Rationality for Mortals

How People Cope with Uncertainty

Gerd Gigerenzer

OXFORD
UNIVERSITY PRESS

2008

CONTENTS

Chapter 1

Bounded and Rational

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, weightlifting, 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 these 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 & Vonèche, 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 (a) constructing *as-if theories of unbounded rationality,*

This chapter is a revised version of G. Gigerenzer, "Bounded and Rational," in *Contemporary Debates in Cognitive Science*, ed. R. J. Stainton (Oxford, UK: Blackwell, 2006), 115–133.

by (b) constructing *as-if theories of optimization under constraints,* by (c) demonstrating *irrational cognitive illusions,* or by (d) studying *ecological rationality.* I have placed my bets on the last of these. Being engaged in the controversy, I am far from dispassionate but will be as impartial as I can.

Four Positions on 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 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 future). The mental ability to deduce the future from perfect knowledge requires *omnipotence*, or *unlimited computational power*. To be able to deduce the future with certainty implies that the structure of the world is *deterministic*. Omniscience, omnipotence, 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 story too literally and, in many a theory, re-created us in proximity to 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 by Daniel Bernoulli to the maximization of expected utility (chap. 10). In unbounded rationality, the three O's reign: *optimization* (such as maximization) replaces determinism, whereas the assumptions of omniscience and omnipotence 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 concerned only 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 those 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 (Gigerenzer, 2008). It underlies 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 model one or several constraints.

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 models 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 (1947) 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 of search 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 many economists and biologists (wrongly) tend to equate optimization under constraints with *bounded rationality,* the latter 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. Many sophisticated Bayesian models in cognitive science are of this kind, sacrificing the goal of modeling cognitive processes for that of applying an optimization model.

Cognitive Illusions: Logical Irrationality

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 & Tversky, 1996; Gilovich, Griffin, & Kahneman, 2002) opposes theories assuming that humans are basically rational. It has two 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. The cognitive processes underlying these errors are called heuristics, and the major three proposed are representativeness, availability, and anchoring and adjustment, with some new additions, including "affect." The program has produced a long list of biases. 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 & Tversky, 1982: 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 a law of logic (the "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 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.

The Science of Heuristics: Ecological Rationality

The starting point for the study of heuristics is the relation between mind and environment rather than between mind and logic (Gigerenzer, Todd, & the ABC Research Group, 1999; Gigerenzer & Selten, 2001a). Humans have evolved in natural environments, both social and physical. To survive and reproduce, 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 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" (Simon, 1990: 7). Just as one cannot understand how scissors cut by looking only at one blade, one will not understand human behavior by studying either 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, using heuristics rather than optimization models, 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 often unknown or *computationally intractable*.

A problem is computationally intractable if no mind or machine can find the optimal solution in reasonable time, such as a lifetime or a millennium. The game of chess is one example, where no computer or mind can determine the best sequence of moves. In order to be able to compute the optimal strategy, one could trim down the 8×8 board to a 4×4 one and reduce the number of pieces accordingly. Whether this result tells us much about the real game, however, is questionable.

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."

The science of heuristics has three goals, the first descriptive, the second normative, and the third of design.

The adaptive toolbox. The goal is to analyze the adaptive toolbox, that is, the heuristics, their building blocks, and the evolved capacities exploited by the building blocks. Heuristics should be specified in the form of computational models. This analysis includes the phylogenetic and ontogenetic development of the toolbox as well as cultural and individual differences.

Ecological rationality. The goal is to determine the environmental structures in which a given heuristic is successful, that is, the match between mind and environment (physical and social). This analysis includes the coevolution between heuristics and environments.

Design. The goal is to use the results of the study of the adaptive toolbox and ecological rationality to design heuristics and/or environments for improving decision making in applied fields such as health care, law, and management.

To see how this program differs from the cognitive illusions program, consider four general beliefs about heuristics that are assumed to be true in the cognitive illusions program but that turn out to be misconceptions from the point of view of the ecological rationality program (table 1.1). First, heuristics are seen as second-best approximations to the "correct" strategy defined by an optimization model; second and third, their use is attributed either to our cognitive limitations or to the fact that the problem at hand is not important; and finally, it is assumed that more information and more computation is always better if they are free of charge. I use an asset-allocation problem to demonstrate that, as a general truth, each of these beliefs is mistaken. Rather, one has to measure heuristics and optimization models with the same yardstick—neither is better per se in the real world.

