



# Heuristics

RALPH HERTWIG and PETER M. TODD

*Max Planck Institute for Human Development, Berlin*

- I. A Short History of the Concept "Heuristic"
- II. Unbounded Rationality versus the Bounded Reality of Human Decision Making
- III. Heuristics for Human Judgment, Choice, and Search Behavior
- IV. How to Measure a Heuristic's Success
- V. Conclusion

## GLOSSARY

**algorithm** A strategy for solving a problem that guarantees solution in a finite number of steps if the problem has at least one solution. An example of a very simple algorithm is that for obtaining temperature on the Fahrenheit scale when the value for the centigrade scale is known: Multiply the known value by 1.8 and add 32.

**bounded rationality** Principles underlying nonoptimizing adaptive behavior of real information processing systems working under conditions of limited time, information, and computational capacities. Among those principles specified by Herbert Simon is satisficing.

**cognitive illusions or biases** Systematic deviations of human judgments from rules of probability theory and logic. The occurrence of cognitive illusions is attributed to some judgment heuristics that are often useful but sometimes lead to predictable errors.

**heuristic** An approximate strategy or "rule of thumb" for problem solving and decision making that does not guarantee a correct solution but that typically yields a reasonable solution quickly or brings one closer to hand.

**noncompensatory heuristic** A heuristic is noncompensatory if, once it has used a piece of information to make a decision, no further information in any combination can undo or compensate for the effect of the original information. In contrast, a heuristic is compensatory if there is at least one piece of information that can

be outweighed by other pieces of information. A compensatory strategy integrates at least some of the available information and makes trade-offs between the relevant pieces of information to form an overall evaluation of each of the available alternatives or options.

**satisficing** According to Herbert Simon, satisficing is using experience to construct an expectation (or aspiration level) of a reasonable solution to some problem and stopping the search for solutions as soon as one is found that meets the expectation.

**unbounded rationality** Decision-making strategies that have no regard for constraints of time, knowledge, or computational capacities. Modern mainstream economic theory is largely based on unbounded rationality models that portray economic agents as fully rational Bayesian maximizers of subjective utility.

**Many decisions faced by people cannot be made in an optimal way because optimal solutions may take too much computation to find or may not even exist. Instead, real decision makers must often take shortcuts and use heuristics that yield reasonable solutions in a reasonable amount of time, even if they do not guarantee always reaching a good decision. These heuristics are thus an essential aspect of human intelligence, leading to adaptive behavior despite the challenging conditions of limited time, knowledge, and computational capacity under which people have to solve problems. Heuristics are most commonly studied in psychology, particularly within the domains of judgment and decision making, and in computer-based applications in artificial intelligence and operations research. This article focuses on research in psychology that has proposed heuristic models of how people search for information and make decisions and choices.**

## I. A SHORT HISTORY OF THE CONCEPT "HEURISTIC"

Recent research on decision heuristics descends from earlier schools of thought. For this reason, understanding current thought can be aided by first considering the history of the concept. In 1905, the 26-year-old Albert Einstein published his first fundamental paper in quantum physics, titled "On a Heuristic Point of View Concerning the Production and Transformation of Light." In that Nobel prize-winning paper, Einstein used the term heuristic to indicate that he considered the view he presented therein as incomplete, even false, but still useful. Einstein could not wholeheartedly accept the quantum view of light that he started to develop in this paper, but he believed that it was of great transitory use on the way to building a more correct theory. As used by Einstein, a heuristic (a term of Greek origin meaning "serving to find out or discover") is an approach to a problem that is necessarily incomplete given the knowledge available, and hence unavoidably false, but that is useful nonetheless for guiding thinking in appropriate directions.

A few decades later, Max Wertheimer (a close friend of Einstein's), Karl Duncker, and other Gestalt psychologists spoke of heuristic reasoning, but with a slightly different meaning from that of Einstein. Gestalt psychologists conceptualized thinking as an interaction between inner mental processes and external problem structure. In this view, heuristic methods such as "looking around" and "inspecting the problem" are first used to guide the search for appropriate information in the environment, which is then restructured or reformulated by inner processes.

Heuristic methods also play a prominent role in George Pólya's approach to mathematical problem solving. According to the Hungarian mathematician, effective problem solving consists of four main phases—understanding the problem, devising a plan, carrying out the plan, and looking back—all of which can incorporate heuristics. Devising a plan, for instance, can include heuristic methods such as "examine a simpler or special case of the problem to gain insight into the solution of the original problem," "work backward," or "identify a subgoal."

In the 1950s and 1960s, Herbert Simon and Allen Newell started to develop heuristics for searching for solutions to problems. They replaced the somewhat vague notion of heuristic reasoning of the Gestalt school and of Pólya's with much more precise computer-based models (e.g., in the General Problem

Solver system) of human problem solving and reasoning largely based on the means–ends analysis heuristic. This heuristic found some way to reduce the distance between the current state and the goal state. With the advent of information processing theory in cognitive psychology, a heuristic came to mean a useful shortcut, an approximation, or a rule of thumb for guiding search through a space of possible solutions, such as a strategy that a chess master uses to explore the enormous number of possible moves at each point in a game.

Such general-purpose or "weak" methods as the means–ends analysis heuristic, however, proved insufficient to deal with problems other than artificial and well-defined mathematical problems or the games of chess and cryptoarithmetic that Newell and Simon investigated. As a consequence, research in artificial intelligence (AI) in the 1970s turned to collecting domain-specific rules of thumb from specialists in a particular field and incorporating these into expert systems. At approximately the same time, mathematicians working in operations research began dealing with new results from computational complexity theory indicating that efficient algorithmic solutions to many classes of challenging combinatorial problems (such as the traveling salesman problem) might not be found; as a consequence, they too turned to the search for problem-specific heuristics, although through invention rather than behavioral observation.

