarth 1975).

Noisy information. The larger the noise in a set of training data, the better the accuracy of a simple strategy relative to a complex strategy (where simpler means fewer free parameters) when making predictions about new data (Akaike 1973; Forster and Sober 1994).

J-shaped Distributions. Many distributions are not normal, but J-shaped, that is, most objects have small values and only a few have large values (consider the population of cities, or number of publications per person). If the distribution of objects is J-shaped on a criterion, a simple heuristic, QuickEst, can exploit this structure to make fast and frugal quantitative estimates of individual objects on the criterion that are highly accurate (Hertwig et al. 1999).

This is an incomplete catalogue of when and why simple heuristics work. We agree that their secret is in the environment.

WHAT WE DO NOT YET UNDERSTAND

As the various examples in this chapter show, simple strategies can be quite effective in the right environments. However, there is work to be done. We still need a conceptual language to measure and communicate the structure of environments. We still need precise models of heuristics built with respect to the cognitive architecture of organisms. We still need to understand how the two blades of the scissors fit together, i.e., which heuristics are suited to which environments. Finally, there is the large question which kept us arguing days and nights, lunches and dinners, coffees and teas: which homunculus selects among heuristics, or is there none?

REFERENCES

- Akaike, H. 1973. Information theory and an extension of the maximum likelihood principle. In: Second International Symposium on Information Theory, ed. B.N. Petrov and F. Csaki, pp. 267–281. Budapest: Akademiai Kiado.
- Anderson, J.R. 1983. *The Architecture of Cognition*. Cambridge, MA: Harvard Univ. Press.
- Bateson, M. 1993. *Currencies for Decision Making: The Foraging Starling as a Model Animal*. Ph.D. diss., Oxford Univ.
- Bateson, M., and A. Kacelnik. 1996. Rate currencies and the foraging starling: The fallacy of the averages revisited. *Behav. Ecol.* 7:341–352.

- Bateson, M., and A. Kacelnik. 1997. Starlings' preferences for predictable and unpredictable delays to food. *Anim. Behav.* 53(6):1129–1142.
- Campbell, D.T. 1959. Systematic error on the part of human links in communication systems. *Inf. Cont.* 1:334–369.
- Cosmides, L., and J. Tooby. 1997. Dissecting the computational architecture of social inference mechanisms. In: *Characterizing Human Psychological Adaptations*, ed. G. Bock and G. Cardew, vol. 208, pp. 132–158. Chichester: Wiley.
- Czerlinski, J., G. Gigerenzer, and D.G. Goldstein. 1999. How good are simple heuristics? In: *Simple Heuristics That Make Us Smart*, G. Gigerenzer, P.M. Todd, and the ABC Research Group, pp. 97–118. New York: Oxford Univ. Press.
- Dawes, R.M. 1979. The robust beauty of improper linear models. *Am. Psych.* 34:571–582.
- Einhorn, H.J., and R.M. Hogarth. 1975. Unit weighting schemes for decision making. *Org. Behav. Human Perf.* 13:171–192.
- Einhorn, H.J., and W. McCoach. 1977. A simple multiattribute procedure for evaluation. *Behav. Sci.* 22:270–282.
- Forster, M., and E. Sober. 1994. How to tell when simpler, more unified, or less ad hoc theories will provide more accurate predictions. *Brit. J. Phil. Sci.* 45:1–35.
- Fraser, J.M., P.J. Smith, and J.W. Smith. 1992. A catalog of errors. *Intl. J. Man-Machine St.* 37:265–307.
- Gagné, R.M. 1984. Learning outcomes and their effects: Useful categories of human performance. *Am. Psych.* 39:377–385.
- Garner, R. 1987. *Metacognition and Reading Comprehension*. Norwood, NJ: Ablex.
- Gibbon, J., R.M. Church, et al. 1988. Scalar expectancy theory and choice between delayed rewards. *Psych. Rev.* 95:102–114.
- Gigerenzer, G., and D.G. Goldstein. 1996. Reasoning the fast and frugal way: Models of bounded rationality. *Psych. Rev.* 103:650–669.
- Gigerenzer, G., P.M. Todd, and the ABC Research Group. 1999. *Simple Heuristics That Make Us Smart*. New York: Oxford Univ. Press.
- Gilliam, J.F., R.F. Green, et al. 1982. The fallacy of the traffic policeman: A response to Templeton and Lawlor. *Am. Nat.* 119:875–878.
- Goldstein, D.G., and G. Gigerenzer. 1999. The recognition heuristic: How ignorance makes us smart. In: *Simple Heuristics That Make Us Smart*, G. Gigerenzer, P.M. Todd, and the ABC Research Group, pp. 75–96. New York: Oxford Univ. Press.
- Hayes-Roth, B., and F. Hayes-Roth. 1979. A cognitive model of planning. *Cog. Sci.* 3:275–310.
- Henrich, J. 1999. Rationality, Cultural Transmission, and Adaptation: The problem of culture and decision-making in anthropology. Univ. of California, www.bol.ucla.edu/~henrich/.
- Hertwig, R., U. Hoffrage, and L. Martignon. 1999. Quick estimation: Letting the environment do some of the work. In: *Simple Heuristics That Make Us Smart*, G. Gigerenzer, P.M. Todd, and the ABC Research Group, pp. 209–234. New York: Oxford Univ. Press.
- Heyes, C.M. 1993. Imitation, culture and cognition. *Anim. Behav.* 46:999–1010.
- Heyes, C.M., G.R. Dawson, et al. 1992. Imitation in rats: Initial responding and transfer evidence. *Q. J. Exp. Psych.* 45B:229–240.
- Hogarth, R.M. 1978. A note on aggregating opinions. *Org. Behav. Human Perf.* 21:40–46.