Table 1.1: Four common but erroneous beliefs about heuristics

Misconception	Clarification
1. Heuristics produce second-best results; optimization is always better.	Optimization is not always the better solution, for instance, when it is computationally intractable or lacks robustness due to estimation errors.
2. Our minds rely on heuristics only because of our cognitive limitations.	We rely on heuristics for reasons that have to do with the structure of the problem, including computational intractability, robustness, and speed of action.
3. People rely or should rely on heuristics only in routine decisions of little importance.	People rely on heuristics for decisions of low and high importance, and this it not necessarily an error.
4. More information and computation is always better.	Good decision making in a partly uncertain world requires ignoring part of the available information and, as a consequence, performing less complex estimations because of the robustness problem. See investment example.

Investment Behavior

In 1990, Harry Markowitz received the Nobel Prize in Economics for his theoretical work on optimal asset allocation. He addressed a vital investment problem that everyone faces in some form or other, be it saving for retirement or earning money on the stock market: how best to invest your money in N assets. Markowitz proved that there is an optimal portfolio that maximizes the return and minimizes the risk. One might assume that when he made his own retirement investments he relied on his award-winning optimization strategy. But he did not. Instead he relied on a simple heuristic, the $1/N$ rule:

> *Allocate your money equally to each of* N *funds.*

There is considerable empirical evidence for this heuristic: About 50 percent of people studied rely on it, and most consider only about 3 or 4 funds to invest in. Researchers in behavioral finance have criticized this behavior as naïve. But how much better is optimizing than $1/N$? A recent study compared twelve optimal asset-allocation policies (including that of Markowitz) with the $1/N$ rule in seven allocation problems, such as allocating one's money to ten American industry portfolios. The twelve policies included Bayesian and non-Bayesian models of optimal choice. Despite their complexity, none could consistently beat the heuristic on various financial measures (DeMiguel, Garlappi, & Uppal, 2006).

How can a heuristic strategy be better than an optimizing one? At issue is not computational intractability, but robustness. The optimization models performed better than the simple heuristic in data fitting (adjusting their parameters to the data of the past ten years) but worse in predicting the future. Similar to the results that will be reported in the following chapters (figures 2.6 and 3.1), they thus overfitted the past data. The $1/N$ heuristic, in contrast, does not estimate any parameter and consequently cannot overfit.

Note that $1/N$ is not always superior to optimization. The important question of when in fact it predicts better can be answered by studying the rule's *ecological rationality*. Three relevant environmental features for the performance of $1/N$ and the optimizing models are:

(i) the predictive uncertainty of the problem,
(ii) the number N of assets, and
(iii) the size of the learning sample.

Typically, the larger the uncertainty and the number of assets and the smaller the learning sample, the greater the advantage of the heuristic. Since the uncertainty of funds is large and cannot be changed, we focus on the learning sample, which comprised 10 years of data in the above study. When would the optimization models begin to outperform the heuristic? The authors report that with 50 assets to allocate one's wealth to, the optimization policies would need a window of 500 years before it eventually outperformed the $1/N$ rule.

Note that $1/N$ is not only an investment heuristic. Its range is broader. For instance, $1/N$ is employed to achieve fairness in sharing among children and adults (dividing a cake equally), where it is known as the equality rule; it is the voting rule in democracies, where each citizen's vote has the same weight; it represents the modal offer in the ultimatum game; and it is a sibling of the tallying rules that will be introduced in chapter 2, where each reason is given the same weight. $1/N$ can achieve quite different goals, from making money to creating a sense of fairness and trust.

Markowitz's use of $1/N$ illustrates how each of the four general beliefs in table 1.1 can be wrong. First, the $1/N$ heuristic was better than the optimization models. Second, Markowitz relied on the heuristic not because of his cognitive limitations. Rather, as we have seen, his choice can be justified because of the structure of the problem. Third, asset allocations, such as retirement investments, are some of the most consequential financial decisions in one's life. Finally, the optimization models relied on more information and more computation than $1/N$, but that did not lead to better decisions.

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 independent of its content. For instance, the truth table of the material conditional *if P then Q* is defined independent 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?) of the problem. The program of studying whether people's judgments deviate from content-blind norms proceeds in four steps:

Syntax first. Start with a law of logic or probability.
Add semantics and pragmatics. Replace the logical terms (e.g., material conditional, mathematical probability) by English terms (e.g., if…then; probable), add content, and define the problem to be solved.
Content-blind norm. Use the syntax to define the "rational" answer to the problem. Ignore semantics and pragmatics.
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 conditional as a norm for reasoning about any content and the set-inclusion or "conjunction rule" (chap. 4). 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 how we should think (Piatelli-Palmarini, 1994:158). Consider the material conditional.

