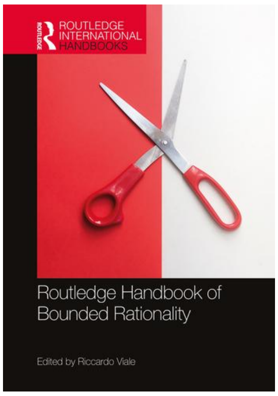

Routledge Handbook of Bounded Rationality

Riccardo Viale

What is bounded rationality?

Publication details
https://test.routledgehandbooks.com/doi/10.4324/9781315658353-2
Gerd Gigerenzer
Published online on: 29 Oct 2020

2
WHAT IS BOUNDED RATIONALITY?

Gerd Gigerenzer

Bounded rationality has, surprisingly, three faces. The original is by Herbert A. Simon, who coined the term. The other two arrived later and bent Simon's term into meanings that could hardly be more different. Equally puzzling, the two new meanings contradict each other, each having absorbed the term into their own framework. This double take-over has been so subtle that few people notice that, when talking about bounded rationality, they are often talking about different things. The three faces of bounded rationality are not simply a matter of terminology; they reflect fundamentally dissimilar research programs and visions about the nature of human rationality. In this chapter, I will describe the three faces, explain how we got there, and outline how to develop Simon's original program.

Simon's bounded rationality

Herbert A. Simon (1916–2001) coined the term *bounded rationality*. While still a student, fresh out of a price theory class at the University of Chicago, he tried to apply the perspective of utility maximization to budget decision problems in his native Milwaukee's recreation department in the mid-1930s. To his surprise, he learned that managers did not try to compare the marginal utility of a proposed expenditure with its marginal costs, but simply added incremental changes to last year's budget and engaged in other rules of thumb. During that concrete encounter, Simon realized what managers could and could not measure, and concluded that the framework of utility maximization "was hopeless" (Simon 1988, p. 286). This discrepancy between theory and reality marked the beginning of what he later called the study of bounded rationality: "Now I had a new research problem: How do human beings reason when the conditions for rationality postulated by the model of neoclassical economics are not met?" (Simon, 1989, p. 377).

What is rationality in neoclassical economics? Although not all economists agree, they typically refer to three pillars: consistency, maximization of expected utility, and—if learning is involved—Bayesian updating of probabilities. Leonard J. Savage (1954), known as the father of modern Bayesian decision theory, defined two conditions necessary for these three pillars of rationality:

1. *Perfect Foresight of Future States*: The agent knows the exhaustive and mutually exclusive set S of future states of the world.
2. *Perfect Foresight of Consequences*: The agent knows the exhaustive and mutually exclusive set C of consequences of each of his or her actions, given a state.

Savage called the pair (S, C) a *small world*. States and consequences must necessarily be described at some limited level of detail—hence the adjective *small*. The prototype of a small world is a lottery where all possible future states (the tickets) along with all possible outcomes (the payoffs) and their probabilities are known for certain, or a game such as roulette where all states (in roulette: the 36 numbers plus zero), consequences, and probabilities are known. In a small world, nothing new or unexpected can ever happen. These conditions have been variously called unbounded rationality, full rationality, constructivist rationality, or the Bayesian rationality approach. What Simon noticed, however, was that the managers he observed and humans in general mostly have to deal with situations that are unlike small worlds with perfect foresight and that do not meet the assumptions for expected utility maximization and Bayesian updating. For instance, if the exhaustive and mutually exclusive set of future states of the world and their consequences is not known, one cannot maximize expected utility or assign prior probabilities to states that add up to one, which makes Bayesian updating unfeasible. Simon called for an empirical study of *how* humans reason when perfect foresight is not possible. In his lifetime, most economists showed little interest in Simon's question and preferred theories assuming perfect foresight.

Risk *and* uncertainty

I will use the term *risk* for situations in which agents have perfect foresight of future states and their consequences, as defined above, and also certain knowledge about the probabilities of the states. For the many situations in which this is not the case, I will use the term *uncertainty*. This distinction between risk and uncertainty goes back to Frank Knight (1921). *Uncertainty* is sometimes used to denote a small world without known probabilities (Luce & Raiffa, 1957), a situation known as *ambiguity*. Yet ambiguity contains only a minor degree of uncertainty; it still assumes perfect foresight of future states and consequences. When speaking of uncertainty, I also refer to situations that do not meet the definition of small worlds. Uncertainty is sometimes called *fundamental uncertainty* or *radical uncertainty*. It includes ambiguity and intractability (such as in chess), ill-defined problems such as budget decisions where optimal solutions cannot be known because perfect foresight is an illusion, and, in general, all future events that are not perfectly foreseeable. In short, uncertainty includes most of the interesting problems that humans face in real life.

Now we can formulate the first principle of Simon's bounded rationality program: *to study how human beings make decisions under uncertainty, not only under risk.*

Decision making under uncertainty is obviously relevant for understanding how managers, business people, or other human beings make choices. In contrast, decision making under risk is rare, mostly restricted to human-made environments such as lotteries, games, and gambling. Nevertheless, lotteries remain the stock-in-trade of decision research, defining risk aversion, regret, and loss aversion, and are the domain of modifications of expected utility theory, such as prospect theory.

