Cognitive Unconscious and Human Rationality

edited by Laura Macchi, Maria Bagassi, and Riccardo Viale

The MIT Press
Cambridge, Massachusetts
London, England

This book was set in Stone Sans Std and Stone Serif Std by Toppan Best-set Premedia Limited. Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Names: Macchi, Laura, 1961– editor. | Bagassi, Maria, editor. | Viale, Riccardo, editor.
Title: Cognitive unconscious and human rationality / Laura Macchi, Maria Bagassi, and Riccardo Viale, eds. ; foreword by Keith Frankish.
Description: Cambridge, MA : MIT Press, [2015] | Includes bibliographical references and index.
Identifiers: LCCN 2015038315 | ISBN 9780262034081 (hardcover : alk. paper)
Subjects: LCSH: Reasoning (Psychology) | Cognition. | Subconsciousness.
Classification: LCC BF442 .C64 2015 | DDC 154.2–dc23 LC record available at http://lccn.loc.gov/2015038315

10 9 8 7 6 5 4 3 2 1

1 Rationality without Optimization: Bounded Rationality

Gerd Gigerenzer

The study of bounded rationality is not the study of optimization in relation to task environments.

—Herbert A. Simon (1991)

What is bounded rationality? The answer to this question seems obvious. After all, numerous books, articles, and encyclopedia entries have been written and conferences held on this topic ever since Herbert A. Simon coined the term in the mid-20th century. Yet when Kenneth Arrow, Daniel Kahneman, and Reinhard Selten—all three, like Simon, winners of the Bank of Sweden Prize in Economic Sciences in Memory of Alfred Nobel—write about bounded rationality, each of them has a strikingly different concept in mind. Nonetheless, each links his own particular concept to the same person, Herbert Simon, so that there appears to be a single definition of bounded rationality. In this chapter I explain the three different faces of this concept and what Herbert Simon actually meant. Let us begin with Simon himself.

How do people make decisions in a world of uncertainty? In the mid-1930s young Herbert Simon, fresh from a class on price theory at the University of Chicago, tried to apply the perspective of utility maximization to budget decision problems in his native Milwaukee recreation department. To his surprise, he learned that managers did not compare the marginal utility of a proposed expenditure with its marginal costs. Instead, they simply added incremental changes to last year's budget, engaged in habits and bargaining, or voted on the basis of their identification with organizations. Simon 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)

Let us call this *Simon's question*. Budget problems are not the only ones for which Simon's question must be posed. It is equally relevant for financial investment, personnel selection, strategic decisions, and in general, in situations where the future cannot be known with certainty and surprises happen. What Simon means is that in such situations the neoclassical theory is of little help either *as a model of how decisions are actually made* or *as a prescription for how to make better decisions*. The hypothesis that managers behave *as if* they maximized their expected utility is an essentially different claim (see below). What are these conditions for rationality to which Simon refers?

There are two main approaches to defining rationality in economics and beyond (Sen, 1987). The first emphasizes internal *consistency*. In other words, choices must obey certain rules such as "if *x* is chosen from a set containing *y*, then *y* will not be chosen from any set containing *x*," which is known as the *Weak Axiom of Revealed Preferences* (Samuelson, 1938). A variety of further conditions of internal consistency have been proposed in the literature. The second approach emphasizes the reasoned pursuit of *self-interest*. The view that human beings relentlessly maximize self-interest is often attributed to Adam Smith, although Smith considered many other motivations as well; nevertheless, self-interest has shaped the understanding of individual motivation in economics for a long time.

These two approaches are not the same; a person can be consistent without maximizing self-interest by maximizing something else. As a consequence, some consider the maximization of self-interest too narrow a conception of rationality, and various proposals have been made to extend the motivational structure of *Homo economicus* to "other-regarding" motivations (e.g., kinship, reciprocal altruism, kindness, politeness) while maintaining the maximization framework. To others, the consistency framework is too general; for instance, acting consistently against one's own interests would be an odd conception of rationality.

What Is Radical about Simon's Question?

Simon's question is more radical than the question of how to add other-regarding motives into the utility function. He argued that most of the time, motivated by either self-interest or other-regarding interest, humans do not maximize at all but instead "satisfice" (Simon, 1955, 1979). Why would humans not optimize? As a polymath with a strong interest in artificial intelligence, Simon was aware that most important problems are computationally intractable, that is, no computer or mind can calculate the optimal course of action. This was not a new insight. Jimmy Savage, the father of modern Bayesian decision theory, had earlier pointed out how "utterly ridiculous" it would be to "plan a picnic or to play a game of chess in accordance with the principle," that is, with the expected utility maximization framework (Savage, 1954, p. 16). In chess the insurmountable problem is tractability: although the game is well defined, the optimal sequence of moves cannot be calculated by a human or machine. For planning a picnic, the problem is even deeper. Here all possible moves (what can happen) cannot be known ahead, and surprises may occur; thus, the problem space is ill defined.

In the following, I use the term *uncertainty* to refer to problems that are well defined but computationally intractable or ill defined, that is, where some alternatives, consequences, or probability distributions are either not known or knowable. My use of the term extends its earlier use by Knight (1921), who focused mainly on unknown probability distributions. In an uncertain world, by definition, we cannot calculate the "optimal" course of action (although, after the fact, many tend to exhibit supernatural powers of hindsight). We can now see that Simon's question involves two objectives.

