In R.J. Stainton (Ed.), Contemporary debates in cognitive science. Oxford, UK: Blackwell. © 2006

CHAPTER SEVEN

Bounded and Rational

Gerd Gigerenzer

At first glance, *Homo sapiens* is an unlikely contestant for taking over the world. "Man the wise" would not likely win an Olympic medal against animals in wrestling, weight lifting, jumping, swimming, or running. The fossil record suggests that *Homo sapiens* is perhaps 400,000 years old, and is currently the only existing species of the genus *Homo*. Unlike our ancestor, *Homo erectus*, we are not named after our bipedal stance, nor are we named after our abilities to laugh, weep, and joke. Our family name refers to our wisdom and rationality. Yet what is the nature of that wisdom? Are we natural philosophers equipped with logic in search of truth? Or are we intuitive economists who maximize our expected utilities? Or perhaps moral utilitarians, optimizing happiness for everyone?

Why should we care about this question? There is little choice, I believe. The nature of *sapiens* is a no-escape issue. As with moral values, it can be ignored yet will nonetheless be acted upon. When psychologists maintain that people are unreasonably overconfident and fall prey to the base rate fallacy or to a litany of other reasoning errors, each of their claims is based on an assumption about the nature of *sapiens* – as are entire theories of mind. For instance, virtually everything that Jean Piaget examined, the development of perception, memory, and thinking, is depicted as a change in logical structure (Gruber and Voneche, 1977). Piaget's ideal image of *sapiens* was logic. It is not mine.

Disputes about the nature of human rationality are as old as the concept of rationality itself, which emerged during the Enlightenment (Daston, 1988). These controversies are about norms, that is, the evaluation of moral, social, and intellectual judgment (e.g., Cohen, 1981; Lopes, 1991). The most recent debate involves four sets of scholars, who think that one can understand the nature of *sapiens* by (i) constructing *as-if theories of unbounded rationality*, (ii) constructing *as-if theories of optimization under constraints*, (iii) demonstrating *irrational cognitive illusions*, or (iv) studying *ecological rationality*. Being engaged in this controversy, I am far from dispassionate, and have placed my bets on ecological rationality. Yet I promise that I will try to be as impartial as I can.

Four Visions of Human Rationality

The heavenly ideal of perfect knowledge, impossible on earth, provides the gold standard for many ideals of rationality. From antiquity to the Enlightenment, knowledge – as opposed to opinion – was thought to require certainty. Such certainty was promised by Christianity, but began to be eroded by the events surrounding the reformation and counter-reformation. The French astronomer and physicist Pierre-Simon Laplace (1749–1827), who made seminal contributions to probability theory and was one of the most influential scientists ever, created a fictional being known as Laplace's superintelligence or demon. The demon, a secularized version of God, knows everything about the past and present, and can deduce the future with certitude. This ideal underlies the first three of the four positions on rationality, even though they seem to be directly opposed to one another. The first two picture human behavior as an approximation to the demon, while the third blames humans for failing to reach this ideal.

I will use the term *omniscience* to refer to this ideal of perfect knowledge (of past and present, not of the future). The mental ability to deduce the future from perfect knowledge requires *unlimited computational power*. To be able to deduce the future with certainty implies that the structure of the world is *deterministic*. Omniscience, unbounded computational abilities, and determinism are ideals that have shaped many theories of rationality. Laplace's demon is fascinating precisely because he is so unlike us. Yet as the Bible tells us, God created humans in his own image. In my opinion, social science took this model too seriously and, in many a theory, recreated us in close proximity of that image.

Unbounded rationality

The demon's nearest relative is a being with "unbounded rationality" or "full rationality." For an unboundedly rational person, the world is no longer fully predictable, that is, the experienced world is not deterministic. Unlike the demon, unboundedly rational beings make errors. Yet it is assumed that they can find the *optimal* (best) strategy – that is, the one that maximizes some criterion (such as correct predictions, monetary gains, or happiness) and minimizes error. The seventeenth-century French mathematicians Blaise Pascal and Pierre Fermat have been credited with this more modest view of rationality, defined as the maximization of the expected value, later changed to the maximization of expected utility by Daniel Bernoulli (Hacking, 1975; Gigerenzer et al., 1989). In unbounded rationality, *optimization* (such as maximization) replaces determinism, whereas the assumptions of omniscience and unlimited computational power are maintained. I will use the term "optimization" in the following way:

> *Optimization* refers to a *strategy* for solving a problem, not to an *outcome*. An optimal strategy is the *best* for a given class of problems (but not necessarily

a perfect one, for it can lead to errors). To refer to a strategy as optimal, one must be able to prove that there is no better strategy (although there can be equally good ones).

Because of their lack of psychological realism, theories that assume unbounded rationality are often called *as-if* theories. They do not aim at describing the actual cognitive processes, but are only concerned with predicting behavior. In this program of research, the question is: If people were omniscient and had all the necessary time and computational power to optimize, how would they behave? The preference for unbounded rationality is widespread. This is illustrated by some consequentionalist theories of moral action, which assume that people consider (or should consider) the consequences of all possible actions for all other people before choosing the action with the best consequences for the largest number of people (Williams, 1988). It is illustrated by theories of cognitive consistency, which assume that our minds check each new belief for consistency with all previous beliefs encountered and perfectly memorized; theories of optimal foraging, which assume that animals have perfect knowledge of the distribution of food and of competitors; and economic theories that assume that actors or firms know all relevant options, consequences, benefits, costs, and probabilities.

Optimization under constraints

Unbounded rationality ignores the constraints imposed on human beings. A *constraint* refers to a limited mental or environmental resource. Limited memory span is a constraint of the mind, and information cost is a constraint on the environment. The term *optimization under constraints* refers to a class of theories that assume that one or several constraints exist.

Lack of omniscience – together with its consequence, the need to search for information – is the key issue in optimization under constraints, whereas the absence of search is a defining feature of theories of unbounded rationality. Models of search specify a searching direction (where to look for information) and a stopping rule (when to stop search). The prototype is Wald's (1950) sequential decision theory. In Stigler's (1961) classical example, a customer wants to buy a used car. He continues to visit used car dealers until the expected costs of further search exceed its expected benefits. Here, search takes place in the environment. Similarly, in Anderson's (1990) rational theory of memory, search for an item in memory continues until the expected costs of further search exceed the expected benefits. Here, search occurs inside the mind. In each case, omniscience is dropped but optimization is retained: the stopping point is the optimal cost-benefit trade-off.

