Published in: *Swiss Journal of Economics and Statistics, 133* (2/2), 1997, 201–218.

Bounded Rationality: Models of Fast and Frugal Inference

Gerd Gigerenzer[1]
Max Planck Institute for Psychological Research, Munich, Germany

Humans and other animals need to make inferences about their environment under constraints of limited time, knowledge, and computational capacities. However, most theories of inductive inferences model the human mind as a supercomputer like a Laplacean demon, equipped with unlimited time, knowledge, and computational capacities. In this article I review models of fast and frugal inference, that is, satisficing strategies whose task is to infer unknown states of the world (without relying on computationaly expensive procedures such as multiple regression). Fast and frugal inference is a form of bounded rationality (Simon, 1982). I begin by explaining what bounded rationality in human inference is not.

1. Bounded Rationality is Not Irrationality

In his chapter in John Kagel and Alvin Roth's *Handbook of Experimental Economics* (1995), Colin Camerer explains that "most research on individual decision making has taken normative theories of judgment and choice (typically probability rules and utility theories) as null hypotheses about behavior," and has labeled systematic deviations from these norms "cognitive illusions" (p. 588). Camerer continues, "The most fruitful, popular alternative theories spring from the idea that limits on computational ability force people to use simplified procedures or 'heuristics' that cause systematic mistakes (biases) in problem solving, judgment, and choice. The roots of this approach are in Simon's (1955) distinction between *substantive* rationality (the result of normative maximizing models) and *procedural* rationality." (p. 588) In the preface to their anthology, Daniel Kahneman, Paul Slovic, and Amos Tversky (1982) relate their heuristics-and-biases program to "Simon's treatment of heuristics of reasoning and of bounded rationality" (p. xii). Richard Thaler (1991) explains that Kahneman and Tversky have shown that "mental illusions should be considered the rule rather than the exception. Systematic, predictable differences between normative models of behavior and actual behavior occur because of what Herbert Simson [sic!] (1957, p. 198) called 'bounded rationality'." (p. 4)

My first point is to disentangle the confusion between bounded rationality (or procedural rationality) and irrationality inherent in these statements—a confusion which has been repeated many times (e.g., Oaksford & Chater, 1992; see Lopes, 1992). I use the term "irrationality" as a shorthand for the various "errors" and "fallacies" in statistical and probabilistic judgment which Camerer lists, such as the conjunction fallacy, the base rate fallacy, and the overconfidence bias. In each of these alleged demonstrations of irrationality, the assumption is made that it is crystal-

1 I am grateful for comments on earlier versions of this paper by Bernhard Borges, Ralph Hertwig, Ulrich Hoffrage, Timothy Ketelaar, and Laura Martignon.

clear what the correct judgment is. Sound reasoning is reduced to applying a simple rule such as the conjunction rule or Bayes' rule, without even looking at the content and context of the task (Gigerenzer, 1996a; Gigerenzer & Murray, 1987). Systematic deviations of human judgment from these norms (the "null hypotheses") are called "biases" or "errors" and attributed to crude "heuristics"—representativeness, availability, and anchoring. What do these "heuristics and biases" have to do with bounded rationality?

To start with, most of the heuristics and biases in statistical and probabilistic judgment that Camerer lists stem from the anthology by Kahneman et al. (1982), in which, as mentioned before, the link to bounded rationality is made in the preface. This anthology contains all of Tversky and Kahneman's major papers since the early 1970s, none of which has a single citation to Simon. Given the normal care that Tversky and Kahneman take in crediting others, it is unlikely that their research actually had its roots in Simon's concept of bounded rationality (Lopes, 1992). This leaves us with the possibility that there is, nevertheless, a deep link between the two programs which has just gone unnoticed for a decade or so. So let us examine how the heuristics and biases actually relate to bounded rationality.

For that we need some criteria for bounded rationality. To find specific criteria turns out to be harder than it seems: initially, the concept of bounded rationality was only vaguely defined, and one could "fit a lot of things into it by foresight and hindsight" (Simon, 1992, p. 18). I introduce four general requirements (rather than specific ones such as explicit stopping rules, see below), two related to each "blade" of the "scissors" that shape bounded rationality: "the structure of task environments and the computational capabilities of the actor" (Simon, 1990, p. 7).

(1) *The task is too hard to compute an exact solution.* In Simon's (1979) words, "Satisficing [is] aiming at the good when the best is incalculable" (p. 3). To use one of his favorite examples, "If the game of chess, limited to its 64 squares and six kinds of pieces, is beyond exact computation, then we may expect the same of almost any real-world problem …" (Simon, 1990, p. 6). There are two readings of the term "incalculable." First, the task is too hard for the computational power humankind has available today, in minds and machines, such as in the case of chess. Second, the task is too hard for the limited computational resources the average mind has available.

(2) *The task environment needs to be studied.* The moral from Simon's two-bladed scissors analogy is "that one must consider both the task environment and the limits upon the adaptive powers of the system" (Simon, 1991, p. 36).

(3) *Limited cognitive resources.* A person has to make an inference under limited time, limited knowledge, limited computational capacities, and limited resources for obtaining further information. These resources are insufficient to compute the exact solution.

(4) *A satisficing strategy is specified.* This condition requires some precisely formulated strategy, which is proposed as a model of bounded rationality. This satisficing strategy computes a judgment or decision from the analysis of the task environment.

