

Striking a Blow for Sanity in Theories of Rationality

Gerd Gigerenzer

I took the title of this chapter from an email Herbert A. Simon sent me in May 1999. In this email, he wrote a statement for the back cover of *Simple Heuristics That Make Us Smart* in which he commented: “I think the book strikes a great blow for sanity in the approach to rationality [and shows] why more than minor tampering with existing optimization theory is called for.” But Herb wouldn’t be Herb if he hadn’t added “and you wouldn’t believe I had ever skimmed the volume if I didn’t find SOMETHING to disagree with.” And so he continued, pointing out that he hadn’t found the expert/novice topic treated, that scientific discovery would have been a great example for ill-structured domains ...

Bringing sanity into theories of rationality was a major guideline in Herbert Simon’s scientific life. However, as he himself was prepared to admit, sanity in rationality entered his thinking as a negatively defined concept, a kind of black box that contained everything that was not optimization. What he opposed has various names: full rationality, substantial rationality, maximization of expected utility, *Homo economicus*, or simply optimization. What he proposed had its seeds in his revised dissertation, *Administrative Behavior* (1947), and eventually became termed bounded rationality, satisficing, or procedural rationality. Because of its initial vague definition, bounded rationality, however, came to mean many things to many people.

Bounded Rationality in the Plural

Simon’s (1955, 1956) concept of bounded rationality has been claimed by three different programs. One of these Simon opposed, one he tolerated, and one he embraced. I will call the three programs optimization under constraints, cognitive illusions and anomalies, and ecological rationality, although I am not sure that Herb always wanted to distinguish the latter two programs the way I do.

What Bounded Rationality is Not: Optimization under Constraints

In models of full rationality, all relevant information is assumed to be available to *Homo economicus* at no cost. This classical version of *Homo economicus* has a distinctive Christian flavor: He is created in the image of an omniscient God. Real humans, however, need to search for information first. In an attempt to render economic theory more realistic, Stigler (1961) introduced constraints on full rationality, such as costs of information search. The idea of optimization under constraints is to propose one or a few constraints (too many would make the mathematics too hard or even intractable) while retaining the ideal of optimization. In this view, for instance, a person who wants to buy a used car of a certain brand stops search when the costs of further

search—direct costs and opportunity costs—exceed those of its benefits. Introducing real constraints does make the approach more realistic, but maintaining the ideal of optimization, that is, calculating an optimal stopping point, does not. Such an ideal of optimization invokes new kinds of omniscience, that is, foreseeing the benefits and costs of further information search (Conlisk, 1996). This is little evidence that humans make decisions this way.

Lack of psychological reality was an objection Herb made time and again. The argument against his and others' concern with omniscience and psychological evidence has been the "as if" conjecture: The question is not whether people actually optimize, with or without constraints, but whether they *act as if* they were doing so. As long as optimization predicts behavior, one need not be concerned with the actual motivations, emotions, and reasoning of people (Friedman, 1953). In this view, the bounds in bounded rationality are just another name for constraints, and bounded rationality is merely a case of optimizing under constraints. Despite Herb's vehement protests, this message has become the doctrine. But this doctrine comes at a price: Retaining the ideal of optimization can make models of optimization under constraints more demanding than models of full rationality, both mathematically and psychologically. In the words of Thomas Sargent, a proponent of the view that bounded rationality means optimization under constraints: "Ironically, when we economists make the people in our models more 'bounded' in their rationality ... *we* must be smarter, because our models become larger and more demanding mathematically and econometrically." (Sargent, 1993, p. 2) In optimization under constraints, agents are recreated in the image of econometricians, one step above the gods. In personal conversation, Herb once remarked with a mixture of humor and anger that he had considered suing those authors who misuse his term of bounded rationality to construct ever more psychologically unrealistic models of human decision making.

Optimization, with or without constraints, has also spread beyond economics. Psychologists often propose models of cognition that assume almost unlimited memory, storage capacities, and computational power. That is, many psychologists also build "as if" models of behavior. Over lunch, I once asked Herb about his impact on psychology and got one of his straightforward responses:

GG: Do you think you had much effect on psychologists with "bounded rationality"?

HS: Yes. There is an abundant literature on recognition and search, for instance, in Newell and Simon's *Human Problem Solving*.

GG: But what about exemplar models of categorization and the many other Laplacean demon models of cognition?

HS: Oh, these are of no relevance.

Why should we listen to Herb rather than building "as if" models of optimization under constraints? Is there a problem with the program of making right predictions from wrong assumptions?