Optimization and realism can inhibit one another, with a paradoxical consequence. Each new realistic constraint makes optimization calculations more difficult, and eventually impossible. The ideal of optimization, in turn, can undermine the attempt to make a theory more realistic by demanding new unrealistic assumptions – such as the knowledge concerning cost and benefits necessary for estimating the optimal stopping point. As a consequence, models of optimization under constraints tend to be more complex than models of unbounded rationality, depicting people in the image of econometricians (Sargent, 1993). This unresolved paradox is one reason why constraints are often ignored and theories of unbounded rationality preferred. Since economists and biologists (wrongly) tend to equate optimization under constraints with *bounded rationality*, bounded rationality is often dismissed as an unpromisingly complicated enterprise and ultimately nothing but full rationality in disguise (Arrow, 2004). Theories of optimization under constraints tend to be presented as as-if theories, with the goal of predicting behavior but not the mental process - just as models of unbounded rationality do.

Heuristics and biases: cognitive illusions

Unbounded rationality and optimization under constraints conceive of humans as essentially rational. This is sometimes justified by the regulating forces of the market, by natural selection, or by legal institutions that eliminate irrational behavior. The "heuristics and biases" or "cognitive illusions" program (Kahneman et al., 1982; Gilovich et al., 2002) opposes theories assuming that humans are basically rational. It has two stated goals. The main goal is to understand the cognitive processes that produce both valid and invalid judgments. Its second goal (or method to achieve the first one) is to demonstrate errors of judgment, that is, systematic deviations from rationality also known as cognitive illusions (Kahneman and Tversky, 1996, p. 582). The cognitive processes underlying these errors are called heuristics, and the major three proposed are representativeness, availability, and anchoring and adjustment. The program has produced a long list of biases (see Krueger and Funder, 2004). It has shaped many fields, such as social psychology and behavioral decision making, and helped to create new fields such as behavioral economics and behavioral law and economics.

Although the heuristics-and-biases program disagrees with rational theories on whether or not people follow some norm of rationality, it does not question the norms themselves. Rather, it retains the norms and interprets deviations from these norms as cognitive illusions: "The presence of an error of judgment is demonstrated by comparing people's responses either with an established fact . . . or with an accepted rule of arithmetic, logic, or statistics" (Kahneman and Tversky, 1982, p. 493). For instance, when Wason and Johnson-Laird (1972) criticized Piaget's logical theory of thinking as descriptively incorrect, they nevertheless retained the same logical standards as normatively correct for the behavior studied. When Tversky and Kahneman (1983) reported that people's reasoning violated the laws of logic ("conjunction rule"), they nevertheless retained logic as the norm for rational judgment.

The heuristics-and-biases program correctly argues that people's judgments do in fact systematically deviate from the content-blind laws of logic or optimization. But it has hesitated to take two necessary further steps: to rethink the norms, and to provide testable theories of heuristics. The laws of logic and probability are neither necessary nor sufficient for rational behavior in the real world (see below), and mere verbal labels for heuristics can be used post hoc to "explain" almost everything.

The term "bounded rationality" has been used both by proponents of optimization under constraints, emphasizing rationality, and by the heuristics-and-biases program, emphasizing irrationality. Even more confusing is the fact that the term was coined by Herbert A. Simon, who was not referring to optimization or irrationality, but to an ecological view of rationality (see next section), which was revolutionary in thinking about norms, not just behavior (Simon, 1956; Selten, 2001; Gigerenzer, 2004b).

The science of heuristics: ecological rationality

The starting point for the science of heuristics is the relation between mind and environment, rather than between mind and logic (Gigerenzer et al., 1999; Gigerenzer and Selten, 2001). Humans have evolved in natural environments, both social and physical. To survive, reproduce, and evolve, the task is to adapt to these environments, or else to change them. Piaget called these two fundamental processes assimilation and accommodation, but he continued to focus on logic. The structure of natural environments, however, is ecological rather than logical. In Simon's (1990) words: "Human rational behavior is shaped by a scissors whose two blades are the structure of task environments and the computational capabilities of the actor" (p. 7). Just as one cannot understand how scissors cut by looking only at one blade, one will not understand human behavior by studying cognition or the environment alone.

The two key concepts are *adaptive toolbox* and *ecological rationality*. The analysis of the adaptive toolbox is descriptive, whereas that of ecological rationality is normative. The adaptive toolbox contains the *building blocks* for *fast and frugal heuristics*. A heuristic is fast if it can solve a problem in little time, and frugal if it can solve it with little information. Unlike as-if optimization models, heuristics can find good solutions independent of whether an optimal solution exists. As a consequence, one does not need to "edit" a real world problem in order to make it accessible to the optimization calculus (e.g., by limiting the number of competitors and choice alternatives, by providing quantitative probabilities and utilities, or by ignoring constraints). Heuristics work in real-world environments of natural complexity, where an optimal strategy is typically unknown.

The study of ecological rationality answers the question: In what environments will a given heuristic work? Where will it fail? Note that this normative question can only be answered if there is a process model of the heuristic in the first place, and the results are gained by proof or simulation. As mentioned beforehand, the ecological rationality of a verbal label such as representativeness cannot be determined. At most one can say that representativeness is sometimes good and sometimes bad – without being able to explicate the "sometimes."

Let me illustrate the difference between the adaptive toolbox program and the previous visions of rationality with the example of problem solving in baseball and cricket.

