

EVOLUTION AND COGNITION

RATIONALITY *for* MORTALS

HOW PEOPLE COPE WITH UNCERTAINTY



Gerd Gigerenzer

Press

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior permission of Oxford University Press.

Library of Congress Cataloging-in-Publication Data
Gigerenzer, Gerd.

Rationality for mortals : how people cope with uncertainty / Gerd Gigerenzer.

p. cm. — (Evolution and cognition)

Includes bibliographical references and index.

ISBN 978-0-19-974709-2

1. Reasoning (Psychology) 2. Thought and thinking.

I. Title.

BF442.G54 2008

153.4'3—dc22 2007029644

9 8 7 6 5 4 3 2 1

Printed in the United States of America

CONTENTS

- [1. Bounded and Rational](#)
- [2. Fast and Frugal Heuristics](#)
- [3. Rules of Thumb in Animals and Humans](#)
- [4. I Think, Therefore I Err](#)
- [5. Striking a Blow for Sanity in Theories of Rationality](#)
- [6. Out of the Frying Pan into the Fire](#)
- [7. What's in a Sample? A Manual for Building Cognitive Theories](#)
- [8. A 30 Percent Chance of Rain Tomorrow](#)
- [9. Understanding Risks in Health Care](#)
- [10. The Evolution of Statistical Thinking](#)
- [11. Mindless Statistics](#)
- [12. Children Can Solve Bayesian Problems](#)
- [13. In the Year 2054: Innumeracy Defeated](#)

[Bibliography](#)

[Name Index](#)

[General Index](#)

Rationality for Mortals

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*, 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.

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.

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 consequentialist 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 is 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 $P \wedge \text{not-Q}$.

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 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).

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).
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.

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.

Chapter 2

Fast and Frugal Heuristics

If you open a book on judgment and decision making, chances are that you will stumble over the following moral: Good reasoning must adhere to the laws of logic, the calculus of probability, or the maximization of expected utility; if not, there must be a cognitive or motivational flaw. Don't be taken in by this fable. Logic and probability are mathematically beautiful and elegant systems. But they do not always describe how actual people—including the authors of books on decision making—solve problems, as the subsequent story highlights. A decision theorist from Columbia University was struggling whether to accept an offer from a rival university or to stay. His colleague took him aside and said, “Just maximize your expected utility—you always write about doing this.” Exasperated, the decision theorist responded, “Come on, this is serious.”

The study of heuristics investigates how people actually make judgments and decisions in everyday life, generally without calculating probabilities and utilities. The term *heuristic* is of Greek origin and means “serving to find out or discover.” In the title of his Nobel Prize–winning paper of 1905, Albert Einstein used the term *heuristic* to indicate an idea that he considered incomplete, due to the limits of our knowledge, but useful (Holton, 1988). For the Stanford mathematician George Polya (1954), heuristic thinking was as indispensable as analytical thinking for problems that cannot be solved by the calculus or probability theory—for instance, how

to find a mathematical proof. The advent of computer programming gave heuristics a new prominence. It became clear that most problems of any importance are computationally intractable; that is, we know neither the optimal solution nor a method to find it. This holds even for well-defined problems, such as chess, the classic computer game Tetris, and the traveling salesman problem (Michalewicz & Fogel, 2000). It also holds for less well-structured problems, such as which job offer to accept, what stocks to invest in, and whom to marry. When optimal solutions are out of reach, we are not paralyzed to inaction or doomed to failure. We can use heuristics to discover good solutions.

This chapter is a revised version of G. Gigerenzer, “Fast and Frugal Heuristics: The Tools of Bounded Rationality,” in *Blackwell Handbook of Judgment and Decision Making*, ed. D. J. Koehler and N. Harvey (Oxford, UK: Blackwell, 2004), 62–88.

What Is a Heuristic?

How does a baseball outfielder catch a fly ball? He might compute the trajectory of the ball and run to the point where it is supposed to land. How else could he do it? In Richard Dawkins’s words (1976/1989: 96):

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. He may neither know nor care what a differential equation is, but this does not affect his skill with the ball. At some subconscious level, something functionally equivalent to the mathematical calculations is going on.

