(2004). In J. P. Leighton & R. J. Sternberg (Eds.), *The nature of reasoning* (pp. 273–287). New York: Cambridge University Press.

10

Cognitive Heuristics

Reasoning the Fast and Frugal Way

Barnaby Marsh, Peter M. Todd, and Gerd Gigerenzer

It is commonly assumed that more information can lead to more precision and better decisions. Consider for instance, the task of catching a baseball. In order to catch the ball, the ballplayer must be at the appropriate spot on the field, where the incoming ball lands. How could this spot be determined? One way to arrive at a precise prediction of the landing point would be to find the ball's launch angle and initial velocity. However, this would lead to an incorrect prediction unless other considerations, such as wind speed, humidity, and the spin of the ball were taken into account. Additional considerations such as differences in wind speed at different parts of the ball's flight could be used to further refine the estimate. Of course, any of these considerations might not be easily or perfectly knowable, and there is a chance that an important consideration would be overlooked. Given that there are typically only a few seconds to gather information and take action, it is likely that by the time that all necessary calculations have been made, the ball would have already landed.

What is the player to do? One way to solve the problem effectively is to forget all the calculations and use the *gaze heuristic*. Following this heuristic, the player fixates visually on the ball and starts running in the general direction where the ball is likely to land, adjusting running speed so that the angle of gaze to the ball remains constant (see McLeod & Dienes, 1996). If the player can run fast enough to maintain the specific angle of gaze, the player will be in the right place to catch the ball as it nears the ground. As this example illustrates, finding a workable solution to a problem does not need to depend on taking all causally relevant information into account (in fact, for many tasks in everyday life, this is not feasible anyway). In contrast, a simple rule, or heuristic can be used. Heuristics have specific qualities that are often absent in other solutions: First, they are *fast*, meaning that they can be used in cases when time constraints would greatly restrict the use of other methods. Second, they are *frugal*, meaning that they make use of only a few pieces of information, rather than taking all available information into account. Surprisingly, heuristics can often compete successfully against standard problem-solving methods that use more information and which require greater computational power. Additionally, the decisions that heuristics produce can be more accurate and robust across different problem situations than those made by comparative formal methods.

In the past few years, a new research approach has emerged on how and when specific heuristics can work, when people use them, and how they compare with more complex decision mechanisms (Gigerenzer, Todd, & the ABC group, 1999; Todd & Gigerenzer, 2000; Gigerenzer & Selten, 2001). This chapter gives an overview of some of the specific heuristics that have been investigated within this research program, showing how heuristics can be used in a variety of decision problems involving choice, categorization, estimation, elimination, and other tasks faced by real decision makers. However, before discussing specific examples, we will first briefly outline the conceptual background that has motivated and guided our search for fast and frugal heuristics.

The term *heuristic* has acquired a range of meanings since its introduction to English in the early 1800s. For instance, in his 1905 paper in quantum physics entitled, "On a Heuristic Point of View Concerning the Generation and Transformation of Light," Einstein used the term *heuristic* to indicate a view that was incomplete and unconfirmed, but nonetheless useful. In early psychological contexts, *heuristics* have come to refer to useful mental shortcuts, approximations, or rules of thumb used for guiding search and making decisions (Duncker, 1945; Polya, 1954). Since the 1970s, much attention has focused on ways that heuristics can lead to errors or bias in human judgment (Camerer, 1995; Kahneman, Slovic, & Tversky, 1982; Kahneman & Tversky, 2000; Nisbett & Ross, 1980). Our use of the term heuristic stands in contrast to this more recent fashion and is more closely related to earlier psychology traditions in that we place our emphasis on heuristics as adaptive tools for making decisions, given real constraints. Specifically, we are interested in the relationship between heuristics and the structure of the environment, and how heuristics might be able to provide tenable solutions to commonly faced problems.