Optimization is impossible in most natural situations. The ideals of "as if" optimization is obviously limited because, in most natural situations, optimization is computationally intractable in any implementation, whether machine or neural (Michalewicz & Fogel, 2000). In computer science, these situations are called NP-complete; that is, the solution cannot be computed in polynomial time. For instance, no mind or computer can apply Bayes's rule to a large number of variables that are mutually dependent, because the number of computations increases exponentially with the number of variables. In such situations, a fully rational Bayesian cannot exist. Even for games with simple and well-defined rules, such as chess and Go, we do not know the optimal strategy.

Nevertheless, we know what a good outcome is. In these situations, “as if” optimization can only be achieved when the real situation is changed and simplified in a mathematically convenient way so that optimization is possible. Thus, the choice is between finding a good heuristic solution for a game where no optimal one is known, and finding an optimal solution for a game with modified rules. That may mean abandoning our study of chess in favor of tic-tac-toe.

Optimization is unfeasible when problems are unfamiliar and time is scarce. In situations where optimization is in principle possible (unlike those under the first point), a practical issue remains. Selten (2001) distinguishes between familiar and unfamiliar problems. In the case of a familiar problem, the decision maker knows the optimal solution. This may be due to prior training or because the problem is simple enough. In the case of an unfamiliar problem, however, the decision maker cannot simply execute a known method that leads to the best result because that method must first be discovered. In other words, the agent has to solve two tasks: level 1, executing a method that leads to a solution, and level 2, finding this method. Thus, two questions arise. What is the optimal method to be chosen? And what is the optimal approach to discovering that method? (There may be an infinite regress: level 3, finding a method for level 2, and so on.) At each level, time must be spent. Although Selten’s argument concerning unfamiliar problems has not yet been cast into mathematical form, as the previous issue of combinatorial explosion has been, it strongly suggests that an optimizing approach to unfamiliar problems is rarely feasible when decision time is scarce.

Optimization does not imply an optimal outcome. Some economists, biologists, and cognitive scientists seem to believe that a theory of bounded rationality must rely on optimization in order to promise optimal decisions. No optimization, no good decision. But this does not follow. Optimization needs to be distinguished from an optimal outcome. Note that the term optimization refers to a mathematical process—computing the maximum or minimum of a function—which does *not* guarantee optimal outcomes in the real world. The reason is that one has to make assumptions about the world in order to be able to optimize. These assumptions are typically selected by mathematical convenience, based on simplifications, and rarely grounded in psychological reality. If they are wrong, one has built the optimization castle on sand, and optimization will not necessarily lead to optimal results. This is one reason why models of bounded rationality that do not involve optimization often can make predictions at least as good as those made by models that do involve optimization (Gigerenzer & Selten, 2001a; March, 1978; Selten, 1998).

A good fit, per se, is not an empirical validation of the model. Friedman’s (1953) argument in favor of “as if” models was this: What counts is not descriptive accuracy, that is, the psychological validity of the axioms and assumptions, but rather the accuracy of the predictions a model makes. Despite Friedman’s introductory example of the law of falling bodies, this explicit disinterest in a proper description of the underlying mechanisms would be unusual in physics, molecular biology, or genetics. (This does not mean that “as if” models are never used; optimal foraging models in animal biology are an example.) The point I want to make here is that one needs to be careful in distinguishing between two kinds of statistical tests that have both been labeled “predictions.” One is *data fitting*, that is, “explanations” of existing data; the other is *ex ante prediction*, that is, predictions of *new* observations.

In cognitive science and economics, the validity of a model is often reported in terms of its fit with *given* observations, such as what proportion of the variance a model explains. However, the belief that a good fit between model and data would provide empirical evidence for the

model is unfounded if the model has numerous free parameters (Roberts & Pashler, 2000). For instance, introducing more and more relevant constraints into models of optimization increases the number of adjustable parameters, which can make the resulting model too “powerful” to allow for falsification by empirical data. In these situations, a model can fit almost all possible data, including data produced by two logically inconsistent theories. Here, a good fit is a mathematical truism, not an empirical result. Utility maximization models often have many adjustable parameters—such as the utilities and utility functions in each particular case (Arrow, 1986; Simon, 1986). This problem of “overfitting” is not specific to optimization models, but rather occurs in any statistical model that has a relatively large number of adjustable parameters, including neural networks (Geman, Bienenstock, & Doursat, 1992). If smart enough, one can always find parameters so that the model fits a given situation. The problem of overfitting becomes particularly stringent in the “as if” philosophy because the only empirical test for a model concerns its predictive power. Models of bounded rationality such as fast and frugal heuristics (Gigerenzer & Selten, 2001a) that dispense with optimization and also, for the most part, with utilities and probabilities, reduce this validation problem in two ways. First, they model the underlying mechanisms of choice and inference, and thereby provide a second source for testing their validity (process and outcome). Second, they are simple and robust so that their predictions show consistently less overfitting than optimizing models (Gigerenzer, Todd, & ABC Research Group, 1999).