The research program of bounded rationality is about:

1. Decision making under *uncertainty*—as opposed to programs about calculable *risk* only—that is, situations in which optimizing (determining the *best* course of action) is not feasible for the reasons mentioned above.
2. The process by which individuals and institutions actually reach decisions—as opposed to *as-if* models—including their search, stopping, and decision rules.

In Simon's words, the program's first objective is to model "rationality without optimization" (1990, p. 8). Note that the reason is not people's irrationality but the very nature of ill-defined or intractable problems. As we will see, some psychologists missed this point, blaming the human mind for not being able to optimize all the time. The second objective is to model the decision processes rather than construct as-if models, as do neoclassical economics and parts of psychology as well. Ever since Milton Friedman declared that the realism of the assumptions is of little relevance in economic models and that only the quality of the predictions matters, the as-if doctrine has successfully eliminated concern with psychological processes.

The reader might ask: What is so radical about Simon's question? Let me consider two groups of researchers whom Simon addressed, economists and psychologists. Many neoclassical economists would consider a model that does not involve optimization as something alien to their field. As we will see, even behavioral economists have continued on this path. The resulting optimizing models are typically as-if. Rationality without some form of optimization appears to be unthinkable. Similarly, many psychologists define rational behavior by the laws of probability or logic and confront their experimental subjects with problems of risk rather of uncertainty. Standards of this kind of research are choices among lotteries, Bayesian problems, trolley problems, and experimental games in which the optimal course of action can be calculated. Both the problems and the theories of rationality are reduced to situations of calculable risks.

As a consequence, research devoted to Simon's question has been the exception rather than the rule.

Why Is There No General Theory of Bounded Rationality?

In *Administrative Behavior* (Simon, 1947), bounded rationality still referred to everything outside of full rationality or optimizing. Later, Simon (1955, 1956) became more specific by proposing concepts such as aspiration levels, recognition, limited search, and satisficing. Yet in the half-century after these beginnings, relatively little conceptual and empirical work has been done. As a consequence, we lack an overarching theory of bounded rationality. One reason is that Simon's question has been systematically avoided by misrepresenting the study of bounded rationality as either the study of optimization or deviations from optimization. Simon himself would be the first to admit that he left us with an unfinished project and that his own thinking developed gradually from an optimizing version of satisficing to the clear statement of his question in later years. In an email dated September 2, 1999, he wrote:

Dear Gerd,

I have never thought of either bounded rationality or satisficing as precisely defined technical terms, but rather as signals to economists that they needed to pay attention to reality, and a suggestion of some ways in which they might. ...

I guess a major reason for my using somewhat vague terms—like bounded rationality—is that I did not want to give the impression that I thought I had "solved" the problem of creating an empirically grounded theory of economic phenomena. ... There still lies before us an enormous job of studying the actual decision processes that take place in corporations and other economic settings. ... End of sermon—which you and Reinhard don't need. I am preaching to believers.

Cordially,

Herb

In signing off, he refers to the work of Reinhard Selten, who devoted part of his professional life to answering Simon's question (e.g., Gigerenzer & Selten, 2001; Selten, 1998, 2001). Selten's work on decision making under uncertainty, together with the work of my research group and others around the world (e.g., Gigerenzer, Todd, & the ABC Research Group, 1999; Todd, Gigerenzer, & the ABC Research Group, 2012), indicates that the tide is changing and that we are finally beginning to know how to answer Simon's question, including its extension to a normative one (see below). In this chapter, I do not cover the hundreds of experiments and analyses that have been performed to answer Simon's question in the last two decades (for an overview, see Gigerenzer, Hertwig, & Pachur, 2011; Hertwig, Hoffrage, & the ABC Research Group, 2013). Instead, I focus on the conceptual issues crucial to Simon's question.

I argue that there are three quite different programs of bounded rationality, two of them explicitly ignoring Simon's question. I then explain how these programs differ from what Simon had in mind and thus provide one answer to the question why there is no general theory of bounded rationality.

Three Programs of "Bounded Rationality"

The first two programs represent what most economists and psychologists, respectively, think bounded rationality is about. Each of these two programs ignores Simon's question by assuming that the conditions of rationality are met. Nevertheless, both misleadingly refer to Simon's concept of bounded rationality. In that way, the revolutionary potential of Simon's question has been largely defused.

As-If Rationality: Optimization under Constraints Most neoclassical economists argue that bounded rationality is nothing but full rationality, taking costs of information and deliberation into account. As Arrow (2004, p. 48) asserts:

> ... 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.

As an aside, this statement is from Arrow's contribution to the memorial volume for the late Simon. In this view, bounded rationality is nothing but constrained optimization in disguise. Stigler's (1961) model of the purchase of a second-hand car is a classic example, where the buyer is assumed to stop search when the costs of further search exceed its benefits. But in Simon's own words that is by no means what it is about: "The study of bounded rationality is not the study of optimization in relation to task environments" (Simon, 1991).

In personal conversation Simon once wryly told me that he wanted to sue authors who misused his concept to describe another form of optimization (Gigerenzer, 2004). In doing so they ignored both objectives of Simon's bounded rationality program: to study behavior under uncertainty and the actual decision processes themselves.

Moreover, the optimization-under-constraints version of bounded rationality leads to a paradoxical result that provides a good reason not to pursue this program any further. The more realistic constraints are introduced into an optimization model, the more unrealistic knowledge and computations it needs to assume. Accordingly, in his *Bounded Rationality in Macroeconomics*, Sargent (1993, p. 22) describes bounded rationality as a "research program to build models populated by agents who behave like working economists or econometricians." As a consequence, the program of incorporating realistic constraints creates highly complex models that are unattractive compared to the easier models of full rationality, leaving us back at square one.