Do "heuristics and biases" satisfy these four general requirements? Consider first a concrete example, one of the most celebrated cognitive illusions: the "conjunction fallacy" in the Linda problem. Imagine you are a participant in an experiment; in front of you is a text problem and you begin to read (Tversky & Kahneman, 1983):

Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations. Which is more probable:

(a) Linda is a bank teller,

(b) Linda is a bank teller and active in the feminist movement.

Assume you chose (b), just as some 80%–90% of the participants in previous experiments did. Tversky & Kahneman (1983) argued that this judgment is a reasoning fallacy, because it violates the conjunction rule:

$$p(A \cap B) \leq p(A), \text{ and } p(A \cap B) \leq p(B).$$

In words, the probability of the conjunction of two events A and B cannot be larger than the probability of either of the two events. This alleged demonstration of human irrationality has been widely accepted and publicized. Stephen J. Gould (1992) puts the message clearly: "Tversky & Kahneman argue, correctly I think, that our minds are not built (for whatever reason) to work by the rules of probability." (p. 469) The conjunction fallacy has been suggested as a cause of many human misfortunes and disasters, such as US security policy (Kanwisher, 1989) and people's assessment of the chance of nuclear reactor failures (Stich, 1985). Let us now see what this alleged demonstration of human irrationality has to do with bounded rationality.

(1) *Is the task too hard?* No. Different from chess and real-world situations, one does not even need a pocket calculator to compute (what is considered to be) the "correct" solution to the Linda problem.
(2) *Is the task environment studied?* No. No analysis of the situation is needed for the "correct" solution. One does not even need to read the description of Linda. Tversky and Kahneman (1983) assume that all that matters for sound reasoning is to map the term "probable" into mathematical probability, and the term "and" into logical AND. That's all that is needed for the conjunction rule. The norm is content-blind, therefore the environment does not matter (Gigerenzer, 1996a). Any knowledge, such as about bank tellers and feminists, is considered irrelevant for sound reasoning. As a consequence, the participants' assumptions about what the experimenter wants them to do are not analyzed either.
(3) *Do limited cognitive resources apply?* Limited cognitive resources are not an issue in the Linda problem. For finding the "correct" solution, absolutely no knowledge about the environment is needed and no resources for obtaining further information are required; thus the issue of limited knowledge does not apply. Similarly, little if any computational capacities are needed, and time constraints or information costs are of no relevance.
(4) *Is a satisficing strategy specified?* The standard explanation for the "conjunction fallacy" is that people do not reason according to the laws of probability, but use a heuristic called "representativeness": The description of Linda is more representative of a feminist bank teller than of a bank teller. The term "representative" seems to mean "similar." But which of the many different strategies for computing similarity is meant by this word? The strategy for computing representativeness has not yet been specified.

I conclude that the conjunction fallacy and its proposed explanation, the representativeness heuristic, satisfy none of these four general criteria of bounded rationality. This result holds more generally for the "heuristics and biases" in statistical and probabilistic reasoning (though occasionally one of the criteria may be satisfied). First, what is considered to be the "correct" solution can almost always be computed with a few keystrokes on a cheap calculator (overconfidence bias is one exception). Second, the norms are content-blind, therefore any analysis of the task environment is assumed to be unnecessary in the first place. Third, for the same reason—content-blind norms—knowledge and information search play little if any role, and nor do limits of memory and attention. Fourth, none of the three heuristics proposed in the early 1970s—representativeness, availability, and anchoring—has ever been turned into a precise model. They have remained one-word explanations with the virtue of Rorschach inkblots. Every researcher can read into them what he or she wishes. The reluctance to specify precise and falsifiable process models, to clarify the antecedent conditions that elicit various heuristics, and to work out the relationship between heuristics has been repeatedly pointed out (e.g., Einhorn & Hogarth, 1981; Lopes, 1991; Shanteau, 1989; Wallsten, 1983). However, Kahneman and Tversky (1996, p. 585) still continued to defend undefined "heuristics" in reaction to critique (see Gigerenzer, 1993, 1994, 1996a). Thus, by all four criteria, the heuristics-and-biases program has little to do with models of bounded rationality.

However, one could argue that at least the first criterion holds in its weak version, which says that the solution is incalculable for the average person. Even if the task does not look difficult, it is so for human minds, so the argument goes. And this result allegedly holds fairly stable across variations, as we learn from Camerer (1995). But this picture is misleading; it ignores the recent demonstrations of how to make the conjunction fallacy largely disappear (Fiedler, 1988; Gigerenzer, 1991; Hertwig & Gigerenzer, 1996). Consider the simple fact that the term "probable," attached to a single event (such as that Linda is a bank teller), has several legitimate meanings besides mathematical probability. Some examples are "plausible," "credible" as in "a credible story," and "that may in view of present evidence be reasonably expected to happen," (see e.g., the Oxford English Dictionary). Similarly, statisticians of the frequentist school would not accept that single-event "probabilities" as in the Linda problem have anything to do with the mathematical theory of probability (Gigerenzer, 1994). Thus, for both psychological and statistical reasons, an adequate test of the human capacity to reason according to the conjunction rule is to state the problem in frequencies rather than in ambiguous single-event probabilities. Ralph Hertwig and I have formulated the Linda problem in terms of frequencies (Hertwig & Gigerenzer, 1996). Everything was left constant except that we replaced the ambiguous phrase "Which is more probable?" by a frequency judgment: "There are 100 women like Linda. How many of them are (a) bank tellers, (b) bank tellers and active in the feminist movement?" In a series of experiments, conjunction violations dropped from almost 90% in the original probability version to as low as 0% in the frequency version. Fiedler (1988) had earlier shown similar results: Violations of the conjunction rule dropped from 80% to 90% in probability judgments to about 20% in frequency judgments.[2] Thus, the task is not too hard for most people, once it is clarified that it is about mathematical probability and not about something else.[3]

---

2 Note that Tversky and Kahneman (1983) had reported an effect of frequency for a different problem, but did not pay much attention to it.