How does an outfielder catch a flyball? In Richard Dawkins' words: "When a man throws a ball high in the air and catches it again, he behaves as if he had solved a set of differential equations in predicting the trajectory of the ball" (1989, p. 96). Note that Dawkins invokes a form of omniscience. The trajectory of a ball is determined by a number of causal factors, including the ball's initial velocity, angle, and distance from the player; the air resistance; the speed and direction of the wind at each point of the trajectory, and spin. An unboundedly rational player would measure all causally relevant factors for computing the ball's trajectory in no time, would know the optimal formula to compute the trajectory, and run to the point where the ball lands. A player who optimizes under constraints would not be able to measure all factors given the time constraint of a few seconds, but would know the optimal formula to compute the trajectory given these constraints. These are as-if theories. Real humans perform poorly in estimating the location where the ball will strike the ground (Saxberg, 1987; Babler and Dannemiller, 1993). Note also that as-if theories have limited practical use for instructing novice players or building a robot player. The heuristics-and-biases program might respond with an experiment demonstrating that experienced players make systematic errors in estimating the point where the ball lands, such as underestimating the distance to the player. Tentatively, these errors might be attributed to player's overconfidence bias or optimism bias (underestimating the distance creates a false certainty that one can actually catch a ball). The demonstration of a discrepancy between judgment and norm, however, does not lead to an explanation of how players actually catch a ball. One reason for this is that the critique is merely descriptive, that is, what players can do, and does not extend to the norm, that is, what players should do. The norm is the same – to compute the trajectory correctly, requiring knowledge of the causal factors.

Yet it is time to rethink the norms, such as the ideal of omniscience. The normative challenge is that real humans do not need a full representation and unlimited computational power. In my view, humans have an adaptive toolbox at their disposal, which may contain heuristics that can catch balls in a fast and frugal way. Thus, the research question is: Is there a fast and frugal heuristic that can solve the problem? Experiments have shown that experienced players in fact use several heuristics (e.g., McLeod and Dienes, 1996). One of these is the *gaze heuristic*, which works only when the ball is already high up in the air:

> *Gaze heuristic*: Fixate your gaze on the ball, start running, and adjust your speed so that the angle of gaze remains constant.

The angle of gaze is between the eye and the ball, relative to the ground. The gaze heuristic ignores all causally relevant factors when estimating the ball's trajectory. It attends to only one variable: the angle of gaze. This heuristic belongs to the class of *one-reason decision making heuristics*. A player relying on this heuristic cannot predict where the ball will land, but the heuristic will lead him to that spot. In other words, computing the trajectory is not necessary; it is not an appropriate norm. The use of heuristics crosses species borders. People rely on the gaze heuristic and related heuristics in sailing and flying to avoid collisions with other boats or planes; and bats, birds, and flies use the same heuristics for predation and pursuit (Shaffer et al., 2004).

To repeat, the gaze heuristic consists of three building blocks: fixate your gaze on the ball, start running, and adjust running speed. These building blocks can be part of other heuristics, too.

> *Definition*: A fast and frugal heuristic is a (conscious or unconscious) strategy that searches for minimal information and consists of building blocks that exploit evolved abilities and environmental structures.

Heuristics can be highly effective because they are anchored in the evolved brain and in the external environment. Let me explain.

Embodiment: heuristics exploit evolved abilities. For instance, the first building block "fixate your gaze on the ball" exploits the evolved ability of *object tracking*, in this case, to track a moving object against a noisy background. It is easy for humans to do this; 3-month-old babies can already hold their gaze on moving targets (Rosander and Hofsten, 2002). Tracking objects, however, is difficult for a robot; a computer program that can track objects as well as a human mind can does not yet exist. Thus, the gaze heuristic is simple for humans but not for today's generation of robots. The standard definition of optimization as computing the maximum or minimum of a function, however, ignores the "hardware" of the human brain. In contrast, a heuristic exploits hard-wired or learned cognitive and motor processes, and these abilities make it simple. This is the first reason why fast and frugal heuristics can, in the real world, be superior to some optimization strategy.

Situatedness: heuristics exploit structures of environments. The rationality of heuristics is not logical, but ecological. The study of the ecological rationality of a heuristic answers the normative question concerning the environments in which a heuristic will succeed and in which it will fail. It specifies the class of problems a given heuristic can solve (Martignon and Hoffrage, 1999; Goldstein et al., 2001). Ecological rationality implies that a heuristic is not good or bad, rational or irrational per se, but only relative to an environment. It can exploit certain structures of environments or change them. For instance, the gaze heuristic transforms the complex trajectory of the ball in the environment into a straight line.

Note that as-if optimization theories – because they ignore the human mind – are formulated more or less independently of the hardware of the brain. Any computer can compute the maximum of a function. Heuristics, in contrast, exploit the specific hardware and are dependent on it. Social heuristics, for instance, exploit the evolved or learned abilities of humans for cooperation, reciprocal altruism, and identification (Laland, 2001). The principles of embodiment and situatedness are also central for "New AI" (Brooks, 2002). For an introduction to the study of fast and frugal heuristics, see Payne et al., 1993; Gigerenzer et al., 1999; and Gigerenzer and Selten 2001.

The Problem With Content-Blind Norms

In the heuristics-and-biases program, a norm is typically a law (axiom, rule) of logic or probability rather than a full optimization model. A law of logic or probability is used as a *content-blind norm* for a problem if the "rational" solution is determined independently of its content. For instance, the truth table of the material implication *if P then Q* is defined independently of the content of the Ps and Qs. The definition is in terms of a specific syntax. By content, I mean the semantics (what are the Ps and Qs?) and the pragmatics (what is the goal?). The heuristics-and-biases program of studying whether people's judgments deviate from content-blind norms proceeds in four steps:

1 *Syntax first*: Start with a law of logic or probability.
2 *Add semantics and pragmatics*: Replace the logical terms (e.g., material implication, mathematical probability) by English terms (e.g., if . . . then; probable), add content, and define a problem to be solved.
3 *Content-blind norm*: Use the syntax to define the "rational" answer to the problem. Ignore semantics and pragmatics.

4 *Cognitive illusion*: If people's judgments deviate from the "rational" answer, call the discrepancy a cognitive illusion. Attribute it to some deficit in the human mind (not to your norms).

Content-blind norms derive from an internalist conception of rationality. Examples are the use of the material implication as a norm for reasoning about any content (Wason and Johnson-Laird, 1972), the set-inclusion or "conjunction rule" (Tversky and Kahneman, 1983), and Bayes's rule (Kahneman and Frederick, 2002; see also Matheson, chapter 8, BOUNDED RATIONALITY AND THE ENLIGHTENMENT PICTURE OF COGNITIVE VIRTUE). Proponents of content-blind norms do not use this term, but instead speak of "universal principles of logic, arithmetic, and probability calculus" that tell us what we should think (Piatelli-Palmarini, 1994, p. 158). Consider the material implication.