Note that Dawkins carefully inserts the qualifier “as if.” To compute the trajectory is no simple feat; no computer program or robot to date can compute it in real time. What about an experienced player? First, we might assume that the player intuitively knows the family of parabolas, because, in theory, balls have parabolic trajectories. In order to select the right parabola, the player needs to be equipped with sensory organs that can measure the ball’s initial distance, initial velocity, and projection angle. Yet in the real world, influenced by air resistance, wind, and spin, balls do not fly in parabolas. Thus, the player would further need to be capable of estimating the speed and direction of the wind at each point of the ball’s flight, in order to compute the resulting path and the point where the ball will land, and to then run there. All this would have to be completed within a few seconds—the time a ball is in the air. This explanation is based on the ideals of *omniscience* and *omnipotence*: To solve a complex problem, a person constructs a complete representation of its environment and relies on the most sophisticated computational machinery.

An alternative vision exists, which does not aim at complete representation and information. It poses the question: Is there a smart heuristic that can solve the problem? One way to discover heuristics is to study experienced players. Experimental studies have shown that players actually use several heuristics (e.g., McLeod & Dienes, 1996). The simplest one is the *gaze heuristic*, which works if the ball is already high up in the air:

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

The angle of gaze is the angle between the eye and the ball, relative to the ground (figure 2.1). A player who uses this heuristic does not need to estimate wind, air resistance, spin, or the other causal variables. He can get away with ignoring every piece of causal information. All the relevant information is contained in one variable: the angle of gaze. Note that a player using the gaze heuristic is not able to compute the point at which the ball will land. But the player will be there where the ball lands.

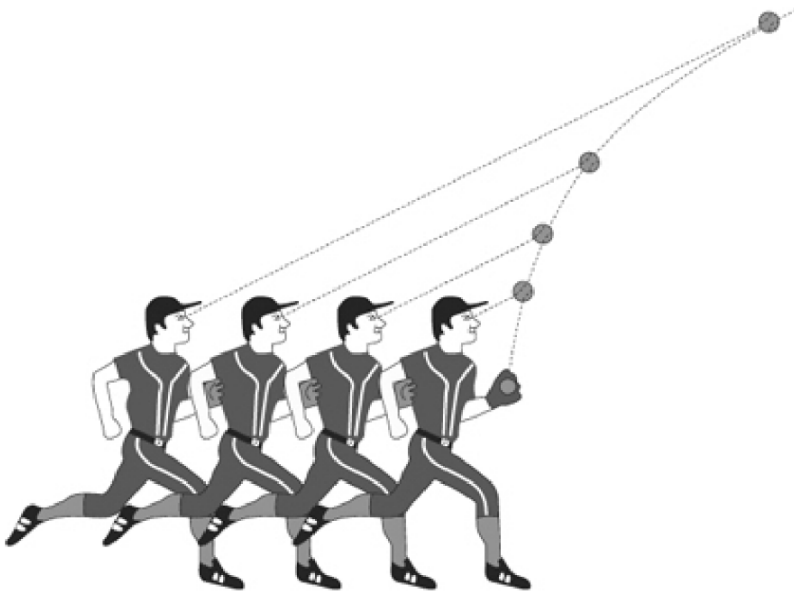


Figure 2.1: How to catch a fly ball? Players rely on unconscious rules of thumb. When a ball comes in high, a player fixates his gaze on the ball, starts running, and adjusts the speed so that the angle of gaze remains constant.

The gaze heuristic is a fast and frugal heuristic. It is fast because it can solve the problem within a few seconds, and it is frugal because it requires little information, just the angle of gaze. The heuristic consists of three building blocks: fixate your gaze on the ball, start running, and adjust your running speed. These building

blocks can be part of other heuristics, too.

Definition: A fast and frugal heuristic is a strategy, conscious or unconscious, that searches for minimal information and consists of building blocks that exploit evolved capacities and environmental structures.