3 The reason why most people chose to interpret "Which is more probable?" other than in terms of mathematical probability seems to be that the latter would imply that the description of Linda is irrelevant for the task, which in turn would imply that the experimenter violates Grice's (1975) conversational maxim of "relevance." This conclusion is supported by the task analysis, experiments, and paraphrasing tasks reported in Hertwig and Gigerenzer (1996). In other words, people's judgments reflect social rationality, not mental inability. In defense against my critique (e.g., Gigerenzer, 1991), Kahneman and Tversky (1996) constructed a between-subjects design for the Linda problem, and claimed that at least in this special situation the conjunction fallacy is obtained even with frequency judgments. They asked one group "Suppose there are 1,000 women who fit this description. How many of them are (a) high school teachers? (b) bank tellers?" and a second group "How many of them are (a) high school teachers, and (c) bank tellers and active feminists." (p. 587) The estimate of (c) was higher than that of (b), which they took as a violation of the conjunction rule. Note that Kahneman and Tversky (1996) had changed in this experiment the original conjunction "bank teller and active in the feminist movement" into "bank tellers and feminists," that is, a noun-active/adjective to a noun-noun combination. They did not point out this change. Hertwig (1997) has provided evidence that people actually tend to read the new formulation as a disjunction (rather than a conjunction), and that the "conjunction fallacy" in Kahneman and Tversky's (1996) between-subjects design disappeared when this misleading formulation is replaced by the original conjunction. The problem with Kahneman and Tversky's (1996) defense is the same as with their original analysis of the Linda problem: A content-blind norm is applied, and how people actually understand the task environment is not analyzed.

The conjunction fallacy is not the only so-called cognitive illusion that largely disappears when probabilities are replaced by frequencies. Gigerenzer, Hoffrage and Kleinbölting (1991) showed that overconfidence bias completely disappeared when participants were asked "How many of the last 50 questions did you get correct?" instead of "What is the probability that your answer to this question is correct?" (see also May, 1987; Sniezek & Buckley, 1993). Lay persons' reasoning followed Bayes' rule about three times as often when the information was in a frequency format rather than in a probability format (Cosmides & Tooby, 1996; Gigerenzer & Hoffrage, 1995). Physicians' diagnostic inferences followed Bayes' rule four times as often with frequency formats than with probabilities (Gigerenzer, 1996b; Hoffrage & Gigerenzer, 1996). Teigen (1974) reported that overestimation of probabilities (e.g., What is the probability that a randomly chosen female student at the University of Bergen is above 160 cm tall?) changed into more realistic estimates when subjects were given the opportunity to estimate frequencies (e.g., If we measure 500 female students, how many of them will be above 160 cm tall?). The difference between single events and repeated events also makes the "illusion of control" (Langer, 1975) largely disappear (Budescu & Bruderman, 1995; Koehler, Gibbs, & Hogarth, 1994), makes the certainty effect and the possibility effect (Kahneman & Tversky, 1979) largely disappear (Keren, 1991; Keren & Wagenaar, 1987), and reduces preference reversals (Wedell & Böckenholt, 1990). A review of the effects of frequency is in Gigerenzer (1991, 1994), and theoretical explanations in Gigerenzer and Hoffrage (1995) and Gigerenzer et al. (1991). For a different view see Kahneman and Tversky (1996), and for my response, Gigerenzer (1996a).

"Cognitive illusions" have been presented in the last three decades as hard facts similar to "visual illusions"—stubborn, largely ineradicable, genuine illusions, to which laymen and experts fall prey. The fact that one-and-the-same factor, frequencies versus probabilities, can make such a broad spectrum of alleged cognitive illusions largely disappear suggests that the tasks are not too hard, and the fault is not simply in the human mind. These results should not be read to imply that frequency judgments are always correct. There exist theories of cognitive processes that predict when they are and when not (Gigerenzer et al., 1991; Gigerenzer & Hoffrage, 1995). But it should be clear that the single most trenchant conclusion reached by the heuristics-and-biases program, namely that people are all too bad at reasoning, is itself, to a large degree, an illusion fostered by all-too-narrow norms of sound reasoning.

To summarize: The study of bounded rationality has been recently associated with the search for biases, defined as systematic discrepancies from some rule of probability. I have stated four general requirements for bounded rationality and concluded that the heuristics-and-biases approach to human judgment has little to do with studying bounded rationality. The stock-in-trade biases tend to disappear largely when the problems are formulated in terms of frequencies rather than probabilities.