BENCHMARKS FOR HEURISTIC TOOLS

A starting point in the study of heuristics from an adaptive point of view rests in selecting the proper standard by which they can be evaluated. For instance, many behavioral regularities uncovered by research on heuristics and biases have been found to violate standard models of rational choice. However, this does not mean that such behaviors are "irrational" in wider contexts, because standard rational choice models typically do not allow for natural constraints that affect real decision makers, such as limitations in time, knowledge, and computational capacity. As such, these methods cannot be confidently used to assess the true goodness of observed decision strategies.

Rather, the ultimate test of the "rationality" of a heuristic can be found in its fitness consequences relative to real constraints and real environmental structures. To illustrate, findings such as a systematic avoidance of variability (risk aversion) may seem puzzling from a normative profit-maximizing perspective, but may make more sense when viewed in terms of the costs of extra time generally needed to learn about the nature of variance in real environments and nonlinearity in payoff value (such as diminishing marginal utility for larger amounts; see Bernoulli, 1738/1954). In this light, the observation of aversion to variance may not seem so irrational after all – it can in fact be ecologically rational in the right circumstances. Similar analyses can be applied to other biases (Gigerenzer, 1991).

The study of ecological rationality involves analyzing the structure of commonly encountered environments, the structure of specific cognitive heuristics, and the fit between the two. Incorporating a perspective of ecological rationality into the study of behavior enables a clearer understanding of why humans and animals solve problems as they do when faced by different conditions. Our approach in exploring heuristics assumes that organisms make inferences on the basis of stable but possibly uncertain or noisy cues. Furthermore, various stable patterns of environmental features have existed over long periods of time, allowing selective pressures to operate on the design of decision-making processes matched to those environmental structures. In particular, we expect cognitive mechanisms to have been selected to be both quick, by needing to gather little information, and effective, by using the most appropriate cues (Todd, 2001). Our working hypothesis is that people use a range of heuristics depending on the particular details of the environmental context. In this sense, the evolved mind functions as an adaptive tool box, providing specific heuristic tools to solve particular types of problems that are commonly faced in real environments.

UNPACKING THE TOOL BOX: FAST AND FRUGAL HEURISTICS

Just like tools from any tool box, heuristic tools can range from the quite specific to the more general. It is also often the case that solving a problem may require several tools to be used, rather than just one. We have investigated a range of heuristics that can all be defined in precise computational terms, allowing us to evaluate and analyze them rigorously. Most of the heuristics can be broken down into a set of simple building blocks, that control their search for information, how the search is stopped, and what is done with the results of the search. By developing a better understanding of the psychological building blocks from which heuristics can be composed, we can get a bottom-up handle on what tools fill the mind's adaptive tool box.

In general, then, our approach to studying heuristic tools involves (a) using knowledge of behavior and psychologically plausible building blocks to design computational models of candidate heuristics, (b) analyzing the environmental structures in which these heuristics are expected to perform well, (c) testing their performance in real-world environments, and (d) determining whether and when people (and other animals) really use these heuristics. The results of the investigatory stages (b), (c), and (d) can be used to revise the next round of theorizing in stage (a). We now turn to several specific fast and frugal heuristics that we have identified and tested. These heuristics fall into four distinct classes, depending on the structure of the problem situation that is being addressed.

Ignorance-Based Decision Making

One of the most common and fundamental problems that decision makers face is to infer which of two options scores higher on some criterion or more generally to choose one option from two possibilities. Sometimes, the only information available is whether or not the option has ever been encountered before. If one option is recognized and the other is not, then the decision maker can use recognition as a cue in making the decision. This is known as the recognition heuristic (Goldstein & Gigerenzer, 1999, 2002). In computational terms, the recognition heuristic can be formulated as follows (for simplicity, we assume here that there is a positive relation between recognition and the choice criterion):

The Recognition Heuristic

Given: Two options that may or may not have been encountered before.
Find: Which of the two options has the higher value on some criterion. If one option is recognized and the other is not, *then* select the recognized object, otherwise choose at random.
To illustrate: Imagine that you are traveling in an exotic country, and you are offered ham and eggs. Suppose, further, that you are offered the choice between eggs that have either green or yellow yolks. Using the recognition heuristic, you would select the yellow-yoked eggs because they are recognized.