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

Heuristics exploit evolved capacities. A heuristic is *simple* because it can take advantage of the evolved or learned capacities of an organism. For example, it is easy for humans to track a moving object against a noisy background; three-month-old babies can already hold their gaze on moving targets (Rosander & von Hofsten, 2002). Tracking objects, however, is difficult for a robot; a computer program as capable as a human mind of solving this problem does not yet exist. Similarly, in contrast to robots, humans are able to run. Thus, the gaze heuristic is simple for humans but not for robots. Simplicity is not only a characteristic of beauty; it also enables *fast, frugal, transparent, and robust* judgments. The gaze heuristic, like all heuristics, is transparent in the sense that it can be easily understood and taught to a novice, and the term *robust* refers to the ability of heuristics to generalize to new situations (see below). To summarize, a heuristic exploits hard-wired or learned cognitive and motor processes, and these features make it simple.

Heuristics exploit structures of environments. The rationality of heuristics is not logical, but ecological. Ecological rationality implies that a heuristic is not good or bad, rational or irrational per se, only relative to an environment. It can exploit particular environmental structures or change an environment. For instance,

the gaze heuristic transforms the complex trajectory of the ball in the environment into a straight line. All heuristics are to some degree domain-specific; they are designed to solve specific classes of problems. The gaze heuristic can solve problems that involve the interception of moving objects. If you learn to fly an airplane, you will be taught a version of it: When another plane is approaching, and you fear a collision, then look at a scratch in your windshield and observe whether the other plane moves relative to that scratch. If it does not, dive away quickly. For the pilot, the goal is to avoid a collision, whereas for the outfielder, the goal is to produce a collision. The nature of the heuristic is the same. To summarize, evolved capacities can make a heuristic simple, while the structure of the environment can make it smart.

Heuristics are distinct from as-if optimization models. The idea of calculating the ball's trajectory by solving differential equations is a form of optimization. When optimization is proposed to explain human behavior (as opposed to building artificial systems), this is called *as-if optimization*. As-if optimization models are silent about the actual process, although it is sometimes suggested that the measurements and calculations might happen unconsciously. The gaze heuristic, however, illustrates that the logic of a heuristic, conscious or unconscious, can be strikingly distinct from as-if optimization. This yields an advantage. With a good model of a heuristic, one can deduce predictions that cannot be obtained from an as-if optimization model. The gaze heuristic, for instance predicts that players catch the ball while running, which follows from the fact that the player must move to keep the angle of gaze constant. Similarly, when the ball is thrown to the side of the player, one can predict that the player will run a slight arc, as can be observed in baseball outfielders and in dogs who catch Frisbees

(e.g., Shaffer & McBeath, 2002). In summary, a model of a heuristic is a rule whose purpose is to describe the actual process—not merely the outcome—of problem solving.

Models of Heuristics

A model of a heuristic specifies (i) a process rule, (ii) the capacities that the rule exploits to be simple, and (iii) the kinds of problems the heuristic can solve, that is, the structures of environments in which it is successful.

Models of heuristics need to be distinguished from mere labels. For instance, terms such as *representativeness* and *availability* are commonsense labels without specification of a process and the conditions under which a heuristic succeeds and fails. These need to be developed into testable models; otherwise they can account for almost everything post hoc.

There already exist a number of testable models for heuristics, such as satisficing (Selten, 2001; Simon, 1982), elimination-by-aspects (Tversky, 1972), and various heuristics for multiattribute choice discussed in Payne, Bettman, & Johnson 1993. Much of this earlier work addressed heuristics for preferences, not for inferences, that is, for problems where no single external criterion of success exists. Criteria for the accuracy of heuristics were typically internal, such as whether they used all of the information or how closely they mimicked the gold standard of a weighted additive model. Because there were no external criteria for accuracy, the true power of heuristics could not be fully demonstrated.

I focus on heuristics for inferences—such as comparative judgments, classification, and estimation. From the seminal work on heuristics with simple unit weights (+1 and -1; see Dawes, 1979),

we know that the predictive accuracy of simple heuristics can be as high as or higher than that of the gold standard of weighing and adding. For instance, unit weights matched multiple regression in predicting the academic performance of students (Dawes & Corrigan, 1974), and the take-the-best heuristic was as successful as Bayes's rule at predicting the outcomes of basketball games in the 1996 NBA season, but it did so faster and with less information (Todorov, 2002). Models of heuristics for classification, estimation, comparative judgments, and choice are discussed in Gigerenzer, Todd, and the ABC Research Group 1999, and Gigerenzer and Selten 2001b. In what follows, I will select a few heuristics and discuss their ecological rationality and the empirical evidence.