The purpose of models of bounded rationality cannot be to explain deviations of human judgment from rules of probability. In the situations in which bounded rationality applies, the solution cannot be reduced to one of these rules. The purpose is to explain how people can do better than chance, that is, to explain deviations from random performance in the direction of successful performance.

2. Models of Satisficing Inference

How can a mind infer unknown properties of its environment on the basis of limited knowledge about that environment? How can these inferences be modeled, assuming the constraints of limited time and computational capacities? I will consider models of satisficing inference that embody, in addition to the general criteria listed above, the following specific criteria:

– Step-by-step procedures
– Limited search (simple stopping rules)
– One-reason decision making (non-compensatory strategies)
– Exploitation of a lack of knowledge (how to make positive use of one's ignorance)
– Exploitation of structures of information (structures of environments).

This paper deals with the following type of inference: Which of two objects scores higher on a criterion? This inference is a special case of the more general problem of infering which object in a class of M objects has the highest value on a criterion, but I will consider here only the case of $M = 2$. Examples are treatment allocation (e.g., which of two patients to treat first in the emergency room, with life expectancy after treatment as criterion), financial investment (e.g., which of two options to buy, with profit as criterion), and demographic predictions (e.g., which of two places has higher pollution, mortality rates, and so on).

Take The Best

Take The Best is a satisficing algorithm designed for problems of this kind, that is, for situations in which fast inferences have to be made about which of two objects (patients, alternatives) scores higher on some criterion (Gigerenzer & Goldstein, 1996). The general situation is illustrated in Figure 1. There are N objects (*a, b, c,* …) and a number of predictors that have binary values (the situation can be generalized to continuous predictors, e.g., by dichotomizing). I explain the step-by-step algorithm of Take The Best with a demographic problem that we originally used to

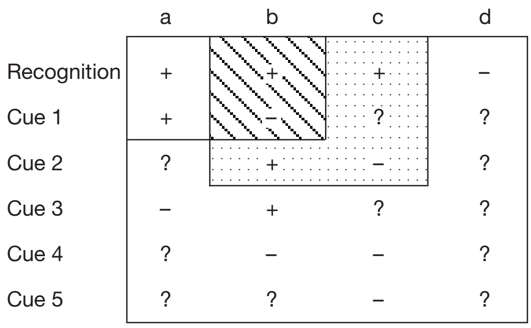

Objects *a, b,* and *c* are recognized, *d* is not. Predictor values are positive (+) or negative (–); missing knowledge is shown by a question mark. Predictors are ordered by their validities. To infer whether *a* > *b,* the Take The Best algorithm looks up only the values in the striped space. To infer whether *b* > *c* search is bounded to the dotted space. The other predictor values are not looked up (Gigerenzer & Goldstein, 1996).

Figure 1. Illustration of bounded search through limited knowledge.

study its performance: Which of two cities has a larger population? Here, *a* and *b* are two German cities, say Bremen and Heidelberg. Examples of predictors that indicate higher population are soccer team (whether or not a city has a team in the major soccer league) and state capital (whether or not a city is a state capital). The predictors are ordered according to their (perceived) validity, with Predictor 1 at the top. The predictor values can be positive (a city has a soccer team, which indicates larger population), negative (has no soccer team), or unknown (the person has no information). The task is to infer which city, *a* or *b,* has a larger population. In addition to these ecological predictors, there is a subjective cue, recognition (whether or not the person has heard of the city). Recognition only plays a role when it is correlated with the criterion, as it is with population.

Step-by-step procedure. Take The Best looks up in memory, step-by-step, information concerning predictors, until a predictor is found that discriminates. Discrimination occurs when one object has a positive value and the other has no positive value (negative or unknown). How does Take The Best infer which of two cities, Bremen *(a)* and Heidelberg *(b),* has the larger population, given the limited knowledge in Figure 1? First, the recognition values are looked up, which in this case do not discriminate, because both are positive. Next, the values on the top-ranking ecological predictor, the soccer team cue (Predictor 1) are searched. Bremen has a soccer team in the major league, but Heidelberg does not. Search in memory is terminated, and the inference is made that Bremen has the larger population. No other predictor values are looked up in memory. Thus, only 4 out of 12 values in Figure 1 (striped area) are looked up. None are integrated. Consider now the inference, which of *b* and *c* has a higher population. The values for recognition and Predictor 1 do not discriminate, but those of Predictor 2 do. Thus 6 values are looked up (dotted area in Figure 1) before search is terminated and the inference is made that *b* has the higher population. Finally, consider the inference, which of *c* and *d* is larger? Object *c* is recognized, *d* is not; that's it. The inference is made that *c* has the larger population. The Take The Best algorithm is shown in the form of a flow chart in Figure 2.

Limited search. Take The Best operates by limited search with an explicit stopping (discrimination) rule. Its motto is "take the best, ignore the rest." In contrast, "rational" inference, as traditionally conceived, needs to look up all available information. Take Figure 2.

The Best violates this tenet of classical rationality. The stopping rule makes the algorithm fast (search is quickly terminated) and frugal (only a few predictor values are used for the inference).

One-reason decision making. The inference is made by one predictor only; there is no integration and compensation of predictors. Take The Best is non-compensatory. For instance, the positive values of object *b* in Predictors 2 and 3 (Figure 1) cannot reverse the decision made solely on the basis of the higher ranking Predictor 1. In contrast, "rational" inference, as traditionally conceived, integrates all available information in some optimal way. Take The Best violates this maxim. One-reason decision making makes the algorithm computationally simple, if computation is sequential.