Savage himself explicitly limited his theory, currently known as *Bayesian decision theory*, to small worlds only. For instance, he pointed out that his theory does not apply to planning a picnic or playing a game of chess (Savage, 1954, p. 16). When planning a picnic, one cannot know ahead all future states that might happen. And when playing chess, where a finite number of states (sequences of moves) indeed exists, the problem is intractability because no human or machine can enumerate the exhaustive set of possible moves and determine the optimal one. As a consequence, one cannot assign probabilities to these states that add up to one. To illustrate, chess has approximately 10^{120} different unique sequences of moves or games, which is a number greater than the estimated number of atoms in the universe. Savage believed it would be "utterly ridiculous" to apply rational choice theory to situations without perfect foresight (1954, p. 16). Both Simon and Savage were fully in agreement that expected utility maximization is a local but not a universal tool.

As-if *and* real decision-making processes

Simon's question is about the *process* of decision making: "How do human beings reason?" Simon criticized the fact that neo-classical economists showed surprisingly little interest in studying how people actually make decisions, but instead engaged in armchair speculation and relied mainly on their intuition. When confronted with the same kind of evidence as Simon once was, namely, that business professionals did not behave as postulated by utility theory, Milton Friedman reacted differently: He dismissed the empirical evidence as irrelevant. In his *as-if* defense, Friedman (1953) famously argued that the goal of "positive economics" is prediction, not psychological realism. The theory says that people behave *as if* they maximized their expected utility, not that they actually do. All that counts are sufficiently good predictions. I do not doubt that expected utility theory and its variants can explain behavior *after* the fact, due to its great flexibility; whether it can actually predict behavior is less clear. True predictions must be out-of-sample, not by fitting a model to known data. After reviewing 50 years of research for evidence about how well utility functions, such as utility of income functions, utility of wealth functions, and the value function in prospect theory actually predict behavior, D. Friedman, Isaac, James, and Sunder (2014) concluded: "Their power to predict out-of-sample is in the poor-to-nonexistent range, and we have seen no convincing victories over naïve alternatives" (p. 3).

Now we can formulate the second principle of Simon's bounded rationality program: *To study how human beings make decisions, as opposed to relying on as-if expected utility theories.*

In my opinion, Simon's insistence on studying the process and Friedman's goal of prediction are not incompatible opponents; they can be reconciled. My hypothesis is that by constructing realistic process models, one can, on average, make better predictions than by using as-if models. I will provide evidence for this argument below using the example of the prediction of customer purchases.

But what are the tools humans use to make good decisions under uncertainty? Knight (1921) had proposed "judgment" and "experience," but not much more. In his *General Theory*, John Maynard Keynes (1936) proposed "animal spirits," without telling us what these exactly are. In their book *Animal Spirits*, Akerlof and Shiller (2009) remedied this by moving a step forward and distinguishing five such spirits—confidence, fairness, corruption, money illusion, and stories. The authors argue that these "restless and inconsistent" elements were the principal reasons for the financial crisis of 2008 (pp. vii, 4).

Simon (1955, 1956) made a more concrete start, proposing that humans make decisions under uncertainty by relying on aspiration levels, limited search, and heuristics such as *satisficing*. In its simplest version, the satisficing rule is:

Step 1: Set an aspiration level α.
Step 2: Choose the first option that satisfies α.

To illustrate, consider a study of how real estate entrepreneurs decide in which location to invest in order to develop a new commercial high-rise or a residential area. Berg (2014) reports that of 49 professional investors, each and every one relied on a version of satisficing: *If I believe I can get at least* x *return within* y *years, then I take the option*. Here, "*x* return in *y* years" is the aspiration level. In general, an aspiration is a goal, and an aspiration level is a goal value. The psychologist Kurt Lewin (1935), who promoted the concept of aspiration, considered successful people to be those who set goals that are within their ability to reach. Using a satisficing heuristic, people can deal with uncertainty because the rule does not require perfect knowledge about all options, such as possible locations, marriage partners, or professions. How did Simon decide to study political science and economics? In his own account, he "simply picked the first profession that sounded fascinating" (1978, p. 1).

In general, satisficing can handle situations where the exhaustive and mutually exclusive set of options is not knowable, and where options need to be searched one by one. In such situations, optimizing is not an option.

Behavior = f (cognition, environment)

The third principle that defines Simon's bounded rationality is captured by his scissors analogy: behavior is a function of both cognition and environment (Simon, 1990, p. 7): "Human rational behavior (and the rational behavior of all physical symbol systems) is shaped by a scissors whose two blades are the structure of task environments and the computational capabilities of the actor."

Simon did not flesh out the scissors principle in detail; I will provide an example below in the section on ecological rationality. In general, it requires an analysis of both the heuristics people use and the structure of the environments in which they act. Economic theory analyzes the incentive structure in environments but not the actual decision strategies, assuming as-if utility maximization. Thus, it deals with only one blade. Similarly, much of psychological theorizing restricts itself to cognitive processes and is mute about environmental structure. Thus, it deals solely with the other blade. Debates in psychology tend to center around "internal" questions: whether the mind is a Bayesian or not, whether information is integrated in an additive or multiplicative way, whether the mind works by mental models or mental logic, or whether a person is risk averse or risk seeking. By analyzing only one blade of Simon's scissors, one cannot understand the ecological rationality of a heuristic, a belief, or other strategies.

To summarize, there are three principles that define Simons' program of bounded rationality:

1. *Uncertainty.* To study decision making under uncertainty, not only risk.
2. *Process*. To study the actual process of decision making, as opposed to as-if expected utility maximization.
3. *Scissors.* To study how the structure of an environment, together with the cognitive process, produces the resulting behavior.

For many of Simon's contemporaries, this program was asking too much. It clashed with the practice of reducing all uncertainties to risks by assuming perfect foresight, such as by converting real-world problems into lotteries. It also clashed with the ideal of explaining all behavior after the fact by saying that it maximizes some utility, regardless of Savage's small-world restriction. Needless to say, Simon's revolution did not happen during his lifetime.

