universität wien

MASTERARBEIT | MASTER'S THESIS

Titel | Title

Deliberating in the world.

Prescriptive arguments in bounded rationality from heuristics to habits.

verfasst von | submitted by

Thijmen Abel Adrian Koomen

angestrebter akademischer Grad | in partial fulfilment of the requirements for the degree of

Master of Arts (MA)

Wien | Vienna, 2024

Studienkennzahl lt. Studienblatt | Degree programme code as it appears on the student record sheet:	UA 066 642
Studienrichtung lt. Studienblatt | Degree programme as it appears on the student record sheet:	Masterstudium Philosophy and Economics
Betreut von | Supervisor:	Univ.-Prof. Dr. Hans Bernhard Schmid

Abstract

This thesis examines prescriptive arguments put forward in the tradition of bounded rationality. Focusing on Herbert Simon's critique of subjective expected utility theory and Gerd Gigerenzer's program of ecological rationality, the first two chapters argue that the bounded rationality literature often overlooks the prescriptive assumptions it is built on. Specifically, both Simon and Gigerenzer posit computational efficiency as an intuitive criterion of rationality without further explanation. This leads to spurious conclusions, like Gigerenzer's claim that ecological rationality bridges the historic is-ought gap in philosophy. Drawing inspiration from Jennifer Morton's ecological theory of rationality, Chapter 3 argues that a more fruitful approach would be to present the norms of rationality as habits. In such an account, boundedly rational strategies are deemed rational when they are generally conducive to reaching the agent's ends. This approach avoids strong assumptions about the nature of rationality, other than that it is instrumental. An important insight that results from this approach is that increasing descriptive adequacy is by itself not enough to generate better prescriptions.

Zusammenfassung

Diese Arbeit untersucht die präskriptiven Argumente, die in der Tradition der begrenzten Rationalität vorgebracht werden. In den ersten beiden Kapiteln wird anhand von Herbert Simons Kritik an der subjektiven Wert-Erwartungstheorie und Gerd Gigerenzers Programm der ökologischen Rationalität argumentiert, dass die Literatur zur begrenzten Rationalität häufig die präskriptiven Annahmen übersieht, auf denen sie beruht. Insbesondere gehen sowohl Simon als auch Gigerenzer davon aus, dass Recheneffizienz ein intuitives Kriterium für Rationalität ist, ohne dies näher zu erläutern. Dies führt zu falschen Schlussfolgerungen, wie z. B. Gigerenzers Behauptung, dass ökologische Rationalität die historische Ist-Soll-Lücke in der Philosophie überbrückt. In Anlehnung an Jennifer Mortons ökologische Theorie der Rationalität wird in Kapitel 3 argumentiert, dass ein fruchtbarerer Ansatz darin bestünde, die Normen der Rationalität als Gewohnheiten darzustellen. In einem solchen Ansatz werden begrenzt rationale Strategien als rational angesehen, wenn sie im Allgemeinen dazu beitragen, die Ziele des Handelnden zu erreichen. Dieser Ansatz vermeidet starke Annahmen über die Natur der Rationalität, außer dass sie instrumentell ist. Eine wichtige Einsicht, die sich aus diesem Ansatz ergibt, ist, dass eine Erhöhung der beschreibenden Angemessenheit allein nicht ausreicht, um bessere Präskriptionen zu erstellen.