The recognition heuristic has been investigated in a range of different contexts. Test results indicate both that the recognition heuristic is often used, and that it is efficacious in many environments (Goldstein & Gigerenzer, 1999, 2002). For instance, common wisdom has it that people should invest in companies that they know. One can test whether using the recognition heuristic in the construction of investment portfolios is actually a good idea – that is, whether or not it will make money in the complex environment of the stock market. In a series of experiments, we addressed this question by forming real investment portfolios based on firms most recognized by German and American pedestrians. Surprisingly, we found that by forming portfolios by recognition alone, it was possible to match and often beat experienced investment professionals, despite the fact that this latter group has extensive information and a formidable range of tools at its disposal (see Borges, Goldstein, Ortmann, & Gigerenzer, 1999 for details).

A counterintuitive consequence of the use of the recognition heuristic is what we have called the *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* the decision-making performance (Goldstein & Gigerenzer, 1999, 2002). Thus, for instance, performance on a multiple-question exam about relative sizes of cities or heights of mountains or wealth of people can go down if the exam is given repeatedly over a number of weeks, because any difference in recognition knowledge that exam takers could bring to bear the first time they saw the questions (and hence use via the recognition heuristic) is erased by the repeated exposure to exam items, rendering everything recognized.

One-Reason Decision Making

For most of the decisions we make, however, we have more information than just recognition to go on. In these cases, other heuristics can be employed. When multiple cues are available for guiding decisions, what methods might be used to make a choice? The most frugal approach is to use one-reason decision making (Gigerenzer & Goldstein, 1996, 1999), choosing between options based on the first cue found that favors one over the others:

One-Reason Decision-Making Heuristics

Given: Two objects and their corresponding values on cue dimensions that can be used to infer their relative value on some criterion.

Find: The object with a higher value on the decision criterion.

a. Select any cue dimension.

b. Look for the corresponding cue value for each object.

c. If they differ, then stop and choose the object with the cue indicating a greater criterion value.

d. If they do not differ, then return to the beginning of this loop (a) to look for another cue dimension.

To illustrate: Suppose that you are shopping for fresh pasta and the two products available differ in terms of, color, moistness, and smell. Using one-reason decision making, if your criterion was freshness, you would select any one cue (for instance, color), and compare the two options on this single cue. You would then select the product having the most appealing color. If there is a tie, then you would select another cue (such as moistness), and make your decision on the basis of that cue, and so on.

The 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). This class of heuristics does not need to compute an optimal cost-benefit tradeoff, as in optimization under constraints. In fact, it need not compute any costs or benefits at all. Specific one-reason decision heuristics can be created by using particular search rules that dictate how the appropriate cue dimensions are selected (step a); we have explored three such heuristics that differ only in their search rule (see Gigerenzer & Goldstein, 1996). The *take the best heuristic* searches for cues in the order of their validity – that is, how often each cue has indicated the correct versus incorrect option (when it discriminated between them). The *take the last heuristic* 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. These three heuristics can be understood conceptually as follows:

The Take The Best Heuristic

Given: A pair of objects to choose between according to some criterion, each with associated binary cue values (some of which may be unknown) where value 1 is associated with higher criterion values.
Find: The object with the higher value on the criterion.
a. If applicable, use the recognition heuristic (see above): If one object is recognized, predict that object as having a higher value on the criterion; if neither is recognized, simply guess randomly; if both are recognized:
Then
b. Ordered search: Choose the cue with the highest validity that has not yet been tried in the current decision. Look up the values of the two objects.
c. If one object has a cue value = 1 and the other does not (i.e., value = 0 or unknown) then stop search and go to step "d." Otherwise, go back to step "b" and search for another cue. If no further cue is found, then guess.
d. Predict that the object with the cue value = 1 has the higher value on the criterion.
To illustrate: Suppose that you are choosing between two holiday packages – one to Bali, the other to Tahiti, and that your most important vacation criterion is the potential for active outdoor leisure time. Since you recognize both destinations, you look to the cue that you feel is most important – whether it is likely to be sunny at the time of year when you are wanting to travel. You find that both destinations have acceptable weather during the month of anticipated travel (e.g., the options are tied on this cue), and so you then turn to the next most important cue – whether or not the hotel happens to be on the beach. Upon looking up this cue, you discover that both hotels are on beachfronts, and so you look to the next most important cue-whether boat rental is available. You find that boats are easily available in Tahiti but not near the hotel in Bali. On the basis of this cue, you select Tahiti as your holiday destination for this year.