In 1966, the British psychologist Peter Wason invented the *selection task*, also known as the *four-card problem*, to study reasoning about conditional statements. This was to become the most frequently studied task in the psychology of reasoning. Wason's starting point was the material implication $P \rightarrow Q$, as defined by the truth table in elementary logic. In the second step, the Ps and Qs are substituted by some content, such as "numbers" (odd/even) and "letters" (consonants/vowels). The material implication "$\rightarrow$" is replaced by the English terms "if... then," and a rule is introduced: "If there is an even number on one side of the card, there is a consonant on the other." Four cards are placed on the table, showing an even number, an odd number, a consonant, and a vowel on the surface side. People are asked which cards need to be turned around in order to see whether the rule has been violated. In the third step, the "correct" answer is defined by the truth table: to turn around the P and the not-Q card, and nothing else, because the material conditional is false if and only if $P \cap \neg Q$. However, in a series of experiments, most people picked other combinations of cards, which was evaluated as a reasoning error, due to some cognitive illusion. In subsequent experiments, it was found that the cards picked depended on the content of the Ps and Qs, and this was labeled the "content effect." Taken together, these results were interpreted as a demonstration of human irrationality and a refutation of Piaget's theory of operational thinking. Ironically, as mentioned before, Wason and Johnson-Laird (1972) and their followers held up truth-table logic as normative even after they criticized it as descriptively false.

Are content-blind norms reasonable norms? Should one's reasoning always follow truth-table logic, the conjunction rule, Bayes's rule, the law of large numbers, or some other syntactic law, irrespective of the content of the problem? The answer is no, and for several reasons. A most elementary point is that English terms such as "if... then" are not identical to logical terms such as the material conditional "$\rightarrow$" (Fillenbaum, 1977). This confusion is sufficient to reject logic as a content-blind norm. More interesting, adaptive behavior has other goals than logical truth or consistency, such as dealing intelligently with other people. For instance, according to Trivers' (1971, 2002) theory of reciprocal altruism, each human possesses altruistic and cheating tendencies. Therefore, one goal in a social contract is to search for information revealing whether one has been cheated by the other party (Cosmides, 1989). Note that the perspective is essential: you want to find out whether you were cheated by the other party, not whether you cheated the other. Logic, in contrast, is without perspective. Consider a four-cards task whose content is a social contract between an employer and an employee (Gigerenzer and Hug, 1992):

> If a previous employee gets a pension from the firm, then that person must have worked for the firm for at least ten years.

The four cards read: got a pension, worked ten years for the firm, did not get a pension, worked eight years for the firm. One group of participants was cued into the role of the employer, and asked to check those cards (representing files of previous employees) that could reveal whether the rule was violated. The far majority picked "got a pension" and "worked for eight years." Note that this choice is consistent with both the laws of the truth table and the goal of cheater detection. Proponents of content-blind norms (mis-)interpreted this and similar results as indicating that social contracts somehow facilitated logical reasoning. But when we cued the participants into the role of an employee, the far majority picked "did not get a pension" and "worked for ten years." (In contrast, in the employer's group, no participant had checked this information.) Now the result was inconsistent with the truth table, but, from the employee's perspective, again consistent with the goal of not being cheated. Search for information was Machiavellian: to avoid being cheated oneself, not avoiding cheating others.

The perspective experiment clearly demonstrates that logical thinking is not central to human reasoning about these problems, as well as that truth-table logic is an inappropriate norm (Cosmides and Tooby, 1992; Gigerenzer, 2000). Yet several decades and hundreds of thousands of dollars of grant money have been wasted trying to show that human thinking violates the laws of logic. We have learned next to nothing about the nature of thinking or other cognitive processes. The same holds for research on other content-blind norms (Gigerenzer, 1996, 2001). Inappropriate norms tend to suggest wrong questions, and the answers to these generate more confusion than insight into the nature of human judgment (Gigerenzer, 2000). My point is not new. Wilhelm Wundt (1912/1973), known as the father of experimental psychology, concluded that logical norms have little to do with thought processes, and that attempts to apply them to learn about psychological processes have been absolutely fruitless. But at least some psychologists have learned. For instance, Lance Rips, who had argued that deductive logic might play a central rule in cognitive architecture (Rips, 1994), declared that he would not defend this "imperialist" theory anymore (Rips, 2002).

The Ecological Rationality of Heuristics

Is there an alternative to optimization and content-blind norms? The alternatives are external forms of rationality, where reasonableness is measured by the actual success in solving problems. Success has several aspects, depending on the problem. These include predictive accuracy of a strategy, how much information it needs to search, and how fast it leads to a decision – in some situations, to act now is better than to wait until the best action is found. Ecological rationality can be expressed in comparative terms: a given heuristic performs better in environment A than in B. For instance, imitation of others' successful behavior (as opposed to individual learning) works in environments that change slowly, but not in environments under rapid change (Boyd and Richerson, 2001). Ecological rationality can also be analyzed in quantitative terms: a heuristic will make 80 percent correct predictions in environment E, and requires only 30 percent of the information.

The ecological rationality of a heuristic is conditional on an environment. Content-blind norms, in contrast, are defined without consideration of any environment. Ecological rationality is comparative or quantitative; it is not necessarily about the best strategy. This provides it with an advantage: ecological rationality can be determined in all situations where optimization is out of reach, and one does not need to edit the problem so that optimization can be applied. In what follows, I describe the ecological rationality of three classes of heuristics. For a more extensive analysis, see Gigerenzer et al., 1999; Goldstein et al., 2001; and Smith, 2003.

Recognition heuristic

Daniel Goldstein and I asked American and German students the following question (Goldstein and Gigerenzer, 2002):

> Which city has more inhabitants: San Diego or San Antonio?

Sixty-two percent of Americans answered correctly: San Diego. The Germans knew little of San Diego, and many had never heard of San Antonio. What percentage of the more ignorant Germans found the right answer? One hundred percent. How can people who know less make more correct inferences? The answer is that the Germans used the *recognition heuristic*:

> If you recognize one city, but not the other, infer that it has the larger population.