These four issues highlight some limits of “as if” optimization. There are other well-known problems, such as the “infinite regress” problem of determining how much information to gather in order to determine the cost of information. These issues indicate that optimization under constraints is not the last word, despite its mathematical beauty. Bounded rationality needs a different intellectual home. But which?

What Bounded Rationality is Not: Cognitive Illusions and Anomalies

Optimization with decision costs taken into account is one misreading of Herb’s concept of bounded rationality. It is not the only one. Many psychologists and some economists assume that the study of bounded rationality is the study of cognitive limitations and systematic errors in judgment and decision making (e.g., Camerer, 1998; Rabin, 1998). Surprisingly, this second meaning amounts to something like the converse of the first. The cognitive illusions program aims at demonstrating that people’s judgments and decisions do not follow the predictions of “as if” optimization.

For instance, in his article “Bounded rationality in individual decision making,” Camerer (1998, p. 179) reviews anomalies in decisions and errors in judgments and calls this the “exploration of procedural (bounded) rationality of individuals.” Kaufman (1999, p. 141) gives the example of a gay man who practiced unsafe sex with multiple partners and “is now HIV positive and admits to his bounded rationality.” This view that the study of bounded rationality is the study of systematic errors in judgment and decision making has spread from psychology into economics and law, shaping new research areas such as behavioral economics (e.g., Camerer, 1995) and behavioral law and economics (e.g., Jolis, Sunstein, & Thaler, 1998). In Camerer’s words, “the goal is to test whether normative rules are *systematically* violated and to propose alternative theories to explain any observed violations” (p. 588). The products of this research are well known: a list of cognitive biases such as base rate neglect, overconfidence bias, and the sunk-cost effect.

This program assumes that the meaning of bounded rationality is that humans have cognitive limitations, which express themselves in errors in judgment and decision making; therefore, the study of errors is the study of bounded rationality. Compared to optimization under constraints, this second interpretation of bounded rationality is a relatively new one. The origin of this interpretation seems to be a mentioning of Simon's work on bounded rationality in the preface of Kahneman, Slovic, and Tversky's 1982 anthology. Since there are no citations at all to Simon in the early influential papers of Kahneman and Tversky, which are reprinted in this anthology, this mentioning was probably more an acknowledgement to a distinguished figure than an intellectual debt (Lopes, 1992). Nevertheless, the notion that bounded rationality is the study of cognitive illusions has since become widespread.

Herb applauded the demonstrations of systematic deviations from expected utility by Kahneman, Tversky, and others. But what did he think when the followers of Kahneman and Tversky labeled these demonstrations the study of "bounded rationality?" I asked him once, and his response was "That's rhetoric. But Kahneman and Tversky have decisively disproved economists' rationality model." Herb was surprised to hear that I held their notion of cognitive illusions and biases to be inconsistent with his concept of bounded rationality. I think he liked their results so much that he tended to overlook that these experimenters accepted as normative the very optimization theories that Herb so fought against, at least when the results were interpreted as cognitive illusions. A true theory of bounded rationality does not rely on optimization theories, neither as descriptions nor as norms of behavior. (I gave reasons for the normative limits in the previous section.) We once discussed this issue on a walk through the beautiful Carnegie Mellon campus in the spring of 1997. A systematic deviation from an "insane" standard should not automatically be called a judgmental error, should it? "I hadn't thought about it in this way," Herb replied.

Why is bounded rationality not the same as irrationality? Herb has given the answer in the form of an analogy. Bounded rationality is like a pair of scissors: The mind is one blade, and the structure of the environment is the other. To understand behavior, one has to look at both, at how they fit. In other words, to evaluate cognitive strategies as rational or irrational, one also needs to analyze the environment, because a strategy is rational or irrational only with respect to an environment, physical or social (Simon, 1990). The study of cognitive illusions and errors, however, studies only the cognitive blade, and compares it with laws of probability rather than with the structure of the environment. One blade alone does not cut well.