Recognition Heuristic

Imagine you are a contestant in a TV game show and face the \$1 million question: "Which city has more inhabitants, Detroit or Milwaukee?"

What is your answer? If you are American, then your chances of finding the right answer, Detroit, are not bad. Some 60 percent of undergraduates at the University of Chicago did (Goldstein & Gigerenzer, 1999). If, however, you are German, your prospects look dismal because most Germans know little about Detroit, and many have not even heard of Milwaukee. How many correct inferences did the less knowledgeable German group that we tested make? Despite a considerable lack of knowledge, virtually all of the Germans answered the question correctly. How can people who know less about a subject nevertheless make more correct inferences? The answer is that the Germans used a fast and frugal heuristic, the recognition heuristic: If you recognize the name of one city but not the other, then infer that the recognized city has

the larger population. The Americans could not use the heuristic, because they had heard of both cities. They knew too much.

The recognition heuristic is useful when there is a strong correlation—in either direction—between recognition and criterion. For simplicity, let us assume that the correlation is positive. For two-alternative choice tasks, the heuristic can be stated as follows:

Recognition heuristic: If one of two objects is recognized and the other is not, then infer that the recognized object has the higher value with respect to the criterion.

The recognition heuristic builds on an evolved capacity for recognition—such as face, voice, and name recognition. No computer program yet exists that can perform face recognition as well as a human child. Note that the capacity for recognition is different from that for recall. For instance, one may recognize a face but not recall anything about who that person is. If people use the recognition heuristic in an adaptive way, they will rely on it in situations where it is ecologically rational.

Ecological Rationality: The recognition heuristic is successful when ignorance is systematic rather than random, that is, when recognition is strongly correlated with the criterion.

The direction of the correlation between recognition and the criterion can be learned from experience, or it can be genetically coded. Substantial correlations exist in competitive situations, such as between name recognition and the excellence of colleges, the value of companies' products, and the quality of sports teams. One way to measure the degree of ecological rationality of the

recognition heuristic (the correlation between recognition and criterion) is the *recognition validity* α , which is the proportion of times a recognized object has a higher criterion value than an unrecognized object in a reference class, such as cities, companies, or sports teams:

$$\alpha = R / (R + W), \tag{2.1}$$

where R is the number of correct (right) inferences the recognition heuristic would achieve, computed across all pairs in which one object is recognized and the other is not, and W is the number of incorrect (wrong) inferences, computed under the same circumstances.

The research summarized in [Table 2.1](#) suggests that people use the recognition heuristic in a relatively adaptive way, that is, it is followed most consistently when the recognition validity is high. For instance, Pohl (2006, Exp. 1) used the 20 largest Swiss cities and asked one group of participants to judge which of two cities has the larger population and another group to judge which of the two cities was located farther from the Swiss city Interlaken. Recognition is valid for inferring population ($\alpha = .86$) but not for inferring distance ($\alpha = .51$). Participants were not informed about validity. Nevertheless, they intuitively followed the recognition heuristic for inferring population in 89 percent of all cases, compared to in only 54 percent for inferring distance, which is almost at chance level. The correlation between recognition validity and proportion of judgments consistent with the recognition heuristic in [Table 2.1](#) is $r = .55$.

Table 2.1: People tend to rely on the recognition heuristic in situations where the recognition validity is substantially above

chance (.5).

Task	Reference	Recognition validity α	Consistent with recognition heuristic
Career point total of NHL hockey players	Snook & Cullen 2006	.87	96%
Winner of 2003 Wimbledon tennis matches	Serwe & Frings 2006	.73 (amateurs) .67 (laypeople)	93% 88%
Winner of 2005 Wimbledon tennis matches	Scheibehenne & Bröder 2007	.71 (amateurs) .69 (laypeople)	89% 79%
Winner of 2004 European Soccer Championship matches	Pachur & Biele 2007	.71	91%
Population of German cities	Goldstein & Gigerenzer 2002	.80	90%
Population of Swiss cities	Pohl 2006, Exp. 1 and 2	.86 (Exp. 1) .72 (Exp. 2)	89% 75%
Distance of Swiss cities from Interlaken	Pohl 2006, Exp. 1	.51	54%
Population of European cities	Pohl 2006, Exp. 3	.89 (Italian) .82 (Belgian)	88% 89%
Largest mountains, rivers, and islands	Pohl 2006, Exp. 4	.49 (mountains) .74 (rivers) .85 (islands)	89% 94% 81%
Population of European cities	Volz et al. 2006	.63	84%
Prevalence of infectious diseases	Pachur & Hertwig 2006	.62 (Study 1) .62 (Study 2)	62% 69%