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 one of the most frequently studied tasks in the psychology of reasoning. Wason's starting point was the material conditional $P \rightarrow Q$, 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 conditional "$\rightarrow$" 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$\cap$not-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? My 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 "$\rightarrow$". 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's (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-card task whose content is a social contract between an employer and an employee (Gigerenzer & Hug, 1992):

> If a previous employee gets a pension from the firm, then that person must have worked for the firm for at least 10 years.

The four cards read: got a pension, worked 10 years for the firm, did not get a pension, worked 8 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 8 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 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 10 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 to avoid 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 here. 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 from these studies.The same holds for research on other content-blind norms (Gigerenzer, 2001). Inappropriate norms tend to suggest wrong questions, and the answers to these generate more confusion than insight into the nature of human judgment. 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 psychologists do learn. 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).

Rethinking Cognitive Biases

The above selection task illustrates the limits of logical norms for understanding good thinking. That is not to say that logic is never an appropriate norm, but rather that, like other analytical and heuristic tools, its domain is restricted. Violations of logical reasoning were previously interpreted as cognitive fallacies, yet what appears to be a fallacy can often also be seen as adaptive behavior, if one is willing to rethink the norms. More recently, a reevaluation of so-called cognitive biases that takes into account the structure of the environment and the goals of the decision maker has finally taken place. Table 1.2 illustrates a dozen cognitive illusions that are under debate. What unites these examples is the fact that as soon as researchers began to study the structure of information in the environment, an apparently dull cognitive illusion often took on the form of a sharp pair of scissors.

Consider the first item in the list, overconfidence bias, as an illustration. In a series of experiments, participants answered general-knowledge questions, such as:

Which city is farther north—New York or Rome?

How confident are you that your answer is correct?

50 percent / 60 percent / 70 percent / 80 percent / 90 percent / 100 percent

The typical finding was that when participants were 100 percent confident of giving a correct answer, the average proportion correct was lower, such as 80 percent; when they said they were 90 percent confident, the average proportion correct was 75 percent, and so on. This "miscalibration" phenomenon was labeled *overconfidence bias* and interpreted as a cognitive illusion. The explanation was sought in the minds of people who participated in the experiments, not in the environment. It was attributed to a confirmation bias in memory search: People first choose an answer, then search for confirming evidence only and grow overly confident. Yet Koriat, Lichtenstein, and Fischhoff's (1980) experiments showed only small or nonsignificant effects that disappeared in a replication (Fischhoff & MacGregor, 1982). Others proposed that people are victims of insufficient cognitive processing or suffer from self-serving motivational biases or from fear of invalidity. No explanation could be verified. In a social psychology textbook, the student was told: "Overconfidence is an accepted fact of psychology. The issue is what produces it. Why does experience not lead us to a more realistic self-appraisal?" (Myers, 1993: 50). Overconfidence bias was taken as the explanation for various kinds of personal and economic disasters, such as the large proportion of start-ups that quickly go out of business. As Griffin and Tversky (1992: 432) explained, "The significance of overconfidence to the conduct of human affairs can hardly be overstated." Finally, in a Nobel laureate's words, "some basic tendency toward overconfidence appears to be a robust human character trait" (Shiller, 2000: 142).

Eventually several researchers realized independent of each other that this phenomenon is a direct reflection of the *unsystematic* variability in the environment (Erev, Wallsten, & Budescu, 1994; Pfeiffer, 1994; Juslin, Winman, & Olsson, 2000). The large unsystematic variability of confidence judgments leads, *in the absence of any overconfidence bias,* to regression toward the mean, that is, the average number correct is always lower than a high confidence level. When one plots the data the other way round, the same unsystematic variability produces a pattern that looks like *underconfidence:* When participants answered 100 percent correctly, their mean confidence was lower, such as 80 percent, and so on (Dawes & Mulford, 1996). The phenomenon seems less a result of systematic cognitive bias and more a consequence of task environments with unsystematic error. Every unbiased mind and machine exhibits it.