The rejection of novel ideas is nothing new in the history of science. However, the reaction to Simon's ideas was rather unusual. Instead of being rejected as worthless or unimportant, his concept of *bounded rationality* was hijacked and radically reframed, making his revolutionary ideas no longer recognizable. Neo-classical economists absorbed the term into the orthodoxy of perfect foresight and concluded that satisficing is quintessentially optimizing. Proponents of the heuristics-and-biases program in psychology, in contrast, appropriated the term for their own focus on human lack of rationality. This is why bounded rationality now has three faces that could hardly be more dissimilar.

Bounded rationality as optimization under constraints

Satisficing involves search. Simon (1955, 1956) was one of the first to raise the topic of search in the economics literature. Note that in situations of risk, search is not an issue, given that all alternatives are already known. However, it was Stigler (1961) who popularized the topic by changing satisficing search to optimal search. In his classical example of the purchase of a second-hand car, the buyer is assumed to stop searching when the costs of further search exceed its benefits. Yet to factor costs and benefits into the equation, the buyer needs to know all alternatives in the first place, in addition to the costs of search. Stigler's approach poured search theory back into the old bottle of expected utility maximization, introducing perfect foresight once again, and reducing uncertainty to risk. With this move, satisficing is equivalent to optimizing with the constraint of search costs. In fact, Simon (1955) had ventured in the same direction earlier. There now exists a large literature in economics in which satisficing is modeled as optimization under constraints, mostly in the expected utility maximization tradition, whereas models of satisficing under uncertainty are rare. In a memorial book for the late Simon, Arrow (2004, p. 48) argued that "boundedly rational procedures are in fact fully optimal procedures when one takes account of the cost of computation in addition to the benefits and costs inherent in the problem as originally posed." Bounded rationality is thereby seen as nothing but constrained optimization in disguise.

In 1993, Thomas Sargent published his book *Bounded Rationality in Macroeconomics*, where he re-interprets "the idea of bounded rationality as a research program to build models populated by agents who behave like working economists or econometricians" (1993, p. 22). As he points out, by adding more and more constraints to optimization, models of bounded rationality become larger and mathematically more demanding, which is why there is no rush among econometricians to implement such models. Econometricians try to reduce the number of parameters, and that "is not what bounded rationality promises" (p. 5). In this view, bounded rationality makes optimization more difficult, and unbounded rationality ends up as the simpler and tractable alternative. This interpretation turned the idea of simple heuristics such as satisficing upside down. In personal conversation, Simon playfully remarked that he had considered suing authors who misused the term *bounded rationality* to mean optimization under constraints (Gigerenzer, 2004). In 1998, Ariel Rubinstein published *Modeling Bounded Rationality* and admirably included some of Simon's comments (Rubinstein, 1998). Simon criticized Rubinstein for showing no awareness of psychological research apart from that of Kahneman and Tversky. Rubinstein responded that, unlike some other economists, he was concerned neither with predicting behavior nor with giving normative advice. Rather, "by modeling bounded rationality, we try to examine the logic of a variety of principles that guide decision makers, especially within interactive systems (markets and games)" (1998, p. 191). But these logical armchair analyses were exactly what Simon was battling.

During his lifetime, Simon was highly influential in many fields, from artificial intelligence to psychology to political science. The major exception was in economics, the field in which he was awarded the Nobel Memorial Prize. In his own assessment, his program of bounded rationality was received with "something less than unbounded enthusiasm" and "largely ignored as irrelevant for economics" (Simon, 1997, p. 269). I believe that economics has missed an important call to extend its theoretical and methodological toolbox.

Bounded rationality as irrationality

Beginning in the 1970s, psychologists within the heuristics-and-biases program ventured to show that the assumptions of neo-classical economic theory are descriptively incorrect: the consistency axioms, the maximization of expected utility, and the updating of probabilities by Bayes' rule (e.g., Tversky & Kahneman, 1974). These assumptions are also known as rational choice theory. This program produced a long list of discrepancies between theory and actual behavior. A discrepancy could be due to flaws either in the theory or in behavior. Unlike Simon, the program maintained rational choice theory as normatively correct and blamed the discrepancies on people's lack of rationality. Its proponents have even defended utility theory as a universal norm against critics such as Lola Lopes (e.g., Tversky & Bar-Hillel, 1983, p. 713). Accepting rational choice theory as a universal norm and attributing deviations from this norm to people's lack of rationality is clearly not what Simon had in mind. "Bounded rationality is not irrationality" (Simon, 1985, p. 297).

Nevertheless, the heuristics-and-biases program also appropriated the term *bounded rationality* for itself. "Our research attempted to obtain a map of bounded rationality, by exploring the systematic biases that separate the beliefs that people have and the choices they make from the optimal beliefs and choices assumed in rational agent models" (Kahneman, 2003, p. 1449). Now we have a third meaning of the term: human errors, defined as whenever human judgment *deviates from rational choice theory*. That contradicts both Simon's original meaning and economists' re-interpretation of the term as economic rationality. Moreover, these alleged biases were attributed to people's use of heuristics. As a consequence, the term *heuristic* became associated with bias, contrary to a long tradition in psychology and computer science and to the definitions listed in the *Oxford English Dictionary* (n.d.).

The observation that people sometimes violate rational choice theory eventually led to *dual process theories* that "explain" this apparent flaw by assuming an error-prone "System 1" that intuitively uses heuristics and a "System 2" that operates according to rational choice theory. So far, these systems have been characterized by lists of general dichotomies (heuristic versus logical, unconscious versus conscious, and so on), without precise models of the processes, such as formal models of heuristics. Moreover, aligning "heuristic" with "unconscious" and "biases" is not even correct; every heuristic I have studied can be used both consciously and unconsciously, and can lead to better or worse decisions than what would be considered "rational" (Kruglanski & Gigerenzer, 2011).