Deviations from Optimization: The Cognitive Illusions Program Whereas most economists argue that bounded rationality is about rational behavior, most psychologists argue that it is about systematic biases and other forms of irrational behavior. The *cognitive illusions program*, also called the *heuristics-and-biases program*, takes this stance. Although the two interpretations appear to be diametrically opposed, a closer look shows that they differ only in their descriptive claims about people's rationality. In fact, the heuristics-and-biases program accepts the economic optimizing models as the universal norm for people's reasoning. In Kahneman's (2003, p. 1449) words:

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.

In this view, optimization is descriptively incorrect as a model of human behavior but remains normatively correct. Accordingly, the experimental studies conducted in this program rely on content-free principles of logic, probability theory, or expected utility maximization in situations where the optimal belief or choice is known—or assumed to be known. Any discrepancy between human judgment and the proposed norm is then blamed on the humans rather than the models.

Let us compare the two programs. As-if rationality ignores both research goals in Simon's question. Instead, it models situations in which the optimal answer can be calculated and favors as-if models rather than models of the decision process. Although the cognitive illusions program also ignores Simon's first objective, it does accept his second objective, repeatedly stating its aim to uncover the underlying cognitive processes, called heuristics. However, unlike what Simon had in mind, its proponents propose no testable formal models of heuristics. With the exception of the earlier seminal work by Tversky (1972) on elimination by aspects, virtually no models of heuristics have been specified algorithmically. Instead, vague labels such as availability and representativeness were introduced after the fact (Gigerenzer, 1996).

By now it should be clear that Kahneman's use of the term *bounded rationality* is not Simon's. As Simon (1985, p. 297) said, "Bounded rationality is not irrationality." One explanation for this discrepancy is that Kahneman and Tversky never set out to find an answer to Simon's question. As Lopes (1992) pointed out, they adopted the fashionable term *bounded rationality* as an afterthought and as an acknowledgment to a distinguished figure. This hypothesis is supported by the fact that although bounded rationality is mentioned in the preface of Kahneman, Slovic, and Tversky's 1982 anthology, their influential papers reprinted in the same anthology neither mention it nor cite Simon at all.

Whereas economists have tried to reduce bounded rationality to neoclassical economics, Kahneman, Tversky, and their followers have tried to reduce it to their study of cognitive illusions. In this way, Simon's question has been abandoned (Callebaut, 2007; Gigerenzer, 2004), and the role of psychology has been reduced to documenting errors. A distinguished economist made this point crystal clear, concluding a discussion with the dictum "either reasoning is *rational* or it's *psychological*" (Gigerenzer, 2000, p. vii). By contrast, the study of ecological rationality dispenses with this supposed opposition between the rational and the psychological. To analyze ecological rationality, one needs to understand the heuristics that people use in the first place.

Ecological Rationality: The Adaptive Toolbox A third perspective, the program of ecological rationality (Gigerenzer et al., 1999; Gigerenzer & Selten, 2001), addresses both objectives in Simon's question. Looking beyond expected utility maximization and Bayesian decision theory, it seeks to answer the question of how people make rational decisions under uncertainty, that is, when optimization is out of reach (Gigerenzer & Selten, 2001, p. 4):

> Models of bounded rationality describe how a judgment or decision is reached (that is, the heuristic processes or proximal mechanisms) rather than merely the outcome of the decision, and they describe the class of environments in which these heuristics will succeed or fail.

This program has both a descriptive and a normative goal. Its descriptive goal is to model the heuristics, their building blocks, and the core cognitive capacities these exploit. This is called the study of the *adaptive toolbox* of an individual or an institution. Although the heuristics-and-biases program has made a first step in this direction, the program goes one step further and provides mathematical and computational models for heuristics that are testable. The normative study of heuristics addresses the question when people *should* rely on a simple heuristic rather than on a competing, more complex strategy. It is called the study of the ecological rationality of heuristics.

The Adaptive Toolbox

Simon had proposed the satisficing heuristic as a model of decision making under uncertainty. The task is to choose a satisfactory alternative when the set of alternatives is not known in advance. That kind of uncertainty is typical for many problems, such as whom to marry, where to invest one's money, or which house to buy. Here, the "optimal" stopping rule—the point in time where the costs of further search exceed its benefits—cannot be calculated, and a reduction to optimization of constraints would mean creating an illusion of certainty. The basic version of satisficing with a fixed aspiration level has three building blocks: a search rule, a stopping rule, and a decision rule (Simon, 1955):

1. Set an aspiration level α and sequentially search for objects.
2. Stop search after the first object x is found with $x \geq \alpha$.
3. Choose object *x*.

By adapting building blocks to the problem at hand, one can modify heuristics. For instance, aspiration levels can be constant over time or adjusted to reflect a learning process (Selten, 1998, 2001). By replacing the stopping rule (2) with (2′), we can incorporate such a learning process and attain an adjustable aspiration level:

2′. Stop search after the first object is found with $x \geq \alpha$. If no such object is found within a time interval β, lower α by a factor γ, and continue search.