In psychology after 1970, researchers became increasingly interested in how people reason about unknown or uncertain aspects of real-world environments. The research program that spurred this interest was the heuristics-and-biases program initiated by Amos Tversky and Daniel Kahneman. This program's research strategy has been to measure human decision making against various normative standards taken from probability theory and statistics. Based on this strategy two major results about people's reasoning under uncertainty emerged: a collection of violations of the normative standards (that in analogy to perceptual illusions are often called "cognitive illusions" or "biases") and explanations of these illusions in terms of a small number of cognitive heuristics. According to Kahneman and Tversky, people rely on a limited number of heuristics—most prominently representativeness, availability, and anchoring and adjustment—that often yield reasonable judgments but sometimes lead to severe and systematic biases. Diverging from earlier usage, the term heuristics now gained a different connotation: fallible cognitive shortcuts that people often use when faced with

uncertainty and that can lead to systematic biases and lapses of reasoning indicating human irrationality. This more negative view of heuristics—and of the people who use them as “cognitive misers” using little information or cognition to reach biased conclusions—has spread to many other fields, including law, economics, medical decision making, sociology, and political science.

Recently, however, a new appreciation is emerging that heuristics may be the only available approach to decision making in the many problems for which optimal logical solutions do not exist (as researchers in operations research realized). Moreover, even when exact solutions do exist, domain-specific decision heuristics may be more effective than domain-general logical approaches, which are often computationally infeasible (as AI found). This has led to research programs such as the study of ecological rationality by Gigerenzer, Todd, and colleagues. Their program focuses on precisely specified computational models of heuristics and how they are matched to the ecological structure of particular decision environments. It also explores the ways that learning and evolution can achieve this match in human behavior, something that has already been widely accepted for other animals in research on rules of thumb in behavioral ecology.

In the following sections, the focus is on research in psychology exploring heuristics proposed to model how people search for information and make decisions and choices. Researchers such as Payne, Bettman, and Johnson and Svenson have been concerned with psychological heuristics for preferences, but here we are mostly concerned with inference heuristics. Heuristics and shortcuts are also important in human perception and higher order reasoning processes (e.g., hypothesis testing), planning and problem solving, as well as in computer applications in these domains.

## II. UNBOUNDED RATIONALITY VERSUS THE BOUNDED REALITY OF HUMAN DECISION MAKING

Both the heuristics-and-biases program and the recently emerging work on ecologically rational heuristics have been linked to Herbert Simon's notion of bounded rationality. This concept can be understood by contrasting it to the traditional decision-making approach embodied in unbounded rationality, illu-

strated by the following example. Imagine being faced with the decision of whether or not to marry. How can this decision be made in a rational way? Assume that you attempted to resolve this question by maximizing your subjective expected utility. To compute your personal expected utility for marrying, you would have to determine all the possible consequences that marriage could bring (e.g., children, companionship, and countless further consequences), attach quantitative probabilities to each of these consequences, estimate the subjective utility of each, multiply each utility by its associated probability, and finally add all these numbers. The same procedure would have to be repeated for the alternative “not marry.” Finally, you would have to choose the alternative with the higher total expected utility.

Maximization of expected utility in this way is probably the best known realization of the prominent vision of unbounded rationality. Models of unbounded rationality have been criticized for having little or no regard for the constraints of time, knowledge, and computational capacities that real humans face. For instance, while you are deliberating about whether marrying is the right choice, considering each of the myriad conceivable consequences and assigning probabilities to each, any potential partner will probably have married someone else. To this criticism proponents of unbounded rationality generally concede that their models assume unrealistic mental abilities, but they nevertheless defend them by arguing that humans act as if they were unboundedly rational. In this interpretation, the models of unbounded rationality do not describe the process but merely the outcome of reasoning.

If the lofty ideals of human reasoning do not capture the processes of how real people make decisions in the real world, what then are those processes? In other words, what models take into account the challenging conditions under which people have to solve problems, including limited time, knowledge, and computational capacity? Herbert Simon proposed that these constraints force humans to use “approximate methods” (heuristics) to handle most tasks. These approximate methods form the basis of bounded rationality.

Simon's vision of bounded rationality has two interlocking components that act like a pair of scissors to shape human rational behavior. The two blades in this metaphor are the computational capabilities of the actor and the structure of task environments. First, the computational capability blade implies that models of human judgment and decision making should be built on what we actually know about the mind's limitations

rather than on fictitious competencies assumed in models of unbounded rationality. There are two key limitations central to bounded rationality. First, contrary to models of unbounded rationality, humans cannot search for information for all of eternity. In computationally realistic models, search must be limited because real decision makers have only a finite amount of time, knowledge, attention, or money to apply to a particular decision. Limited search requires rules to specify what information to seek and in what order (i.e., an information search rule) and a way to decide when to stop looking for information (i.e., a stopping rule).

Another key limitation of the human mind is that the pieces of information uncovered by the search process are not likely to be processed in an overly complex way. In contrast, most traditional models of inference, from linear multiple regression models to Bayesian models to neural networks, try to find some optimal integration of all information available: Every bit of information is taken into account, weighted, and combined in some more or less computationally expensive way. Models of bounded rationality instead rely on processing steps that are computationally bounded. For instance, a bounded decision or inference can be based on only one or a few pieces of information, whatever the total amount of information found during search. The simple decision rule used to process this limited knowledge need not weigh or combine pieces of information, thus eliminating the need to convert different types of information into a single common currency (e.g., utilities). Note that decision rules and information search and stopping rules are connected. For instance, when a heuristic searches for only one (discriminating) cue, this largely constrains the possible decision rules to those that do not integrate information. On the other hand, if search extends to many cues, the decision rule will be less constrained. The cues may then be weighted and integrated, or only the best of them may determine the decision.