How could bounded rationality come to be equated with people's irrationality? A key to finding an answer is found in the classical papers of Kahneman and Tversky, which are reprinted in the anthology by Kahneman, Slovic, and Tversky (1982). In these articles, neither Simon nor bounded rationality is cited. But in the Preface to the anthology, Simon is briefly mentioned, apparently more as a nod to a distinguished figure than an acknowledgment of a significant intellectual link (Lopes, 1992). Thus, the re-interpretation of bounded rationality as irrationality likely occurred as an afterthought.

Does the heuristics-and-biases research implement Simon's three principles? First, like neo-classical theory, the program does not make the distinction between risk and uncertainty. For most problems investigated, it assumes that logic or probability theory provides the correct answer. With respect to the study of the process of decision making, it has made some contributions in its work on heuristics, but these are all associated with biases. Moreover, the heuristics proposed in the 1970s have never been defined, so that their predictions could be tested but instead were used to explain deviations between theory and behavior after the fact, as with the availability heuristic in the original letter-R study (Tversky & Kahneman, 1973). When availability was later more clearly defined and its predictions tested for the very same study, it was shown not to predict people's judgment (Sedlmeier, Hertwig, & Gigerenzer, 1998). Hindsight is easier than foresight. The use of labels instead of models differs from the earlier work by Tversky (1972) on elimination-by-aspect. In terms of the scissors analogy, the heuristics-and-biases program focused solely on internal explanations of behavior, without an analysis of the environmental conditions under which a heuristic is likely to succeed in reaching a goal or not. In this program, the use of heuristics is the problem and the supposed cause of errors. For decisions under uncertainty, however, heuristics are the solution, not the problem.

Homo heuristicus, *Homo economicus*, and Homer Simpson

The three faces of bounded rationality are not complementary, but conflicting. Let us give each of the three faces a name. *Homo heuristicus* personifies Simon's original program to study decision making under uncertainty. *Homo economicus* stands for the neo-classical program assuming situations of risk. Finally, Homer Simpson is an apt name for the program studying the deviations of people's behavior from the ideal of *Homo economicus*, attributing these to flaws in people rather in the theory (the name is from Thaler & Sunstein, 2008).

These three meanings of bounded rationality are not the only ones in circulation. Every time I enter the US to give a talk, the immigration officer asks me about the topic of my talk. Usually I respond by giving a two-sentence summary of bounded rationality. One officer nodded and remarked: "Oh, that's when you get old!" Another one got the point and said: "Oh, uncertainty, that's what I do, trying to make sense of people. Would you please write down the title of your book?" Most of the time, however, explaining bounded rationality just speeds up the process: "Oh, thanks, just go through."

I end this chapter with a short introduction to the program of ecological rationality, the modern extended version of Simon's bounded rationality.

The ecological rationality program: *Homo heuristicus*

Simon's question was descriptive: What are the tools humans use to make decisions under uncertainty? The program of ecological rationality starts from that question and extends it to three altogether, a descriptive, prescriptive, and engineering one (Gigerenzer, Hertwig, & Pachur, 2011; Gigerenzer & Selten, 2001). The descriptive question concerns the repertoire of tools that individuals and institutions rely on to make decisions under uncertainty. Simon suggested one such tool, satisficing. This repertoire is called the *adaptive toolbox*. The term *toolbox* signifies that there is more than one tool, and *adaptive* that the tools are adapted to specific classes of problems, just as a hammer is intended for nails, and a screwdriver for screws. The prescriptive questions fleshes out the scissors analogy: What environmental structures can a heuristic exploit in order to make better decisions? This question is addressed using analysis and computer simulation. The results to the first two questions provide answers to the engineering question: How can simple, intuitive, and efficient decision systems be designed? This is the study of *intuitive design* (not covered here; see Gigerenzer et al., 2011).

The ecological rationality program is based on three methodological principles:

1. *Algorithmic models of heuristics.* Early candidates of heuristics such as availability and representativeness were largely undefined. Algorithmic models of heuristics such as satisficing replace vague labels.
2. *Competitive testing*. Models of heuristics should be tested against the best competing models, not against a null hypothesis of no effect.
3. *Test of predictions, not of data fitting.* Prediction is about the future, or about data that have not yet been observed. Fitting takes place when the data have already been observed and the parameters of a model are chosen so that they maximize the fit (such as R^2). Fitting alone is not a proper test of a model: The more free parameters are added to a model, the better it fits the data, but this does not necessarily hold for prediction, which involves uncertainty. In fact, simple heuristics can outperform more complex strategies in prediction.

The adaptive toolbox

The study of the adaptive toolbox of an individual, an organization, or a species proceeds by observation and experiment. It aims at algorithmic models of the heuristics that people actually use under uncertainty rather than aiming at as-if models of expected utility maximization. Models of bounded rationality describe not merely the outcome of a judgment or decision but also how it is reached (that is, the heuristic processes or proximal mechanisms). They also describe the building blocks of heuristics, which allow heuristics to be combined or extended. Consider once again the satisficing heuristic. In situations where one is unsure about the proper aspiration level, the basic satisficing rule can be extended to *satisficing with adaptation*:

Step 1: Set an aspiration level α.
Step 2: Choose the first option that satisfies α.
Step 3: If after time β no option has satisfied α, then change α by an amount γ and continue until an option is found.