As a consequence, apparently stable cognitive illusions can be made to disappear and reappear by varying crucial structures of the environment. For instance, Gigerenzer, Hoffrage, and Kleinbölting (1991) theoretically derived and experimentally demonstrated that two apparently stable cognitive illusions, the overconfidence bias and the hard-easy effect, disappear when the underlying questions are randomly sampled from an environment rather than systematically selected. Juslin, Winman, and Olsson (2000) confirmed this initial demonstration in a quantitative analysis of all 130 extant studies. Other research has pointed to the ecological reasons for people's difficulty in following Bayes's rule when reasoning with probabilities. Bayesian reasoning can be strongly improved when the information is not presented in probabilities but rather in natural frequencies, which correspond to the environmental input that humans have received during most of their evolution (Gigerenzer & Hoffrage, 1995, 1999). Meanwhile, this method has been proven to help physicians and patients, judges and law students alike to understand the uncertainties in HIV tests, DNA fingerprinting, and other technologies (e.g., Gigerenzer, 2002; Hoffrage & Gigerenzer, 1998; Hoffrage, Lindsay, Hertwig, & Gigerenzer, 2000; Koehler, 1996). An analysis of environmental structures—such as sampling type, sample size, representation,

degree of predictability, and social rules—helps to understand how the cognitive blade and the environmental blade work together (Fiedler, 2000; Gigerenzer, 2000).

The confusion between bounded rationality and the study of cognitive illusions and irrationality is the second misunderstanding of Herb's idea of sanity in theories of rationality. Bounded rationality is not merely an attack on the assumptions underlying optimizing models; it is a positive program to replace optimization—when it is unrealistic—with something better.

What Bounded Rationality is: Ecological Rationality (the Adaptive Toolbox)

Let me illustrate a fast and frugal heuristic from the adaptive toolbox with an example from sports (Gigerenzer & Selten, 2001b). In cricket, baseball, and soccer, players need to catch balls that come in high. A company wants to design a robot that can catch the ball. This is a thought experiment; no such robot exists as yet. For the sake of simplicity, we consider only the case where a ball comes in high, behind or in front of a player, but not to his left or right.

One team of engineers, which I call the optimizing team, proceeds by programming the family of parabolas into the robot's brain (in theory, balls fly in parabolas). To select the proper parabola, the robot needs to be equipped with instruments that can measure the distance from where the ball was thrown or shot, as well as its initial velocity and projection angle. Yet in a real game, due to air resistance and wind, balls do not fly in parabolas. Thus, the robot would need further instruments that can measure the speed and direction of the wind at each point of the ball's flight in order to compute the resulting path. In addition, there are further factors, such as spin, that affect the flight of the ball. The optimizing team eventually succeeds in producing a lengthy equation that specifies the trajectory of the flight and the spot where the ball will land, given all these measurements. Note that this equation is an "as if" model—the team is not concerned with the actual mechanisms that real players or robots use—and, consequently, the equation does not inform us how to actually build a robot. The optimizing team responds that their task is not "robot psychology," that is, to understand how a robot actually does, or could do, the job. Their claim is that the model will predict the point to which real players and robots would run to catch the ball.

A subgroup within the optimizing team objects that there is no time for the robot to make the proposed measurements and computations, given that the ball is only in the air for a few seconds. A fully rational robot would just sit on the field, measuring and calculating, and thus missing every ball. The more precise the measurements are, the longer they take, and the less time the robot has left to run to the spot where the ball is supposed to land. The real constraint, they argue, is not money but time. Instead of trying to model an omniscient robot, the team proposes to build one that optimizes under constraints. After some deliberation, the subgroup puts forward a number of constraints concerning the robot's ability for information search. Finally, the members of this subgroup design a sophisticated formula that optimizes the outcome under the given constraints, a mathematical masterpiece. The hard-core members of the optimizing team, however, object that this formula is even more complex than the first one, and that the robot will sit even longer on the field measuring and calculating. So what's the point?

A second team enters the field and argues that you need to understand the cognitive processes in order to find out what players do in order to eventually design a robot. One should experimentally study real players, and create situations in which they systematically demonstrate judgmental errors, that is, deviations from the optimizing model. These errors will be the window for the underlying cognitive processes. This team calls itself the "cognitive illusions team." After

a phase of trial and error, the team has found a task in which the players show a bias. A player is positioned on a fixed point in the field, a ball is shot in high, and the player is asked to predict how many yards in front of or behind his position the ball will hit the ground. The surprising result is that the players don't predict very well: They consistently underestimate the distance between their position and the point where the ball will land. This systematic error is called the "optimistic bias" in baseball, cricket, and soccer—because underestimating the distance suggests to players that they might actually get the ball even when they can't. A debate opens on whether this judgmental error could be adaptive, because not trying to run for a ball that actually could have been reached is a more costly error in a real game than trying without success. The cognitive illusions team claims that the model of the optimizing team has been descriptively disproved; actual players show systematic errors whereas the optimizing model doesn't predict any. The optimization team responds that they will nevertheless maintain their model, because a model that can at least approximately predict the data is better than no model. After all, they argue, the notion of an "optimistic bias" is only a redescription of the data; the bias team hasn't put forward any alternative model of the underlying cognitive mechanisms, nor of how to build the robot.