The Take The Last Heuristic

a. If applicable, use the recognition heuristic (see above): If one object is recognized, predict that object as having a higher value on the criterion; if neither is recognized, simply guess randomly; if both are recognized:

Then

b. Memory search: if there is a record of which cues stopped search during previous decisions, choose the cue that stopped search on the most recent decision but which has not yet been tried on the current one. Look up the cue values of the two objects. If there is no record, try a random cue.
c. If one object has a cue value = 1 and the other does not (i.e., value = 0 or unknown) then stop search and go to step "d." Otherwise, go back to step "b" and search for another cue. If no further cue is found, then guess.
d. Predict that the object with the cue value = 1 has the higher value on the criterion.
To illustrate: Suppose that that you are once again looking for a beach vacation package. This year you are choosing between the Bahamas or Jamaica – both destinations that you recognize. You then perform a memory search and recall that last year whether or not the hotel had easy access to boat rental was the decisive inference cue. You then use this cue in current inference. You discover that only the hotel in the Bahamas has easy access to boat rental; on the basis of this cue, you decide on the package to the Bahamas.

The Minimalist Heuristic

a. If applicable, use the recognition heuristic (see above), if one object is recognized, predict that object as having a higher value on the criterion; if neither is recognized, simply guess randomly; if both are recognized:

Then

b. Random search: Draw a cue randomly (without replacement) and look up the cue value of the two objects.
c. If one object has a cue value = 1 and the other does not (i.e., value = 0 or unknown) then stop search and go to step "d." Otherwise, go back to step "b" and search for another cue. If no further cue is found, then guess.
d. Predict that the object with the cue value = 1 has the higher value on the criterion.
To illustrate: Suppose once again that you are faced with a choice between two vacation packages-one to the Bahamas and the other to Jamaica. Your primary criterion is to be someplace that permits ample outdoor leisure time. Since you recognize both, you next select any cue at random that may relate to your criterion. For instance, you look at whether the hotel has a pool. Since the hotel in Jamaica has a pool but the hotel in the Bahamas does not, you select the package to Jamaica.

What we found when we tested the performance of these one-reason decision-making heuristics was surprising. Despite their simplicity, they still made very accurate choices. When we compared these heuristics against a range of more traditional information-combining methods such as multiple regression, we found that the simple heuristics always came close to, and in the case of take the best even often exceeded, the efficacy of

TABLE 10.1. *Performance of Different Decision Strategies Across 20 Data Sets*

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

Note: The performance of two fast and frugal heuristics (Minimalist, Take the Best) and two linear strategies (Dawes's Rule, Multiple Regression) is shown for a range of 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).

the traditional methods (see Czerlinski, Gigerenzer, & Goldstein, 1999; see Table 10.1). Thus, making good decisions need not rely on the standard approach of collecting all available information and combining it according to the relative importance of each cue; simply betting on one good reason can be sufficient.

Those who are used to working with formal models might be suspicious of the prospect of using only a few cues instead of many. How can we expect simple domain-specific heuristics to be as accurate as complex general strategies that work with many free parameters? Perhaps the key answer lies in their robustness in providing accurate answers across a range of environments. More complex general strategies that work by making use of a large number of free parameters, such as multiple linear regression, can suffer from trying to make sense of every piece of information they encounter. This failure of generalization, known as *overfitting* (Geman, Bienenstock, & Doursat, 1992; Massaro, 1988), stems from assuming that every detail is of relevance. Fast and frugal heuristics can reduce overfitting by ignoring the noise inherent in many cues and looking instead for the most important cues. Thus, simply using only one or a few of the most useful cues can automatically yield robustness.