The first two steps are identical to the basic form of satisficing, but the third allows the aspiration level to be adapted if no option is found after a certain amount of time.

To illustrate, consider the question of how car dealers should price used cars. There is no unique answer in economic theories. Building on Stigler's (1961) optimization under constraint theory, some theories assume that prices should be fine-tuned to the changes in market conditions, such as supply and demand; others assume two kinds of customers, informed shoppers who seek the lowest price and naïve customers who go for a brand name, reputation, or other surrogates for quality. To address this variability among customers, dealers should respond with mixed strategies where prices are randomized, just as some supermarkets change prices in a quasi-random way. An analysis of 628 German car dealers and over 16,000 cars showed, however, that none of these theories predicted the actual prices (Artinger & Gigerenzer, 2016). There was virtually no fine-tuning, despite drastic changes in supply, and randomizing prices was also absent.

How then do dealers set prices? In Artinger and Gigerenzer's (2016) study, almost all of the dealers (97 percent) used a satisficing heuristic. Nineteen percent of the dealers relied on the basic satisficing heuristic: They chose an initial price (the aspiration level), typically below the average price on the market, and kept it constant until the car was sold. The majority (78 percent) of dealers relied on satisficing with adaptation. Typically, they chose an initial price in the middle of the price range on the market and, if the car was not sold after about four weeks, lowered the price, and so on (Artinger & Gigerenzer, 2016). Do dealers leave money on the table by relying on satisficing? A comparison with the economic models mentioned above indicated that both versions of satisficing used by the dealers led to higher, in fact, more than double the profit. Heuristics can perform well under uncertainty.

The study of the adaptive toolbox has documented several classes of heuristics besides satisficing. One class consists of one-reason heuristics, which rely on a single powerful cue and ignore all others (see below). Another class contains lexicographic heuristics, including fast-and-frugal trees and take-the-best. Lexicographic rules also base the final decision on one reason only, but unlike one-reason heuristics, they may search through several reasons before finding the decisive one. A fourth class consists of social heuristics, designed for coordination, competition, and co-operation. An overview of these heuristics can be found in Gigerenzer and Gaissmaier (2011) and Gigerenzer et al. (2011)

Ecological rationality of heuristics

The goal of the study of ecological rationality is to determine the match between heuristics and environment, that is, the structure of environments that a given class of heuristics can exploit (Todd, Gigerenzer, & the ABC Research Group, 2012; Gigerenzer & Selten, 2001). As a first approximation, "a heuristic is ecologically rational to the degree it is adapted to the structure of the environment" (Gigerenzer & Todd, 1999, p. 13). In his Nobel lecture, Smith (2003) used this definition and generalized it from heuristics to markets. I will illustrate the study of ecological rationality with the prediction of customer behavior.

A large apparel retailer sends special offers and targeted information to customers. The retailer's goal is to target offers at "active" customers who are likely to make purchases in the future. The overall aim is to remove inactive, unprofitable customers from the customer base; identify profitable, inactive customers who should be reactivated; and determine active customers who should be targeted with regular marketing activities. But how can the marketing department predict which of their previous customers will make a purchase in the future? In the tradition of optimizing models, stochastic customer base models such as the Pareto/NBD model (NBD stands for "negative binomial distribution") have been proposed as models for repeat purchase behavior. These models try to fine-tune predictions by making complex assumptions about individual behavior and the heterogeneity across customers, assuming that customers buy at a steady albeit stochastic rate until they eventually become inactive (see Wübben & von Wangenheim, 2008). When experienced managers from a large apparel retailer were studied, however, it was found that they did not use these complex models. Instead, the managers relied on a simple rule:

> Hiatus heuristic: *If a customer has not made a purchase for nine months or longer, classify them as inactive, otherwise as active.*

The hiatus can vary, depending on the kind of business, but the decision is based on a single reason, the recency of purchase. All other cues, such as the frequency of purchase, are ignored. The hiatus heuristic is an instance of the class of *one-reason heuristics* (Gigerenzer & Gaissmaier, 2011).

Why would experienced managers rely on only a single reason? The traditional heuristics-and-biases view might conclude that relying on one reason is naïve and may be due to managers' limited cognitive capacities. From this viewpoint, managers should better use stochastic customer base models such as the Pareto/NBD model. But note that the managers are in a situation of uncertainty, not in a small world of risk. The methodology outlined above—algorithmic models of heuristics, competitive testing, and tests of predictions—explains how to perform a competitive test between the complex Pareto/NBD model and the simple hiatus heuristic. The result of such a test showed that the Pareto/NBD model did predict future purchases for 75 percent of the customers correctly, but the hiatus heuristic did so for 83 percent (Wübben & von Wangenheim, 2008). Is this apparel company an exception? Subsequent tests showed otherwise: In 60 further tests, the hiatus heuristic was, on average, better at predicting behavior than were several stochastic customer base models. Or are these stochastic customer models substandard? In the same 60 tests, the hiatus heuristic on average also better predicted behavior in comparison with logistic regression and with random forests, one of the most powerful machine learning techniques (Artinger, Kozodi, Wangenheim, & Gigerenzer, 2018). Under uncertainty, a single powerful reason can lead to more accurate, more transparent, and less costly predictions than those made by highly complex machine learning algorithms.

This existence proof of the predictive power of the heuristic does not yet tell us *when* the heuristic will succeed or not. After all, it predicted better on average, but not in every individual case. What environmental structures does the hiatus heuristic exploit when making accurate predictions with little effort? The study of the ecological rationality of heuristics provides answers to such questions (Gigerenzer et al., 2011). To illustrate, I briefly describe a condition for the ecological rationality of the hiatus heuristic, which also holds for similar one-reason heuristics.