A third team is called in. This team agrees with the second team that humans may not be able to compute the point where the ball will land. However, they argue that the negative goal of disproving the optimization team's predictions does not directly lead to positive models of the underlying cognitive processes. For instance, the notion of an "optimistic bias" does not describe how a player actually catches a ball, but only how his judgment deviates from the actual landing point. The third team proposes to unpack what they call the adaptive toolbox of human behavior and cognition, that is, the smart cognitive and social heuristics as well as the building blocks from which the heuristics are constructed. What do players actually do in order to catch a ball, given that they do not seem to perform the measurements and calculations that the optimization team proposes? Are there fast and frugal heuristics that players use? Experiments show that there are. One of these is the "gaze heuristic." When a ball is high in the air, an experienced player visually fixates the ball and starts running. The heuristic is to adjust the running speed so that the angle of gaze—the angle between the eye and the ball—remains constant (or within a certain range—see McLeod & Dienes, 1996). Assuming the player can run fast enough, this gaze heuristic will take him to the spot where the ball will land. The gaze heuristic is fast and frugal: It pays attention to only one cue, the angle of gaze. It does not attempt to acquire the relevant information concerning initial distance, velocity, wind, spin, or other causal variables, nor does it try to calculate the resulting path given this information. All the relevant information is contained in one variable: the angle of gaze. I call this *one-reason decision making*. The gaze heuristic creates an environment that the player can easily exploit. In place of the complicated trajectory of the ball's flight—which the optimizing team has worked out—the gaze heuristic creates a linear relation between the player's eye and the ball. The rationale behind the gaze heuristic is an instance of ecological rationality.

The optimizing team responds that it may be very interesting to know how actual players do the job, but that it is not really relevant. The successful player will run exactly to the same point that can be calculated from our equations, the team maintains, and so the player acts "as if" he were optimizing. Not exactly, the adaptive toolbox team replies, for there are two advantages to realistic process models. First, the omniscient, optimizing player exists only "as if" and therefore does not lead to instructions on how to build a robot, or how to teach human beginners. The information and computation necessary will likely lead to computational explosion, that is, the model is impossible to implement in any hardware, human or computer. In contrast, the gaze heuristic can be taught to inexperienced players, and we may eventually succeed in building a

robot that uses it. Second, with a good model of the heuristic a person uses, one can make much more precise predictions than with an “as if” model, including behaviors that an “as if” model cannot foresee. For instance, the gaze heuristic predicts that the player will catch the ball *while running*. This is an experimentally testable prediction, and in fact, players do not move to the spot where they think the ball will land and then wait there; rather, they catch the ball while running. The reason is in the environment: The trajectory of the ball is not linear, and thus the player has to move in order to keep the angle of gaze constant. Optimization models would not predict this behavior. Similarly, knowing the heuristics players use helps to predict what players cannot do. Remember that the cognitive illusions team had shown that even experienced players cannot correctly predict where the ball will land. Knowing the heuristic, we understand why this is in fact not necessary for successfully catching a ball. The gaze heuristic succeeds without this ability. A player using the gaze heuristic does not calculate where the ball will land; the heuristic takes him where the ball lands. Thus, what looks like a serious judgmental bias in need of debiasing turns out to be irrelevant for good ball catching.

Homo Heuristicus

This thought experiment illustrates the program of ecological rationality: to study (i) the heuristics people actually use to solve a class of tasks, (ii) the structure of the task environment, and (iii) what environmental structure a heuristic can exploit. The corresponding methodologies used to investigate these issues are experimental research, analytical proofs, and computer simulations. The aim of such research is to establish a “periodic system” of heuristics and their building blocks, as well as a conceptual language to describe the structures of relevant real-world environments. This program develops Simon’s ideas: to study the rational principles that underlie the behavior of real people, who do not optimize and, for the most part, do not calculate utilities and probabilities. Recall that this program differs from the optimizing program in that it analyzes the actual process—the heuristics—rather than constructing “as if” models based on a convenient mathematical structure. Unlike the cognitive illusions program, it directly analyzes the decision process rather than trying to demonstrate violations of the assumptions underlying “as if” models. I have used a heuristic from sports, because sports are a “neutral” topic for most economists and cognitive scientists and are unlikely to invoke strong emotional a priori beliefs. Most of the heuristics studied, however, deal with consumer choice, treatment allocation, risk estimation, social games, and other forms of judgment under limited time and knowledge. The various heuristics in the adaptive toolbox consist of a small number of building blocks, including rules for information search, stopping, and decision. The program is described in detail in Gigerenzer et al. (1999) and in Gigerenzer and Selten (2001a).