These two key limitations, limited information search and limited processing of information, can be instantiated into models of heuristics. The limitations help explain how heuristics achieve one of their most important advantages, namely, speed. In fact, for much decision making in the real world—the stock broker who decides within seconds to keep or sell a stock, or the captain of the firefighter squad who within a few moments must predict how a fire will progress and whether or not to pull out the squad—speed is often the crucial objective.

The second blade in Simon's scissors metaphor, operating in tandem with computational capability, is environmental structure. This blade is of crucial importance in shaping bounded rationality because it can explain when and why heuristics perform well, namely, if the structure of the heuristic is adapted to the structure of the environment (i.e., if the heuristic is ecologically, rather than logically, rational). Simon's classic example concerns foraging organisms that have a single need—food. An organism living in an environment in which little heaps of food are randomly distributed can survive with a simple heuristic: Run around randomly until a heap of food is found. For this, the organism needs some capacity for movement, but it does not need a capacity for inference or learning. For an organism in an environment in which food is distributed not randomly but in patches whose locations can be inferred from cues, more sophisticated strategies are possible. For instance, it could learn the association between cues and food and store this information in memory. The general point is that in order to understand which heuristic an organism employs, and when and why the heuristic works well, one needs to examine the structure of the information in the environment.

### III. HEURISTICS FOR HUMAN JUDGMENT, CHOICE, AND SEARCH BEHAVIOR

The models of heuristics for human judgment, choice, and decision making that have been proposed in psychology since 1970 can be linked to two traditions of heuristics with earlier beginnings as described previously. These two traditions have employed different levels of description. First, following the line of the heuristic methods studied by the Gestalt psychologists, one class of heuristics consists of psychological principles that are verbally described. These models are only relatively loosely specified and usually do not explicate all the processes they involve (e.g., in terms of information search, stopping, and decision rules). Second, following from Simon and Newell's computer-based models of human decision making, another class of heuristics has been formulated as process models, with explicit specification of the processes involved. Because of this explication, the latter heuristics can be both mathematically analyzed and tested with the help of computer simulations. We consider each class of heuristics in turn.

### A. Heuristics for the Judgment of Probability and Frequencies: Availability, Representativeness, and Anchoring and Adjustment

The heuristics most widely studied within psychology are those that people use to make judgments or estimates of probabilities and frequencies in situations of uncertainty (i.e., in situations in which people lack exact knowledge). Most prominent among these are the availability, representativeness, and anchoring and adjustment heuristics.

The availability heuristic leads one to assess the frequency of a class or the probability of an event by the number of instances or occurrences that can be brought to mind or by how easy it seems to call up those instances. For instance, which class of words is more common: seven-letter English words of the form “\_\_\_\_\_n\_” or the form “\_\_\_\_\_i n g”? According to the availability heuristic, to estimate the frequency of occurrences people draw a sample of the events in question from memory. Specifically, for this case they retrieve words ending in -ing (e.g., “jumping”) and retrieve words with “n” in the sixth position (e.g., “raisins”) and then count the number of words retrieved in some period or assess the ease with which such words could be retrieved. They then answer that the more numerous or easier class of words is more common. Because people find it easier to think of words ending with -ing than to think of words with the letter “n” in the next-to-last position, they usually estimate the class “\_\_\_\_\_i n g” to be more common. This judgment, however, is wrong because all words ending with -ing also have “n” in the sixth position; in addition, there are seven-letter words with “n” the sixth position that do not end in -ing.

The availability heuristic has been suggested to underlie diverse judgment errors, ranging from the tendency to overestimate how many people die from some specific causes of death (e.g., tornado) and underestimate the death toll of others causes (e.g., diabetes) to why people’s answers to life satisfaction questions (“How happy are you?”) may be overly influenced by events that are especially memorable.

The representativeness heuristic has been proposed as a means to assess the probability that an object A belongs to a class B (e.g., that a person described as meek is a pilot) or that an event A is generated by a process B (e.g., that the sequence HTHTHT was generated by randomly throwing a fair coin). This heuristic produces probability judgments according to the extent that object A is representative of or similar

to the class or process B (e.g., meekness is not representative of pilots, so a meek person is judged as having a low probability of being a pilot). This heuristic can lead to errors because similarity or representativeness judgments are not always influenced by factors that should affect judgments of probability, such as base rates. The representativeness heuristic has also been evoked to explain numerous judgment phenomena, including “hot hand” observations in basketball (the belief that a player is more likely to score again after he or she already scored successfully than after missing a shot) and the gambler’s fallacy (the belief that a successful outcome is due after a run of bad luck).

Another heuristic, anchoring and adjustment, produces estimates of quantities by starting with a particular value (the anchor) and adjusting upward or downward from it. For instance, people asked to quickly estimate the product of either  $8 \times 7 \times 6 \times 5 \times 4 \times 3 \times 2 \times 1$  or  $1 \times 2 \times 3 \times 4 \times 5 \times 6 \times 7 \times 8$  give a higher value in the former case. According to the anchoring and adjustment heuristic, this happens because the first few numbers presented are multiplied together to create a higher or lower anchor, which is then adjusted upwards in both cases, yielding a higher final estimate for the first product.

