



Axiomatic rationality and ecological rationality

Gerd Gigerenzer¹

Received: 18 June 2018 / Accepted: 14 June 2019 / Published online: 9 July 2019
© The Author(s) 2019

Abstract

Axiomatic rationality is defined in terms of conformity to abstract axioms. Savage (The foundations of statistics, Wiley, New York, 1954) limited axiomatic rationality to *small worlds* (S, C), that is, situations in which the exhaustive and mutually exclusive set of future states S and their consequences C are known. Others have interpreted axiomatic rationality as a categorical norm for how human beings should reason, arguing in addition that violations would lead to real costs such as money pumps. Yet a review of the literature shows little evidence that violations are actually associated with any measurable costs. Limiting axiomatic rationality to small worlds, I propose a naturalized version of rationality for situations of intractability and uncertainty (as opposed to risk), all of which are not in (S, C) . In these situations, humans can achieve their goals by relying on heuristics that may violate axiomatic rationality. The study of ecological rationality requires formal models of heuristics and an analysis of the structures of environments these can exploit. It lays the foundation of a moderate naturalism in epistemology, providing statements about heuristics we *should* use in a given situation. Unlike axiomatic rationality, ecological rationality can explain less-is-more effects (when using less information can be expected to generate more accurate predictions), formalize when one should move from ‘is’ to ‘ought,’ and be evaluated by goals beyond coherence, such as predictive accuracy, frugality, and efficiency. Ecological rationality can be seen as a formalization of means–end instrumentalist rationality, based on Herbert Simon’s insight that rational behavior is a function of the mind and its environment.

Keywords Axiomatic rationality · Ecological rationality · Heuristics · Instrumentalism · Less-is-more effect · Naturalism · Uncertainty

Epistemology is often seen as a strictly normative discipline and psychology as a purely descriptive one. Epistemology tells us how we ought to think, psychology how we actually think. According to this divide, psychology has nothing to offer for under-

✉ Gerd Gigerenzer
gigerenzer@mpib-berlin.mpg.de

¹ Max Planck Institute for Human Development, Lentzeallee 94, 14195 Berlin, Germany

standing the nature of rationality. This view can be traced from Frege through Russel, Wittgenstein, and Carnap to the post-World War II analytical philosophy in North America, Britain, and elsewhere. Axiomatic rationality is a case in point. Abstract axioms have been interpreted as categorical norms we should follow, even as universal prescriptions without specified limits. When experiments showed that humans systematically deviate from these norms, the discrepancy was attributed to flaws in humans rather than in the norms. In this article, I will argue that axiomatic rationality can, if at all, provide norms only in *small worlds*, exemplified by lotteries, where the exhaustive and mutually exclusive set of future states of the worlds and consequences is known or knowable, and when these norms are limited to logical coherence. Outside small worlds, I propose a naturalized version of rationality for situations of uncertainty (as opposed to risk) and intractability (Gigerenzer and Sturm 2012). In these situations, humans can achieve their goals by relying on fast-and-frugal heuristics that may violate axiomatic rationality. These goals go beyond coherence to include predictive accuracy, frugality, and efficiency. In general terms, a heuristic is ecologically rational to the degree that it is adapted to the structure of an environment. The study of the ecological rationality requires formal models of heuristics and an analysis of the structures of environments these can exploit. It lays the foundation of a moderate naturalism in epistemology, providing statements about heuristics we *should* use in a given situation.

1 Axiomatic rationality and risk

Von Neumann and Morgenstern are credited with having formulated the first set of choice axioms. However, a normative interpretation of the axioms or the maximization of expected utility is absent in the three editions of their *Theory of Games and Economic Behavior* (1944, 1947, 1953). Their great contribution was to prove that if an individual satisfies a set of axioms, then choices can be represented by a utility function, similar to the axioms that guarantee that elements can be represented on a number line. By itself, a representation theorem does not imply a prescription of what people should do.

When Savage (1954), one of the founders of modern Bayesian decision theory, laid out his own set of axioms, he attached a normative interpretation. But he also stated limits to his theory using two specific examples, not general principles. The examples were *playing chess* and *planning a picnic* (Savage 1954, p. 16). To such situations, axiomatic rationality does not apply.

What are the general principles behind these two examples? We can work these out by considering Savage's central concept of a *small world*. A small world consists of a set S of mutually exclusive and exhaustive future states of the world and a set C of mutually exclusive and exhaustive consequences of one's actions if a particular state occurs. Actions are defined as mappings from states of the world to consequences. States of the world must necessarily be described at some limited level of detail, hence the qualifier *small*. Technically, a small world is described by the pair (S, C) .

Playing chess lies outside of Savage's small worlds because no human or machine can determine the exhaustive set S of all possible states—here, all sequences of possible moves—and choose the optimal one. To understand the order of magnitude of

this limitation, note that chess has approximately 10^{120} unique sequences of moves or games, a number greater than the estimated number of atoms in the universe. In computer science, such problems are called *computationally intractable* (with subdivisions into *NP-hard*, *NP-complete*, etc.). An intractable problem is defined as one for which no efficient (i.e., polynomial-time) algorithm to solve it exists. Thus, the first general principle that limits axiomatic rationality is *intractability*.

Planning a picnic, by contrast, is an ill-defined problem. A problem can be ill-defined in several respects: the set S may not be known because the future is uncertain, or the set C may not be known because of unexpected events and accidents, or because the problem is unfamiliar and decision time is scarce. For Savage (1954), the proverbs “Look before you leap” and “You can cross that bridge when you come to it” mark the demarcation line between the narrow domain of axiomatic rationality and the world beyond:

Carried to its logical extreme, the “Look before you leap” principle demands that one envisage every conceivable policy for the government of his whole life (at least from now on) in its most minute details, in the light of the vast number of unknown states of the world, and decide here and now on one policy. This is utterly ridiculous ... It is even utterly beyond our power to plan a picnic or to play a game of chess in accordance with the principle. (p. 16)

The ‘look before you leap’ principle exemplifies the three pillars of Savage’s decision theory: a set of choice axioms, the maximization of subjective expected utility, and Bayesian updating of probabilities. ‘You can cross that bridge when you come to it,’ in contrast, represents situations where (S, C) is not known or knowable.