Table of Contents

Introduction

The study of decision making is often divided into two camps. On the one hand, there is the descriptive project, which tries to explain and predict the decisions humans make. On the other hand, there is the prescriptive project, which tries to assess what decisions are rational. To an important extent, this distinction can be traced back to David Hume, who in his *Treatise of Human Nature* (1739) observed that statements about what people *ought* to do cannot be deduced from premises that describe the world as it *is* alone.[1] Nonetheless, when a theory tells an agent that she ought to do something, this prescription can be followed only if it lies within her natural abilities. This means that prescriptive theories may benefit from an accurate picture of the capacities of the human agent. This intuition has led to the emergence of bounded rationality, one of the most outspokenly interdisciplinary approaches to rational decision making.[2] Pioneered by the economist Herbert Simon, bounded rationality seeks to improve decision making through better observations, thus bringing description and prescriptions closer together.[3]

Nevertheless, even Simon conceded that there can be "no *ought*'s from *is*'s alone".[4] The question therefore arises how prescriptive arguments are possible in the descriptively informed program of bounded rationality to begin with. This thesis addresses this question based on three accounts put forward in the bounded rationality literature. The first of these is Simon's own procedural approach, which draws prescriptions from the assumption that it is rational to use time and cognitive effort efficiently. The second is Gigerenzer's ecological rationality, which tries to abolish the division between *is* and *ought* through the study of so-called heuristics. The third is Jennifer Morton's ecological theory of rationality, which suggests that the requirements of rationality can be conceived as habits and thereby evaluated based on their outcomes.

---

[1] David Hume, *A Treatise of Human Nature Volume 1*, eds. David Fate Norton and Mary J. Norton (Oxford: Clarendon Press), 302.

[2] David Bell et al. distinguish not two, but three functions of decision theory: the normative, the descriptive and the prescriptive. In this typology, normative models deal with the behavior of ideally rational agents (*ought*-questions), while descriptive models deal with observed behavior (*is*-questions). Differently, prescriptive models use observation to help agents improve their decision making, so that they fall outside the *is-ought* dichotomy. However, authors in the bounded rationality tradition tend to treat *ought*-statements as significant only when real human agents can follow them. This means that in this typology, *ought*-questions would belong to the prescriptive, not the normative domain. See David E. Bell et al., "Descriptive, Normative, and Prescriptive Interactions in Decision Making," in *Decision Making: Descriptive, Normative, and Prescriptive Interactions*, eds. David Bell et al. (Cambridge: Cambridge University Press, 1988), 9.

[3] Herbert A. Simon. *Reason in Human Affairs* (Stanford: Stanford University Press, 1983), 19-23.

[4] Ibid., 7.

Before turning to these approaches in more detail, however, it is helpful to illustrate the complex relation between *is* and *ought* through an example. Imagine an agent – let's call him Daniel – who wants to buy a sweatshirt at his favorite store. For simplicity, let's assume that he only cares about two criteria: softness and warmth. A sensible recommendation for Daniel might be that he ought not to spend money on an item that scores worse on these criteria than another available alternative. In practice, this means that he should settle on a budget and pick the item that scores best on both counts within this budget. In case no item is both softest *and* warmest, Daniel could reflect on the relative importance of the two criteria and weigh his alternatives accordingly. In this way, he can still select an alternative that is best all-things-considered. This rough sketch of a prescriptive theory might be called Simple Maximizing.

In reality, however, it might turn out that Daniel gets overwhelmed by the number of sweatshirts in the store. As a result, comparing all items to each other would require much time and a fair amount of memory retention. What is more, he might find it impossible to rank his alternatives from worst to best. When he finds a sweatshirt that is soft and thin, for example, it might be hard to compare it to an alternative that is tough and warm, because the qualities weighed against each other are so different. As a result of these complexities, Daniel opts simply to buy the first item he finds that matches both his criteria satisfactorily. This procedure, which can be called Simple Satisficing, leads to a satisfactory decision using little time and effort. The downside is that it risks selecting an item that is worse than another available alternative, which violates Simple Maximizing.