For instance, when judging which of two NHL hockey players has the higher career point total ($\alpha = .87$), participants followed the recognition heuristic in 96% of the cases. All studies tested inferences from memory (not from givens) with recognition that is ecologically valid (rather than defined as objects presented in a previous experimental session, as in many memory tasks).

The adaptive toolbox perspective implies two processes that precede the use of the recognition heuristic: recognition and evaluation. To be able to use the heuristic, one alternative must be recognized and the other not. To decide whether to use the heuristic (as opposed to another strategy) in a given situation, an evaluation process is enacted. Consistent with this hypothesis, a neuroimaging study (Volz et al., 2006) showed that different brain

regions were activated for mere recognition judgments (“Which of the two cities do you recognize?”) compared to tasks that allow the recognition heuristic to be used (“Which city has the larger population?”). Specifically, a deactivation was observed within the anterior frontomedian cortex (aFMC) when people did not follow the recognition heuristic. Since the aFMC has been previously associated with self-referential judgments, this result indicates that the recognition heuristic is the default, but that contradicting source knowledge can inhibit its use.

Knowledge that seems to inhibit the use of the recognition heuristic includes (i) low recognition validity (see above); (ii) recognizing individual objects for reasons that have nothing to do with the criterion, such as recognizing Chernobyl because of its nuclear power accident (Oppenheimer, 2003); and (iii) direct criterion knowledge for the recognized object, such as when comparing the population of the town around the corner with a known small population to that of an unknown city. Reliance on the recognition heuristic seems to be largely maintained even in the presence of contradicting cue information. For instance, the median participant in Richter and Späth (2006, Exp. 3) judged a recognized city as larger than an unrecognized one in 100 percent of the cases when they were told that the recognized city had an international airport, 97 percent when they were given no information, and 97 percent when they were told that the recognized city had no such airport. Contradicting cue information has an effect on some participants, but the majority in this experiment abided by the default: “go with what you know.”

The recognition heuristic should not be confused with *availability* (Tversky & Kahneman, 1974). Availability refers to ease of recall, not recognition. The recognition heuristic implies several

counterintuitive phenomena that cannot be deduced from any other theory I am aware of. As mentioned before, recognition information tends to dominate contradictory clues, in rats as well as in people, even if there is conflicting evidence (Pachur, Bröder, & Marewski, in press). Next, I will deduce a counterintuitive phenomenon, the *less-is-more effect*, and the conditions under which it occurs.

The Less-Is-More Effect

Equation 2.2 specifies the proportion of correct answers c on an exhaustive test of all pairs of N objects (such as cities, soccer teams) for a person who recognizes n of these objects.

$$c = \frac{2n(N-n)}{N(N-1)}\alpha + \frac{(N-n)(N-n-1)}{N(N-1)}\frac{1}{2} + \frac{n(n-1)}{N(N-1)}\beta \quad (2.2)$$

The three terms on the right side of the equation correspond to the three possibilities: A person recognizes one of the two objects, none, or both. The first term accounts for the correct inferences made by the recognition heuristic, the second term for guessing, and the third term equals the proportion of correct inferences made when knowledge beyond recognition is used. The *knowledge validity* b is the relative frequency of getting a correct answer when both objects are recognized, which is computed like the recognition validity. All parameters in equation 2.2 can be independently measured.