The Americans could not use this heuristic. They knew too much. The Americans had heard of both cities, and had to rely on their recall knowledge. Exploiting the wisdom in partial ignorance, the recognition heuristic is an example of ignorance-based decision making. It guides behavior in a large variety of situations: rats choose food they recognize on the breath of a fellow rat and tend to avoid novel food; children tend to approach people they recognize and avoid those they don't; teenagers tend to buy CDs of bands whose name they have heard of; adults tend to buy products whose brand name they recognize; participants in large conferences tend to watch out for faces they recognize; university departments sometimes hire professors by name recognition; and institutions, colleges, and companies compete for a place in the public's recognition memory through advertisement.

Like all heuristics, the recognition heuristic works better in certain environments than in others. The question of when it works is the question of its ecological rationality:

> The recognition heuristic is ecologically rational in environments where the recognition validity α is larger than chance: $\alpha > 0.5$.

The validity α is defined as the proportion of cases where a recognized object has a higher criterion value (such as population) than the unrecognized object, for a given set of objects. This provides a quantitative measure for ecological rationality.

For instance, α is typically around 0.8 for inferring population (Goldstein and Gigerenzer, 2002), 0.7 for inferring who will win a Grand Slam tennis match (Serwe and Frings, 2004), and 0.6 for inferring disease prevalence (Pachur and Hertwig, 2004).

Take The Best

As a second illustration of ecological rationality, consider a heuristic from the family of one-reason decision-making. The task is to infer which of two objects has a higher value on a criterion, based on binary cues. The heuristic is called Take The Best, because it relies on only one cue to make this inference, the cue with the highest validity on which the objects differ. The rest of the cues are ignored. Take The Best has three building blocks, a search rule, stopping rule, and decision rule:

1. *Search rule*: Search through cues in order of their validity. Look up the cue values of the cue with the highest validity first.
2. *Stopping rule*: If one object has a positive cue value and the other does not (or is unknown), then stop search and proceed to Step 3. Otherwise exclude this cue and return to Step 1. If no more cues are found, guess.
3. *Decision rule*: Predict that the object with the positive cue value has the higher value on the criterion.

The validity of a cue i is defined as $v_i = R_i/P_i$, where R_i = number of correct predictions by cue i, and P_i = number of pairs where the values of cue i differ between objects.

When is relying on only one cue rational? Consider an environment with M binary cues ordered according to their weights W_j (such as beta weights), with $1 \leq j \leq M$. A set of cue weights is noncompensatory if for every weight:

$$W_j > \sum_{k>j} W_k$$

I refer to this environmental structure as *noncompensatory information*. An example is the set of weights 1, $^1/_2$, $^1/_4$, $^1/_8$, and so on. The sum of the cue weights to the right of a cue can never be larger than this cue's weight – the sum cannot compensate for a cue with a higher weight. Here, Take The Best makes the same inferences as any linear strategy (with the same order of weights). Thus, we get the following result (Martignon and Hoffrage, 1999):

> In an environment with non-compensatory information, no linear strategy can outperform the faster and more frugal Take The Best heuristic.

A second environmental structure that Take The Best can exploit is scarce information:

$$M < \log_2 N,$$

where M and N are the number of cues and objects, respectively (Martignon and Hoffrage, 2002). An example of scarce information is a sample with 30 objects measured on 5 cues or predictors ($\log_2 30 < 5$). Related environmental structures are discussed in Hogarth and Karelaia, 2005. Consistent with these results, Take The Best and other one-reason decision making heuristics have been proven to be, on average, more accurate than multiple regression in making various economic, demographic, environmental, and health forecasts, as well as in the prediction of heart attacks (Green and Mehr, 1997; Czerlinski et al., 1999). Todorov (2003) showed that Take The Best predicted the outcomes of basketball games during the 1996/97 NBA season as accurately as Bayes's rule did, but with less information. Chater et al. (2003) demonstrated that Take The Best matched highly complex computational methods such as a three-layer feedforward connectionist network, Quinlan's (1993) decision three algorithm, and two exemplar-based models, Nearest Neighbor and Nosofsky's (1990) Generalized Context Model. When the environment had scarce information, specifically when the training set included less than 40 percent of all objects, Take The Best was more accurate than any of these computationally expensive strategies. The effectiveness of one-reason decision making has been demonstrated, among others, in the diagnosis of heart attacks (Green and Mehr, 1997) and in the simple heuristic for prescribing antibiotics to children (Fischer et al., 2002).

Follow the majority

Social heuristics exploit the capacity of humans for social learning and imitation (imitation need not result in learning), which is unmatched among the animal species. For instance, just like the recognition heuristic, the following social heuristic allows an individual to act with only a surface analysis of the situation:

> *Do-what-the-majority-do heuristic*: If you see the majority of your peers display a behavior, engage in the same behavior.

This simple social heuristic seems to cause a broad range of adaptive behaviors (Laland, 2001). It saves an organism from having to extract information anew from the environment, and hence starting from scratch. It facilitates the spread of cultural behavior, from religious ceremonies to rules of conduct to riots against minorities. Imitating the majority – as opposed to imitating one skilled individual – is characteristic of adaptive behaviors related to social issues, such as moral and political actions (Gigerenzer, 2004a). Studies have reported behavior copying in animals and humans. Dugatkin (1992) argued that female guppies choose between males by copying the mate preferences of other females. In modern human societies, teenagers admire a movie star because everyone else in their peer group adulates that person. Advertisement exploits this heuristic by portraying a product surrounded by many admirers (not just one). People may display disgust for members of a minority because they notice that most of their peers do the same. A man may start thinking of marriage at a time when most other men in his social group do, say, around age 30. Copying the behavior of one's peers is a most frugal heuristic, for it almost guarantees the peer group's approbation, is sometimes even a condition for peer acceptance, and one does not need to consider the pros and cons of one's behavior.

Do-what-the-majority-do tends to be ecologically rational in situations where:

1. the observer and the demonstrators of the behavior are exposed to similar environments, such as social systems;
2. the environment is stable or changing slowly rather than quickly;
3. the environment is noisy and consequences are not immediate, that is, it is hard or time-consuming to figure out whether a choice is good or bad, such as which political or moral system is preferable (Boyd and Richerson, 1985; Goldstein et al., 2001).