In this example, therefore, prescription and description come apart. Yet this does not mean that Simple Maximizing has no prescriptive value at all. In fact, the view that theories that describe behavior badly can still be prescriptively valid is commonly held among decision theorists. At the same time, the case of Daniel shows that prescriptions may fall on flat ears if they do not respect the agent's capacities. Indeed, it is possible that Daniel violates Simple Maximizing not because he disagrees with it in principle, but because it would result in endless agonizing in practice. Differently, Simple Satisficing helps him make good, if not perfect, decisions that save him time and effort. If Daniel would rather obtain a satisfactory outcome than endlessly deliberate, Simple Satisficing seems to serve his interest better than Simple Maximizing. As a result, Simple Satisficing might be attractive not just from a descriptive, but also from a prescriptive standpoint.

When Simon advanced bounded rationality, he relied on a similar observation. He argued that many theories of rationality present formal rules and axioms without factoring in the cognitive abilities needed to use these rules in practice. When the reality of deliberation is taken into account, however, theories that seem prescriptively valid in principle often turn out to be so computationally complex that they are hardly practicable for real-world agents. One of the theories he criticized in this way was subjective expected utility theory and its use in economics. In the real world, Simon suggested, agents make better decisions when they abstain from using this theory and use their time and cognitive capacities more efficiently. For example, like Daniel, agents may be better off when they simply choose the first alternative they find that is satisfactory by their own standards. To describe this decision strategy, Simon coined the term satisficing, which he presented both as a descriptive and a prescriptive model of rational choice because of its relative efficiency.[5]

More recently, the bounded project of connecting the descriptive and the prescriptive sides of decision making has been taken up by the psychologist Gerd Gigerenzer. In his program of ecological rationality, Gigerenzer studies the cognitive algorithms that the human mind uses to deal with complex decision problems, called heuristics, in order to determine the conditions under which these heuristics work well. For example, an agent in Daniel's situation could choose to do what most of the surrounding people do, which Gigerenzer calls the imitate-the-majority heuristic. In situations where this heuristic leads to more satisfactory results in a way that is also faster and requires less information than alternative decision strategies, Gigerenzer argues, it would be irrational not to use it. In this way, ecological rationality purports to *derive* the norms of rationality from the observed success of heuristics. This means that it attempts to bridge the divide between *is* and *ought* in an important number of cases.[6]

A third, philosophically motivated perspective on bounded rationality is offered by Morton. In her ecological theory of rationality, Morton interprets the requirements of rationality as relatively stable habits that function at the background of human reasoning. In this picture, these requirements should be evaluated based on how well they serve the agent's instrumental interests in a global sense, rather than in some individual decision

[5] Herbert A. Simon, "From substantive to procedural rationality," in *Method and appraisal in economics*, ed. Spiro Latsis (Cambridge: Cambridge University Press, 1976), 133-144."

[6] Gerd Gigerenzer and Thomas Sturm, "How (far) can rationality be naturalized?" *Synthese* 187 (2012): 244.

problem. If a norm of rationality like Simple Satisficing reliably helps an agent achieve her ends, it is rational to be in a habit of using it.[7] In this way, Morton's approach embraces the intuition that descriptive considerations are important for prescriptive theories, without purporting to derive *ought*-statements from *is*-statements.

In short, all three authors approach the prescriptive side of bounded rationality in a unique way. This thesis analyzes these approaches to gauge if and how prescriptions are possible in bounded rationality. To set the stage, Chapter 1 starts with a representative "unbounded" account of rationality, namely Savage's subjective expected utility theory. It then presents Simon's critique of this theory and the satisficing conception of rationality he introduced instead. In this chapter, it is argued that despite his valuable observations about Savage's theory, Simon struggled to elaborate his efficiency-based approach into a full prescriptive argument. Subsequently, Chapter 2 analyzes Gigerenzer's attempt to "naturalize" bounded rationality by deriving *ought* from *is*. To do this, it first presents the heuristics that form the basis for his program. It then discusses his argument that the rationality of heuristics can be derived from their accuracy, speed and informational frugality. The chapter assesses this prescriptive claim, concluding that it relies on strong and implausible assumptions about the nature of instrumental rationality. As a result, Gigerenzer fails to bridge *is* and *ought*.