To return to the initial question, which city is in fact farther north, New York or Rome? Temperature is a very good cue for latitude, but not a certain one. The correct answer is Rome. When researchers predominantly select

Table 1.2: Twelve examples of phenomena that were first interpreted as cognitive illusions (left) but later revalued as reasonable judgments given the environmental structure (right)

Is a phenomenon due to a "cognitive illusion"...	...or to an environmental structure plus an unbiased mind?
Overconfidence bias (defined as miscalibration)	"Miscalibration" can be deduced from an unbiased mind in an environment with unsystematic error, causing regression toward the mean (Dawes & Mulford, 1996; Erev et al., 1994).
Overconfidence bias (defined as mean confidence minus proportion correct)	"Overconfidence bias" can be deduced from an unbiased mind in an environment with unrepresentative sampling of questions; disappears largely with random sampling (Gigerenzer et al., 1991; Juslin et al., 2000).
Hard-easy effect	"Hard-easy effect" can be deduced from an unbiased mind in an environment with unsystematic error, causing regression toward the mean (Juslin et al., 2000).

Table 1.2 (continued)

Overestimation of low risks and underestimation of high risks	This classical phenomenon can be deduced from an unbiased mind in an environment with unsystematic error, causing regression toward the mean (Hertwig, Pachur, & Kurzenhäuser, 2005).
Contingency illusion	"Contingency illusion" can be deduced from an unbiased mind performing significance tests on samples with unequal sizes, such as minorities and majorities (Fiedler, Walther, & Nickel, 1999).
Most drivers say they drive more safely than average	The distribution of the actual number of accidents is highly skewed, which results in the fact that most drivers (80% in one U.S. study) have fewer than the average number of accidents (Lopes, 1992; Gigerenzer, 2002a).
Availability bias (letter "R" study)	"Availability bias" largely disappears when the stimuli (letters) are representatively sampled rather than selected (Sedlmeier, Hertwig, & Gigerenzer, 1998).
Preference reversals	Consistent social values (e.g., don't take the largest slice; don't be the first to cross a picket line) can create what look like preference reversals (Sen, 2002).
Probability matching	Probability matching is suboptimal for an individual studied in isolation but not necessarily for individuals in an environment of social competition (Gallistel, 1990).
Conjunction fallacy	"Conjunction fallacy" can be deduced from the human capacity for semantic inference in social situations (Hertwig & Gigerenzer, 1999).
False consensus effect	This "egocentric bias" can be deduced from Bayes's rule for situations where a person has no knowledge about prior probabilities (Dawes & Mulford, 1996).
Violations of logical reasoning	A number of apparent "logical fallacies" can be deduced from Bayesian statistics for environments where the empirical distribution of the events (e.g., P, Q, and their negations) is highly skewed (McKenzie & Amin, 2002; Oaksford & Chater, 1994) and from the logic of social contracts (Cosmides & Tooby, 1992).

The general argument is that an unbiased (not omniscient) mind plus a specific environmental structure (such as unsystematic error, unequal sample sizes, skewed distributions) is *sufficient* to produce the phenomenon. Note that other factors can also contribute to these phenomena. The moral is not that people would never err but that in order to understand good and bad judgments, one needs to analyze the structure of the problem or of the natural environment.

questions where a reliable cue fails (but do not inform experiment participants), the mean proportion correct will be lower than the mean confidence. This difference has also been called overconfidence, the second item in table 1.2, and attributed to people's mental flaws rather than to researchers' unrepresentative sampling. When researchers began to sample questions randomly from the real world (e.g., comparing all metropolises on latitude), this alleged cognitive illusion largely disappeared (see chap. 7).

Cognitive Luck

Matheson (2006) discusses the study of ecological rationality as a way to overcome the epistemic internalism of the Enlightenment tradition. But he raises a concern: "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." His 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 not always correct 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 a fundamental problem in our health systems, namely that a large number of physicians are innumerate (Gigerenzer, 2002a), as illustrated by screening for breast cancer. A woman with a positive mammogram asks the physician what the probability is that she actually has cancer. What do physicians tell that worried woman? In 2007, I asked 160 experienced gynecologists this question. To help them out, I gave them the relevant information, in the form of *conditional probabilities* (expressed as percentages).