Instead of the hiatus heuristic, consider now a linear strategy that uses and combines more cues. It uses *n* binary cues $x_1, \ldots, x_n$, with values of either +1 or -1, where the positive value indicates future purchases. The weights of the cues are $w_1, \ldots, w_n$, all of which are positive:

$$y = w_1x_1 + w_2x_2 + w_3x_3 + \ldots + w_nx_n$$

The linear rule makes the inference "active" if $y > 0$, otherwise "inactive." Linear algorithms are standard approaches to prediction, which I use as reference point for a comparative analysis with one-reason heuristics. Let us give the single cue used by the heuristic the number 1, and the remaining cues the numbers 2, …., n. Like the beta weights in a logistic regression, the weights of each of the remaining cues reflect their *additional* contribution to the higher ranked cues. If the following *dominant cue* condition holds, one can show that the linear rule cannot lead to more accurate inferences than a one-reason heuristic:

> *Dominant Cue Condition.* The weights $w_1, w_2, w_3, \ldots w_n$ form a dominant cue structure if they satisfy the inequality constraint:

$$w_1 > \sum_{i=2}^{n} w_i$$

In words, the condition is that the weight of the first cue is larger than the sum of the weights of all other cues. The weights 1, ½, ¼, and ⅛ are an example. If this condition holds, the linear rule will always make the same inference as the hiatus heuristic. For instance, if the dominant cue has a positive value (such as "made a purchase in the last nine months"), the hiatus heuristic leads to the inference that the customer is "active," as does the linear algorithm, whatever the values on the remaining cues are—even if all are negative. The reason is that a dominant cue cannot be overturned or compensated by the sum of all lower-ranking cues. In this situation, despite more effort, the linear strategy cannot be more accurate than the heuristic. If instead the weights were, say, 1, ¾, ½, and ¼, the dominant cue condition would not hold. For a more systematic treatment of the ecological rationality of heuristics, see Gigerenzer (2016), Martignon and Hoffrage (2002), Martignon, Katsikopoulos, and Woike (2008) and Şimşek (2013).

This analysis makes clear that the rationality of a heuristic cannot be evaluated by looking at the heuristic and concluding that it is overly simple, but only by looking at its match with the environment. Dominant cues and even noncompensatory cues (a stronger condition where dominance must hold for every cue relative to the lower-ranked ones, not only for the top cue) appear to be the rule rather than the exception in many real-world problems (Şimşek, 2013). The dominant cue condition explains when it is rational to rely on a single cue and ignore the rest. At a more general level, which I cannot cover here, the bias–variance dilemma explains in what situations simple heuristics can predict better than more complex strategies, and vice versa (Gigerenzer & Brighton, 2009).

Ecological rationality of beliefs

The study of ecological rationality is equally relevant when evaluating the rationality of beliefs. Beliefs have often been evaluated using only logic or probability theory as the gold standard, assuming situations of risk instead of uncertainty. As a result, beliefs that are correct under uncertainty have been mistaken as systematic biases and cognitive illusions (Gigerenzer, 2015; 2018; Gigerenzer, Fiedler, & Olsson, 2012). Consider intuitions about chance, where one tradition of experimental studies concludes that people in general have good intuitions (e.g., Kareev, 1992; Piaget & Inhelder, 1975) and where researchers in the heuristics-and-bias program have diagnosed systematic misconceptions. Consider a classical example.

> A fair coin is thrown four times. Which of the two strings of heads (H) and tails (T) is more likely to be encountered?
> HHH
> HHT

Most people think that HHT is more likely. This belief has been deemed a fallacy, based on the argument that each of the two strings has the same probability of occurring. The alleged fallacy was attributed to people's illusory belief in the "law of small numbers," that is, that people expect the equal probability of H and T to hold in small samples as well: HHT is more "representative" than HHH (Kahneman & Tversky, 1972).

Surprisingly, however, this popular belief is actually correct if a person observes a sequence of n coin tosses that is longer than the length k of the string (here, $k = 3$). For instance, for $n = 4$ four tosses, there are 16 possible sequences, each equally likely. Yet 4 of those contain at least one HHT, and only 3 an HHH (see Figure 2.1; Hahn & Warren, 2009). Similarly, the expected waiting time for encountering HHT is 8 tosses of a coin, compared with 14 tosses for HHH. Now we can specify the general condition under which people's belief is *ecologically rational*:

> If $k < n$, a string of Hs with a single alternation such as HHT is more likely to be encountered than a pure string such as HHH.

The belief is only a fallacy in the special case of $k = n$, or with an infinite sample, which corresponds to the population probability (Hahn & Warren, 2009). Moreover, the ecological analysis

H	H	H	H	H	H	H	H	T	T	T	T	T	T	T	T
H	H	H	H	T	T	T	T	H	H	H	H	T	T	T	T
H	H	T	T	H	H	T	T	H	H	T	T	H	H	T	T
H	T	H	T	H	T	H	T	H	T	H	T	H	T	H	T
✓	✓+	+	+	–	–	–	–	✓	+	–	–	–	–	–	–

Figure 2.1 A fair coin is flipped four times, and the result is recorded. H = heads, T = Tails. Which string of flips is more likely: at least one HHH or HHT? Most people believe that HHT is more likely, whereas some researchers say that both strings are equally likely and that people's intuition are therefore biased. In fact, to encounter HHT is more likely than HHH. There are 16 possible outcomes, each equally likely; four of these have strings of HHT ("cross"), but only three strings of HHH ("check mark"). The general principle is: If the number n of tosses (here n = 4) is higher than the size k of the string (k = 3), then pure strings of heads (or tails) are less likely than those with an alteration. In this situation, people's intuition is correct (Hahn & Warren, 2009).

clarifies that the belief that HHT is more likely than HHH is not a case of the gambler's fallacy, as often assumed. The gambler's fallacy refers to the intuition that after witnessing a string of, say, two heads, one expects that the next outcome will be more likely tail than head. This would be a true fallacy because it corresponds to the condition $k = n$. In other words, a total of three throws is considered, either HHH or HHT, and there is no sample k with the property $k < n$ (Gigerenzer, 2018).