While this seemingly leaves bounded rationality with little prescriptive force, Chapter 3 argues that a more modest argument is possible based on Morton's ecological theory of rationality. The chapter first describes the central characteristics of Morton's habitual approach. It then assesses its compatibility with bounded rationality more broadly, concluding that for agents who face scarcity of time and cognitive resources, it may be rational to grow a habit of using heuristics. The final section of this chapter reflects on what such a habit may look like. A striking conclusion of this perspective is that from an habitual point of view, it may be cognitively taxing to be a perfect heuristic user. As a result, bounded rationality may at places have to sacrifice descriptive precision to yield realistic prescriptions.

[7] Jennifer M. Morton, "Reasoning Under Scarcity," *Australasian Journal of Philosophy* 95, No. 3 (2017): 554.

Chapter I. Simon and subjective expected utility theory

Introduction

One of the most influential approaches to prescriptive decision theory is Savage's subjective expected utility theory.[8] This theory presents a schema that agents can use to approach decisions in a way that is rationally consistent, namely by structuring their binary preferences over available alternatives using seven postulates. Savage presented these postulates as intuitively compelling, similarly to the rules of logic.[9] However, Simon put forward forceful criticism of this approach, arguing that its postulates can only be put in practice when agents possess large amounts of information and extensive cognitive capacities. In the real world, therefore, agents are often better off when they use a more efficient procedure Simon called satisficing, which means simply choosing the first alternative that they find satisfactory. Since satisficing can lead to good results using little time and energy, Simon deemed it a rational procedure for real-world agents, even if it sometimes fails to select alternatives that would be better by the agent's own light.[10]

This chapter traces Simon's prescriptive argument that it can be rational to satisfice. To do this, the first section gives an overview of the central tenets of subjective expected utility theory, based on its seminal formulation by Savage. The second section covers Simon's critique of this theory and the so-called procedural conception of rationality he proposed instead, which forfeits the idea of maximization. The third section engages with two critical attempts to present satisficing as a hidden form of optimization and argues that these approaches cannot solve the procedural problem Simon addressed. It then presents one of the sources of value in Simon's own account, which is the concept of relative efficiency. Ultimately, it is argued, Simon fails to explain why efficiency is an adequate replacement of optimization. Therefore, his procedural account lacks a solid prescriptive basis.

[8] Savage used the label "normative" to describe his theory. However, given that he presented subjective expected utility theory as a manual for agents to improve their actual decision making, it will be called "prescriptive" in the context of this thesis. See note 2 and Leonard J. Savage, *The Foundations of Statistics. Second Revised Edition* (New York: Dover Publications, 1972), 19-20.

[9] Ibid., 20-21.

[10] Simon, "Substantive," 135-137.

1. Savage's theory

Imagine Bea, who wants to go to a new bakery before work next morning. She has heard amazing stories about the bakery's pastries and she would love to start her day off with one of them. This means that she would have to wake up extra early, say around 5:30, and take a detour compared to her normal commute. Unfortunately, however, she does not know at what time the bakery opens. If she would wake up early only to be confronted with a closed shop, the frustration would ruin her morning. In fact, this would be even worse than not going at all. In this situation, what makes Bea's deliberation about what to do difficult is that the consequences of her action depend on an uncertain factor, namely whether the bakery is open or closed. It is in situations like these that Savage's subjective expected utility theory can offer some guidance. When an agent structures her preferences along the theory's guidelines, this guarantees that there exists at least one most preferred option, even in the face of uncertainty.[11] This section spells out the central tenets of this theory and the way it can help agents like Bea tackle the decisions they face.