Although it has been pointed out that availability, representativeness, and anchoring and adjustment are quite useful heuristics (because they often lead to good judgments without much time or mental effort), most of the large body of evidence amassed that is consistent with the use of these heuristics comes from studies showing where they break down and lead to cognitive illusions or biases (i.e., deviations from some normative standards). This heuristics-and-biases research program has caught the attention of numerous social scientists, including economists and legal scholars. There are good reasons for this attention, since systematic biases question the empirical validity of classic rational choice models (i.e., models of unbounded rationality) and may have important economic, legal, and other implications.

However, the exclusive focus on cognitive illusions has evoked the criticism that research in the heuristics-and-biases tradition equates the notion of bounded rationality with human irrationality and portrays the human mind in an overly negative light, with some researchers even arguing that cognitive illusions are the rule rather than the exception. It has also been criticized that, to date, the cognitive heuristics posited have not been precisely formalized such that one could either simulate or mathematically analyze their

behavior, leaving them free to account for all kinds of experimental performance in a post hoc fashion. For instance, it is still an open question of how people assess similarity to make probability judgments with the representativeness heuristic or how many items (e.g., words ending with -ing) the availability heuristic retrieves before it affords a frequency estimate of a class of object (albeit theoretical progress has been made, for instance, by testing whether availability works in terms of ease of recall or number of items recalled). Moreover, the heuristics-and-biases program focuses on human computational capabilities (the first blade of Simon's scissors), largely ignoring the role of the environment by not specifying how such heuristics capitalize on information structure to make inferences. Finally, this program appears to consider heuristics as dispensable mechanisms (that would not be needed if people had the right tools of probability and logic to call on), in contrast to Simon's view of indispensable heuristics as the only available tools for solving many real-world problems.

Kahneman and Tversky have countered some of this critique by drawing a parallel between their heuristic principles and the qualitative principles of Gestalt psychology—the latter being still valuable despite not being precisely specified. Irrespective of the various criticisms, the heuristic and biases program has undoubtedly led to a tremendous amount of research into the idea that people rely on cognitive heuristics made up of simple psychological processes rather than on complex procedures to make inferences about an uncertain world. As a result, this insight has been firmly established as a central topic of psychology.

## B. Fast and Frugal Choice Heuristics

More precisely specified models of heuristics have been studied by another research program that emerged in psychology in the 1990s. This new program considers fast and frugal heuristics for making decisions as the way the human mind can take advantage of the structure of information in the environment to arrive at reasonable decisions. Thus, it focuses on how mental capabilities and structured environments together can lead to accurate and useful inferences rather than focusing on the cases in which heuristics may account for poor reasoning. Most of the fast and frugal heuristics that Gigerenzer, Todd, and colleagues have proposed model the way humans make choices rather than probability judgments (a few others deal with additional tasks, such as estimation and classification).

Many of the choices humans make involve an inference or prediction about which of two objects will score higher on a criterion: Which soccer team will win? Which of two cities has a higher crime rate or higher cost of living? Which of two applicants will do a better job? When making such choices, we may have different amounts of information available. In the most limited case, if the only information available is whether or not each option has been encountered before, the decision maker can do little better than rely on his or her own partial ignorance, for instance, choosing recognized options over unrecognized ones. This may not sound like much for a decision maker to go on, but there is often information implicit in the failure to recognize something, and this failure can be exploited.

This kind of "ignorance-based" decision making is embodied in the recognition heuristic. This heuristic states that, when choosing between two objects (according to some criterion), if one is recognized and the other is not, then select the former. For instance, if predicting whether Manchester United or Bayer Leverkusen will win the European Soccer Champions League, this heuristic would lead most of us to bet on Manchester United. Why? European soccer teams are often named after European cities (e.g., Arsenal London or AC Milano), and people who are ignorant of the quality of European soccer teams can still use city recognition as a cue for soccer team performance. Cities with successful soccer teams are likely to be large, and large cities are likely to be recognized; hence, Manchester, which is more than two times as large as Leverkusen, is also more likely to be recognized and thus chosen as the winner.

The recognition heuristic will yield good choices more often than would random choice in those decision environments in which exposure to different possibilities is positively correlated with their ranking along the decision criterion being used. Animals also behave as if they apply similar rules: Norway rats, for instance, prefer to eat things they recognize through past experience with other rats (e.g., items they have smelled on the breath of others) over novel items.

Employing the recognition heuristic can lead to a surprising phenomenon called the less-is-more effect. This is the analytical and empirical observation that 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. A context in which this effect appears in the reasoning of people is judgments about

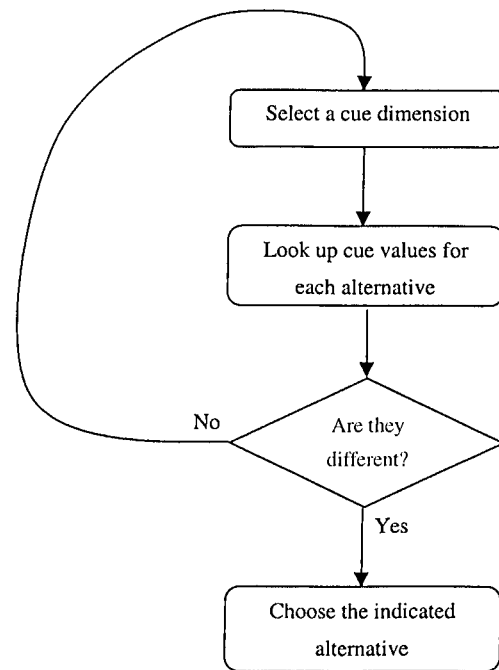
demographics. When American students were asked to pick the larger of two cities, they scored a median 71% correct inferences when city pairs were randomly constructed from the 22 largest U.S. cities and 73% when city pairs were from the 22 largest cities in Germany. The result is counterintuitive when viewed from the premise that more knowledge is always better: The students knew a lifetime of facts about U.S. cities that could be useful for inferring population, but they knew little or nothing beyond mere recognition about the German cities—and they did not even recognize about half of them. The latter fact, however, allowed them to employ the recognition heuristic and to pick German cities that they recognized as larger than those they did not. This heuristic could not be applied for choosing between U.S. cities, however, because the students recognized all of them and thus had to rely on additional retrievable information.