Thus, the basic version has one individual parameter (the aspiration level α), whereas the adjustable version has three (aspiration level α, time interval β, and size of adjustment γ). A study of the pricing strategies for used BMWs (models 320 and 730) by 748 dealers selling 24,482 cars showed that 95% of these dealers used a satisficing heuristic: 15% relied on fixed aspiration levels, and 80% relied on an adjustable aspiration level (Artinger & Gigerenzer, 2014). Typically, α was set above the mean market price and adjusted downward after around 1 month. Simon (1955) referred to pricing in the housing market, but given the lack of research on satisficing, we do not know whether housing prices, like used car prices, are also shaped by aspiration levels. Yet that is also quite likely. For instance, Levitt and Syverson (2008) report that when selling on their own behalf, real estate agents take considerably longer until they sell than when selling for their clients. Satisficing has also been investigated as an alternative to optimizing models in mate choice (Todd & Miller, 1999) and to explain aggregate age-at-marriage patterns (Todd, Billari, & Simao, 2005).

Research on the adaptive toolbox has identified a number of other heuristics, many of which can be decomposed into the three common building blocks:

1. *Search* rules specify what information is considered and in which order or direction information search proceeds.
2. *Stopping* rules specify when the search for information is terminated.
3. *Decision* rules specify how the final decision is reached.

Although there are many heuristics, they consist of a small number of building blocks, similar to how the chemical elements in the periodic table are built from a small number of particles.

Ecological Rationality

The traditional explanation for why people rely on heuristics is that heuristics reduce effort at the cost of accuracy. This hypothesis is called the *accuracy–effort trade-off*. In this view, heuristics are always second-best—in terms of accuracy, not effort. One of the first discoveries in our research was that this accuracy–effort trade-off is not generally true in situations of uncertainty (Goldstein & Gigerenzer, 2002; Gigerenzer & Brighton, 2009). For illustration, across 20 real-world problems, a simple sequential heuristic called take-the-best made on average more accurate predictions than did a multiple regression, and with less effort (Czerlinski, Gigerenzer, & Goldstein, 1999). We called the reversal of the accuracy–effort trade-off a *less-is-more effect*; determining the exact conditions under which less-is-more occurs is one part of the normative study of ecological rationality.

The normative study of heuristics is new. It was unthinkable as long as researchers believed in the generality of the accuracy–effort trade-off: that logic and probability theory are always optimal and heuristics always second best. This may be correct in situations of risk but not in those of uncertainty. The study of ecological rationality is a normative extension of Simon's descriptive question:

Ecological Rationality Question: How *Should* Human Beings Reason in Situations of Uncertainty?

To answer this question, it is no longer sufficient to rely on optimization. But what are the alternatives? Simon provided a clue with his scissors analogy: "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, 1990, p. 7). By looking at only one blade, cognition, it is impossible to understand when and why a behavior succeeds or fails. As a consequence, in situations under uncertainty, rationality can be defined as the "match" between a heuristic and the structure of the environment: "A heuristic is ecologically rational to the degree it is adapted to the structure of the environment" (Gigerenzer & Todd, 1999, p. 13).

The study of the ecological rationality of a heuristic is conducted with the help of mathematical analysis and computer simulation. It aims at identifying the environmental structures that a heuristic exploits. These structures include:

1. Predictability: how well a criterion can be predicted.
2. Sample size: number of observations (relative to number of cues).
3. Number *N* of alternatives.
4. Redundancy: the correlation between cues.
5. Variability in weights: the distribution of the cue weights (e.g., skewed or uniform).

For instance, (1) the smaller the sample size, (2) the higher the redundancy, and (3) the higher the variability in weights are, the more accurate the predictions tend to be of sequential search heuristics such as take-the-best in comparison to linear regression models (Martignon & Hoffrage, 2002). Given these insights, normative statements can be made about when a heuristic will likely be better at making predictions than an alternative strategy such as a weighted linear model. The conditions of the ecological rationality of some heuristics have been well identified (e.g., Baucells, Carrasco, & Hogarth, 2008; Hogarth & Karelaia, 2006, 2007; Katsikopoulos, 2013; Katsikopoulos & Martignon, 2006; Simşek, 2013).

A complementary analysis of the ecological rationality of heuristics is based on the bias-variance dilemma (Geman, Bienenstock, & Doursat, 1992). In earlier research on cognitive illusions, the term *bias* refers to ignoring part of the information, as in the *base rate fallacy*. This can be captured in the equation:

$$\text{Error} = \text{bias} + \varepsilon \tag{1.1}$$

where ε is an irreducible random error. In the heuristics-and-biases view, if the bias is eliminated, good inferences are obtained. That is true in a world of known risks. In an uncertain world, however, there are three sources of errors:

$$\text{Error} = \text{bias}^2 + \text{variance} + \varepsilon \tag{1.2}$$

where *error* is the prediction error (the sum of squared error) and *bias* refers to a systematic deviation between a model and the true state, as in equation 1.1. To define the meaning of *variance*, consider 100 people who rely on the same strategy but have different samples of observations from the same population. Because of sampling error, the 100 inferences may not be the same. Across samples, *variance* is the expected squared deviation around their mean; *bias* is the difference between the mean prediction and the true state of nature. Thus, to minimize total error, a certain amount of bias is needed to counteract the error due to oversensitivity to sampling noise (see Gigerenzer, 2016, for formal definitions).

Variance decreases with increasing sample size but also with simpler strategies that have fewer free parameters (and less flexible functional forms) (Pitt, Myung, & Zhang, 2002). Thus, a cognitive system needs to draw a balance between being biased and flexible (variance) rather than simply trying to eliminate bias. Heuristics can be fast, frugal, and accurate by exploiting the structure of information in environments, by being robust, and by striking a good balance between bias and variance. This bias-variance dilemma helps to clarify the rationality of simple heuristics and why less can be more (Brighton & Gigerenzer, 2008).