Exploitation of a lack of knowledge. Take The Best operates with the *recognition principle: If one of the two alternatives is recognized, and the other not, then choose the recognized object.* Note that this principle is non-compensatory. For instance, the three negative predictor values of object *c* do not reverse the inference that *c* is larger than *d* (Figure 1). The recognition principle can only be used when a person has a lack of knowledge (i.e., does not recognize one of the alternatives) and exploits this lack in environments where recognition is not random but correlated with the criterion. The recognition principle is the most frugal satisficing principle, because it feeds on a lack of knowledge rather than just limited knowledge. Its surprising power can lead to the counterintuitive less-is-more effect, that is, that inferences based on less knowledge can be systemati-

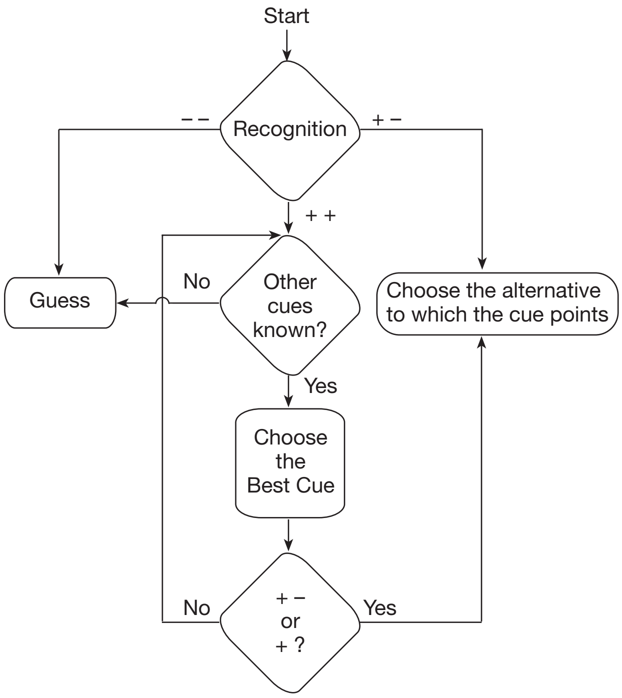

Figure 2. Flow diagram of the Take The Best algorithm (Gigerenzer & Goldstein, 1996).

cally better than inferences based on more knowledge. The structures of environments in which less-is-more effects occur are described in Goldstein and Gigerenzer (1997).

Exploitation of structures of information (environments). The recognition principle can exploit certain structures of information (recognition correlated with the criterion). Similarly, Take The Best can exploit certain structures of information, such as exponentially decreasing weights of binary predictors (see below), which allows high levels of accuracy (Martignon, Hoffrage, & Kriegeskorte, 1997).

A Competition

Although Take The Best seems to reflect what people actually do in many situations under constraints of limited time and knowledge, its simplicity raises the suspicion that it will dismally fail when making inferences about unknown features of real environments. For instance, when Keeney and Raiffa (1993) discussed the lexicographic ordering procedure—a procedure related to Take The Best—they concluded that this procedure "is naively simple" and "will rarely pass a test of 'reasonableness'" (p. 78). How could an inference based on only one predictor compete with one based on an integration of all information available? In order to test how accurate Take The Best is, Daniel Goldstein and I set up a competition between Take The Best and five linear integration algorithms, including multiple regression (Gigerenzer & Goldstein, 1996). The task was to infer which of two cities has the larger population, as described above, for all German cities with more than 100,000 inhabitants (83 cities) with nine ecological predictors. In order to simulate limited knowledge, we created millions of hypothetical subjects, each of whom had a different amount of knowledge, by replacing actual predictor values with unknown values. For each of these subjects, the proportion of correct inferences (whether Heidelberg is really larger than Bonn) in all possible tests (83 x 82/2 pairs of cities) was determined using Take The Best. Similarly, the proportion of correct inferences was determined using each of the five linear integration algorithms. Competitors such as multiple regression computed inferences with the beta weights. The linear algorithms always based each inference on all information (predictor values), whereas Take The Best used, on the average, only less than one third of this information. The counterintuitive result was that Take The Best matched every one of the competing algorithms in accuracy, including multiple regression, and performed better than some (Gigerenzer & Goldstein, 1996). Figure 3 illustrates this result for the special case in which every algorithm performs best, that is, when the simulated persons have complete knowledge of predictor values for each city they recognize. Limited recognition is shown on the x-axis, from 0 to all cities recognized.