In choosing one of two options, most of the time we have more information than just a vague memory of recognition to go on. The situation of American students comparing American cities is just one example. When multiple pieces of information (or “cues”) are available for guiding decisions, how can a heuristic that limits search and processing of information proceed? A decision maker following the dictums of unbounded rationality would collect all the available information, weight it appropriately, and combine it optimally before making a choice. A more 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. One such approach is to rely on “one-reason decision making”: Stop looking for cues as soon as one is found that differentiates between the two options being considered. This allows the decision maker to follow a simple loop, depicted in Fig. 1: (i) Select a cue dimension; (ii) search for the corresponding cue values of each alternative; (iii) if they differ, then stop and choose the alternative for which the cue indicates a greater value on the choice criterion; (iv) if they do not differ, then return to the beginning of this loop to search for another cue dimension.

This four-step loop specifies a stopping rule (stopping after a single cue is found that enables a choice between the two options) and a decision rule (deciding on the option to which the one cue points). Depending on how cue dimensions are searched for in the first step, (i.e., depending on what kind of specific information search rule the heuristic uses), different one-reason decision-making heuristics can be formed. The Take The Best heuristic searches for cues in the order

of their ecological validity (i.e., their correlation with the decision criterion). Take The Last searches for cues in the order determined by their past success so that the cue that was used for the most recent previous decision is checked first during the next decision. Finally, the Minimalist heuristic selects cues in a random order.

For example, consider the task of inferring which of two cities in the United States has a higher homelessness rate. Assume that possible cues are “rent control,” “temperature,” “unemployment,” and “public housing,” each turned into a binary value (0 or 1) according to whether the actual value is below or above the median for U.S. cities. Rent control has the highest ecological validity, temperature has the second highest, and so on. The Minimalist heuristic only needs to know in which direction a cue “points.” For instance, the heuristic needs to know only whether warmer or cooler weather indicates a city with a higher rate of homelessness (in the United States, warmer weather is indeed associated more often with higher homelessness rates than with lower rates). The strategy of Minimalist is to search for cues in random order,



**Figure 1** A flowchart of one-reason decision making. First, select a cue dimension and ascertain the corresponding cue values for each alternative; next, check whether the values for that cue discriminate between the alternatives. If so, then choose the indicated alternative; if not, select another cue dimension and repeat this process. Random choice can be used if no more cues are available.

stop cue search when a cue is found that discriminates between the two cities, and then choose the city that has the cue value 1 when the other city has cue value 0. For instance, when inferring whether New York or Los Angeles has a higher homelessness rate, the unemployment cue might be the first cue randomly selected, and the cue values are found to be 1 for both cities. Because this cue does not discriminate between the cities, search is continued, the public housing cue is randomly selected, and the cues values are 0 for New York and 1 for Los Angeles. Search is stopped at this discriminating cue and the inference is made that Los Angeles has a higher homelessness rate, as it indeed does.

The Take The Best heuristic is exactly like Minimalist except that it considers cues in order of their validity from highest to lowest. If the highest validity cue does not discriminate, the next best cue is tried, and so forth. Thus, Take The Best differs from Minimalist only in the information search rule, but it has the same stopping and decision rule. Take The Best (unlike the Minimalist, Take The Last, and recognition heuristics) is an instance of the class of lexicographic decision strategies. This term signifies that the cues are looked up in a fixed order of validity, and the first cue where choices differ is used alone to make the decision, like the alphabetic order used to arrange words in a dictionary. The Arabic number system is also lexicographic. To determine which of two numbers with equal digit length is larger, one has to start by examining the first (leftmost) digit: If this digit is larger in one of the numbers, the whole number is larger. If they are equal, one has to examine the second digit, and so on (a simple method that is not possible for Roman numbers). There is growing empirical evidence that people actually use lexicographic heuristics such as Take The Best, particularly when time is limited.

How well do these one-reason decision heuristics perform? Table I compares the performance of three fast and frugal heuristics (Minimalist, Take The Best, and Take The Last) to that of multiple regression and Dawes's and Franklin's rule. Unlike the heuristics, multiple regression is a computationally expensive linear strategy that calculates weights that reflect the covariances between predictors or cues. When the task is merely fitting the given data set, multiple regression is the most accurate strategy, by two percentage points, followed by Take The Best. However, when the task is to generalize from a training set to a test set, a simple heuristic such as Take The Best can outperform multiple regression (note that multiple regression has

**Table I**  
Performance of Three Fast and Frugal Heuristics (Take The Best, Minimalist, and Take The Last) and Three Linear Strategies (Dawes's rule, Franklin's rule, and multiple regression) Averaged across 20 Empirical Data Sets<sup>a</sup>

| Heuristic/strategy  | Frugality | Accuracy               |                               |
|---------------------|-----------|------------------------|-------------------------------|
|                     |           | Fitting<br>(% correct) | Generalization<br>(% correct) |
| Take The Best       | 2.4       | 75                     | 71                            |
| Minimalist          | 2.2       | 69                     | 65                            |
| Take The Last       | 2.1       | 70                     | 65                            |
| Franklin's rule     | 7.7       | 75                     | 71                            |
| Dawes's rule        | 7.7       | 73                     | 69                            |
| Multiple regression | 7.7       | 77                     | 68                            |

<sup>a</sup>The average number of predictions available in the 20 data sets was 7.7. Frugality indicates the mean number of cues actually used by each strategy. Accuracy indicates the percentage of correct answers achieved by the heuristics and strategies when fitting data (i.e., fit a strategy to a given set of data) and when generalizing to new data (i.e., use a strategy to predict new data).

all the information Take The Best uses and more). The reason is that by being simple, the heuristics can avoid being too matched to any particular environment—that is, they can escape that curse of overfitting.