> Assume that you screen women in a particular region for breast cancer with mammography. You know the following about women in this region:
>
> > The probability that a woman has breast cancer is 1 percent (prevalence).
> >
> > If a woman has breast cancer, the probability is 90 percent that she will have a positive mammogram (sensitivity).
> >
> > If a woman does not have breast cancer, the probability is 9 percent that she will still have a positive mammogram (false positive rate).
>
> A woman who tested positive asks if she really has breast cancer or what the probability is that she actually has breast cancer.
> What is the best answer?
>
> > (1) "It is not certain that you have breast cancer, yet the probability is about 81 percent." [14]

(2) "Out of 10 women who test positive as you did, about 9 have breast cancer." [47]
(3) "Out of 10 women who test positive as you did, only about 1 has breast cancer." [20]
(4) "The chance that you have breast cancer is about 1 percent." [19]

Note that the gynecologists' answers ranged between 1 percent and 9 out of 10 (90 percent)! The best answer is 1 out of 10, which only 20 percent of them gave. (The numbers in brackets give the percentage of gynecologists [out of 160] who chose each answer.) The most frequent answer was 9 out of 10. Consider for a moment the undue anxiety and panic women with positive mammograms have been caused by such physicians who do not understand the medical evidence.

In an earlier study with 48 physicians from various specialized fields (Hoffrage & Gigerenzer, 1998), we asked for numerical estimates (rather than multiple-choice selection), with similar results. Once again, the estimates ranged between 1 percent and 90 percent. One-third of the physicians thought the answer was 90 percent, one-third gave estimates between 50 percent and 80 percent, and one-third between 1 percent and 10 percent. Physicians' intuitions could hardly vary more—a worrying state of affairs.

This result illustrates a larger problem: When physicians try to draw a conclusion from conditional probabilities, their minds tend to cloud over (chap. 9). What can be done to correct this? From an internalist perspective, one might recommend training physicians how to insert the probabilities into Bayes's rule. Yet this proposal is doomed to failure. When we taught students statistics in this way, their performance dropped by 50 percent just one week after they successfully passed the exam and continued to fade away week by week (Sedlmeier & Gigerenzer, 2001). Moreover, the chance of convincing physicians to take a statistics course in the first place is almost nil; most have no time, little motivation, or believe they are incurably innumerate. Are innumerate physicians then inevitable? No. In the ecological view, thinking does not happen simply in the mind, but in interaction between the mind and its environment. This opens up a second and more efficient way to solve the problem: to change the environment. The relevant part of the environment is the representation of the information, because the representation does part of the Bayesian computation. Natural (nonnormalized) frequencies are such an efficient representation; they mimic the way information was encountered before the advent of writing and statistics, throughout most of human evolution. Here is the same information as above, now in *natural frequencies:*

10 out of every 1,000 women have breast cancer.

Of these 10 women, we expect that 9 will have a positive mammogram.

Of the remaining 990 women without breast cancer, some 89 will still have a positive mammogram.

Imagine a sample of women who have positive mammograms. How many of these women actually have cancer? _____ out of _____ .

When I presented the numerical information in natural frequencies, the confusion in most physicians' minds disappeared; 87 percent of the gynecologists chose "1 out of 10." Most realized that out of some 98 [89 + 9] women who test positive, only 9 are likely to have cancer. Thus, the chances of having breast cancer based on a positive screening mammogram are less than 10 percent, or about 1 in 10. Proper representation of information, such as natural frequencies, helps physicians to understand the outcomes of medical tests and treatments (see also Elmore & Gigerenzer, 2005) and prevents needless shocks to wrongly informed patients. In 2006, this program of teaching transparent risk communication became part of continuing education for gynecologists in Germany; I myself have trained some one thousand physicians in using representations that turn innumeracy into insight (see chap. 9).

Similarly, by changing the environment, we can make many so-called cognitive illusions largely disappear, enable fifth and sixth graders to solve Bayesian problems before they even heard of probabilities (chap. 12), and help judges and law students understand DNA evidence (Hoffrage, Lindsey et al., 2000). Thus, an ecological view actually extends the possibilities to improve judgment, whereas an internalist view limits the chances. To summarize, worrying about "cognitive luck" is bound to an internalist view, where enablers outside the mind are considered suspicious. From an ecological view, environmental structures, not luck, naturally and inevitably influence the mind and can be designed to enable insight. Cognitive virtue is, in my view, a relation between a mind and its environment, very much like the notion of ecological rationality.

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 even the cognitive sciences. Both have an antipsychological bias: They try to define rational behavior without cognitive psychology, promoting as-if theories, which illustrates that "black box" behaviorism is still alive. 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, omnipotence, 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 conclusion has been that we are mostly or sometimes 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 lies in the middle and that we are a little of both, so there is no real disagreement. For instance, the debate between Kahneman and Tversky (1996) and myself (Gigerenzer, 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, Stich, and Bishop (2004: 264) 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" (but see Bishop, 2000). 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 or have not violated the truth-table logic or some other logical law in an experimental task. 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 that 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 artificial 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 counterintuitive 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.