In environments where these conditions do not hold, copying the behavior of the majority can lead to disaster. For instance, copying the production and distribution systems of traditional firms can be detrimental when an economy changes from local to globalized.

Cognitive Luck

In this volume, Matheson discusses the study of ecological (bounded) rationality as a way to overcome the epistemic internalism of the Enlightenment tradition. But he raises a worry:

> If cognitive virtue is located outside the mind in the way that the Post-Enlightenment Picture suggests, then it turns out to be something bestowed on us by features of the world not under our control: it involves an intolerable degree of something analogous to what theoretical ethicists call "moral luck"... [cf. Williams, 1981; Nagel, 1993] – "cognitive luck," we might say. (Matheson, chapter 8, p. 143)

This worry is based on the assumption that internal ways to improve cognition are under our control, whereas the external ones are not.

This assumption, however, is often incorrect, and reveals a limit of an internalist view of cognitive virtue. I conjecture that changing environments can in fact be easier than changing minds. Consider the serious problem of innumerate physicians, as illustrated by screening for colorectal cancer. A man tests positive on the FOB (fecal occult blood) test and asks the physician what the probability is that he actually has cancer. What do physicians tell that worried man? We (Hoffrage and Gigerenzer, 1998) gave exper-ienced physicians the best estimates of base rate (0.3 percent), sensitivity (50 percent), and false positive rate (3 percent), and asked them to estimate the probability of colorectal cancer given a positive test. Their estimates ranged between 1 percent and 99 percent. If patients knew about this variability, they would be rightly scared.

This result illustrates a larger problem: When physicians try to draw a conclusion from probabilities, their minds typically cloud over (Gigerenzer, 2002). What can be done to correct this? An internalist might recommend training physicians to use Bayes's rule in order to compute the posterior probability. In theory, this training should work wonders, but in reality, it does not. One week after students successfully passed such a course, for instance, their performance was already down by 50 percent, and it continued to fade away week by week (Sedlmeier and Gigerenzer, 2001). Moreover, the chance of convincing physicians to take a statistics course in the first place is almost nil; most have no time or little motivation, while others believe they are incurably innumerate. Are we stuck for eternity with innumerate physicians? No. In the ecological view, thinking does not happen simply in the mind, but in interaction between the mind and its environment. This adds a second, and more efficient, way to improve the situation: to edit the environment. The relevant part of the environment is the representation of the information, because the representation does part of the Bayesian computation. Natural (non-normalized) frequencies are such an efficient representation; they mimic the way information has been encountered before the advent of writing and statistics, throughout most of human evolution. For instance: 30 out of every 10,000 people have colorectal cancer, 15 of these will have a positive test; of the remaining people without colorectal cancer, 300 will still have a positive test. When we presented the numerical information in natural frequencies as opposed to conditional probabilities, then the huge variability in physicians' judgments disappeared. They all gave reasonable estimates, with the majority hitting exactly on the Bayesian posterior of about 5 percent, or 1 in 20.

Similarly, by changing the environment, we can make many so-called cognitive illusions largely disappear (Gigerenzer, 2000), enable fifth and sixth graders to solve Bayesian problems before they even heard of probabilities (Zhu and Gigerenzer, forthcoming), and help judges and law students understand DNA evidence (Hoffrage et al., 2000). Thus, an ecological view actually extends the possibilities to improve judgment, whereas an internal view limits the chances. To summarize, the term "cognitive luck" only makes sense from an internalist view, where luck in fact refers to the theory's ignorance concerning the environment, including the social environment. From an ecological view, environmental structures, not luck, directly influence cognition and can be designed to improve it. Cognitive virtue is, in my view, a relation between a mind and its environment, very much like the notion of ecological rationality (see also Bishop, 2000).

What Is the Rationality of Homo Sapiens?

What makes us so smart? I have discussed four answers. The first is that we are smart because we behave as if we were omniscient and had unlimited computational power to find the optimal strategy for each problem. This is the beautiful fiction of unbounded rationality. The second is a modification of the first that diminishes omniscience by introducing the need for searching for information and the resulting costs, but insists on the ideal of optimization. These two programs define the theories in much of economics, biology, philosophy, and the social sciences. Both have an anti-psychological bias: they try to define rational behavior without psychology, promoting as-if theories. The assumption is that one can predict behavior while ignoring what we know about the human mind, an assumption that is not always true. In the image of Laplace's demon, *Homo economicus* has defined *Homo sapiens*: we are basically rational beings, and the nature of our rationality can be understood through the fictions of omniscience and optimization. The heuristics-and-biases program has attacked that position, but only on the descriptive level, using content-blind norms as the yardstick to diagnose human irrationality. The result has been that we are mostly or sometimes – the quantifiers keep changing – irrational, committing systematic errors of reasoning.

There is now a literature that tries to determine which of these positions is correct. Are we rational or irrational? Or perhaps 80 percent rational and 20 percent irrational? Some blessed peacemakers propose that the truth is in the middle and we are a little of both, so there is no real disagreement. For instance, the debate between Kahneman and Tversky and myself (e.g., Gigerenzer, 1996; Kahneman and Tversky, 1996) has been sometimes misunderstood as concerning the question of *how much* rationality or irrationality people have. In this view, rationality is like a glass of water, and Kahneman and Tversky see the glass as half empty, whereas I see it as half full. For instance, Samuels et al. conclude their call for "ending the rationality war" with the assertion that the two parties "do not have any deep disagreement over the extent of human rationality" (2004, p. 264). However, the issue is not quantity, but quality: *what* exactly rationality and irrationality are in the first place. We can easily agree how often experiment participants have violated the truth-table logic or some other logical law in an experimental task, and how often not. But proponents of the heuristics-and-biases program count the first as human irrationality, and the second as rationality. I do not. I believe that we need a better understanding of human rationality than relative to content-blind norms. These were of little relevance for *Homo sapiens*, who had to adapt to a social and physical world, not to systems with content-free syntax, such as the laws of logic.