In *Bounded Rationality: The Adaptive Toolbox*, Reinhard Selten and I start out with this goal: “to promote bounded rationality as the key to understanding how actual people make decisions without utilities and probabilities.” The adaptive toolbox signifies a radical departure from the classical “repair program” of adjusting theories of rational behavior, where one variable, such as regret, is added to the expected utility calculus, or where one tinkers with the functions for probabilities or utilities, as in prospect theory. We start, in contrast, from the empirically rooted knowledge about the human mind and its capabilities (Todd & Gigerenzer, 2000). Quantitative probabilities, utilities, and optimization appear to play little role in the actual capabilities of the human mind, whereas fast and frugal processes, such as name recognition, aspiration levels, imitation learning, sequential search, stopping rules, and one-reason decision making, do. The

models of heuristics typically have zero adjustable parameters, which makes it easier to empirically test and falsify them. In statistical terms, heuristics err on the side of “bias” rather than “variance” (Geman et al., 1992). This work on the adaptive toolbox and on ecological rationality will, I hope, provide a positive alternative to the investigation of rational choice: the study of how *Homo heuristicus* makes decisions in an uncertain world.

The question of the rationality of *Homo heuristicus* concerns the question of ecological rationality. A heuristic is not good or bad, rational or irrational, in itself, but only relative to an environment, just as adaptations are context bound. Heuristics can exploit regularities in the environment—this is the meaning of ecological rationality. In the ball example, the regularity is that a constant angle of gaze will cause a collision between the ball and the player, or between any two moving objects. Just like human organs, heuristics are domain-specific, that is, designed for a class of problems, rather than for general strategies. For instance, the gaze heuristic is useful in situations where one wants to generate a collision between two moving objects, as described before, but also for avoiding collisions. If you learn to fly a plane, for instance, you will be taught a version of the gaze heuristic: When another plane is approaching, look at a scratch in your windshield and see whether the other plane moves relative to that scratch. If it does not, dive away quickly.

The study of heuristics illustrates that, contrary to conventional wisdom, limitations of knowledge, memory, and computational capability need *not* be a disadvantage. The gaze heuristic, for instance, ignores all the causally relevant information; it cannot predict where the ball will land but solves the problem of catching the ball anyway. More information is not always better. Goldstein and Gigerenzer (2002) specify the conditions under which intermediate levels of knowledge lead to systematically more correct predictions than higher levels of knowledge—the “less is more” effect. Computer simulations indicate that memory limitations actually enable a child to learn its first language, while a fully developed memory would in fact prevent language learning (Elman, 1993). Relying on only one good reason can lead to better predictions of demographic and economic variables—such as homeless and school drop-out rates—than can regression analysis with many variables (Czerlinski, Gigerenzer, & Goldstein, 1999). The working memory limitation of “7 plus minus 2” seems to improve the detection of covariances in the environment (Kareev, 2000). Investment portfolios based on pedestrians’ brand name recognition have outperformed the market and major mutual funds (Borges, Goldstein, Ortmann, & Gigerenzer, 1999). For further examples of situations in which limited knowledge or cognitive capacities can speed up learning and promote successful problem solving, see Bjorklund and Green (1992) and Todd (2001). This is not to say that heuristics are foolproof or that limitations are always good, for again the interesting question concerns ecological rationality: to specify the tasks or environmental structures that a heuristic can exploit and those where it will fail.

The term “adaptive toolbox” is not Herb’s, although it is in his spirit. The rationality of the adaptive toolbox is not logical, but ecological. It refers to the match between a heuristic and the structure of an environment, which is the essence of Simon’s analogy of a pair of scissors. For instance, one can specify a class of environmental structures and mathematically prove that a heuristic that simply relies on the best reason while ignoring the rest—such as Take The Best (Goldstein & Gigerenzer, 1999)—is at least as accurate as any linear model with any number of predictors (Martignon & Hoffrage, 1999). In these circumstances, heuristics are not just faster and more frugal than optimization models, but they are also at least as accurate. (Here I mean accuracy in prediction, not in fitting.) When heuristics can exploit the structure of environments, they can avoid a tradeoff between accuracy and effort. Whereas in Stigler’s classical example the used-car buyer typically does not get the best buy, because information search costs limit his at-

tempt at an exhaustive search, there exist situations in which a smart heuristic can solve a task perfectly, as illustrated by the gaze heuristic. Studying ecological rationality—the match between heuristics and environments—is important for freeing the concept of heuristics from the flavor of always being the second-best solution. The issue can be posed with a different twist. If one has a good model of the mechanism, what is the additional value of an “as if” model? “As if” may well turn out to be the second-best solution.