For the general principle underlying the example of planning a picnic, Binmore (2008) speaks of *large worlds*; I will use the term *uncertainty*. In doing so, I connect Savage’s two examples with the distinction between *risk* and *uncertainty*, as proposed by Knight (1921). For Knight, risk refers to situations where the probabilities are known, either by design or from relative frequencies in the long run. Although Knight does not explicitly refer to unknown state spaces, knowledge of the latter is a prerequisite for knowing the probability distribution. For the purpose of this article, I will thus treat situations of risk and small worlds (S, C) as identical.

Conjecture 1 *The normative power of axiomatic rationality is limited to small worlds.*

This conjecture is consistent with my reading of Savage (his writing is unfortunately not known for having the same clarity as his axioms). In plain words, axiomatic rationality cannot prescribe how we should make decisions outside small worlds, whether or not one accepts coherence as the only norm.

Nevertheless, axiomatic decision theory became interpreted as a theory of rationality without specified bounds, ignoring the concerns of Savage, Allais, Ellsberg, and others. This intuition-based, categorical interpretation is remarkable, given that the ideal of a universal calculus proposed by Leibniz had been ridiculed for centuries. Erickson et al. (2013) argued that the interpretation was partly motivated by the cold war, with its threat of mutual nuclear destruction in the 1960s and 1970s. Abstract rationality, utility theory, and game theory embodied the hope that reason could over-

come the emotions of a Khrushchev or Kennedy. For instance, Searle (2001, p. 6) reported about a friend who was a high official at the Pentagon:

He went to the blackboard and drew the curves of traditional microeconomic analysis; and then said, “Where these two curves intersect, the marginal utility of resisting is equal to the marginal disutility of being bombed. At that point, they have to give up. All we are assuming is that they are rational. All we are assuming is that the enemy is rational!”

For cold war rationality, the true enemy was uncertainty and intractability.

In summary, my first and modest proposal is to limit the normative force of axiomatic rationality to small worlds. The proposal is modest because it refrains, for the purpose of this article, from further questioning the normative value of axiomatic rationality in small worlds, as Allais, Ellsberg, and others have done, and from arguing that axiomatic rationality is merely a description of what Savage or others intuitively believe we should believe (Bishop and Trout 2005).

1.1 What is the probability that a problem is intractable or uncertain?

To answer this question, it would be necessary to know the set of all problems a person could encounter, which would entail knowing the exhaustive and mutually exclusive set of all possible situations. This requirement makes a picnic pale in comparison with regard to uncertainty. But we can get an idea of the commonness of intractability and uncertainty by means of several examples.

Consider intractability first. The Travelling Salesperson Problem is likely the best-known scheduling problem that is nondeterministic polynomial (NP) hard. To illustrate, consider a politician who runs for president in the US and plans to tour the country’s 50 largest cities, starting and ending in the same city. How can she determine the shortest route among the approximately 3×10^{62} routes? Not even the fastest computers are able to check this many possibilities in the candidate’s lifetime, let alone before the campaign. A review of scheduling problems reported that 84% were shown to be intractable, 9% were tractable, and 7% had unknown status (Lawler et al. 1993). More generally, computer scientists have argued that most interesting problems are computationally intractable in any implementation, be it neural or machine (Tsotsos 1991).

However, despite intractability, humans can identify ‘near-optimal’ solutions almost effortlessly using heuristics. One simple rule is the *nearest-neighbour heuristic*: “Start with your home city; find the nearest unvisited city and go there; continue until all cities have been visited.” It can provide a useful strategy when the optimal one is unattainable.

Intractability is also a serious limit for the epistemic responsibility of evaluating whether one’s beliefs are coherent, that is, conform to the axioms. Consider the completeness axiom, one of the choice axioms that are necessary so that preferences can be represented by a utility function.

Completeness: $A \succeq B$ or $B \succeq A$.

Completeness means that one prefers either A weakly over B or the opposite. Everything else is excluded, such as not having any preference or not making a choice. This axiom appears almost trivial to satisfy, yet it is not. Consider choosing which websites to visit and in which order. According to Internet Live Stats, 10 websites existed on the Internet in 1992. To order these according to preference, one had to make 45 ($10 \times 9/2$) binary choices. At that time, checking for completeness was tractable. In the year 2016, the number of websites had increased to about 1,085,628,900, which would require in the order of 10^{18} checks. This is no longer tractable, neither for humans nor machines. And without being able to check for completeness, one cannot check for transitivity. Similarly, checking for consistency in probabilistic inferences in Bayesian belief networks is NP-hard; the same holds for approximations (Cooper 1990; Dagum and Luby 1993). In general, the more beliefs a person holds, the more likely that checking coherence becomes intractable. Intractability entails that the principle ‘ought implies can’ is invalid because no mind or machine can do what one ought to do according to axiomatic rationality.

A similar argument can be made for the frequency of encountering uncertainty. The axioms refer to well-defined beliefs or options that are mutually exclusive, like the p 's and q 's in truth tables, whereas many nontrivial political, philosophical, or moral beliefs tend to have fuzzy borders, which makes the task of deciding whether these beliefs conform to a set of axioms fraught with uncertainty. In general, the set of future states S and their consequences C is typically unknown in human affairs, medicine, or finance, all of which are at least as uncertain as when planning a picnic. The prevalence of uncertainty in important affairs can also be inferred from the notable absence of detailed real-world examples in writings about axiomatic rationality.

Savage was relatively open about the limits of axiomatic rationality and pointed out that these already affected his conception of small worlds. A person can always consider a more refined small world, an analysis of which may not agree with the original unrefined small world. The ultimate refinement would be what Savage called the *grand world*. In his book, Savage (1954, p. 467) provided a single example of a small world and none of a grand world.