Note that the performance exhibits a less-is-more effect. For instance, the simulated person who recognizes all cities and has complete information about all values of 83 cities in 9 predictors (at the very right of Figure 3) would make more accurate inferences if she had less complete information, such as information about the values of only 60 cities. The reason is the power

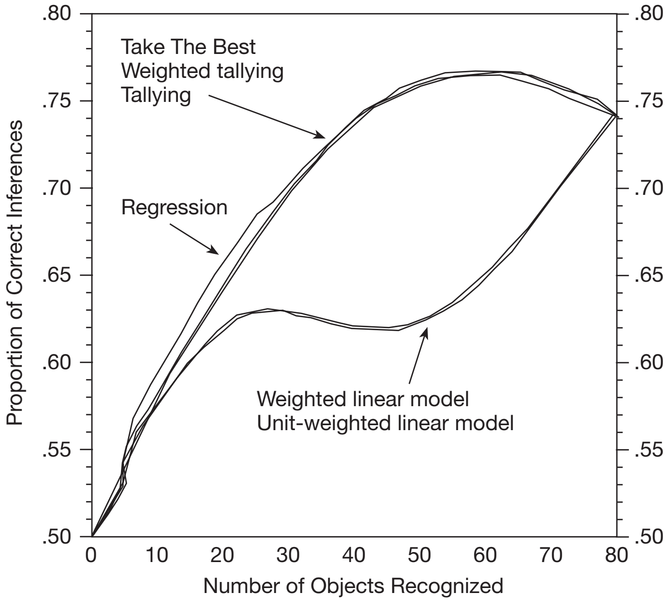

The x-axis shows 84 types of simulated subjects who recognize between 0 and 83 (i.e., all) cities. Tallying counts the number of positive predictor values; weighted tallying weights these values with the validities of each predictor; the unit-weight linear model computes the sum of positive values minus the sum of negative values; the weighted linear model weights these values with the validities of each predictor; and the multiple regression model computes the beta weights. For details of the simulation see Gigerenzer and Goldstein (1996). Copyright 1996 by APA.

Figure 3. Results of the competition between Take The Best and five linear algorithms.

of the recognition principle (which can no longer be applied when all objects are recognized), which is explicit in Take The Best and implicit in some of the linear algorithms (Gigerenzer & Goldstein, 1996).

This result is an existence proof that fast and frugal inference can be as accurate as computationally expensive algorithms that use more knowledge and time. But does this result generalize to other situations, or is there something peculiar with the population demographics of German cities? What is the structure of information in natural environments that Take The Best can exploit, and when would it fail? How do variants of Take The Best that are faster and more frugal perform?

Does the Performance of Take The Best Generalize to Other Environments?

We have simulated the performance of Take The Best in eight task environments, and compared it to the most powerful linear competitor, multiple regression (Czerlinski, Goldstein, & Gigerenzer, 1997). The tasks included predicting the mortality rates in 20 Los Angeles districts from 15 indicators of pollution and demographic information; dropout rates in 57 Chicago high schools based on 18 indicators such as the average salary of the teachers and the proportion of white students; and attractiveness ratings of prominent men and women, based on three cues. These competitions were performed for the case of complete knowledge, that is, where the recognition principle (which exploits a lack of knowledge) could not help Take The Best. In four of the eight environments, the proportion of accurate inferences was the same for multiple regression and Take The Best, in two others multiple regression performed slightly better (1 or 2 percentage points), and in only two environments there was a clear advantage of multiple regression (6 and 9 percentage points). The total proportions of correct inferences ranged between 65% and 84%. Thus, the striking performance of Take The Best did generalize. Equally important, there were systematic differences that provide clues for understanding why and when Take The Best performs so well.

What Structures of Environments Allow Take The Best to Perform So Well?

Martignon et al. (1997) have proven conditions under which Take The Best can and cannot be outperformed by a weighted linear model (with predictor-criterion correlations as weights). I summarize here the gist of their proofs. In environments with *abundant information* (i.e., where the number of cues is very large compared to $\log_2 N$, where N is the number of objects), weighted linear models perform better. Consistent with this proof, the two environments in which multiple regression had a clear edge in performance in the simulations were those where the number of predictors was large relative to the number of objects (such as 15 cues for 20 objects). In environments with *scarce information* (where the number of cues is small relative to the number of objects, defined as less than or equal than $\log_2 N$), Take The Best performs better on average. Finally, when the weights of binary predictors are *exponentially decreasing,* such as 1/2,1/4,1/8, and therefore are non-compensatory, no weighted linear model, including multiple regression, can outperform the faster and more frugal Take The Best.

In many situations humans must make inferences on the basis of scarce information. Environments with strictly non-compensatory cue weights are also not uncommon. For instance, in a study about people's reactions to their experience with police officers and judges, Tyler (1997, Figure 2) reported the beta weights of three predictors (fairness of the procedure, fairness of the outcome, and favorability of the outcome) for each of three criteria: people's respect of the law, their evaluation of the legal authority involved, and their personal feelings following the experience. For the first two criteria, the sets of beta weights were strictly non-compensatory, and for the third, approximately so. Among the eight data sets that Czerlinski et al. (1997) analyzed there were three with strictly non-compensatory weights. Scarce information as well as non-compensatory information is where Take The Best flourishes.

Can Satisficing Inferences Get by With Even Less Knowledge?

Take The Best uses information about the rank order of the validity of the predictors (as opposed to weighted linear models which use information about the quantitative validities of predictors). Assume that this rank order is not known, only the direction into which each of the predictors points (whether a predictor signals a higher or a lower value on the criterion). Two variants of Take The Best operate with this reduced information. They differ from Take The Best only in which predictors they look up first. "Take The Last" tries first the predictor that discriminated the last time; if it does not discriminate, then the predictor that worked the next to last time is examined, and so on. Take The Last works by a well-known psychological principle, the "*Einstellung* effect" (Luchins & Luchins, 1994) of Gestalt psychology. By contrast, the "Minimalist" just tries predictors in random order. Neither of these two algorithms needs information about which predictors are better than others. How accurate are the inferences that these satisficing algorithms draw? For population sizes, Gigerenzer and Goldstein (1996) showed that the accuracy of these two algorithms was, on average, only about 1 percentage point less than that of Take The Best, and still higher than some of the linear models. Each of them stopped earlier than Take The Best, that is, searched for less information. The performance of these two satisficing algorithms was striking.