In the scenarios above, we assume that different options can be selected or rejected on the basis of differences in the presence or absence of specific cues. What would happen if there are more than two options to choose from, but all share the same set of cues? As a test of this question, we investigated the performance of one-reason decision making in the domain of parental investment (Davis & Todd 1999; see also Rieskamp & Hoffrage, 1999, for a different test of this sort). Specifically, we modeled the food provisioning task faced by parent birds: after arriving at the nest with a recently caught insect, the parent must decide which nestling to feed. Cues that may be considered include the size of each chick, variations in vocal begging intensity between different chicks (e.g., hunger), and the position of each chick in the nest. Several cues can be used, or just one. It was found that one-cue feeding rules can perform significantly better (in terms of total chick growth) than traditional rules that combine information from multiple cues.

Elimination Heuristics for Multiple-Option Choices

In situations where there are more available options than values in each available cue dimension, one-reason decision making will usually not suffice because a single cue will be unable to distinguish between all of the alternatives. 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 alternatives until a single option can be selected (see Tversky, 1972). Such rules can be more effective when they rely on some knowledge of the validities of different cues.

One fast and frugal heuristic that works this way is the *QuickEst heuristic* (Hertwig, Hoffrage, and Martignon, 1999), which can be applied to estimation tasks. The QuickEst heuristic can exploit aspects of the commonly encountered J-shaped distribution of features and objects in the environment in order to make quick and useful estimates. J-shaped distributions reflect the tendency in many domains for objects to be unevenly distributed, and highly clustered at one end of a given criterion range. For example, if we ask 1,000 random people to name characters in a Shakespeare play, there will be many more people who are able to name a small number of characters than there are who come up with a long list. If we plot these responses by the number of characters identified, then the resulting plot is likely to look like a "J" tipped on its side, with many people who can name 0–10 characters, and only few in the range of naming 40–50. Similarly, if we plot the number of a nation's cities in different size ranges, there is expected to be a much greater number of cities with populations between 10,000–100,000 than there are with populations between 1,000,000 and 1,090,000. To exploit the common J-shaped environmental structure, QuickEst works in the following way:

The QuickEst Heuristic

Given: An object and a set of corresponding cue values.
Find: An estimate of the object's value on some criterion.
a. Identify a cue that is expected to separate the most common objects from all of the others.

b. Look at the next cue that is expected to separate the remaining common objects from the rest of the distribution tail. Continue looking for cues in this way until a cue is found that no longer places the object with the most common objects.

Then

c. Assign a value to the object based on the value expected based on the attributes of the cue it stopped with.

To illustrate: Imagine that you manage a company and need to estimate the size of a competitor's advertising budget. Since the size of the budget is likely to be related to advertising channels, you look to find a channel that the competitor uses that separates it from most others. For instance, many firms may use local newspapers, but relatively few advertise in larger city or regional papers. You find that your competitor advertises in these larger papers, so you then turn to the next cue: whether the competitor uses spots on national radio. Again, you find that your competitor uses this medium in advertising, so you next check television. Here, you find that your competitor does not use television, and so you stop your search. You then make your estimate on the basis of the expected budget size of a firm that uses newspaper and radio, but not television.

Note that when using QuickEst, no cue combination is necessary at any point. This eliminates complexity problems associated with the integration of a potentially large number of different cue values as found in information-and-computation-intensive methods such as multiple regression and estimation trees. Yet, in a test of this heuristic against these methods, QuickEst consistently matches or outperforms them (Hertwig et al., 1999).

Another elimination-based heuristic is *categorization by elimination* (Berretty, Todd, & Blythe, 1997). This heuristic can be used when the task is to select a single category from several possible categories that a given object might best fall into. The heuristic can be understood as follows:

The Categorization by Elimination Heuristic

Given: An object to be categorized, a set of cue values for the object, a set of possible categories, and which cue values are associated with what categories.