To summarize, intractability implies that the empirical validity of choice axioms cannot be verified, and uncertainty implies that the assumption of a small world (S, C) underlying the choice axioms is invalid in the first place. Although the probability of encountering a situation of uncertainty or intractability is undefined, one might conclude that these situations are the rule rather than the exception. That result implies that a categorical interpretation of choice axioms as a universal norm for behavior is impossible and thus invalid. If that is true and Conjecture 1 is true, then this also questions the ‘instrumental’ justification of categorical norms, namely the claim that violations of logical axioms are associated with substantial costs. Rather, these violations should have little impact on whether or not people reach their goals in everyday life.

1.2 How bad is incoherence?

Many of Savage's followers have ignored intractability and uncertainty, and interpreted violations of logical coherence as signs of human irrationality. In this research, the term *coherence* not only refers to conformity with choice axioms such as transitivity but also includes truth-table logic and rules of probability (e.g., Tversky and Kahneman 1974). I will use the term *logical rationality* for this broader set of rules that includes axiomatic rationality. A fundamental problem with logical rationality (as opposed to axiomatic rationality) is that the various logical and statistical rules proposed as norms do not speak unanimously (otherwise, there would not be centuries of debates among statisticians), and, therefore, a judgment that is diagnosed as "irrational" because it violates one rule, such as modus tollens, can be justified as rational because it satisfies another rule, say Bayes' rule (see Gigerenzer et al. 2012). Here, I do not pursue this internal ambiguity of logical rationality but instead ask whether violations of logical rules actually matter in the real world.

An entirely new discipline, behavioral economics, was created in the 1980s and 1990s to pursue the program of identifying systematic deviations from logical rationality, which was made popular by bestsellers proclaiming that "we are not only irrational, but predictably irrational" (Ariely 2008, p. xviii). Wikipedia lists some 175 cognitive illusions, many of which are violations of coherence. Based on the claim that people's irrationality is persistent, *libertarian paternalists* (Thaler and Sunstein 2008) proposed that governments should "nudge" their citizens into better behavior—to protect them not from external enemies, but from themselves. In this view, people who deviate from logical rationality face economically significant losses (e.g. Thaler and Sunstein 2008; Yates 1990). If true, this would provide indirect evidence for the normative force of logical rationality in the real, mostly uncertain world.

To investigate whether incoherence indeed implies costs in the real world, Arkes et al. (2016) conducted a systematic literature search on the evidence for detrimental material consequences such as false beliefs, lower earnings, impaired health, lower happiness, or shorter lives. In the over 100 studies on violations of transitivity that they identified, not a single demonstration was found of a person becoming a *money pump*, that is, being continually exploited due to intransitive choices. In the more than 1000 articles on preference reversals identified, of which only four actually tested whether these turn people into money pumps or otherwise impose any costs, they found that arbitrage or financial feedback made preference reversals and their costs largely disappear. Arkes et al. then analyzed hundreds of studies on the Asian Disease Problem and other framing effects, and found little to no evidence that 'irrational' attention to framing would be costly. The same result was found in the literature for violations of the independence axiom, the Chernoff condition, and other 'fallacies.'

Lack of evidence for costs should not be misinterpreted as evidence for lack of costs. However, this striking absence suggests that the large and ever-growing list of apparent fallacies is a list of "logical bogeymen," as psychologist Lola Lopes once put it, with little measurable economic or psychological consequences (for details see Gigerenzer 2018). Moreover, violations of coherence were found to be beneficial in some studies. For instance, Berg et al. (2011) reported that people who violated time-consistency

and expected utility theory earned higher monetary payoffs than those who did not, while Houston et al. (2007) reported that fitness maximization can imply violations of transitivity. These results were obtained in situations of risk, not uncertainty. When investigating the decision whether to participate in PSA screening for prostate cancer, Berg et al. (2016) reported that among 133 economists, coherent Bayesians had no more true beliefs than did incoherent Bayesians: the correlation between coherence and accuracy of beliefs was zero, even slightly negative. The most consistent economist had the highest number of false beliefs.

Conjecture 2 *There is lack of evidence that violations of logical rationality have detrimental consequences on people's wealth, health, happiness, proportion of true versus false beliefs, or some other measurable cost.*

Note that Allais, Ellsberg, and others have argued that the axioms have questionable normative force even in Savage's small world. My point, however, is an empirical one: Arkes, Gigerenzer, and Hertwig's literature search did not find any studies that systematically showed that violations of coherence actually matter for goals beyond coherence. This notable absence of evidence poses a problem for a consequentialist justification of axiomatic rationality.

My conclusion is that axiomatic rationality needs to be complemented with a normative theory of rationality that can deal with intractability and uncertainty, and with goals beyond coherence.

2 Ecological rationality and uncertainty

How should we deal with uncertainty? Knight (1921) spoke of "intuitive feelings," "judgment," and "experience," without offering a formal theory. In his *General Theory*, Keynes (1936) suggested relying on our animal spirits, on our spontaneous urge to action, optimism, and hope, all of which make the wheels go round. Yet he too offered no formal alternative. The roots for such an alternative can be found in Herbert Simon's (1979) work on bounded rationality and, specifically, in the concept of *heuristics* (Gigerenzer et al. 2011; Gigerenzer and Selten 2001; Todd et al. 2012).

What are heuristics? The term is of Greek origin, meaning "serving to find out or discover." With its introduction into English in the early 1800s, it referred to a useful tool for solving problems that cannot easily be handled by logic or probabilistic inference. George Polya, Max Wertheimer, and Herbert Simon defined heuristics in similar ways, as tools for finding a proof, solving a novel problem, and planning next year's budget. Einstein entitled his fundamental paper on quantum physics from 1905 "On a heuristic point of view concerning the generation and transformation of light." He used the term *heuristic* to indicate that the view presented was incomplete, even false, yet nonetheless useful and of great transitory value on the path to building a more correct theory (Holton 1988, pp. 360–361). This favorable image of heuristics took a negative bend in psychology around 1970, when heuristics became associated with errors in the heuristics-and-biases program (Tversky and Kahneman 1974). In this influential work, deviations between coherence rules and people's judgments became attributed post hoc to heuristics such as *availability*. The relevant point here is that the

heuristics-and-biases program subscribed to the classical view in epistemology that axiomatic rationality is normative and psychology strictly descriptive. Yet that view of psychology is incorrect.