Take The Best is a member of a larger family, the PMM ("Probabilistic Mental Models") family of satisficing algorithms (Gigerenzer, 1994; Gigerenzer et al., 1991). The closest relatives to Take The Best (but not to Take The Last and the Minimalist, which do not order predictors according to their validity) are lexicographic strategies and the classification and regression tree (CART) models (Breiman et al., 1993). Different from lexicographic strategies, however, Take The Best does not produce systematic intransitive inferences.

3. Summing Up

When a person makes inferences about unknown states of the world under constraints of limited knowledge and time, she is typically not in a position to calculate the optimal solution, even if such a solution is attainable. Take The Best and its variants are fast and frugal algorithms that can draw inferences with a minimum of knowledge and computational effort. These algorithms are based on simple psychologically plausible principles. They violate two classical tenets of rationality: They do not look up all available information and they use one-reason decision making. Nevertheless, Take The Best can be as accurate as weighted linear models, and we can specify the structure of environments in which these satisficing algorithms do well. Models of bounded inference do not necessarily have to forsake accuracy for simplicity, nor rationality for psychological plausibility—the mind can have it both ways.

References

Breiman, L., Friedman, J. H., Olshen, R. A., & Stone, C. J. (1993). *Classification and regression trees.* New York: Chapman and Hall.

Budescu, D. V., & Bruderman, M. (1995). The relationship between the illusion of control and the desirability bias. *Journal of Behavioral Decision Making, 8,* 109–125.

Camerer, C. (1995). Individual decision making. In J. H. Kagel & A. E. Roth (Eds.), *Handbook of experimental economics* (pp. 587–703). Princeton, NJ: Princeton University Press.

Cosmides, L., & Tooby, J. (1996). Are humans good intuitive statisticians after all? Rethinking some conclusions from the literature on judgment under uncertainty. *Cognition, 58,* 1–73.

Czerlinski, J., Goldstein, D., & Gigerenzer, G. (1997). *Information environments and algorithms that exploit them.* Manuscript, Max Planck Institute for Psychological Research, Munich.

Einhorn, H. J., & Hogarth, R. M. (1981). Behavioral decision theory: Processes of judgment and choice. *Annual Review of Psychology, 32,* 53–88.

Fiedler, K. (1988). The dependence of the conjunction fallacy on subtle linguistic factors. *Psychological Research, 50,* 123–129.

Gigerenzer, G. (1991). How to make cognitive illusions disappear: Beyond "heuristics and biases". *European Review of Social Psychology, 2,* 83–115.

Gigerenzer, G. (1993). The bounded rationality of probabilistic mental models. In K. I. Manktelow & D. E. Over (Eds.), *Rationality: Psychological and philosophical perspectives* (pp. 284–313). London: Routledge.

Gigerenzer, G. (1994). Why the distinction between single-event probabilities and frequencies is relevant for psychology (and vice versa). In G. Wright & P. Ayton (Eds.), *Subjective probability* (pp. 129–161). New York: Wiley.

Gigerenzer, G. (1996a). On narrow norms and vague heuristics: A reply to Kahneman and Tversky (1996). *Psychological Review, 103,* 592–596.

Gigerenzer, G. (1996b). The psychology of good judgment: Frequency formats and simple algorithms. *Journal of Medical Decision Making, 16,* 273–280.

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

Gigerenzer, G., & Hoffrage, U. (1995). How to improve Bayesian reasoning without instruction: Frequency formats. *Psychological Review, 102,* 684–704.

Gigerenzer, G., Hoffrage, U., & Kleinbölting, H. (1991). Probabilistic mental models: A Brunswikian theory of confidence. *Psychological Review, 98,* 506–528.

Gigerenzer, G., & Murray, D. J. (1987). *Cognition as intuitive statistics.* Hillsdale, NJ: Erlbaum.

Goldstein, D. G., & Gigerenzer, G. (1997). *Recognition: How to exploit a lack of knowledge.* Manuscript submitted for publication.

Gould, S. J. (1992). *Bully for brontosaurus. Further reflections in natural history.* London: Penguin Books.

Grice, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and semantics 3: Speech acts* (pp. 41–58). New York: Academic Press.

Hertwig, R. (1997). *Judgement under uncertainty: Beyond probabilities.* Manuscript submitted for publication.

Hertwig, R., & Gigerenzer, G. (1996). *The "conjunction fallacy" revisited: How intelligent inferences look like reasoning errors.* Manuscript submitted for publication.

Hoffrage, U., & Gigerenzer, G. (1996). The impact of information representation on Bayesian reasoning. In G. Cottrell (Ed.), *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society* (pp. 126–130). Mahwah, NJ: Erlbaum.

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

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47,* 263–291.

Kahneman, D., & Tversky, A. (1996). On the reality of cognitive illusions. *Psychological Review, 103,* 582–591.