Find: The category to which the object belongs, using as few cues as possible.

a. Rank the cues in terms of their validity (i.e., how often each cue alone indicates the correct category across all objects).

b. Select the highest-validity unchecked cue dimension and check the object's value for that cue.

c. Eliminate from further consideration all categories that do not encompass the current cue value.

d. If only one category remains, it is the final choice; otherwise, if cue dimensions remain, return to "step b", and if there are no more cues to check, pick a category randomly from those left.

To illustrate: Suppose that you need to describe a bottle of an alcoholic beverage. Your basic categories are beer (amber brown, beer bottle), red wine (red, wine bottle), white wine (white, wine bottle), rose (pink, wine bottle), sauterne (yellow, wine bottle) or champagne (white, champagne bottle). Cues thus include bottle shape and color. Since in this case color is the higher validity cue (since color alone will correctly identify the type of beverage in four cases, while bottle shape alone will suffice in only two cases), you first look at the color. Since the color is red, you can eliminate all categories except red wine. If needed, one would next consider the cue validity with the next-highest validity, though in the present case this is not necessary – only one candidate happened to be red.

Categorization by elimination thus uses successive cues to reduce the set of possible categories until only a single possible category remains. Its performance comes within a few percentage points of the accuracy of complex categorization algorithms such as exemplar and neural network models, and yet it uses only about a quarter of the information needed by these other models. Thus, it is a robust candidate for categorization particularly in situations where information about objects is costly or difficult to obtain.

Satisficing Heuristics

All of the heuristics that we have discussed so far are for choosing a single option from multiple alternatives, assuming that all alternatives are presently available. But different heuristics are needed when alternatives (as opposed to cue values) appear sequentially over an extended period or spatial region. In this type of choice task, the stopping rule must specify which object stops search, just as in the previously described heuristics the stopping rule specified which cue stops search. An instance of this type of problem is found in the context of individuals who are searching for a mate from a stream of potential candidates that are seen one at a time – when should the searcher stop and stay with the current candidate?

To address such sequential search problems, agents can satisfice (Simon 1955, 1990). Satisficing works by setting an aspiration level and searching until a candidate is found that exceeds that aspiration. Satisficing eliminates the need to compare a large number of possible outcomes with one another, thus saving time and the need to acquire large amounts of information. The question is how to set the aspiration level in the first place. One way is to examine a certain number of alternatives and use the best criterion value seen in that sample as the aspiration level for further search. Consider the problem of finding the single best alternative from a sequence of fixed length drawn from an unknown distribution – an extreme form of sequential search. In this scenario, using an initial sample of 37% of all available alternatives for setting the aspiration level provides the highest likelihood of picking the best (the optimal solution to the so-called secretary or dowry problem – see Ferguson, 1989). However, much less search (e.g., setting an aspiration using 10% of the available alternatives) is required for attaining other more realistic goals such as maximizing the mean criterion value found across multiple searches (Todd & Miller, 1999). Other search rules have also been explored, such as stopping search after encountering a long gap between attractive candidates (Seale & Rapoport, 1997). In a mutual search setting, for instance where both males and females are searching for a suitable mate, heuristics that learn an aspiration level based on the rejections and offers one receives can lead to successful matching of the two populations of searchers, again with relatively little information (Todd & Miller, 1999).

GOING TO WORK: FUTURE TESTS OF HEURISTIC TOOLS

In this chapter, we have introduced several candidate heuristics and heuristic building blocks that can provide good solutions for various kinds of common decision problems without using too much computation, time, or information. We have also shown how to formalize some of these heuristic solutions in terms of their computational building blocks, and how to test them across different environmental structures where they might be used. Overall, we have found evidence that simple heuristics can work remarkably well for many kinds of problems. Heuristics tend to work well for several reasons. First, they benefit explicitly from not taking all information into account, but only the information that is likely to be important in solving the problem at hand. This particularly has advantages in new or changing environments where more complete detailed knowledge might not be helpful but might actually lead the decision maker astray. Furthermore, they are able to provide relatively quick solutions using specialized decision processes that are suited to the specific nature of the problem domain. This stands in contrast to more complex methods that might need to be fine tuned to give optimal solutions for each and every type of decision process. Finally, by using a set of relatively simple processes, heuristics can offer competitive advantages in terms of computation time, which is often an important consideration in many real-life decisions