Overfitting refers to the problem of a model that is closely matched to one situation (set of data) failing to predict accurately in another similar situation (another set of data). This phenomenon can arise from assuming that every detail in a given environment is of great relevance. Consider forecasting of the U.S. presidential elections as an example. Beyond traditional variables such as incumbency and the state of the election-year economy, a plethora of additional variables have been suggested as predictors of recent U.S. presidential elections, including the voting behavior in Okanogan County (a rural stretch of north-central Washington), the rise or fall of women's hemlines, and the height of the candidates. General strategies such as multiple regression can in fact incorporate each of these and many more variables into the unlimited collection of free variables in their forecast models. As accurate as such parameter-laden forecast models may be for describing particular recent presidential elections, their accuracy in predicting other situations (e.g., earlier U.S. presidential elections or elections in other locations) may well be minimal. That is, these models can easily overfit the particular (training) data set and thereby fail to generalize to the new (testing) data set. In contrast, if a forecast model uses many



fewer parameters, for instance, just incumbency and height of the candidates (which predicted the winner of every election since World War II, except in 1976 and 2000), it is likely to avoid overfitting and thereby generalize better to new situations.

Fast and frugal heuristics (like lexicographic strategies) are noncompensatory, meaning that once they have used a single cue to make a decision no further cues in any combination can undo or compensate for that one cue's effect. When the information in the decision environment is structured in a matching noncompensatory fashion (i.e., the importance or validity of cues decreases rapidly such that each weight of a cue is larger than the sum of all weights to come, e.g., one-half, one-fourth, one-eighth, and so on), the Take The Best heuristic can exploit that structure to make correct decisions as often as compensatory rules. Take The Best also performs comparatively well when information is scarce; that is, when there are many more objects than cues to distinguish them. Further research is needed to explore what environment structures can be exploited by different fast and frugal heuristics.

### C. Heuristics for Multiple Alternative Choices

Not all choices in life are presented to us as convenient pairs of alternatives, of course. Often, we must choose between several alternatives, such as which restaurant to go to, which apartment to rent, or which stocks to buy. Table II lists various decision heuristics that have been proposed in the psychological literature for choosing one out of several alternatives, where each alternative is characterized by cue (or attribute) values and where the importance of a cue is specified by its weight (or validity). This collection is not exhaustive, and it only focuses on heuristics for inference rather than preference (albeit some could be applied to preferences as well). However, the heuristics represent a wide range of different information search, stopping, and decision rules. Among the heuristics for multiple alternative choice, lexicographic (LEX), lexicographic semiorder (LEX-Semi), and elimination by aspects (EBA) are noncompensatory, whereas the rest are compensatory heuristics, integrating (at least some of) the available information and making trade-offs

**Table II**  
Description of Various Multiple Alternative Choice Heuristics<sup>a</sup>

| Heuristic                                     | Description   |
|---|---|
| Franklin's rule or weighted additive rule     | Calculates for each alternative the sum of the cue values multiplied by the corresponding cue weights (validities) and selects the alternative with the highest score.  |
| Dawes's rule or additive rule                 | Calculates for each alternative the sum of the cue values (discretized to either 1 or -1) and selects the alternative with the highest score.   |
| Good features                                 | Selects the alternative with the highest number of good features: a good feature is a cue value that exceeds a specified cut-off.   |
| Weighted pros                                 | Selects the alternative with the highest sum of weighted "pros." A cue that has a higher value for one alternative than for the others is considered a pro for this alternative. The weight of each pro is defined by the validity of the particular cue.   |
| LEX or lexicographic                          | Selects the alternative with the highest cue value on the cue with the highest validity. If more than one alternative has the same highest cue value, then for these alternatives the cue with the second highest validity is considered, and so on.  |
| LEX-Semi or lexicographic semiorder           | Works like LEX, with the additional assumption of a just-noticeable difference. Pairs of alternatives with less than a just-noticeable difference between the cue values are not discriminated.   |
| EBA or elimination by aspects                 | Eliminates all alternatives that do not exceed a specified value on the first cue examined. If more than one alternative remains, another cue is selected. This procedure is repeated until only one alternative is left. Each cue is selected with a probability that is proportional to its weight. |
| LEX-Add or lexicographic additive combination | Represents a combination of two strategies. It first uses LEX-Semi to choose two alternatives as favorites and then evaluates them by Dawes's rule and selects the one with the highest sum.  |

<sup>a</sup>Defined in terms of alternatives (options), cues (information), and weights (importance of information).

between the relevant cues to form an overall evaluation of each alternative.

Weighing and summing of all available information has been used to define rational judgment at least since the Enlightenment: The concepts of expected value and utility, Benjamin Franklin's moral algebra, and *Homo economicus* all rely on these two fundamental processes. The heuristics for multiple alternative choices in Table II can be seen as various shortcuts of these two processes. Dawes's rule, for instance, questions the importance of precise weighting. In the 1970s and 1980s, Robyn Dawes and colleagues showed that tallying information (cues) in terms of simple unit weights, such as +1 and -1, typically led to the same predictive accuracy as the "optimal weights" in multiple regression (particularly when generalizing to new data). Thus, in situations in which the task is to predict what is not yet known (rather than to fit what is already known), weighting information does not seem to matter much, as long as one gets the sign right.