Intuitive Design

Alongside Simon's question and its normative extension to the study of ecological rationality, there is a third question: How can one apply the answers to both questions to design heuristics that improve expert decision making? The key idea is to use the research on the adaptive toolbox to identify classes of strategies that are intuitive to the human mind. For instance, unlike logistic regressions and similar models, which are alien to the thinking of, say, most doctors and lawyers, sequential search heuristics such as take-the-best and fast-and-frugal trees are easy to understand and learn. The research on ecological rationality is then used to determine what kind of heuristic is likely to work for the problem at hand. For instance, fast-and-frugal trees that have been designed for coronary care unit allocations are intuitive for doctors, enable fast decisions, and are reported to make fewer errors in predicting heart attacks than standard logistic regression models used in heart attack diagnosis (Green & Mehr, 1997; Martignon, Vitouch, Takezawa, & Forster, 2003). Similarly, predicting bank vulnerability concerns the realm of uncertainty, not of calculable risk, and standard probability models such as value-at-risk have failed to predict every financial crisis and prevented none. My research group is currently cooperating with the Bank of England to develop simple heuristics for a safer world of finance, both for regulation and for predicting bank failures (Aikman et al., 2014; Neth, Meder, Kothiyal, & Gigerenzer, 2014). Under uncertainty, simplicity and bias are not signs of cognitive limitations but are vital to making better inferences.

Three Programs of Behavioral Economics

These three visions of bounded rationality have shaped the research questions posed in quite a few research programs, including social psychology, neuroeconomics, and behavioral economics.

The term *behavioral economics* appears to have been first used by George Katona (1951, 1980) and by Herbert Simon (Sent, 2004). Behavioral economics as we know it today, however, began in 1984, funded by the Alfred P. Sloan Foundation. In its early stages it was understood as an advancement of Simon's question (Heukelom, 2014). Yet Kahneman, Tversky, and Thaler eventually took over and abandoned Simon's revolutionary ideas in favor of a program with better promise to be accepted by neoclassical economists. Contrary to Simon, Kahneman and Tversky argued that there was nothing wrong with expected utility maximization; there was only something wrong with us humans. These researchers eventually determined the agenda of the emerging field of behavioral economics, to the dismay not only of Simon but notably of Vernon Smith. Smith (2003) put the term *ecological rationality* in the title of his Nobel lecture, drawing an explicit link to our work on the adaptive toolbox, and discussed this relation in his subsequent book (Smith, 2007). The history of the struggle over behavioral economics is told by Heukelom (2014). My point is this: as a closer look reveals, not one but three different kinds of behavioral economics exist, corresponding to the three versions of bounded rationality.

As-If Behavioral Economics

Behavioral economists typically justify their approach with the claim that it "increases the explanatory power of economics by providing it with more realistic psychological foundations" (Camerer & Loewenstein, 2004, p. 3). It is thus surprising that the first version of behavioral economics, which I call *as-if behavioral economics*, appears to be indistinguishable from neoclassical economics in its reliance on as-if models (Berg & Gigerenzer, 2010). The program of as-if behavioral economics is to keep the as-if framework of expected utility maximization untouched, and just add a few new "psychological" parameters to it.

Fehr and Schmidt's (1999) *inequity aversion theory* provides an example of as-if behavioral economics. It accounts for the observation that people care about others as well as themselves by adding two free parameters to the expected utility maximization theory, assuming that individuals behave *as if* they maximized a weighted function of various utilities. *Prospect theory* (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992) similarly introduces additional adjustable parameters to the expected utility equation, including nonlinear transformations of probabilities and utilities, resulting in calculations that few would maintain are realistic models of people's decision processes. Laibson's (1997) model of impulsiveness consists, in essence, of adding a parameter to the neoclassical model of maximizing an exponentially weighted sum of instantaneous utilities. The decision maker is assumed to make an exhaustive search of all consumption sequences, compute the weighted sum of utility terms for each of these, and choose the one with the highest weighted utility.

These three prominent theories illustrate the first kind of behavioral economics: as-if models with parameters added, which allow a better fit, not necessarily better predictions (Brandstätter, Gigerenzer, & Hertwig, 2006). This program serves as a "repair program" of neoclassical economics (Berg & Gigerenzer, 2010; Gigerenzer, 2008, p. 90; Güth, 2008). It ignores both objectives formulated in Simon's question.

Anomalies Program

The second program of behavioral economics is based on the cognitive illusions program. In Camerer's (1995) words, "limits on computational ability force people to use simplified procedures or 'heuristics' that cause systematic mistakes (biases) in problem solving, judgment, and choice. The roots of this approach are in Simon's (1955) ..." (p. 588). More bluntly, Diamond (2008) states: "Behavioral economics is the identification of circumstances where people are making 'mistakes.'" Thaler (1991, p. 138) argues that the major contribution of behavioral economics has been the discovery of a collection of cognitive illusions completely analogous to visual illusions. Similarly, Ariely (2008) argues that visual illusions such as Shepard's two-tables illusion show the deficiencies of our brain. Neither Thaler nor Ariely seems to realize that, without useful biases, we would perceive only what is on the retina. To borrow a phrase from Jerome Bruner, cognition has to go beyond the information given. In order to infer a three-dimensional world from a retinal image, the perceptual system makes intelligent bets. Every intelligent system has to make bets and therefore necessarily makes "good" errors (Gigerenzer, 2008). A system that makes no errors is simply not intelligent.