One of the most fundamental building blocks for Savage's theory is the concept of preferences. While there exists extensive literature about the nature and even the reality of preferences, for the purpose of this thesis it is enough to know that Savage imagined them relatively loosely in terms of actual or hypothetical choice behavior. This means that an agent prefers some alternative x over some alternative y if, in case she must decide between x and y with no other alternatives available, she would choose x. The relation thus described is denoted by the weak inequality "≤", which also includes indifference. In other words, the inequality "y ≤ x" represents the case where an agent would choose x over y or is indifferent between the alternatives.[12]

Savage subsequently showed that an agent's preferences over uncertain outcomes contain at least one highest-ranked alternative if they fulfil certain consistency requirements, which take the form of seven postulates. To take an example, the first postulate states that the preference relation must be a simple ordering among alternatives. This means that the agent's preferences over available alternatives must be both transitive and complete. Here, transitivity entails that if x ≤ y and y ≤ z, then x ≤ z, so that preferences cannot exhibit a cycle. Differently, completeness means that for any two alternatives in the relevant set of alternatives, the agent must either prefer one over the

[11] Savage, *Foundations*, 105
[12] Ibid., 17-18.

other or be indifferent between them. These two properties allow for the pairwise comparison of alternatives to determine which of them is ranked highest.[13] Figure 1 represents these requirements in a situation where an agent gets to choose between a pastry, a croissant and a sandwich for breakfast. Here, only the second preference ordering fulfils both requirements.

Preference ordering	**Transitivity**	**Completeness**
Croissant ≤ pastry	Not violated	Violated
Sandwich ≤ croissant ≤ pastry	Not violated	Not violated
Sandwich ≤ croissant ≤ pastry ≤ sandwich	Violated	Not violated

Figure 1: Transitivity and completeness

Two further building blocks of Savage's theory are utility functions and probability measures. Utility functions attach numerical values to alternatives to represent an agent's preferences. When a well-defined utility function exists for a specific set of alternatives, there exists a maximizing outcome, namely the alternative with the highest utility value attached to it. Differently, probability measures attach probability values to all potential states of the world relevant for a specific decision. Agents are represented as weighting the utilities of their choice alternatives with the subjective probability they attach to each relevant state of the world and summing them together. This procedure yields an *expected* utility function with a maximum, even in cases of uncertainty.[14]

	Bakery open (p=0.5)	**Bakery closed (p=0.5)**	**Expected utility**
Go to bakery	10	1	0.5*10 + 0.5*1 = **5.5**
Normal commute	4	4	0.5*4 + 0.5*4 = **4**

Figure 2: Bea's expected utility

Figure 2 represents Bea's preferences within Savage's framework. In this case, the utility value of going to a bakery that is open is 10, the utility value of arriving at a closed shop is only 1, and the utility value of not going to the bakery is 4 in either case. Now, Bea estimates that the chances that the bakery is open are 50-50. The utility values of Bea's

[13] Ibid., 18.
[14] Ibid., 33, 69, 106-108.

alternatives should therefore be weighted with probability 0.5. This means that the expected utility value attached to going is 5.5, while the expected utility of taking her normal commute is only 4. Despite the possibility that the bakery is closed, then, it is fully rational for Bea to prefer to take her chance and go there.

Importantly, Savage's theory does not tell Bea that she *should* go to the bakery. This is because it is merely a means to evaluate the rationality of preferences themselves. Accordingly, Savage called his postulates "logic-like criteria" that agents can use to "police" their own decisions, and he argued that the rationality of this "code of consistency" was intuitive. Indeed, he suggested that an agent whose preferences would violate his postulates would revise them as soon as the inconsistency is pointed out to her.[15] Concerning transitivity, for example, he wrote:

> "I would, in particular, tell [the critic] that, when it is explicitly brought to my attention that I have shown a preference for f as compared with g, for g as compared with h, and for h as compared with f, I feel uncomfortable in much the same way that I do when it is brought to my attention that some of my beliefs are logically contradictory. Whenever I examine such a triple of preferences on my own part, I find that it is not at all difficult to reverse one of them. In fact, I find on contemplating the three alleged preferences side by side that at least one among them is not a preference at all, at any rate not any more."[16]