Epilogue

Herb is no longer among us. But his spirit is. His struggle with the concept of bounded rationality will stay with us, and I believe that if he could see how it is being developed, he would be enthusiastic. Let me end with his own words:

Dear Gerd,

I have never thought of either bounded rationality or satisficing as precisely defined technical terms, but rather as signals to economists that they needed to pay attention to reality, and a suggestion of some ways in which they might. But I do agree that I have used bounded rationality as the generic term, to refer to all of the limits that make a human being's problem spaces something quite different from the corresponding task environments: knowledge limits, computational limits, incomparability of component goals. I have used satisficing to refer to choice of “good enough” alternatives (perhaps defined by an aspiration level mechanism) or “best-so-far” alternatives to terminate selective search among alternatives—the latter usually not being given in advance, but generated sequentially. So one might apply “satisficing” to the “good-enough criterion” or to any heuristic search that uses such a criterion to make its choice.

Final remark on this point, going from most general to most specific, we have bounded rationality, then heuristic search, then satisficing. Further, on the same level as heuristic search, we have a second class of methods, very important in the theory of expertise: problem solution by recognition. Currently, that is my taxonomy of decision and problem solution methods. You can decide better than I can where you want to place fast-and-frugal in this rough classification. I would tend to regard it as a class of heuristics, hence on the same level as satisficing.

I guess a major reason for my using somewhat vague terms—like bounded rationality—is that I did not want to give the impression that I thought I had “solved” the problem of creating an empirically grounded theory of economic phenomena. What I was trying to do was to call attention to the need for such a theory—and the accompanying body of empirical work to establish it—and to provide some examples of a few mechanisms that might appear in it, which already had some evidential base. There still lies before us an enormous job of studying the actual decision making processes that take place in corporations and other economic settings ...

End of sermon—which you and Reinhard [Selten] don't need. I am preaching to believers.

Cordially,
Herb

References

- Arrow, K. J. (1986). Rationality of self and others in an economic system. *Journal of Business*, 59, 385–399.
- Bjorklund, D. F., & Green, B. L. (1992). The adaptive nature of cognitive immaturity. *American Psychologist*, 47, 46–54.
- Borges, B., Goldstein, D. G., Ortmann, A., & Gigerenzer, G. (1999). Can ignorance beat the stock market? In G. Gigerenzer, P. M. Todd, & the ABC Research Group, *Simple heuristics that make us smart* (pp. 59–72). New York: Oxford University Press.
- Camerer, C. (1995). Individual decision making. In J. H. Kagel & A. E. Roth (Eds.), *The handbook of experimental economics* (pp. 587–703). Princeton, NJ: Princeton University Press.
- Camerer, C. (1998). Bounded rationality in individual decision making. *Experimental Economics*, 1, 16–18.
- Conlisk, J. (1996). Why bounded rationality? *Journal of Economic Literature*, 34, 669–700.
- Czerlinski, J., Gigerenzer, G., & Goldstein, D. G. (1999). How good are simple heuristics? In G. Gigerenzer, P. M. Todd, & the ABC Research Group, *Simple heuristics that make us smart* (pp. 97–118). New York: Oxford University Press.
- Elman, J. (1993). Learning and development in neural networks: The importance of starting small. *Cognition*, 48, 71–99.
- Fiedler, K. (2000). Beware of samples! A cognitive-ecological sampling approach to judgmental biases. *Psychological Review*, 107, 659–676.
- Friedman, M. (1953). *Essays in positive economics*. Chicago: University of Chicago Press.
- Geman, S. E., Bienenstock, E., & Doursat, R. (1992). Neural networks and the bias/variance dilemma. *Neural Computation*, 4, 1–58.
- Gigerenzer, G. (2000). *Adaptive thinking: Rationality in the real world*. New York: Oxford University Press.
- Gigerenzer, G. (2002). *Calculated risks: How to know when numbers deceive you*. New York: Simon & Schuster (UK edition: *Reckoning with risk: Learning to live with uncertainty*. London: Penguin Books).
- Gigerenzer, G., & Hoffrage, U. (1995). How to improve Bayesian reasoning without instruction: Frequency formats. *Psychological Review*, 102, 684–704.
- Gigerenzer, G., & Hoffrage, U. (1999). Overcoming difficulties in Bayesian reasoning: A reply to Lewis & Keren and Meilers & McGraw. *Psychological Review*, 106, 425–430.
- Gigerenzer, G., Hoffrage, U., & Kleinbölting, H. (1991). Probabilistic mental models: A Brunswikian theory of confidence. *Psychological Review*, 98, 506–528.
- Gigerenzer, G., & Selten, R. (Eds.) (2001a). *Bounded rationality: The adaptive toolbox*. Cambridge, MA: MIT Press.
- Gigerenzer, G., & Selten, R. (2001b). Rethinking rationality. In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality: The adaptive toolbox* (pp. 1–12). Cambridge, MA: MIT Press.
- Gigerenzer, G., Todd, P. M., & the ABC Research Group. (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.
- Goldstein, D. G., & Gigerenzer, G. (1999). The recognition heuristic: How ignorance makes us smart. In G. Gigerenzer, P. M. Todd, & the ABC Research Group, *Simple heuristics that make us smart*. (pp. 37–58). New York: Oxford University Press.
- Goldstein, D. G., & Gigerenzer, G. (2002). Models of ecological rationality: The recognition heuristics. *Psychological Review*, 109, 75–90.
- Hoffrage, U., & Gigerenzer, G. (1998). Using natural frequencies to improve diagnostic inferences. *Academic Medicine*, 73, 538–540.
- Hoffrage, U., Lindsay, S., Hertwig, R., & Gigerenzer, G. (2000). Communicating statistical information. *Science*, 290, 2261–2262.
- Jolis, C., Sunstein, C. R., & Thaler, R. (1998). A behavioral approach to law and economics. *Stanford Law Review*, 50, 1471–1550.
- Juslin, P., Winman, A., & Olsson, H. (2000). Naive empiricism and dogmatism in confidence research: A critical examination of the hard-easy effect. *Psychological Review*, 107, 384–396.
- Kahneman, D., Slovic, P., & Tversky, A. (Eds.). (1982). *Judgment under uncertainty: Heuristics and biases*. Cambridge, UK: Cambridge University Press.
- Kareev, Y. (2000). Seven (indeed, plus or minus two) and the detection of correlations. *Psychological Review*, 107, 397–402.
- Kaufman, B. E. (1999). Emotional arousal as a source of bounded rationality. *Journal of Economic Behaviour and Organization*, 38, 135–144.
- Koehler, J. J. (1996). On conveying the probative value of DNA evidence: Frequencies, likelihood ratios, and error rates. *University of Colorado Law Review*, 67, 859–886.