- Hogarth, R.M., C. Michaud, and J.-L. Mery. 1980. Decision behavior in urban development: A methodological approach and substantive considerations. *Acta Psych.* 45:95–117.
- Hunt, G.R. 1996. Manufacture and use of hook-tools by New Caledonian crows. *Nature* 379:249–251.
- Kareev, Y. 1995. Through a narrow window: Working memory capacity and the detection of covariation. *Cognition* 56:263–269.
- Kareev, Y., I. Lieberman, and M. Lev. 1997. Through a narrow window: Sample size and the detection of correlation. *J. Exp. Psych.: Gen.* 126:278–287.
- Klein, G. 1998. *Sources of Power: How People Make Decisions*. Cambridge, MA: MIT Press.
- Klein, G., S. Wolf, L. Militello, and C. Zsombok. 1995. Characteristics of skilled option generation in chess. *Org. Behav. Human Decis. Proc.* 62:63–69.
- Lipman, B.L. 1991. How to decide how to decide how to ... Modeling limited rationality. *Econometrica* 59:1105–1125.
- Martignon, L., and U. Hoffrage. 1999. Why does one-reason decision making work? A case study in ecological rationality. In: *Simple Heuristics That Make Us Smart*, G. Gigerenzer, P.M. Todd, and the ABC Research Group, pp. 119–140. New York: Oxford Univ. Press.
- Martignon, L., and K.B. Laskey. 1999. Bayesian benchmarks for fast and frugal heuristics. In: *Simple Heuristics That Make Us Smart*, G. Gigerenzer, P.M. Todd, and the ABC Research Group, pp. 169–188. New York: Oxford Univ. Press.
- Mazur, J.E. 1984. Tests of an equivalence rule for fixed and variable reinforcer delays. *J. Exp. Psych.: Anim. Behav. Proc.* 10:426–436.
- Nagel, R. 1995. Unraveling in guessing games: An experimental study. *Am. Econ. Rev.* 85:1313–1326.
- Nowak, M.A., and K. Sigmund. 1998a. The dynamics of indirect reciprocity. *J. Theor. Biol.* 194:561–574.
- Nowak, M.A., and K. Sigmund. 1998b. Evolution of indirect reciprocity by image scoring. *Nature* 393:573–577.
- Payne, J.W., J.R. Bettman, and E.J. Johnson. 1993. *The Adaptive Decision Maker*. New York: Cambridge Univ. Press.
- Rieskamp, J., and U. Hoffrage. 1999. When do people use simple heuristics, and how can we tell? In: *Simple Heuristics That Make Us Smart*, G. Gigerenzer, P.M. Todd, and the ABC Research Group, pp. 141–167. New York: Oxford Univ. Press.
- Savage, L.J. 1954. *The Foundations of Statistics*. New York: Wiley.
- Schlag, K.H. 1998. Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits. *J. Econ. Theory* 78:130–156.
- Stephens, D.W. 1991. Change, regularity and value in the evolution of animal learning. *Behav. Ecol.* 2:77–89.
- Thaler, R.H. 1991. *Quasi Rational Economics*. New York: Russell Sage.
- Tversky, A. 1969. Intransitivity of preferences. *Psych. Rev.* 76:31–48.
- Tversky, A. 1996. Contrasting rational and psychological principles of choice. In: *Wise Choices: Decisions, Games, and Negotiations*, ed. R.J. Zeckhauser, R.L. Keeney, and J.K. Sebenius, pp. 5–21. Boston: Harvard Business School Press.