On the other hand, LEX (a generalization of Take The Best), LEX-Semi, and EBA do not require summing procedures. All three heuristics use a simple form of weighting by ordering the cues, but they do not sum the cues. Gigerenzer and colleagues collected counterintuitive evidence that this simple weighting without summing (as in the Take The Best heuristic) can be as accurate and in some circumstances (e.g., generalization) even more accurate than complex decision strategies such as multiple regression.

Among the choice heuristics listed in Table II, EBA, proposed by Amos Tversky, is the most widely known elimination model in psychology. In sequential elimination choice models, one alternative is chosen from a set of possibilities by repeatedly eliminating subsets of alternatives from further consideration until only a single choice remains. One of the motivating factors in developing EBA in particular as a descriptive model of choice was that there are often many relevant cues that may be used in choosing among complex alternatives. EBA deals with this challenge by probabilistically considering successive cues (which are chosen with a probability proportional to their importance), selecting one at a time, and eliminating all the alternatives that do not possess this current cue, until a single alternative remains as the final choice. Other elimination heuristics select cues in a different manner (e.g., deterministically or based on validity) or use them to process the alternatives in other ways (e.g., in terms of thresholds rather than presence or absence).

## D. Heuristics for Sequential Search

The heuristics discussed so far for choosing one option from many operate with the assumption that all the possible options (e.g., cities to choose between) are presently available to the decision maker. In many real-world choice problems, though, an agent encounters options in a sequence spread out over time. The options typically appear in random order and are drawn from a distribution with parameters that are only partially known in advance. In this case, the search for possible options, rather than just for information about those already present, becomes central.

The traditional normative approach to such problems is to search until one finds an option below a precalculated reservation price that balances the expected benefit of further search against its cost; this requires full knowledge of the search costs and the distribution of available alternatives. Heuristics that simplify the reservation price calculation (by replacing an integral with a weighted sum) can come very close to normative performance (e.g., at selecting good prices during a shopping trip to several stores). Other heuristics require less knowledge, such as "Keep searching until the total search cost exceeds 7.5% of the best price found." Herbert Simon's bounded rationality principle of satisficing suggests setting an aspiration level equal to an alternative that is good enough for the decision maker's needs (rather than optimal) and searching until that aspiration is met. Exactly how the aspiration level can be set varies with the search setting (e.g., whether it is a one-sided search such as shopping or a two-sided mutual search such as finding a mate). Finally, another type of search heuristic that people use stops search after a particular pattern of alternatives is encountered rather than after some threshold is exceeded (despite the fact that pattern should not matter from a normative perspective). For instance, the "one-bounce" and "two-bounce" rules state that one should keep searching for a low price until prices go up for the last or two last alternatives, respectively.

## E. Social Decision Heuristics

Decision-making mechanisms can exploit the structure of information in the environment to arrive at better outcomes. The most important aspects of an agent's environment are often created by the other agents with which it interacts. Two of the key problems social agents face are the questions of how to (fairly)

divide up resources among one another and how to make cooperative decisions in situations in which the pursuit of self-interest by each agent would lead to a poor outcome for all. We consider each of these problems in turn.

The task of fairly dividing up resources is ubiquitous, ranging from distributing a cake among siblings to dividing an estate among heirs and splitting a fixed budget among a group of faculty members at an academic institution. Although there are a plethora of fair-division procedures, Brams and Taylor classified them according to a few dimensions, such as the number of players to which they are applicable ( $n=2, 3, 4$ , or more), the properties they satisfy (e.g., proportionality, envy-freeness, and efficiency), and whether or not the division has to be exact or only approximate. A simple but well-known decision heuristic that may be familiar to many parents is "one divides, the other chooses." Although this heuristic stipulates division of labor, it does not specify how the person who divides the resource (i.e., a cake) actually does it. If, however, the person who divides the cake understands the strategic interest of the other, the implied division rule is to divide up the resource such that one is indifferent between the two parts; the other person will then choose whatever he or she considers to be the larger piece. This way each person is assured of getting what he or she perceives to be at least half the resource, and neither party thinks that the other received a larger piece of cake. A fair-division procedure with these properties is said to be proportional and envy-free.

An example of a situation in which the pursuit of self-interest by each party leads to a poor outcome for all is that in which two industrialized regions of the world (e.g., the United States and the European Union) have established trade barriers to each other's exports. Because of the mutual advantages of free trade, both regions would be better off if these barriers were eliminated. However, if either region were to unilaterally give up the barriers, it would be faced with terms of trade that hurt its own economy. In fact, no matter what America does, the European Union (EU) is better off retaining its own trade barriers and vice versa. This strategic situation, in which the incentive to retain trade barriers for both regions produces a worse outcome than would have been possible had both decided to cooperate, is known as a prisoner's dilemma game. This game is just one among many situations that game theorists examine in order to analyze and model the strategic interactions of social agents.

If there is some likelihood that the players will encounter each other in the future, as in trade between

the United States and the EU, the interaction become an iterated prisoner's dilemma game. There are a number of possible decision heuristics for this situation. A particularly simple but surprisingly successful decision heuristic is the tit-for-tat heuristic that Anatol Rapoport submitted to Axelrod's famous computer tournament. Given the possibility of cooperating with or defecting against the other player at each time step, tit-for-tat starts with a cooperative choice and thereafter does what the other player did on the previous move. In other words, tit-for-tat searches for a minimal amount of information (the counterpart's behavior in the last round) and cooperates if the last move was cooperative but defects if the last move was defective. Thus, akin to some of the heuristics described previously, tit-for-tat does not have to weigh and combine pieces of information in some more or less computationally expensive way. Many other successful heuristics, such as generous tit-for-tat and win-stay-lose-shift, have also been proposed for iterated prisoner's dilemma and other games.