However, much work remains. Some of the problems to be explored include (see Todd, Gigerenzer, & the ABC Research Group, 2000, for a full development): further examination of where heuristics come from – how they are learned, or socially constructed, or evolved, or culturally transmitted; how heuristics are selected from the adaptive tool box, including the role of the environment in triggering the use of particular heuristics; how environment structure can be characterized in terms of patterns of cues, social relationships, J-shaped distributions of alternatives, and so on; development of methodology for studying which heuristics people use, when, and how effectively; and specification of criteria that can be used to evaluate the performance of heuristics including Bayesian and other benchmarks. We must also develop studies outside the laboratory to explore how people use heuristics to solve problems when time pressure, competition, and emotions are at play (e.g., Klein, 1998). We also can look at the conditions under which the heuristics that we have studied begin to break down. For instance, we know that the recognition heuristic becomes less effective as the number of recognized options increases. If all options are recognized, then this heuristic is no longer able to provide an advantage. We would like to look at how the use of heuristics might be expected to change as the structure of the environment changes.

Despite these multiple challenges, the study of heuristic strategies nonetheless promises to change the way that decision behavior is understood. In the past, cognitive scientists, economists, and biologists have often addressed issues of how people make choices by building elaborate models and endowing organisms with unlimited abilities to know, memorize, and compute. In other cases, when rational models could not account for real human behavior, researchers often resorted to mere descriptions of decisions made. Work on understanding behavior in the context of ecologically rational heuristics provides a third way forward. We can put aside unattainable hyperrational models and mere collections of facts to construct and test models of the psychological mechanisms used to solve the real problems faced by real decision makers.

References

Bernoulli, D. (1738). *Specimen theoriae novae de mensura sortis* [*Exposition of a new theory on the measurement of risk*]. Translated by L. Sommer, *Econometrica* (1954), 22, 23–36.

Berretty, P. M., Todd, P. M., & Blythe, P. W. (1997). Categorization by elimination: A fast and frugal approach to categorization. In M. G. Shafto & P. Langley (Eds.), *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society* (pp. 43–48). Mahwah, NJ: Erlbaum.

Borges, B., Goldstein, D. G., Ortmann, A., & Gigerenzer, G. (1999). Can ignorance beat the stock market? Name recognition as a heuristic for investing. In G. Gigerenzer, P. M. Todd, & the ABC Research Group. *Simple heuristics that make us smart*. New York: Oxford University Press.

Camerer, C. (1995) Individual decision making. In J. H. Kagel & A. E. Roth (Eds.), *The handbook of experimental economics*. Princeton, NJ: Princeton University Press.

Czerlinski, J., Gigerenzer, G., & Goldstein, D. G. (1999). Accuracy and frugality in a tour of environments. In G. Gigerenzer, P. M. Todd, & the ABC Research Group. *Simple heuristics that make us smart*. New York: Oxford University Press.

Davis, J. N. & Todd, P. M. (1999). Simple decision rules for parental investment. In G. Gigerenzer, P. M. Todd, & the ABC Research Group. *Simple heuristics that make us smart*. New York: Oxford University Press.

Duncker, K. (1945). On problem solving (L. S. Lees, Trans.). *Psychological Monographs, 58* (5 Whole no. 270). (Original work published 1935.)

Ferguson, T. S. (1989). Who solved the secretary problem? *Statistical Science, 4*, 282–296.

Geman, S.,Bienenstock, E., & Doursat E., 1992. Neural networks and the bias/variance dilemma. *Neural Computation*, *4*, 1–58.