The concept of ecological rationality is my answer to the question of the nature of *Homo sapiens*. It defines the rationality of heuristics independently of optimization and content-blind norms, by the degree to which they are adapted to environments. The study of ecological rationality facilitates understanding a variety of counter-intuitive phenomena, including when one reason is better than many, when less is more, and when partial ignorance pays. *Homo sapiens* has been characterized as a tool-user. There is some deeper wisdom in that phrase. The tools that make us smart are not bones and stones, but the heuristics in the adaptive toolbox.

References

Anderson, J. R. (1990). *The Adaptive Character of Thought*. Hillsdale, NJ: Erlbaum.

Arrow, K. J. (2004). Is bounded rationality unboundedly rational? Some ruminations. In M. Augier and J. G. March (eds.), *Models of a Man: Essays in Memory of Herbert A. Simon*. Cambridge, MA: MIT Press.

Babler, T. G. and Dannemiller, J. L. (1993). Role of image acceleration in judging landing location of free-falling projectiles. *Journal of Experimental Psychology: Human Perception and Performance*, 19, 15–31.

Bishop, M. A. (2000). In praise of epistemic irresponsibility: How lazy and ignorant can you be? *Synthese*, 122, 179–208.

Boyd, R. and Richerson, P. J. (1985). *Culture and the Evolutionary Process.* Chicago: University of Chicago Press.

— and — (2001). Norms and bounded rationality. In G. Gigerenzer and R. Selten (eds.), *Bounded Rationality: The Adaptive Toolbox.* Cambridge, MA: MIT Press.

Brooks, R. (2002). *Robot: The Future of Flesh and Machines.* London: Penguin Books.

Chater, N., Oaksford, M., Nakisa, R., and Redington, M. (2003). Fast, frugal, and rational: How rational norms explain behavior. *Organizational Behavior and Human Decision Processes,* 90, 63–86.

Cohen, L. J. (1981). Can human irrationality be experimentally demonstrated? *Behavioral and Brain Sciences,* 4, 317–70.

Cosmides, L. (1989). The logic of social exchange: Has natural selection shaped how humans reason? Studies with the Wason selection task. *Cognition,* 31, 187–276.

Cosmides, L. and Tooby, J. (1992). Cognitive adaptions for social exchange. In J. H. Barkow, L. Cosmides, and J. Tooby (eds.), *The Adapted Mind: Evolutionary Psychology and the Generation of Culture.* New York: Oxford University Press.

Czerlinski, J., Gigerenzer, G., and Goldstein, D. G. (1999). How good are simple heuristics? In G. Gigerenzer, P. M. Todd, and the ABC Research Group, *Simple Heuristics that Make Us Smart.* New York: Oxford University Press.

Daston, L. (1988). *Classical Probability in the Enlightenment.* Princeton, NJ: Princeton University Press.

Dawkins, R. (1989). *The Selfish Gene* (2nd edn.). Oxford: Oxford University Press.

Dugatkin, L. A. (1992). Sexual selection and imitation: Females copy the mate choice of others. *The American Naturalist,* 139, 1384–9.

Fillenbaum, S. (1977). Mind your p's and q's: the role of content and context in some uses of and, or, and if. *Psychology of Learning and Motivation,* 11, 41–100.

Fischer, J. E., Steiner, F., Zucol, F., et al. (2002). Use of simple heuristics to target macrolide prescription in children with community-acquired pneumonia. *Archives of Pediatrics and Adolescent Medicine,* 156, 1005–8.

Gigerenzer, G. (1996). On narrow norms and vague heuristics: A reply to Kahneman and Tversky (1996). *Psychological Review,* 103, 592–6.

— (2000). *Adaptive Thinking: Rationality in the Real World.* New York: Oxford University Press.

— (2001). Content-blind norms, no norms, or good norms? A reply to Vranas. *Cognition,* 81, 93–103.

— (2002). *Calculated Risks: How to Know When Numbers Deceive You.* New York: Simon and Schuster. (Published in UK as *Reckoning with Risk: Learning to Live with Uncertainty.* London: Penguin Books.)

— (2004a). Fast and frugal heuristics: The tools of bounded rationality. In D. Koehler and N. Harvey (eds.), *Handbook of Judgment and Decision Making.* Oxford: Blackwell.

— (2004b). Striking a blow for sanity in theories of rationality. In R. B. Augier and J. G. March (eds.), *Models of a Man: Essays in Honor of Herbert A. Simon.* Cambridge, MA: MIT Press.

Gigerenzer, G. and Hug, K. (1992). Domain-specific reasoning: Social contracts, cheating, and perspective change. *Cognition,* 43, 127–71.

Gigerenzer, G. and Selten, R. (eds.) (2001). *Bounded Rationality: The Adaptive Toolbox.* Cambridge, MA: MIT Press.

Gigerenzer, G., Swijtink, Z., Porter, T., Daston, L., Beatty, J., and Krüger, L. (1989). *The Empire of Chance: How Probability Changed Science and Everyday Life.* Cambridge: Cambridge University Press.

Gigerenzer, G., Todd, P. M., and the ABC Research Group (1999). *Simple Heuristics that Make Us Smart.* New York: Oxford University Press.

Gilovich, T., Vallone, R., and Tversky, A. (1985). The hot hand in basketball: On the misconception of random sequences. *Cognitive Psychology,* 17, 295–314.

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

Goldstein, D. G., Gigerenzer, G., Hogarth, R. M., et al. (2001). Group report: Why and when do simple heuristics work? In G. Gigerenzer and R. Selten (eds.), *Bounded Rationality: The Adaptive Toolbox.* Cambridge, MA: MIT Press.

Green, L. and Mehr, D. R. (1997). What alters physicians' decisions to admit to the coronary care unit? *Journal of Family Practice,* 45, 219–26.

Gruber, H. E. and Vonèche, J. J. (1977). *The Essential Piaget.* New York: Basic Books.

Hacking, I. (1975). *The Emergence of Probability.* Cambridge: Cambridge University Press.

Hoffrage, U. and Gigerenzer, G. (1998). Using natural frequencies to improve diagnostic inferences. *Academic Medicine,* 73, 538–40.

Hoffrage, U., Lindsay, S., Hertwig, R., and Gigerenzer, G. (2000). Communicating statistical information. *Science,* 290, 2261–2.

Hogarth, R. M. and Karelaia, N. (2005). Ignoring information in binary choice with continuous variables: When is less "more"? *Journal of Mathematical Psychology,* 49, 115–25.