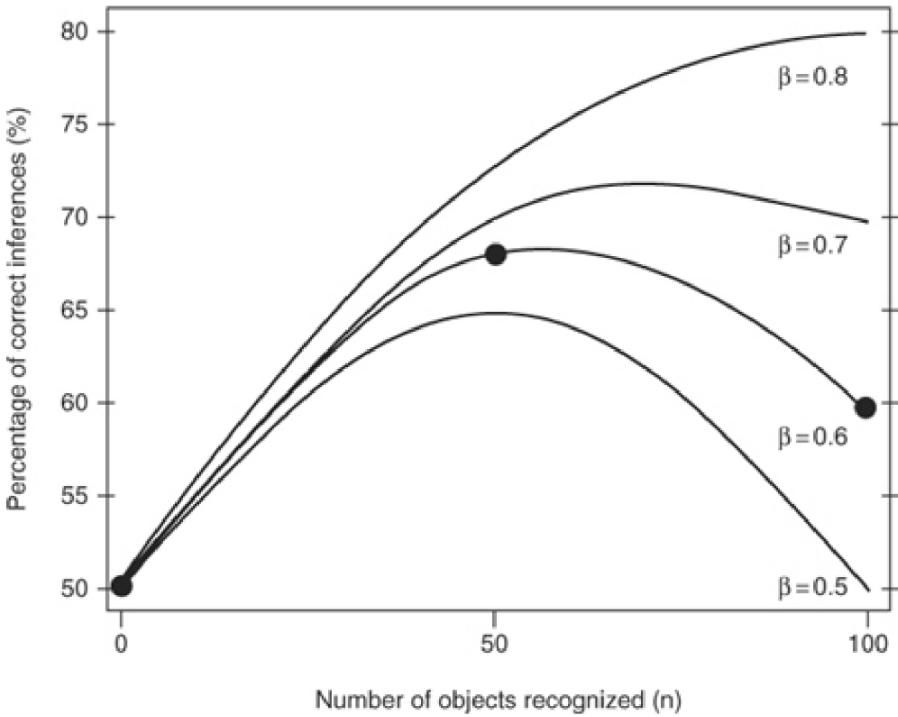


Figure 2.2: The less-is-more effect is a consequence of the recognition heuristic. It occurs when the recognition validity α is larger than the knowledge validity β (and α and β are constant). The curves shown are for $\alpha = .8$. A less-is-more effect can occur between people with the same β , as shown by the middle and right-hand point. It can also occur between people with different knowledge validities. For instance, a person who recognizes only half of the objects ($n = 50$) and has no useful knowledge ($\beta = .5$) will nevertheless make more correct inferences than a person who recognizes all objects ($n = 100$) and has useful knowledge ($b = .6$).

When one plots equation 2.2, a counterintuitive implication can be seen (figure 2.2). Consider first the curve for $\beta = .5$, that is, for people who have no predictive knowledge beyond recognition. A person who has heard of none of the objects will perform at chance level (50 percent, left side). A person who has heard of all objects

will also perform at chance level (50 percent, right side). Only a person who has heard of some but not all objects can use the recognition heuristic, and their accuracy will first increase with n but then decrease again. The reason is that the recognition heuristic can be used most often when about half of the objects are recognized, in comparison to when all or none are recognized. When half of the objects are recognized, a person can use the recognition heuristic about half of the time, which results in some 65 percent (40 percent for $\alpha = .8$ plus 25 percent for guessing) correct inferences, as can be calculated from equation 2.2. The next curve with three dots shows a less-is-more effect in the presence of knowledge beyond mere recognition, for $\beta = .6$. The left dot represents a person who has not heard of any objects, while the dot on the right represents someone who has heard of all objects and has recall knowledge that does better than chance. The middle dot represents a person who recognizes less objects but gets more correct inferences. In general, assuming that α and β are constant, the following result can be proven (Goldstein & Gigerenzer, 2002):

Less-is-more effect: The recognition heuristic will yield a less-is-more effect if $\alpha > \beta$.

A less-is-more effect can emerge in at least three different situations. First, it can occur between two groups of people, when a more knowledgeable group makes worse inferences than a less knowledgeable group in a given domain. An example is the performance of the American and German students on the question of whether Detroit or Milwaukee is larger. Second, a less-is-more effect can occur between domains, that is, when the same group of people achieve higher accuracy in a domain in which they know little than in a domain in which they know a lot. For instance, when

American students were tested on the 22 largest American cities (such as New York versus Chicago) and on the 22 largest German cities (such as Cologne versus Frankfurt), they scored a median 71 percent (mean 71.1 percent) correct on their own cities but slightly higher on the less familiar German cities, with a median of 73 percent correct (mean 71.4 percent). This effect was obtained despite a handicap: Many Americans already knew the three largest U.S. cities in order and did not have to make any inferences. A similar less-is-more effect was demonstrated with Austrian students, whose scores for correct answers were slightly higher for the 75 largest American cities than for the 75 largest German cities (Hoffrage, 1995; see also Gigerenzer, 1993a). Third, a less-is-more effect can occur during knowledge acquisition, that is, when an individual's performance curve first increases but then decreases again.