#### IV. HOW TO MEASURE A HEURISTIC'S SUCCESS

The study of heuristics is a key approach to understanding how real minds make decisions for two main reasons. First, many of life's important problems, from choosing a mate to finding a job, cannot be solved in an optimal way because the space of possibilities that must be taken into account is often unlimited; hence, heuristic shortcuts are called for. Second, even when this space of possible solutions is limited and knowledge is complete, optimization may require unfeasible amounts of computation (as in trying to determine the best next move in chess) so that, again, heuristics will be an appropriate approach for the mind to take.

The fact that there are no optimal strategies for many real-world tasks, however, does not mean that there are no performance criteria. One set of criteria that is often used to evaluate judgments and decisions is their internal coherence, defined as accordance with the laws of probability theory and logic. For instance, if judgments are consistent (e.g., "I always think that event A is more likely than B") and transitive ("I think that A is more likely than B, B is more likely than C, and therefore that A is more likely than C"), this is taken as an indication that the underlying decision strategies are rational. If such criteria are violated, this is typically held to be a sign of irrationality on the part of the decision maker. The heuristics-and-biases

research program has focused on such relatively abstract coherence criteria to indicate when a heuristic produces reasonable or unreasonable decisions.

Alternatively, the success of a heuristic can be measured by comparing its performance with the requirements of its environment, such as accuracy, frugality, and speed. Lexicographic strategies (e.g., Take The Best) are often evaluated via correspondence criteria relating to real-world decision performance (such as how often they correctly choose the larger object in a pair). Comparing heuristics' performance to the requirements of the external world rather than to internal consistency stems from the view that the primary function of heuristics is not to be coherent. Rather, their function is to make reasonable adaptive inferences about the real social and physical world given limited time and knowledge.

The two kinds of criteria, coherence and correspondence, can sometimes be at odds with each other. For instance, in social situations, including some competitive games and predator-prey interactions, it can be advantageous to exhibit inconsistent (and hence noncoherent) behavior in order to maximize adaptive unpredictability (and hence correspondence with real-world goals) and avoid capture or loss. As another example, the Minimalist heuristic violates the coherence criterion of transitivity but nevertheless makes fairly robust and accurate inferences in particular environments. Thus, intransitivity does not necessarily imply high levels of inaccuracy, nor does transitivity guarantee high levels of accuracy: Logic and adaptive behavior are logically distinct.

Finally, it is important to measure the performance of decision mechanisms in terms of how well they make decisions when applied to new data; that is, how they generalize to new situations rather than merely how closely they can be adjusted or fit to a static set of data. In this regard, simple heuristics will often do very well, being about as accurate as complex general strategies that work with many free parameters. The reason is that simple heuristics can avoid being too matched to any particular environment; that is, they can escape the curse of overfitting mentioned earlier. As a consequence, a computationally simple heuristic that uses only some of the available information can be more robust, making more accurate predictions for new data.

## V. CONCLUSION

Simplicity in models has aesthetic appeal. The mechanisms are readily understood and communicated, and they are amenable to step-by-step scrutiny. Furthermore, Popper has argued that simpler models are more falsifiable. However, the idea that humans make many decisions using simple heuristic mechanisms is important not just because the resulting simple models are transparent and easily falsifiable. More important, simple heuristics may be the only approach available for real minds making decisions in the real, uncertain, time-pressured world.

### See Also the Following Articles

ARTIFICIAL INTELLIGENCE • INFORMATION PROCESSING • INTELLIGENCE • LOGIC AND REASONING • NEURAL NETWORKS

### Suggested Reading

- Axelrod, R. (1984). *The Evolution of Cooperation*. Basic Books, New York.
- Brams, S. J., and Taylor, A. D. (1996). *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge Univ. Press, Cambridge, UK.
- Dawes, R. M. (1979). The robust beauty of improper linear models in decision making. *Am. Psychologist* **34**, 571–582.
- Gigerenzer, G., Todd, P. M., and the ABC Research Group (1999). *Simple Heuristics That Make Us Smart*. Oxford Univer. Press, New York.
- Kahneman, D., Slovic, P., and Tversky, A. (Eds.) (1982). *Judgment under Uncertainty: Heuristics and Biases*. Cambridge Univer. Press, Cambridge, UK.
- Newell, A., and Simon, H. A. (1972). *Human Problem Solving*. Prentice Hall, Englewood Cliffs, NJ.
- Payne, J. W., Bettman, J. R., and Johnson, E. J. (1993). *The Adaptive Decision Maker*. Cambridge Univ. Press, New York.
- Pólya, G. (1945). *How to Solve It: A New Aspect of Mathematical Method*. Princeton Univ. Press, Princeton, NJ.
- Simon, H. A. (1955). A behavioral model of rational choice. *Q. J. Econ.* **69**, 99–118.
- Simon, H. A. (1956). Rational choice and the structure of environments. *Psychol. Rev.* **63**, 129–138.
- Simon, H. A. (1990). Invariants of human behavior. *Annu. Rev. Psychol.* **41**, 1–19.
- Svenson, O. (1979). Process descriptions of decision making. *Organizational Behav. Hum. Performance* **23**, 86–112.
- Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychol. Rev.* **79**, 281–299.