Kahneman, D. and Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. In T. Gilovich, D. Griffin, and D. Kahneman (eds.), *Heuristics and Biases: The Psychology of Intuitive Judgment.* New York: Cambridge University Press.

Kahneman, D. and Tversky, A. (1982). On the study of statistical intuitions. In D. Kahneman, P. Slovic, and A. Tversky (eds.), *Judgment Under Uncertainty: Heuristics and Biases.* Cambridge: Cambridge University Press.

— and — (1996). On the reality of cognitive illusions. A reply to Gigerenzer's critique. *Psychological Review,* 103, 582–91.

Kahneman, D., Slovic, P., and Tversky, A. (eds.) (1982). *Judgment Under Uncertainty: Heuristics and Biases.* Cambridge: Cambridge University Press.

Krueger, J. I. and Funder, D. C. (2004). Towards a balanced social psychology: Causes, consequences, and cures for the problem-seeking approach to social behavior and cognition. *Behavioral and Brain Sciences,* 27, 313–27.

Laland, K. (2001). Imitation, social learning, and preparedness as mechanisms of bounded rationality. In G. Gigerenzer and R. Selten (eds.), *Bounded Rationality: The Adaptive Toolbox.* Cambridge, MA: MIT Press.

Lopes, L. L. (1991). The rhetoric of irrationality. *Theory and Psychology,* 1, 65–82.

Martignon, L. and Hoffrage, U. (1999). Why does one-reason decision making work? A case study in ecological rationality. In G. Gigerenzer, P. M. Todd, and the ABC Research Group, *Simple Heuristics that Make Us Smart.* New York: Oxford University Press.

— (2002). Fast, frugal and fit: Lexicographic heuristics for paired comparison. *Theory and Decision,* 52, 29–71.

McLeod, P. and 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, 531–43.

Nagel, T. (1993). Moral luck. In D. Statman (ed.), *Moral Luck.* Albany, NY: State University of New York Press.

Nosofsky, R. M. (1990). Relations between exemplar-similarity and likelihood models of classification. *Journal of Mathematical Psychology,* 34, 393–418.

Pachur, T. and Hertwig, R. (2004). *How Adaptive is the Use of the Recognition Heuristic?* Poster session presented at the Annual Meeting of the Society for Judgment and Decision Making, Minneapolis, MN.
Payne, J. W., Bettman, J. R., and Johnson, E. J. (1993). *The Adaptive Decision Maker.* Cambridge: Cambridge University Press.
Piattelli-Palmarini, M. (1994). *Inevitable Illusions: How Mistakes of Reason Rule our Minds.* New York: Wiley.
Quinlan, J. R. (1993). *C4.5: Programs for Machine Learning.* Los Altos, CA: Morgan Kaufmann.
Rips, L. J. (1994). *The Psychology of Proof: Deductive Reasoning in Human Thinking.* Cambridge, MA: MIT Press.
— (2002). Circular reasoning. *Cognitive Science*, 26, 767–95.
Rosander, K. and Hofsten, C. von (2002). Development of gaze tracking of small and large objects. *Experimental Brain Research*, 146, 257–64.
Samuels, R., Stich, S., and Bishop, M. (2004). Ending the rationality wars: How to make disputes about human rationality disappear. In R. Elio (ed.), *Common Sense, Reasoning and Rationality.* New York: Oxford University Press.
Sargent, T. J. (1993). *Bounded Rationality in Macroeconomics.* New York: Oxford University Press.
Saxberg, B. V. H. (1987). Projected free fall trajectories: I. Theory and simulation. *Biological Cybernetics*, 56, 159–75.
Sedlmeier, P. and Gigerenzer, G. (2001). Teaching Bayesian reasoning in less than two hours. *Journal of Experimental Psychology: General*, 130, 380–400.
Selten, R. (2001). What is bounded rationality? In G. Gigerenzer and R. Selten, *Bounded Rationality: The Adaptive Toolbox.* Cambridge, MA: MIT Press.
Serwe, S. and Frings, C. (2004). *Predicting Wimbledon*: Poster presented at the 4th Summer Institute for Bounded Rationality in Psychology and Economics, Berlin.
Shaffer, D. M., Krauchunas, S. M., Eddy, M., and McBeath, M. K. (2004). How dogs navigate to catch frisbees. *Psychological Science*, 15, 437–41.
Simon, H. A. (1956). Rational choice and the structure of environments. *Psychological Review*, 63, 129–38.
— (1990). Invariants of human behavior. *Annual Review of Psychology*, 41, 1–19.
Smith, V. L. (2003). Constructivist and ecological rationality in economics. *The American Economic Review*, 93, 465–508.
Stigler, G. J. (1961). The economics of information. *Journal of Political Economy*, 69, 213–25.
Todorov, A. (2003). Cognitive procedures for correcting proxy-response biases in surveys. *Applied Cognitive Psychology*, 17, 215–24.
Trivers, R. L. (1971). The evolution of reciprocal altruism. *Quarterly Review of Biology*, 46, 35–57.
— (2002). *Natural Selection and Social Theory: Selected Papers of Robert Trivers.* New York: Oxford University Press.
Tversky, A. and Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, 90, 293–315.
Wald, A. (1950). *Statistical Decision Functions.* New York: Wiley.
Wason, P. C. (1966). Reasoning. In B. M. Foss (ed.), *New Horizons in Psychology.* London: Penguin Books.
Wason, P. C. and Johnson-Laird, P. N. (1972). *Psychology of Reasoning: Structure and Content.* Cambridge, MA: Harvard University Press.
Williams, B. (1981). *Moral Luck.* Cambridge: Cambridge University Press.

— (1988). Consequentialism and integrity. In S. Scheffler (ed.), *Consequentialism and its Critics.* New York: Oxford University Press. (Reprinted from B. Williams and J. J. C. Smart, *Utilitarianism: For and Against.* Cambridge: Cambridge University Press, 1973).

Wundt, W. (1912/1973). *An Introduction to Psychology* (R. Pintner, trans.). New York: Arno.

Zhu, L. and Gigerenzer, G. (forthcoming). Children can solve Bayesian problems: The role of representation in computation. *Cognition.*