Less-Is-More in Groups

Consider now group decision making. Three people sit in front of a computer screen on which such questions as “Which city has more inhabitants: Milan or Modena?” are displayed. The task of the group is to find the correct answer through discussion, and they are free to use whatever means. In this task, the correct solution is difficult to “prove” by an individual group member; thus, one might expect that the majority determine the group decision (the *majority rule*; see Gigone & Hastie, 1997). Consider now the following conflict. Two group members have heard of both cities, and each concluded independently that city *A* is larger. But the third group member has not heard of *A*, only of *B*, and concludes that *B* is larger (relying on the recognition heuristic). After the three members finished their negotiation, what will their consensus be? Given that two members

have at least some knowledge of both cities, one might expect that the consensus is always *A*, which is also what the majority rule predicts. In fact, in more than half of all cases (59 percent), the group voted for *B* (Reimer & Katsikopoulos, 2004). This number rose to 76 percent when two members relied on mere recognition.

That group members let their knowledge be dominated by others' lack of recognition may seem odd. But in fact this apparently irrational decision increased the overall accuracy of the group. This result can be analytically deduced and intuitively seen from [figure 2.2](#). When the recognition heuristic is used in group decisions, a less-is-more effect results if $\alpha > \beta$, just as in [figure 2.2](#), but more strongly. Consistent with the theory, Reimer and Katsikopoulos (2004) observed that when two groups had the same average α and β , the group that recognized *fewer* cities (smaller n) typically had *more* correct answers. For instance, the members of one group recognized on average only 60 percent of the cities, and those in a second group 80 percent; but the first group got 83 percent answers correct in a series of more than one hundred questions, whereas the second got only 75 percent. Thus, group members seem to intuitively trust the recognition heuristic, which can improve accuracy and lead to the counterintuitive less-is-more effect between groups.

Heuristics Based on Reasons and Imitation

When recognition is not valid, or people know too much, heuristics can involve the search for reasons or cues. A few years after his voyage on the *Beagle*, the 29-year-old Charles Darwin divided a scrap of paper (titled “This Is the Question”) into two columns with the headings “Marry” and “Not Marry” and listed supporting reasons for each of the two possible courses of action, such as “nice

soft wife on a sofa with good fire” opposed to “conversation of clever men at clubs.” Darwin concluded that he should marry, writing “Marry—Marry—Marry Q. E. D.” decisively beneath the first column (Darwin, 1887/1969: 232–233). The following year, Darwin married his cousin, Emma Wedgwood, with whom he eventually had 10 children. How did Darwin decide to marry, based on the possible consequences he envisioned—children, loss of time, a constant companion? He did not tell us. But we can use his “Question” as a thought experiment to illustrate various visions of decision making.

Darwin searched his memory for reasons. There are two visions of search: optimizing search and heuristic search. Following Wald’s (1947) optimizing models of sequential analysis, several psychological theorists postulated versions of sequential search and stopping rules (e.g., Busemeyer & Townsend, 1993). In the case of a binary hypothesis (such as to marry or not marry), the basic idea of most sequential models is the following: A threshold is calculated for accepting one of the two hypotheses, based on the costs of the two possible errors, such as wrongly deciding that to marry is the better option. Each reason or observation is then weighted, and the evidence is accumulated until the threshold for one hypothesis is met, at which point the search is stopped and the hypothesis is accepted. If Darwin had followed this procedure, he would have had to estimate, consciously or unconsciously, how many conversations with clever friends are equivalent to having one child and how many hours in a smoky abode can be traded against a lifetime of soft moments on the sofa. Weighting and adding is a mathematically convenient assumption, but it assumes that there is a common currency for all beliefs and desires in terms of quantitative probabilities and utilities. These models are often