Kanwisher, N. (1989). Cognitive heuristics and American security policy. *Journal of Conflict Resolution, 33,* 652–675.

Keeney, R. L., & Raiffa, H. (1993). *Decisions with multiple objectives.* Cambridge, UK: Cambridge University Press.

Keren, G. (1991). Additional tests of utility theory under unique and repeated conditions. *Journal of Behavioral Decision Making, 4,* 297–304.

Keren, G., & Wagenaar, W. A. (1987). Violation of utility theory in unique and repeated gambles. *Journal of Experimental Psychology: Learning, Memory and Cognition, 13,* 387–391.

Koehler, J. J., Gibbs, B. J., & Hogarth, R. M. (1994). Shattering the illusion of control: Multi-shot versus single-shot gambles. *Journal of Behavioral Decision Making, 7,* 183–192.
Langer, E. J. (1975). The illusion of control. *Journal of Personality and Social Psychology, 32,* 311–328.
Lopes, L. L. (1991). The rhetoric of irrationality. *Theory & Psychology, 1,* 65–82.
Lopes, L. L. (1992). Three misleading assumptions in the customary rhetoric of the bias literature. *Theory & Psychology, 2,* 231–236.
Luchins, A. S., & Luchins, E. H. (1994). The water jar experiments and Einstellung effects: I. Early history and surveys of textbook citations. *Gestalt Theory, 16,* 101–121.
Martignon, L., Hoffrage, U., & Kriegeskorte, N. (1997). *Lexicographic comparison under uncertainty: A satisficing algorithm.* Manuscript, Max Planck Institute for Psychological Research, Munich.
May, R. S. (1987). *Realismus von Subjektiven Wahrscheinlichkeiten.* Frankfurt a.M.: Lang.
Oaksford, M., & Chater, N. (1992). Bounded rationality in taking risks and drawing inferences. *Theory & Psychology, 2,* 225–230.
Shanteau, J. (1989). Cognitive heuristics and biases in behavioral auditing: Review, comments and observations. *Accounting Organizations and Society, 14,* 165–177.
Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics, 69,* 99–118.
Simon, H. A. (1979). *Models of thought.* New Haven, CT: Yale University Press.
Simon, H. A. (1982). *Models of bounded rationality* (2 vols.). Cambridge, MA: MIT Press.
Simon, H. A. (1990). Invariants of human behavior. *Annual Review of Psychology, 41,* 1–19.
Simon, H. A. (1991). Cognitive architectures and rational analysis: Comment. In K. Vanlehn (Ed.), *Architectures for intelligence* (pp. 25–39). Hillsdale, NJ: Erlbaum.
Simon, H. A. (1992). *Economics, bounded rationality, and the cognitive revolution.* Aldershot Hants, UK: Elgar.
Sniezek, J. A., & Buckley, T. (1993). Decision errors made by individuals and groups. In N. J. Castellan (Ed.), *Individual and group decision making.* Hillsdale, NJ: Erlbaum.
Stich, S. P. (1985). Could man be an irrational animal? *Synthese, 64,* 115–135.
Teigen, K. H. (1974). Overestimation of subjective probabilities. *Scandinavian Journal of Psychology, 15,* 56–62.
Thaler, R. H. (1991). *Quasi rational economics.* New York: Sage.
Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review, 90,* 293–315.
Tyler, T. R. (1997). Procedural fairness and compliance with the law. *Schweizerische Zeitschrift für Volkswirtschaft und Statistik, 133* (2/2).
Wallsten, T. S. (1983). The theoretical status of judgmental heuristics. In R. W. Scholz (Ed.), *Decision making under uncertainty* (pp. 21–39). Amsterdam: Elsevier.
Wedell, D. H., & Böckenholt, U. (1990). Moderation of preference reversals in the long run. *Journal of Experimental Psychology: Human Perception and Performance, 16,* 429–438.

Summary

I specify general criteria for models of bounded rationality and discuss specific models for satisficing inference. The task of these fast and frugal algorithms is to infer unknown features of their environment under the constraints of limited knowledge, limited time, and limited computational capacities. These algorithms violate fundamental tenets of classical rationality: They neither look up nor integrate all information. I review the performance of the satisficing "Take The Best" algorithm. Despite its frugality, Take The Best can make as many correct inferences as computationally expensive weighted linear models that use and combine all available information. Accurate inferences need not follow the dictates of classical rationality.

Zusammenfassung

Ich formuliere allgemeine Kriterien für Modelle begrenzter Rationalität und diskutiere spezifische "satisficing" Modelle für Inferenz unter Unsicherheit. Die Aufgabe dieser schnellen und

einfachen Algorithmen ist, unbekannte Eigenschaften der Umwelt zu erschließen, und zwar mit begrenztem Wissen, begrenzter Zeit und begrenzter rechnerischer Kapazität. Diese Algorithmen verletzen fundamentale Annahmen klassischer Rationalität: Sie suchen weder alle verfügbare Information, noch integrieren sie Information. Ich berichte über die Leistung des "Take The Best" Algorithmus. Trotz seiner Frugalität kann "Take The Best" genauso viele richtige Inferenzen machen wie rechnerisch aufwendige gewichtete lineare Modelle, welche alle verfügbare Information verwenden und kombinieren. Richtige Inferenzen müssen nicht den Regeln klassischer Rationalität folgen.