2.1 Normative psychology

In *Epistemology Naturalized*, Quine (1969) argued that epistemology is a branch of psychology. His argument was rejected on the grounds that it has a devastating implication: to empty epistemology of its normative character (Bishop 2006). The objection was based on the belief that psychology is a strictly descriptive science.

Conjecture 3 *Parts of psychology are normative. Theoretical and empirical results are used to prescribe what means people should use to achieve goals.*

Consider the following two illustrations, one for situations of risk, the other for uncertainty. It was found that laypeople (Tversky and Kahneman 1980) and physicians (Eddy 1982) had great difficulties making Bayesian inferences from conditional probabilities (such as the sensitivity and specificity of a cancer screening test). Yet when the information was presented in natural frequencies (simple joint frequencies that have not been conditionalized), Bayesian reasoning substantially improved both in undergraduates (Gigerenzer and Hoffrage 1995) and physicians (Gigerenzer 2014). The theoretical explanation for this empirical result is that natural frequencies facilitate the computation of posterior probabilities. This result leads to the prescription:

People *ought* to use natural frequencies to improve Bayesian reasoning.

This prescription amounts to an inference from is to ought. It is valid for situations of risk where the assumptions necessary for Bayes' rule hold (as opposed to situations of uncertainty). Using natural frequencies (as opposed to conditional probabilities) helps to reduce the errors in medical diagnosis and in the evaluation of evidence in court. It leads to further prescriptions such as that natural frequencies should be used to teach Bayesian reasoning, which is currently implemented in both high school textbooks and medical curricula (Gigerenzer 2014).

Consider next a problem companies face: How should managers predict which customers will continue to make purchases in the future? Wübben and von Wangenheim (2008) observed that experienced managers rely on the *hiatus heuristic*: “If the customer has not made a purchase within 9 months, delete from the customer base, otherwise not”. To rely on a single reason—the hiatus—contradicts standard customer base models such as the Pareto/NBD (negative binomial distribution) model, which process more cues and rely on complex stochastic models to estimate for each customer the probability that he or she will make future purchases. Testing both alternatives, they found that the simple heuristic predicted more accurately than the complex model, even though (or because) it did not use the total evidence available. Similarly, heuristics such as *fast-and-frugal trees* and *take-the-best* (see below) also rely on solely one reason to make a prediction, although they may initially search through more cues. Green and Mehr (1997) reported that these one-reason heuristics can predict the risk of ischemic heart disease more accurately and rapidly than standard logistic regression

models. The study of ecological rationality analyzes the conditions E under which relying on one good reason rather than on linear rules that use all cues can be expected to lead to more accurate inferences and is faster and involves less information to boot (see Sects. 2.5–2.9). These results lead to a second prescription:

If conditions E hold, then people *ought* to rely on one-reason heuristics rather than on linear models.

Note that ignoring part of the cues and relying on only one reason has been previously interpreted as a cognitive error in the heuristics-and-biases program and attributed to our cognitive limits. It also appears to conflict with the principle of total evidence. Under uncertainty, however, this can be a rational strategy. Less can be more.

These examples demonstrate that psychology is both a descriptive and a normative discipline. Thus, naturalizing epistemology by taking psychology into account does not imply emptying epistemology of its normative content (Bishop and Trout 2005). What it does do is question the divide of disciplines into strictly normative and descriptive ones, and the associated justification of a priori, intuitive, and categorical norms (Schurz 2014).

2.2 Methodological principles

The study of ecological rationality is characterized by three methodological principles:

1. Formal models of heuristics, as opposed to vague labels.
2. Competitive testing of heuristics, as opposed to null hypothesis tests.
3. Tests of predictive power, such as in out-of-sample prediction, as opposed to data fitting.

The study by Wübben and von Wangenheim (2008) embodies these principles: The hiatus heuristic is a formal model, it is tested against the best competitors in the field, and it predicts future behavior rather than fitting previously known data. I emphasize these principles because they have been largely neglected in past research on heuristics. Heuristics such as *availability* and *representativeness* or terms such as “system 1” (Kahneman 2011) are vague labels, and thus can neither be tested against competitors nor predict behavior. The function of these labels is to ‘explain’ post hoc deviations from axiomatic rationality (Gigerenzer 1996). Post hoc data fitting is a key methodological vice in studies of rationality. Like vague labels, the use of multiple free parameters that are rarely ever fixed allows for fitting any data well without being able to predict well. Consider expected utility theory and its modifications such as cumulative prospect theory, which have up to five adjustable parameters. In a review of half a century of economic literature, Friedman et al. (2014) conclude that their “power to predict out-of-sample is in the poor-to-nonexistent range” (p. 3).

2.3 Epistemic goals

Epistemic rationality is often seen as promoting coherent beliefs rather than incoherent ones, and true beliefs rather than false beliefs. Implicitly, the assumption is often that

coherence is associated with truth. Yet, as we have seen, there is lack of evidence both that coherent beliefs are correlated with true beliefs and that incoherence has substantial costs. Furthermore, in situations of uncertainty, coherence is no longer clearly defined. All of this requires extending epistemic goals from coherence to other goals such as the accuracy of judgment (truth), its speed (how fast a judgment can be made), and frugality (how many cues need to be searched before a judgment can be made). To decide about the ecological rationality of a heuristic, these goals need to be clearly defined. For instance, accuracy can be decomposed into hit rate (such as the proportion of patients who are diagnosed as having a disease among those who actually have the disease) and false alarm rate (the proportion of patients who are diagnosed as having a disease among those who do not have it). Heuristics can be designed to achieve a desired balance between the two rates. For instance, the question of which fast-and-frugal trees are ecologically rational for which balance of hits and misses has been solved (Luan et al. 2011).

The extension of goals of rationality from coherence to performance has been called the “boldest claim” inherent in ecological rationality (Rich 2018, p. 541). Yet this extension has much in common with existing approaches, such as Kitcher’s (1992) naturalism and Goldman’s (1999) epistemological reliabilism.