The anomalies program inherited two problems from the cognitive illusions program: questionable norms for good reasoning and a focus on vague labels instead of precise models of heuristics (e.g., Camerer, 1998; Conlisk, 1996; Kahneman, 2011). What is considered a "cognitive error" is not always so. Consider the first three items in Camerer's list of cognitive errors.

The first apparent illusion is called *miscalibration*. In a typical study people answer a series of trivia questions and provide a confidence judgment that each of their answers is correct. Then the average percentage correct is plotted against each confidence category. Typically, if confidence is 100%, only about 80% of the answers are correct; if confidence is 90%, only 70% are correct; and so on. This discrepancy between confidence and percentage correct, a so-called "miscalibration," has been attributed to systematic mental or motivational flaws in people's minds (e.g., Camerer, 1998; Lichtenstein, Fischhoff, & Phillips, 1982). A closer inspection shows that these "miscalibration" curves look like regression to the mean, and several experiments have shown that this is indeed largely the case. Regression to the mean is a consequence of *unsystematic* error, and if one plots the data the other way around, one gets the other regression curve, which now looks like *underconfidence*. Regression to the mean is generated by unsystematic error, not by a systematic error such as overconfidence. Thus, the problem is largely a systematic error *by researchers*, not one of *ordinary people* (Dawes & Mulford, 1996; Erev, Wallsten, & Budescu, 1994; Gigerenzer, Fiedler, & Olsson, 2012).

The second error on the list is that the confidence intervals people construct are too small. In this research, people are asked to estimate the 95% confidence interval for, say, the length of the river Nile. A person might respond, "between 2,000 and 8,000 kilometers." The typical result reported is that in more than 5% of the cases the true estimate is outside the confidence interval. This phenomenon is attributed to overconfidence. Again, a bit of statistical thinking by the researchers would have helped. Juslin, Winman, and Hansson (2007) noted that when inferences are drawn from samples, statistical theory implies that a systematic error will appear in variance and therefore in the production of confidence intervals, whereas the mean (probability) is an unbiased estimate. In experiments they showed that when people are asked to provide the probability (mean) that the true value lies in the interval, the error largely disappears. In our example that means asking the question, "What is the probability that the length of the river Nile is between 2,000 and 8,000 kilometers?"

The third form of "overconfidence" listed is the observation that average confidence judgment is higher than average percentage correct. Long ago, we (Gigerenzer, Hoffrage, & Kleinbölting, 1991) showed that this difference disappears when experimenters use randomly sampled questions rather than selecting questions with surprising answers (such as that New York is south of Rome). After reviewing 135 studies, Juslin, Winman, and Olsson (2000) found that overconfidence was observed when questions were selected but was on average about zero with random sampling. The bias thus lies in the biased sampling of researchers, not simply in people's minds.

I could continue to review the list of so-called errors, but the point should be sufficiently clear. Quite a few "biases" actually appear to originate from researchers' lack of statistical thinking rather than, as claimed, from people's systematic errors (Gigerenzer, 2001; Gigerenzer, Fiedler, & Olsson, 2012). This is not to say that biases never exist. For instance, most of the research reported in the previous paragraphs and elsewhere contradicting the heuristics-and-biases story is surprisingly absent in the behavioral economics literature from Camerer (1998) to Kahneman (2011). That is known as a confirmation bias.

Toward an Ecological Rationality Program of Behavioral Economics

The third program of behavioral economics is an elaboration of Simon's question. The goal is to enrich economic theory by progressing from as-if models to process models while simultaneously moving from labels for heuristics to testable models and from logical rationality to ecological rationality. As a first step there is a need to develop process models that address systematic violations of expected utility maximization. Tversky's (1972) elimination by aspects described such a process model and was a move away from optimizing models in the very spirit of Simon's question. Unfortunately, he abandoned this program when joining Kahneman and turned to the repair program. Meanwhile, a quite successful process model has been proposed—the priority heuristic (Brandstätter, Gigerenzer, & Hertwig, 2006), which has been shown to outperform cumulative prospect theory in predicting choices between gambles. Moreover, without a single free parameter, it logically implies many of the so-called anomalies, including the Allais paradox and the fourfold pattern of risk attitude (Katsikopoulos & Gigerenzer, 2008). The second goal is to derive better norms, that is, norms that can guide decision making under uncertainty. The adaptation of the ecological rationality program to behavioral economics is practically nonexistent at the moment. As a consequence, behavioral economics bears a much greater similarity to neoclassical economics than to its original goal of a realistic account of decision processes (Berg & Gigerenzer, 2010).

Rationality in the Plural

It is sometimes said that there is only one way to be rational but several ways to be irrational. The only rational way is optimization. I hope to have made clear that this statement is much too narrow. In situations of uncertainty there is no single best answer, but there are several reasonable ways to behave in a rational way. The idea that there is one way to be rational is itself an illusion or, if you like, one of the many ways to be irrational. Simon gave us an idea of how to fare better when confronted with uncertainty as opposed to calculable risk.

References

Aikman, D., Galesic, M., Gigerenzer, G., Kapadia, S., Katsikopoulos, K., Kothiyal, A., et al. (2014, May 2). *Taking uncertainty seriously: Simplicity versus complexity in financial regulation.* Bank of England Financial Stability Paper No. 2028. Available at SSRN: http://ssrn.com/abstract=2432137. doi: 10.2139/ssrn.2432137

Ariely, D. (2008). *Predictably irrational: The hidden forces that shape our decisions.* New York: HarperCollins.