The general point is that the statistics of a sample are not always the same as the probabilities of the population. By the same kind of ecological argument, it could be shown that the belief in the hot hand in basketball is not a fallacy, as claimed by Gilovich, Vallone, and Tversky (1985). Rather, a reanalysis of the original data showed that the belief is correct but that the researchers overlooked the difference between sample statistics and population probabilities and used the wrong gold standard (Miller & Sunjuro, 2016). Labeling a belief a bias even if no bias exists is an example of a broader phenomenon that I call the "bias bias": the tendency to see systematic biases in behavior even when there is no verifiable error at all (Gigerenzer, 2018).

The study of ecological rationality is a general approach for evaluating the rationality of heuristics and beliefs. It differs from logical rationality in that it does not compare behavior with some principle of logic or probability theory that is taken to be universally true, assuming that the situation is always one of risk. Rather, heuristics and beliefs should be evaluated against the structure of the environment, where small samples may systematically differ from a world of risk.

Guidelines for the study of decision making under uncertainty

Let me end with a summary in the form of three theoretical guidelines for studying bounded rationality:

1. *Take uncertainty seriously.* Theories of human behavior should take the distinction between risk and uncertainty seriously. What is rational under risk is not necessarily rational under uncertainty, and what is a biased belief under risk is not necessarily biased under uncertainty. In situations of risk, fine-tuned behavior is likely to be adaptive, such as continuous Bayesian probability updating. In situations of uncertainty, by contrast, simple heuristics are likely to be adaptive.

2. *Take heuristics seriously*. Under uncertainty, heuristics are not the problem but the solution. We need more studies of the adaptive toolbox of individuals and institutions and of their development over time.
3. *Take ecological rationality seriously*. Study the ecological rationality of heuristics, beliefs, and other behavior instead of their logical rationality alone.

References

Akerlof, G. A., & Shiller, R. J. (2009). *Animal spirits*. Princeton, NJ: Princeton University Press.

Arrow, K. J. (2004). Is bounded rationality unbounded rational? Some ruminations. In M. Augier & J. G. March (Eds.), *Models of a man: Essays in memory of Herbert A. Simon* (pp. 47–56). Cambridge, MA: MIT Press.

Artinger, F. M., & Gigerenzer, G. (2016). The cheap twin: From the ecological rationality of heuristic pricing to the aggregate market. In J. Humphreys (Ed.), *Proceedings of the Seventy-sixth Annual Meeting of the Academy of Management*. Online ISSN: 2151–6561. doi:10.5465/ambpp.2016.206.

Artinger, F. M., Kozodi, N., Wangenheim, F., & Gigerenzer, G. (2018). Recency: Prediction with smart data. In J. Goldenberg, J. Laran, & A. Stephen (Eds.), *American Marketing Association Winter Conference Proceedings:* Vol. 29. *Integrating paradigms in a world where marketing is everywhere* (L2–L6). Chicago: American Marketing Association.

Berg, N. (2014), Success from satisficing and imitation: Entrepreneurs' location choice and implications of heuristics for local economic development, *Journal of Business Research*, 67, 1700–1709. doi:10.1016/j.jbusres.2014.02.016.

Friedman, D., Isaac, R. M., James, D., & Sunder, S. (2014). *Risky curves. On the empirical failure of expected utility*. New York: Routledge.

Friedman, M. (1953). *Essays in positive economics*. Chicago: University of Chicago Press.

Gigerenzer, G. (2004). Striking a blow for sanity in theories of rationality. In M. Augier & J. G. March (Eds.), *Models of a man: Essays in memory of Herbert A. Simon* (pp. 389–409). Cambridge, MA: MIT Press.

Gigerenzer, G. (2015). On the supposed evidence for libertarian paternalism. *Review of Philosophy and Psychology*, 6, 361–383.

Gigerenzer, G. (2016). Towards a rational theory of heuristics. In R. Frantz & L. Marsh (Eds.), *Minds, models, and milieux: Commemorating the centennial of the birth of Herbert Simon* (pp. 34–59). New York: Palgrave Macmillan.

Gigerenzer, G. (2018). The bias bias in behavioral economics. *Review of Behavioral Economics*, 5, 303–336. doi:10.1561/105.00000092.

Gigerenzer, G., & Brighton, H. (2009). Homo heuristicus: Why biased minds make better inferences. *Topics in Cognitive Science*, 1, 107–143. doi:10.1111/j.1756-8765.2008.01006.x.

Gigerenzer, G., Fiedler, K., & Olsson, H. (2012). Rethinking cognitive biases as environmental consequences. In P. M. Todd, G. Gigerenzer, & the ABC Research Group, (Eds.), *Ecological rationality: Intelligence in the world* (pp. 80–110). New York: Oxford University Press.

Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology*, 62, 451–482. doi:10.1146/annurev-psych-120709-145346.

Gigerenzer, G., Hertwig, R., & Pachur, T. (Eds.) (2011). *Heuristics: The foundations of adaptive behavior*. New York: Oxford University Press.

Gigerenzer, G., & Selten, R. (Eds.). (2001). *Bounded rationality: The adaptive toolbox*. Cambridge, MA: MIT Press.