2.4 Less-is-more effects

According to one view, heuristics are subject to an *accuracy–effort trade-off* (Kahneman 2011; Shah and Oppenheimer 2008): Heuristics save effort but at the cost of accuracy. In this view, the rationality behind relying on heuristics lies solely in reducing effort, winning time, and avoiding information search. This is probably the dominant interpretation of heuristics in philosophy and the social sciences.

Conjecture 4 *Accuracy–effort trade-offs hold in situations of risk where the optimal course of action is calculable but not necessarily in situations of uncertainty. Under uncertainty, less-is-more effects exist.*

The finding that the hiatus heuristic predicts more accurately than the Pareto/NBD model is a case in point. Although the Pareto/NBD model has all the information the heuristic uses and more, the heuristic generates more accurate predictions (for other less-is-more effects, see Gigerenzer et al. 2011).

Less-is-more effect Assume two strategies, P and T . P uses only a proper subset of the information that T uses. If P makes more accurate predictions, this is called a less-is-more effect.

Less-is-more does not imply a monotonic relationship between decreasing effort and increasing accuracy. Rather, there is generally an inverse U-shaped relationship between effort and accuracy, meaning that after some point on this curve, more information search is not only costly but also reduces accuracy. This phenomenon has been discussed as apparent epistemic irresponsibility (Bishop 2000). Less-is-more effects should not occur if accuracy–effort trade-offs were generally true. To understand why and when less-is-more effects can be expected, I will first introduce the bias–variance dilemma as a general alternative to the accuracy–effort trade-off in situations




Fig. 1 A visual analogy of the two components of prediction error: bias and variance. The bull's eye is the unknown true value μ (here: 0, 0) to be predicted. Each dart represents a predicted value x_m based on a random sample from the population with the true value μ . Bias is zero if the mean prediction 'hits' the target. Left: A systematic bias, whose size is the distance between the mean of the darts thrown and the bull's eye ($\bar{x} - \mu$), and a low variance, that is, the darts are close together. Right: Zero bias ($\bar{x} - \mu$), that is, the darts are lined up exactly around the bull's eye, but considerable variance. The example illustrates situations where biased minds can make more accurate predictions than unbiased ones do (Gigerenzer 2016)

of uncertainty, and thereafter a specific analysis of the ecological rationality for the take-the-best heuristic and similar one-good-reason heuristics.

2.5 The bias–variance dilemma

Consider a minimal form of uncertainty: the problem of estimating the true value μ in a population on the basis of random samples. Each of M samples ($m = 1, \dots, M$) generates an estimate x_m , with \bar{x} as their mean. This situation involves uncertainty because the true value is not known. Yet the uncertainty is minimal because the population is stable and the samples are random. Here, the total error has three sources (Geman et al. 1992):

$$\text{Prediction error} = (\text{bias})^2 + \text{variance} + \epsilon, \quad (1)$$

where ϵ is unsystematic noise (mean zero and uncorrelated with bias), and bias = $\bar{x} - \mu$, that is, the average deviation of the mean of the sample estimates from the true value. For instance, if the true temporal trajectory of a variable is a polynomial of second degree and a linear regression is used to predict the variable, the model has a systematic bias. Variance = $\frac{1}{m} \sum (x_m - \bar{x})^2$, that is, the mean squared deviation of the sample estimates from their mean \bar{x} . The variance component reflects the sensitivity of the predictions to different samples drawn from the same population. Variance decreases with larger sample sizes and increases with the number of free parameters estimated. Figure 1 provides a visual depiction of bias and variance (Brighton and Gigerenzer 2012).

In it, the bull's eye represents the true value, and each dart the estimate from a sample. The darts on the left-hand dartboard show a systematic bias but low variance. In contrast, the darts on the right-hand dartboard are lined up exactly around the bull's eye and show no bias but considerable variance. A moderate bias with low variance (left) may lead to better accuracy than would a zero bias with high variance.

The variance component of the error corresponds to the concept of *overfitting*. A model with many free parameters may fit the data perfectly but predict worse than simpler models (Forster and Sober 1994).

2.6 Heuristics reduce error due to variance

Heuristics can reduce error due to variance in several ways. A total reduction of variance can be achieved by a heuristic that is insensitive to data by using no adjustable parameters. In Fig. 1, this insensitivity would correspond to a set of darts that all end up at the same location, showing no variance due to fluctuations in samples. A hiatus heuristic with a fixed hiatus is an example. Such a heuristic will avoid being overly sensitive to the peculiarities of the sample information available. Another case in point is the $1/N$ heuristic in investment, which divides a sum of money equally over N options. In contrast, Markowitz's mean–variance portfolio—for which Markowitz was awarded the Nobel prize in economics—calculates the 'optimal' weights for each option from the available data. Yet $1/N$ can lead to better returns than the mean–variance portfolio can (DeMiguel et al. 2009). The hiatus heuristic and $1/N$ ignore the total evidence from sample data.

In other situations, however, ignoring the evidence from samples may increase bias to the extent that the total error increases. Heuristics that learn from samples can still reduce error due to variance by (1) ignoring valid predictors, (2) not estimating weights, and (3) not estimating the covariance matrix between cues or reasons and treating the cues as independent. Heuristics such as take-the-best combine all three of these principles.

The purpose of the following case study is to explain in more detail the study of ecological rationality and how it differs from axiomatic rationality.

2.7 A case study of ecological rationality: take-the-best

Consider the task of inferring which of two alternatives has the larger value on some criterion, such as which contestant will win a tennis match or which high school will have more drop-outs. To make this inference, there are n cues or reasons. Experiments showed that people tend not to use all cues, even if each is valid, and often proceed in a lexicographic order. The term *lexicographic* has its origin in the way one looks up a word in a lexicon; first, one searches for the first letter, then the second, and so on. In decision theory, the term refers to a process in which one looks up cues in sequential order and can stop search immediately after the first or a later cue if a stopping rule is satisfied. Because lexicographic choice cannot be mapped into a utility function, and may not conform to the choice axioms, it has been interpreted a priori as irrational.