Arrow, K. J. (2004). Is bounded rationality unboundedly rational? Some ruminations. In M. Augier & J. G. March (Eds.), *Models of a man: Essays in memory of Herbert A. Simon* (pp. 47–55). Cambridge, MA: MIT Press.

Artinger, F., & Gigerenzer, G. (2014). Aspiration-adaptation, price setting, and the used-car market. Unpublished manuscript.

Baucells, M., Carrasco, J. A., & Hogarth, R. M. (2008). Cumulative dominance and heuristic performance in binary multi-attribute choice. *Operations Research, 56,* 1289–1304.

Berg, N., & Gigerenzer, G. (2010). As-if behavioral economics: Neoclassical economics in disguise? *History of Economic Ideas, 18,* 133–165. doi:10.1400/140334.

Brandstätter, E., Gigerenzer, G., & Hertwig, R. (2006). The priority heuristic: Making choices without trade-offs. *Psychological Review, 113,* 409–432. doi:10.1037/0033-295X.113.2.409.

Brighton, H., & Gigerenzer, G. (2008). Bayesian brains and cognitive mechanisms: Harmony or dissonance? In N. Chater & M. Oaksford (Eds.), *The probabilistic mind: Prospects for Bayesian cognitive science* (pp. 189–208). New York: Oxford University Press.

Callebaut, W. (2007). Simon's silent revolution. *Biological Theory, 2,* 76–86. doi:10.1162/biot.2007.2.1.76.

Camerer, C. (1995). Individual decision making. In J. H. Kagel & A. E. Roth (Eds.), *The handbook of experimental economics* (pp. 587–703). Princeton, NJ: Princeton University Press.

Camerer, C. F. (1998). Bounded rationality in individual decision making. *Experimental Economics, 1,* 163–183. doi:10.1007/BF01669302.

Camerer, C., & Loewenstein, G. F. (2004). *Behavioral economics: Past, present, future.* Princeton, NJ: Princeton University Press.

Conlisk, J. (1996). Why bounded rationality? *Journal of Economic Literature, 34,* 669–700.

Czerlinski, J., Gigerenzer, G., & Goldstein, D. G. (1999). How good are simple heuristics? In G. Gigerenzer, P. M. Todd, & the ABC Research Group, *Simple heuristics that make us smart* (pp. 97–118). New York: Oxford University Press.

Dawes, R. M., & Mulford, M. (1996). The false consensus effect and overconfidence: Flaws in judgment, or flaws in how we study judgment? *Organizational Behavior and Human Decision Processes, 65,* 201–211. doi:10.1006/obhd.1996.0020.

Diamond, P. (2008). Behavioral economics. *Journal of Public Economics, 92*(8), 1858–1862.

Erev, I., Wallsten, T. S., & Budescu, D. V. (1994). Simultaneous over- and underconfidence: The role of error in judgment processes. *Psychological Review, 101*, 519–527. doi:10.1037/0033-295X.101.3.519.

Fehr, E., & Schmidt, K. (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics, 114*, 817–868. doi:10.1162/003355399556151.

Geman, S. E., Bienenstock, E., & Doursat, R. (1992). Neural networks and the bias/variance dilemma. *Neural Computation, 4*, 1–58. doi:10.1162/neco.1992.4.1.1.

Gigerenzer, G. (1996). On narrow norms and vague heuristics: A reply to Kahneman and Tversky (1996). *Psychological Review, 103*, 592–596. doi:10.1037/0033-295X.103.3.592.

Gigerenzer, G. (2000). *Adaptive thinking: Rationality in the real world.* New York: Oxford University Press.

Gigerenzer, G. (2001). Content-blind norms, no norms, or good norms? A reply to Vranas. *Cognition, 81*, 93–103. doi:10.1016/S0010-0277(00)00135-9.

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 honor of Herbert A. Simon* (pp. 389–409). Cambridge, MA: MIT Press.

Gigerenzer, G. (2008). *Rationality for mortals.* New York: Oxford University Press.

Gigerenzer, G. (2016). Towards a rational theory of heuristics. In R. Frantz & L. Marsh (Eds.), *Minds, models, and milieu. Commemorating the centenary of Herbert Simon's birth.* Basingstoke, UK: Palgrave Macmillan.

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, *Ecological rationality: Intelligence in the world* (pp. 80–110). New York: Oxford University Press.

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

Gigerenzer, G., Hoffrage, U., & Kleinbölting, H. (1991). Probabilistic mental models: A Brunswikian theory of confidence. *Psychological Review, 98*, 506–528. doi:10.1037/0033-295X.98.4.506.

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, *Simple heuristics that make us smart* (pp. 3–34). New York: Oxford University Press.

Gigerenzer, G., Todd, P. M., & the ABC Research Group. (1999). *Simple heuristics that make us smart.* New York: Oxford University Press.

Goldstein, D. G., & Gigerenzer, G. (2002). Models of ecological rationality: The recognition heuristic. *Psychological Review, 109*, 75–90. doi:10.1037/0033-295X.109.1.75.

Green, L. A., & Mehr, D. R. (1997). What alters physicians' decisions to admit to the coronary care unit? *Journal of Family Practice, 45*, 219–226.

Güth, W. (2008). (Non)behavioral economics: A programmatic assessment. *Zeitschrift für Psychologie. Journal of Psychology, 216*(4), 244.

Hertwig, R., Hoffrage, U., & the ABC Research Group. (2013). *Simple heuristics in a social world.* New York: Oxford University Press.

Heukelom, F. (2014). *Behavioral economics: A history.* Cambridge, UK: Cambridge University Press.