Now, Savage's is not the only version of expected utility theory. The earliest such theory is usually ascribed to Bernoulli, who wrote in the 1700s, while modern versions are often traced back to the work of Von Neumann and Morgenstern in the 20th century. However, the novelty of Savage's theory was that it supplied a way to deal with uncertainty in decision making through subjective beliefs. This made it an attractive model of decision under uncertainty in the behavioral sciences, specifically in economics. Ivan Moscati observes that in this discipline, Savage's subjective version of expected utility theory was an authoritative model until deep in the 20th century. However, unlike Savage himself, many authors accepted the theory as a descriptive account as well. For instance, the economist Milton Friedman argued that Savage's theory had scientific value as an empirical account despite making unrealistic assumptions about the human mind. This is because an economic theory should presumably be judged based on its ability to predict

[15] Ibid., 19-21, 59.
[16] Ibid., 21.

behavior. If a theory does this successfully, humans can be taken to behave *as if* the theory's assumptions are realistic. While Friedman initially formulated this as-if methodology together with Savage, the latter ultimately endorsed subjective expected utility theory on prescriptive grounds alone.[17]

The descriptive-prescriptive divide is not the only fault line in the expected utility theory literature. For example, Samir Okasha details the difference between "behaviorist" and "mentalist" approaches, where the first treat utility functions and probability measures as mere representations or re-descriptions of preferences, while the latter assume that they correspond to real psychological entities.[18] Similarly, Johanna Thoma distinguishes "constructivist" from "realist" interpretations, which correspond to roughly the same divide. In terms of these dichotomies, Savage's approach should be deemed behaviorist and constructivist. This is because it derives utility functions and probability measures from preferences alone. Strictly speaking, then, there can be talk of utilities and probabilities only when an agent's preferences conform to the axioms of the theory in the first place. This also explains why the prescriptions that result from Savage's theory focus exclusively on the consistency of preferences.[19]

To end this section, it is important to note that the prescriptive nature of Savage's theory necessitates a minimal commitment to deliberation, even if in a constructivist interpretation. This is because Savage's postulates are presented as a code of consistency that agents should consciously use to structure their preferences. In other words, agents should actively reflect on their preferences and revise them if they violate one of the theory's postulates. However, this poses a potential problem. Simon, one of the most forceful critics of expected utility theory, observed that it is often impossible for real-world agents to form preference profiles consistent with the theory's axioms. This would make it inadequate as a prescriptive theory of rational decision making altogether. The next section presents this criticism, and the so-called procedural approach to rationality Simon suggested instead.

[17] Ivan Moscati, *The History and Methodology of Expected Utility* (Cambridge: Cambridge University Press, 2023), 2, 45-46.
[18] Samir Okasha, "On the Interpretation of Decision Theory," *Economics and Philosophy* 32, No. 3 (2016): 415.
[19] Johanna Thoma, "Decision Theory," in *The Open Handbook of Formal Epistemology*, eds. Richard Pettigrew and Jonathan Weisberg (The PhilPapers Foundation, 2019), 69, 73.

2. The procedural alternative

Simon objected to the use of subjective expected utility theory in economics on both descriptive and prescriptive grounds. He claimed that descriptively, subjective expected utility theory is a bad predictor of human behavior, specifically in cases where uncertainty prevails. In addition, the theory does not describe the psychological processes that govern the human decision-making process itself, which Simon considered a weakness.[20] Prescriptively, he argued that the theory's demands were so complex that they could not be put in practice in most real-life situations. This means that it could not be used as a behavioral code after all, leaving it with little practical value. For these reasons, Simon sought to construct an account of rationality based on realistic models of choice. The resulting program, called bounded rationality, models how the cognitively limited human mind approaches decision problems, and formulates rationality criteria based on these observations.[21] This section discusses Simon's criticism of the classical theory based on his paradigmatic discussion of a game of chess. It then presents his own *procedural* approach to rationality, which forms the prescriptive core of his bounded program.