The take-the-best heuristic is a model of lexicographic choice (Gigerenzer and Goldstein 1996). Like many heuristics, take-the-best has three building blocks: a search rule, a stopping rule, and a decision rule. For convenience, assume that all cues are binary (0 and 1) and the cue value that signals a higher criterion value is 1.

1. *Search rule* Search cues in order of their validity v .
2. *Stopping rule* Stop search on finding the first cue that discriminates between the alternatives (i.e., cue values are 1 and 0, or 0 and 1).
3. *Decision rule* Infer that the alternative with the positive cue value (1) has the higher criterion value.

The validity v of a cue is given by:

$$v = C/(C + W), \quad (2)$$

where C is the number of correct inferences when a cue discriminates, and W is the number of wrong inferences, all estimated from samples.

Numerous studies have shown that in situations where take-the-best is ecologically rational (see below), a large proportion of people tend to rely on it. This includes student populations (e.g., Bergert and Nosofsky 2007; Bröder 2012), airport customs officers, police officers, and burglars (e.g., Garcia-Retamero and Dhami 2009; Pachur and Marinello 2013). However, people who rely on this heuristic appear to commit several transgressions against logical rationality: First, as mentioned before, lexicographic choices cannot be represented by a utility function. Second, choices appear to violate the principle of total evidence by (1) ignoring all dependencies between cues, that is, the entire covariance matrix, when ordering the cues by v , and (2) ignoring all other cues after the first cue is found that allows for a judgment, leaving valid information on the table. Third, variants of the stopping rule can lead to systematic intransitivity (Arló-Costa and Pedersen 2013). Each of these properties has been interpreted as a cognitive error. In the words of Keeney and Raiffa (1993), for instance, lexicographic rules are “naively simple” and “will rarely pass a test of ‘reasonableness’” (p. 78).

However, Keeney and Raiffa argued solely from their a priori normative view of axiomatic rationality, without any testing. Since the mid-1990s, others have gone on to conduct such tests, showing that take-the-best can not only model people’s choices (the descriptive question) but also predict ‘objective’ criteria as accurately as or better than complex models, including multiple regression and sophisticated machine learning algorithms such as classification and regression trees, support vector machines and random forests (Brighton and Gigerenzer 2015; Czerlinski et al. 1999; Şimşek and Buckmann 2016). Take-the-best corresponds to the left-hand dart board in Fig. 1, with systematic bias but low variance, while complex models with many free parameters correspond to the right-hand dartboard, unless sample sizes are very large.

The amount of bias take-the-best has depends on the structure of the environment. What are the environmental conditions E that take-the-best and similar one-reason heuristics can exploit?

2.8 Ecological axioms

The term *environment* refers to the alternatives, cues, criteria, and other factors relevant for the decision maker. The environment determines the bias of a heuristic and other strategies. Which environmental structures ‘help’ lexicographic heuristics perform well so that they have a small bias (in addition to small variance)? Because the true value may not be accessible and will differ from problem to problem, the question asked is a comparative one: Can we identify general environmental structures in which the bias of take-the-best is *equal* to that of linear models? Answering this question helps identify situations where take-the-best can be expected to be more accurate than linear models, that is, when the bias is the same but the heuristic generates less error by variance, resulting in smaller total error (Eq. 1).

Consider a choice between objects A and B , based on n cues, where the value of the i th cue is represented by x_i and weighted in the linear payoff function by w_i . To simplify, assume that the cues are binary and the weights are nonnegative. We know of three environmental conditions where the bias of take-the-best is the same as that of linear models: noncompensatoriness, dominance, and cumulative dominance (Gigerenzer 2016). These features can be seen as ecological axioms:

Noncompensatoriness. The weights $w_1, w_2, w_3, \dots, w_n$ are noncompensatory if they satisfy the $n - 1$ inequality constraints:

$$w_i > \sum_{j=i+1}^n w_j, \quad i = 1, 2, \dots, n-1. \quad (3)$$

An example is the set of weights $\{1, 1/2, 1/4, 1/8\}$. If the weights are noncompensatory, then a linear rule (with the same order of cues) will always lead to the same choice as a lexicographic rule (Martignon and Hoffrage 2002). Take the example above. If the lexicographic rule yields decisions on the basis of the first cue (with weight 1), every linear rule will match this choice because the sum of all other weights ($1/2 + 1/4 + \dots$) will always be smaller than the weight of the first cue. Thus, if noncompensatoriness holds, then a lexicographic heuristic will have the same bias as any linear model with the same order of cues.

Dominance If alternative A has a value higher than or equal to alternative B on all n cues and a higher value on at least one cue, then alternative A dominates alternative B . Thus, if dominance holds, then a lexicographic heuristic will have the same bias as any linear model.

Cumulative Dominance The *cumulative profile* of an alternative consists of n values, where the i th value is the sum of the first i values. Alternative A cumulatively dominates B if its cumulative profile exceeds or equals the cumulative profile of B in every term and exceeds it in at least one term (Baucells et al. 2008). Dominance implies cumulative dominance, but not vice versa. If cumulative dominance holds, then a linear rule (with the same order of cues) predicts the cumulative dominant object, just as a lexicographic rule does.

In sum, if either noncompensatoriness, dominance, or cumulative dominance holds, then take-the-best or similar lexicographic heuristics will have the same bias as a linear

model that relies on more cues and ‘optimal’ weighting. In that case, a lexicographic heuristic can be said to be ecologically rational relative to any linear model because one can expect (at least) the same accuracy with less effort. The three conditions explain when to expect this. They do not explain when and why heuristics can predict more accurately, which can, however, be understood from Eq. 1: heuristics tend to reduce error by variance, so that the total error of the heuristic can be less than that of a linear model, such as when the bias of the heuristic is not much higher than that of the linear model.

Noncompensatoriness refers to the relative strength of the cues in the environment, while the two dominance conditions refer to the relative quality of alternatives (Katsikopoulos 2011). The result can be generalized from take-the-best to other sequential search (lexicographic) heuristics, with single-cue heuristics as a special case. Now we can make the conditions E in the prescription in Sect. 2.1 explicit:

If conditions E hold—noncompensatoriness, dominance, or cumulative dominance—then people *ought* to rely on take-the-best rather than on linear models.