Hogarth, R. M., & Karelaia, N. (2006). "Take-the-best" and other simple strategies: Why and when they work "well" with binary cues. *Theory and Decision, 61*, 205–249. doi:10.1007/s11238-006-9000-8.

Hogarth, R. M., & Karelaia, N. (2007). Heuristic and linear models of judgment: Matching rules and environments. *Psychological Review, 114*, 733–758. doi:10.1037/0033-295X.114.3.733.

Juslin, P., Winman, A., & Hansson, P. (2007). The naïve intuitive statistician: A naïve sampling model of intuitive confidence intervals. *Psychological Review, 114*, 678–703. doi:10.1037/0033-295X.114.3.678.

Juslin, P., Winman, A., & Olsson, H. (2000). Naive empiricism and dogmatism in confidence research: A critical examination of the hard-easy effect. *Psychological Review, 107*, 384–396. doi:10.1037/0033-295X.107.2.384.

Kahneman, D. (2003). Maps of bounded rationality: A perspective on intuitive judgment and choice. In T. Frangsmyr (Ed.), *Les Prix Nobel: The Nobel Prizes 2002* (pp. 1449–1489). Stockholm: Nobel Foundation.

Kahneman, D. (2011). *Thinking fast and slow.* London: Allen Lane.

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

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47*, 263–291. doi:10.2307/1914185.

Katsikopoulos, K. (2013). Why do simple heuristics perform well in choices with binary attributes? *Decision Analysis, 10*, 327–340.

Katsikopoulos, K. V., & Gigerenzer, G. (2008). One-reason decision-making: Modeling violations of expected utility theory. *Journal of Risk and Uncertainty, 37*, 35–56. doi:10.1007/s11166-008-9042-0.

Katsikopoulos, K. V., & Martignon, L. (2006). Naive heuristics for paired comparisons: Some results on their relative accuracy. *Journal of Mathematical Psychology, 50,* 488–494. doi:10.1016/j.jmp.2006.06.001.

Katona, G. (1951). *Psychological analysis of economic behavior.* New York: McGraw-Hill.

Katona, G. (1980). *Essays on behavioral economics.* Ann Arbor: Survey Research Centers, Institute for Social Research, The University of Michigan.

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

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics, 112*(2), 443–477.

Levitt, S. D., & Syverson, C. (2008). Market distortions when agents are better informed: The value of information in real estate transactions. *Review of Economics and Statistics, 90*(4), 599–611. doi:10.1162/rest.90.4.599.

Lichtenstein, S., Fischhoff, B., & Phillips, L. D. (1982). Calibration of probabilities: The state of the art to 1980. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases* (pp. 306–334). Cambridge, UK: Cambridge University Press.

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

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., Vitouch, O., Takezawa, M., & Forster, M. R. (2003). Naive and yet enlightened: From natural frequencies to fast and frugal decision trees. In D. Hardman & L. Macchi (Eds.), *Thinking: Psychological perspectives on reasoning, judgment and decision making* (pp. 189–211). Chichester, UK: Wiley.

Neth, H., Meder, B., Kothiyal, A., & Gigerenzer, G. (2014). *Homo heuristicus* in the financial world: From risk management to managing uncertainty. *Journal of Risk Management in Financial Institutions, 7*(2), 134–144.

Pitt, M. A., Myung, I. J., & Zhang, S. (2002). Toward a method of selecting among computational models of cognition. *Psychological Review, 109,* 472–491.

Samuelson, P. A. (1938). A note on the pure theory of consumers' behavior. *Economica, 5,* 61–71.

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

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

Selten, R. (1998). Aspiration adaptation theory. *Journal of Mathematical Psychology, 42,* 191–214.

Selten, R. (2001). What is bounded rationality? In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality: The adaptive toolbox* (pp. 13–36). Cambridge, MA: MIT Press.

Sen, A. (1987). Rational behaviour. In J. Eatwell, M. Milgate, P. Newman, & I. Erev (Eds.), *The new Palgrave dictionary of economics* (Vol. 4, pp. 68–76). London: Macmillan.

Sent, E.-M. (2004). Behavioral economics: How psychology made its (limited) way back into economics. *History of Political Economy, 36*(4), 735–760.

Simon, H. A. (1947). *Administrative behavior: A study of decision-making processes in administrative organizations.* New York: Macmillan.

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,* 129–138. doi:10.1037/h0042769.

Simon, H. A. (1979). Rational decision making in business organizations. *American Economic Review, 69,* 493–513.

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: Lawrence Erlbaum Associates.

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. (1991). Cognitive architectures and rational analysis: Comment. In K. V. Lehn (Ed.), *Architectures for intelligence* (pp. 25–39). Hillsdale, NJ: Lawrence Erlbaum Associates.

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. *American Economic Review, 93,* 465–508. doi:10.1257/000282803322156954.

Smith, V. L. (2007). *Rationality in economics: Constructivist and ecological forms.* Cambridge: Cambridge University Press.

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

Thaler, R. H. (1991). *Quasi rational economics.* New York: Russell Sage Foundation.

Todd, P. M., Billari, F. C., & Simao, J. (2005). Aggregate age-at-marriage patterns from individual mate-search heuristics. *Demography, 42,* 559–574.

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

Todd, P. M., & Miller, G. F. (1999). From pride and prejudice to persuasion: Satisficing in mate search. In G. Gigerenzer, P. M. Todd, & the ABC Research Group, *Simple heuristics that make us smart* (pp. 287–308). 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., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty, 5*, 297–323. doi:10.1007/BF00122574.