- Lopes, L. L. (1992). Three misleading assumptions in the customary rhetoric of the bias literature. *Theory and Psychology*, 2, 231–236.
- March, J. G. (1978). Bounded rationality, ambiguity, and the engineering of choice. *Bell Journal of Economics*, 9, 587–608.
- Martignon, L., & Hoffrage, U. (1999). Why does one-reason decision making work? A case study in ecological rationality. In G. Gigerenzer, P. M. Todd, & the ABC Research Group, *Simple heuristics that make us smart* (pp. 119–140). New York: Oxford University Press.
- McLeod, P., & Dienes, Z. (1996). Do fielders know where to go to catch the ball or only how to get there? *Journal of Experimental Psychology: Human Perception and Performance*, 22, 531–543.
- Michalewicz, Z., & Fogel, D. B. (2000). *How to solve it: Modern heuristics*. New York: Springer.
- Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.
- Rabin, M. (1998). Psychology and economics. *Journal of Economic Literature*, 36, 11–46.
- Roberts, S., & Pashler, H. (2000). How persuasive is a good fit? A comment on theory testing. *Psychological Review*, 107, 358–367.
- Sargent, T. J. (1993). *Bounded rationality in macroeconomics*. New York: Oxford University Press.
- Selten, R. (1998). Aspiration adaptation theory. *Journal of Mathematical Psychology*, 42, 191–214.
- Selten, R. (2001). What is bounded rationality. In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality: The adaptive toolbox*. Cambridge, MA: MIT Press.
- Simon, H. A. (1947). *Administrative behavior*. New York: Macmillan.
- Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69, 99–118.
- Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63, 129–138.
- Simon, H. A. (1986). Rationality in psychology and economics. In R. Hogarth & M. Reder (Eds.), *Rational choice: The contrast between economics and psychology* (pp. 25–40). Chicago: University of Chicago Press.
- Simon, H. A. (1990). Invariants of human behavior. *Annual Review of Psychology*, 41, 1–19.
- Stigler, G. J. (1961). The economics of information. *Journal of Political Economy*, 69, 213–225.
- Todd, P. M. (2001). Fast and frugal heuristics for environmentally bounded minds. In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality: The adaptive toolbox* (pp. 51–70). Cambridge, MA: MIT Press.
- Todd, P. M., & Gigerenzer, G. (2000). Precis of simple heuristics that make us smart. *Behavioral and Brain Sciences*, 23, 727–780.
- Tsotsos, J. K. (1991). Computational resources do constrain behavior. *Behavioral and Brain Sciences*, 14, 506–507.