Gigerenzer, G. (1991). How to make cognitive illusions disappear. Beyond heuristics and biases. In W. Stroebe & M. Hewstone (Eds.), *European review of social psychology* (vol. 2, pp. 83–115). Chichester, UK: Wiley.

Gigerenzer, G., & Goldstein, D. G., (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, *103*, 650–669.

Gigerenzer, G., & Goldstein, D.G. (1999). Betting on one good reason: Take the best and its relatives. In G. Gigerenzer, P. M. Todd, & the ABC Research Group (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.

Gigerenzer, G., & Selten, R. (Eds.) (2001). *Bounded rationality: The adaptive toolbox* (Dahlem Workshop Report). Cambridge, MA: MIT Press.

Gigerenzer, G., Todd, P. M., & the ABC Research Group (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.

Goldstein, D. G., & Gigerenzer, G. (1999). The recognition heuristic: How ignorance makes us smart. In G. Gigerenzer, P. M. Todd, & the ABC Research Group. *Simple heuristics that make us smart*. New York: Oxford University Press.

Goldstein, D. G., & Gigerenzer, G. (2002). Models of ecological rationality: The recognition heuristic. *Psychological Review*, *109*, 75–90.

Hertwig, R., Hoffrage, U., & Martignon, L. (1999). Quick estimation: Letting the environment do some of the work. In G. Gigerenzer, P. M. Todd, & the ABC Research Group. *Simple heuristics that make us smart*. New York: Oxford University Press.

Kahneman, D., Slovic, P., & Tversky, A. (1982). *Judgement under uncertainty: Heuristics and biases*. Cambridge: Cambridge University Press.

Kahneman, D., & Tversky, A. (Eds.). (2000). *Choices, values, and frames*. Cambridge: Cambridge University Press

Klein, G. (1998). *Sources of power: How people make decisions*. Cambridge, MA: MIT Press.

Massaro, D. W. (1988). Some criticisms of connectionist models of human performance. *Journal of Memory and Language*, *27*, 213–234.

McLeod, P., & Dienes, Z. (1996). Do fielders know where to go to catch the ball or only how to get there? *Journal of Experimental Psychology: Human perception and performance*, *22*, 531–543.

Nisbett, R. E., & Ross, L. (1980). *Human inference: Strategies and shortcomings of social judgement*. Englewood Cliffs, NJ: Prentice-Hall.

Polya, G. (1954). *Mathematics and plausible reasoning: Vol. 1. Induction and analogy in mathematics*. Princeton, NJ: Princeton University Press.

Rieskamp, J., & Hoffrage, U. (1999). When do people use simple heuristics and how can we tell? In G. Gigerenzer, P. M. Todd, & the ABC Research Group. *Simple heuristics that make us smart*. New York: Oxford University Press.

Seale, D. A., & Rapoport, A. (1997). Sequential decision making with relative ranks: An experimental investigation of the "secretary problem." *Organizational Behavior and Human Decision Processes*, *69*(3), 221–236.

Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, *69*, 99–118.

Simon, H. A. (1990). Invariants of human behavior. *Annual Review of Psychology*, *41*, 1–19.

Todd, P. M. (2001). Fast and frugal heuristics for environmentally bounded minds. In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality: The adaptive toolbox (Dahlem Workshop Report)* (pp. 51–70). Cambridge, MA: MIT Press.

Todd, P. M., & Gigerenzer, G. (2000). Simple heuristics that make us smart. *Behavioral and Brain Sciences*, 23(5), 727–741.

Todd, P. M., Gigerenzer, G., & the ABC Research Group (2000). How can we open up the adaptive toolbox? (Reply to commentaries.) *Behavioral and Brain Sciences*, 23(5), 767–780.

Todd, P., & Miller, G. (1999) From pride and prejudice to persuasion: Satisficing in mate search. In G. Gigerenzer, P. M. Todd, & the ABC Research Group. *Simple heuristics that make us smart*. New York: Oxford University Press.

Tversky A. (1972). Elimination by Aspects: A theory of choice. *Psychological Review*, 79, 281–299.