Gigerenzer, G., & Todd, P. M. (1999). Fast and frugal heuristics: The adaptive toolbox. In G. Gigerenzer, P. M. Todd, & the ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 3–34). New York: Oxford University Press.

Gilovich, T., Vallone, R. & Tversky. A. (1985). The hot hand in basketball: On the misperception of random sequences. *Cognitive Psychology*, 17, 295–314. doi:10.1016/0010-0285(85)90010-6.

Hahn, U., & Warren, P. A. (2009). Perceptions of randomness: Why three heads are better than four. *Psychological Review*, 116(2), 454–461. doi:10.1037/a0015241.

Kahneman, D. (2003). Maps of bounded rationality: Psychology for behavioral economics. *American Economic Review,* 93(5), 1449–1475.

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

Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology*, 3, 430–454.

Kareev, Y. (1992). Not that bad after all: Generation of random sequences. *Journal of Experimental Psychology: Human Perception and Performance*, 18, 1189–1194. doi:10.1037/0096-1523.18.4.1189.

Keynes, J. M. (1936). *The general theory of employment, interest and money*. London: Macmillan.

Knight, F. (1921). *Risk, uncertainty and profit*. Boston, MA: Houghton Mifflin Co.

Kruglanski, A. W., & Gigerenzer, G. (2011). Intuitive and deliberative judgments are based on common principles. *Psychological Review*, 118, 97–109.

Lewin, K. A. (1935). *A dynamic theory of personality*. New York: McGraw-Hill.

Lopes, L. L. (1992). Three misleading assumptions in the customary rhetoric of the bias literature. *Theory & Psychology*, 2, 231–236. doi:10.1177/0959354392022010.

Luce, R. D., & Raiffa, H. (1957). *Games and decisions*. New York: Dover.

Martignon, L., & Hoffrage, U. (2002). Fast, frugal, and fit: Lexicographic heuristics for paired comparison. *Theory and Decision*, 52, 29–71. doi:10.1023/A:1015516217425.

Martignon, L., Katsikopoulos, K. V., & Woike, J. K. (2008). Categorization with limited resources: A family of simply heuristics. *Journal of Mathematical Psychology*, 52(6), 352–361. doi:10.1016/j.jmp.2008.04.003.

Miller, J. B., & Sanjuro, A. (2016). Surprised by the gambler's and hot hand fallacies? A truth in the law of small numbers. IGIER Working Paper No. 552. doi:10.2139/ssrn.2627354.

Oxford English Dictionary (n.d.). Heuristic. In *OED Online*. Available at: www.oed.com/view/Entry/86554?isAdvanced=false&result=1&rskey=C4NVLA& (accessed July 8, 2020).

Piaget, J., & Inhelder, B. (1975). *The origin of the idea of chance in children*. New York: Norton.

Rubinstein, A. (1998). *Modeling bounded rationality*. Cambridge, MA: MIT Press.

Sargent, T. J. (1993). *Bounded rationality in macroeconomics*. Oxford: Clarendon Press.

Savage, L. J. (1954). *The foundations of statistics*. New York: Wiley.

Sedlmeier, P., Hertwig, R., & Gigerenzer, G. (1998). Are judgments of the positional frequencies of letters systematically biased due to availability? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24, 754–770.

Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69, 99–118. doi:10.2307/1884852.

Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63, 129–38. doi:10.1037/h0042769.

Simon, H. A. (1978). Rationality as process and as product of thought. *American Economic Review*, 68, 1–16.

Simon, H. A. (1985). Human nature in politics: The dialogue of psychology and political science. *American Political Science Review*, 79, 293–304.

Simon, H. A. (1988). Nobel laureate Simon "looks back": A low-frequency mode. *Public Administration Quarterly*, 12, 275–300.

Simon, H. A. (1989). The scientist as problem solver. In D. Klahr & K. Kotovsky (Eds.), *Complex information processing: The impact of Herbert A. Simon* (pp. 375–398). Hillsdale, NJ: Erlbaum.

Simon, H. A. (1990). Invariants of human behavior. *Annual Review of Psychology,* 41, 1–19. doi:10.1146/annurev.ps.41.020190.000245

Simon, H. A. (1997). *Models of bounded rationality,* Vol. 3: *Empirically grounded economic reason*. Cambridge, MA: MIT Press.

Şimşek, Ö. (2013). Linear decision rule as aspiration for simple decision heuristics. *Advances in Neural Information Processing Systems*, 26, 2904–2912.

Smith, V. L. (2003). Constructivist and ecological rationality in economics. *The American Economic Review,* 93(3), 465–508.

Stigler, G. J. (1961). The economics of information. *Journal of Political Economy*, 69, 213–225. doi:10.1086/258464.

Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. New Haven, CT: Yale University Press.

Todd, P. M., Gigerenzer, G., & the ABC Research Group (Eds.) (2012). *Ecological rationality: Intelligence in the world*. New York: Oxford University Press.

Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review*, 79, 281–299. doi:10.1037/h0032955.

Tversky, A., & Bar-Hillel, M. (1983). Risk: The long and the short. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 9*(4), 713–717. doi:10.1037/0278-7393.9.4.713.

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5, 207–232.

Tversky, A., & Kahneman, D. (1974). Judgement under uncertainty: Heuristics and biases. *Science*, 185, 1124–1131. doi:10.1126/science.185.4157.1124.

Wübben M, & von Wangenheim F. (2008). Instant customer base analysis: managerial heuristics often "get it right." *Journal of Marketing*, 72, 82–93. doi:10.1509/jmkg.72.3.82.