2.9 How often do these favorable conditions hold?

Şimşek (2013) analyzed 51 natural data sets from online repositories, textbooks, research publications, packages for R statistical software, and individual scientists’ collected field data. The data sets spanned areas as diverse as biology, business, computer science, ecology, economics, education, engineering, and medicine, among others. The number of cues ranged from 3 to 21, which were numeric or binary; the number of objects (alternatives) ranged from 12 to 601, which resulted in a number of possible pairwise comparisons ranging from 66 to 180,300. In each of these comparisons, Şimşek examined how often one or more of the three conditions—noncompensatoriness, dominance, and cumulative dominance—was satisfied. The result was surprising. The median for the 51 data sets was 90% (Şimşek 2013). That is, in half of the data sets, 90% or more of the decisions encountered were such that a lexicographic rule yielded the same prediction as a linear model. When the continuous cues were dichotomized at their medians, that is, transformed into binary cues, this number increased to 97%. This means that in the majority of the cases, the lexicographic heuristics had the same bias as a linear model. Together with their potential for reducing variance, these results explain when and why simple heuristics can outperform linear models in prediction.

3 Rationality under uncertainty and intractability

In this article, I argued that the domain of axiomatic rationality is defined by Savage’s small worlds (S, C) where the exhaustive and mutually exclusive set of future states S of the world and their consequences C is known. My first conjecture is that outside of these stable, well-defined situations, axiomatic rationality has no normative force. My second, related conjecture is that despite the widespread interpretations of violations of logical rationality as signs of human irrationality, there exists little to no empirical

evidence that these violations would incur substantial costs such as diminished health, wealth, or happiness. The third conjecture is that psychology can inform prescriptions for what people should do to achieve a given goal, which is the topic of the theory of ecological rationality. The study of heuristics extends axiomatic rationality from small worlds to situations of uncertainty and intractability, and from coherence to performance goals such as predictive accuracy, speed, and frugality. Developing such a theory entails the descriptive study of how individuals and institutions actually make decisions and the prescriptive study of the ecological rationality of heuristics. It also teaches us that in situations of uncertainty, less information can be more beneficial.

These results conflict with the ideal of an a priori, categorical interpretation of axiomatic rationality and, more generally, logical rationality. The very existence of uncertainty and intractability defies choice axioms or logical rules as universal norms. Ecological rationality, in contrast, emphasizes the adaptive character of rules or heuristics to reach goals. It provides a formal approach to what is called instrumental or practical rationality. Last but not least, it can help to put an end to the idea that psychology has nothing to offer for understanding the nature of rationality.

Acknowledgement Open access funding provided by Max Planck Society. I would like to thank Carl Hoefler and Thomas Sturm for helpful comments.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Ariely, D. (2008). *Predictably irrational*. London: Harper Collins.
- Arkes, H. R., Gigerenzer, G., & Hertwig, R. (2016). How bad is incoherence? *Decision*, 3, 20–39.
- Arló-Costa, H., & Pedersen, A. P. (2013). Fast and frugal heuristics: Rationality and the limits of naturalism. *Synthese*, 190, 831–850.
- Baucells, M., Carrasco, J. A., & Hogarth, R. M. (2008). Cumulative dominance and heuristic performance in binary multiattribute choice. *Operations Research*, 56, 1289–1304.
- Berg, N., Biele, G., & Gigerenzer, G. (2016). Consistent Bayesians are no more accurate than non-Bayesians: Economists surveyed about PSA. *Review of Behavioral Economics*, 3, 189–219.
- Berg, N., Eckel, C., Johnson, C. (2011). *Inconsistency pays? Time-inconsistent subjects and EU violators earn more*. Unpublished Manuscript, University of Texas, Dallas.
- Bergert, F. B., & Nosofsky, R. M. (2007). A response-time approach to comparing generalized rational and take-the best models of decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 33(1), 107–129.
- Binmore, K. (2008). *Rational decisions*. Princeton, NJ: Princeton University Press.
- Bishop, M. A. (2000). In praise of epistemic irresponsibility: How lazy and ignorant can you be? *Synthese*, 122, 179–208.
- Bishop, M. A. (2006). Fast and frugal heuristics. *Philosophy Compass*, 1(2), 201–223.
- Bishop, M. A., & Trout, J. D. (2005). The pathologies of standard analytic epistemology. *NOUS*, 39, 696–714.
- Brighton, H., & Gigerenzer, G. (2012). Are rational actor models “rational” outside small worlds? In S. Okasha & K. Binmore (Eds.), *Evolution and rationality: Decisions, co-operation and strategic behavior* (pp. 84–109). Cambridge: Cambridge University Press.
- Brighton, H., & Gigerenzer, G. (2015). The bias bias. *Journal of Business Research*, 68, 1772–1784.

- Bröder, A. (2012). The quest for take-the-best. In P. M. Todd, G. Gigerenzer, & the ABC Research Group (Eds.), *Ecological rationality: Intelligence in the world* (pp. 216–240). New York: Oxford University Press.
- Cooper, G. F. (1990). The computational complexity of probabilistic inference using Bayesian belief networks. *Artificial Intelligence*, 42, 393–405.
- Czerlinski, J., Gigerenzer, G., & Goldstein, D. G. (1999). How good are simple heuristics? In G. Gigerenzer, P. M. Todd, & the ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 97–118). New York: Oxford University Press.
- Dagum, P., & Luby, M. (1993). Approximating probabilistic inference in Bayesian belief networks is NP-hard. *Artificial Intelligence*, 60, 141–153.
- DeMiguel, V., Garlappi, L., & Uppal, R. (2009). Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy? *Review of Financial Studies*, 22, 1915–1953.
- Eddy, D. M. (1982). Probabilistic reasoning in clinical medicine: Problems and opportunities. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases*. New York: Cambridge University Press.
- Erickson, P., Klein, J. L., Daston, L., Lemov, R., Sturm, T., & Gordin, M. D. (2013). *How reason almost lost its mind: The strange career of cold war rationality*. Chicago: University of Chicago Press.
- Forster, M., & Sober, E. (1994). How to tell when simpler, more unified, and less *ad hoc* theories will provide more accurate predictions. *British Journal of the Philosophy of Science*, 45, 1–35.
- Friedman, D., Isaac, R. M., James, D., & Sunder, S. (2014). *Risky curves. On the empirical failure of expected utility*. New York: Routledge.
- Garcia-Retamero, R., & Dhami, M. K. (2009). Take-the-best in expert-novice decision strategies for residential burglary. *Psychonomic Bulletin & Review*, 16, 163–169.
- Geman, S., Bienenstock, E., & Doursat, R. (1992). Neural networks and the bias/variance dilemma. *Neural Computation*, 4, 1–58.
- Gigerenzer, G. (1996). On narrow norms and vague heuristics: A reply to Kahneman and Tversky. *Psychological Review*, 103, 592–596.
- Gigerenzer, G. (2014). *Risk savvy: How to make good decisions*. New York: Viking.
- 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.
- 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., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103, 650–669.
- Gigerenzer, G., Hertwig, R., & Pachur, T. (Eds.). (2011). *Heuristics: The foundations of adaptive behavior*. New York: Oxford University Press.
- Gigerenzer, G., & Hoffrage, U. (1995). How to improve Bayesian reasoning without instruction: Frequency formats. *Psychological Review*, 102, 684–704.
- Gigerenzer, G., & Selten, R. (Eds.). (2001). *Bounded rationality: The adaptive toolbox*. Cambridge, MA: MIT Press.
- Gigerenzer, G., & Sturm, T. (2012). How (far) can rationality be naturalized? *Synthese*, 187, 243–268.
- Goldman, A. (1999). *Knowledge in a social world*. Oxford: Oxford University Press.
- Green, L., & Mehr, D. R. (1997). What alters physicians' decisions to admit to the coronary care unit? *Journal of Family Practice*, 45, 219–226.
- Holton, G. (1988). *Thematic origins of scientific thought* (2nd ed.). Cambridge, MA: Harvard University Press.
- Houston, A. I., McNamara, J. M., & Steer, M. D. (2007). Violations of transitivity under fitness maximization. *Biology Letters*, 3, 365–367.
- Internet Live Stats. <http://www.internetlivestats.com/total-number-of-websites/>.
- Kahneman, D. (2011). *Thinking, fast and slow*. London: Allen Lane.
- Katsikopoulos, K. V. (2011). Psychological heuristics for making inferences: Definition, performance, and the emerging theory and practice. *Decision Analysis*, 8, 10–29.
- Keeney, R. L., & Raiffa, H. (1993). *Decisions with multiple objectives*. Cambridge: Cambridge University Press.

- Keynes, J. M. (1936). *The general theory of employment, interest and money*. London: Macmillan.
- Kitcher, P. (1992). The naturalists return. *The Philosophical Review*, 101, 53–114.
- Knight, F. (1921). *Risk, uncertainty and profit*. Boston, MA: Houghton Mifflin Co.
- Lawler, E. L., Lenstra, J. K., Kan, A. H. G. R., & Shmoys, D. B. (1993). Sequencing and scheduling: Algorithms and complexity. In S. S. Graves, A. H. G. R. Kan, & P. Zipkin (Eds.), *Handbooks in operations research and management science* (Vol. 4, pp. 445–522)., Logistics of production and inventory Amsterdam: North Holland.
- Luan, S., Schooler, L., & Gigerenzer, G. (2011). A signal detection analysis of fast-and-frugal trees. *Psychological Review*, 118, 316–338. <https://doi.org/10.1037/a0022684>.
- Martignon, L., & Hoffrage, U. (2002). Fast, frugal, and fit: Lexicographic heuristics for paired comparison. *Theory and Decision*, 52, 29–71.
- Pachur, T., & Marinello, G. (2013). Expert intuitions: How to model the decision strategies of airport customs officers? *Acta Psychologica*, 144, 97–103.
- Quine, W. V. O. (1969). *Epistemology naturalized*. New York: Columbia University Press.
- Rich, P. (2018). Comparing the axiomatic and ecological approaches to rationality: Fundamental agreement theorems in SCOP. *Synthese*, 195, 529–547.
- Savage, L. J. (1954). *The foundations of statistics*. New York: Wiley.
- Schurz, G. (2014). Cognitive success: Instrumental justifications of normative systems of reasoning. *Frontiers in Psychology*, 5, 625.
- Searle, J. R. (2001). *Rationality in action*. Cambridge, MA: MIT Press.
- Shah, A. K., & Oppenheimer, D. M. (2008). Heuristics made easy: An effort-reduction framework. *Psychological Bulletin*, 137, 207–222.
- Simon, H. A. (1979). *Models of thought*. New Haven, CT: Yale University Press.
- Şimşek, Ö. (2013). Linear decision rule as aspiration for simple decision heuristics. *Advances in Neural Information Processing Systems*, 26, 2904–2912.
- Şimşek, Ö., & Buckmann, M. (2016). On learning decision heuristics. *Proceedings of Machine Learning Research*, 58, 75–85.
- 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. (2012). *Ecological rationality: Intelligence in the world*. New York: Oxford University Press.
- Tsotsos, J. (1991). Computational resources do constrain behavior. *Behavioral Brain Sciences*, 14, 506–507.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185, 1124–1131.
- Tversky, A., & Kahneman, D. (1980). Causal schemata in judgments under uncertainty. In M. Fishbein (Ed.), *Progress in social psychology* (Vol. 1). Hillsdale, NJ: Erlbaum.
- Von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior* (2nd ed. 1947; 3rd ed. 1953). Princeton, NJ: Princeton University Press.
- Wübben, M., & von Wangenheim, F. (2008). Instant customer base analysis: Managerial heuristics often ‘get it right’. *Journal of Marketing*, 72, 82–93.
- Yates, J. F. (1990). *Judgment and decision making*. Englewood Cliffs, NJ: Prentice-Hall.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.