For Simon, the game of chess shows that expected utility theory in its different versions can offer prescriptions that are impossible to follow in the real world. This is especially striking given that chess is a so-called game of perfect information, meaning that both players know the past and current positions of all pieces on the board. Theoretically, they could use this information to determine the objectively optimal decision path going forward. To do this, the ideally rational player would evaluate each of her potential moves based on what countermoves they may elicit. She would then think further ahead by also evaluating her own potential responses to the opponent's countermoves, her opponent's subsequent responses, and so on until the end of the game. This process would result in a complete picture of all legal conditional sequences of moves that specifies for each whether they end in a win, lose or draw. These game plans could then be compared sequentially to form a complete preference profile, in which the winning strategies are ranked highest.[22]

---

[20] Herbert A. Simon, *Models of Bounded Rationality: Empirically Grounded Economic Reason. Volume 3* (Cambridge: The MIT Press, 1997), 285, 291-293

[21] Herbert A. Simon, "Rationality as Process and as Product of Thought," *The American Economic Review* 68, No. 2 (1978): 12.

[22] Herbert A. Simon, "Theories of Bounded Rationality," in *Decision and Organization: A Volume in Honor of Jacob Marschak*, eds. C. B. McGuire and R. Radner (Amsterdam: North-Holland Publishing Company, 1972), 165-166.

However, while this procedure yields an optimum in theory, in practice it would require that an agent evaluate what Simon estimates are around 10^{120} possible game plans. Seeing as this is incomprehensibly complex for human agents, players cannot realistically be assumed to have a complete picture of their choice alternatives to begin with. Even though an objectively optimal strategy is available, then, human agents are unable to identify and deliberately select it. For Simon, this example shows the limits of theories like expected utility theory, which recommend that agents compare their alternatives and rank them exhaustively. If agents cannot form a full picture of the alternatives they can choose from, it becomes impossible to comparatively rank them, leaving expected utility theory without practical value.[23]

According to Simon, one attempt to solve this problem might be to model the chess player as choosing individual moves instead of complete game plans. To rank the moves by their expected result, the agent would then only have to compare the 30 or so moves open to her at a specific turn, which seems entirely feasible within the bounds of reason. However, Simon observed that even in this model, a player would have to rank her moves based on all the countermoves they may elicit. This analysis would quickly branch out into an extensive move tree that transcends real human computational capacity within reasonable time limits. As a result, when chess players are modeled as move-choosing agents, they might be able to ascertain which alternatives are available to them, but not what the potential consequences of each of their alternatives are. This again makes it impossible to rank the choice alternatives based on their outcomes. In consequence, this model has no practical use for the real-world agent either.[24]

Still a different solution might be to model the chess player's limitations as uncertainty using Savage's probability measures. However, Simon argued that given the knowledge that chess players possess, this route is not available either. This is because the information needed to ascertain the optimal decision path is already present, so that there is no talk of uncertainty in the probabilistic sense. Indeed, finding the optimal path in chess is like proving a theorem: there is a true answer, but the real-world player is cognitively unable to process the available information to find it. As far as there is talk of

[23] Ibid.
[24] Ibid., 169.

uncertainty, then, it refers to a type of ignorance about the decision situation that cannot be resolved through probability measures.[25]

In other words, chess shows that expected utility theory can impose unrealistic requirements in the real world, even in cases where it is possible in theory. This is because it does not account for the computational complexity involved in dealing with concrete decision situations. For this reason, Simon called theories like expected utility theory not only incomplete, but "worse than incomplete". Indeed, he claimed that the requirements of theories that ignore computational complexity can be "seriously misleading by providing 'solutions' to […] questions that are without operational significance". To resolve this issue, Simon proposed that accounts of rationality be based on how the human mind deals with complex decisions in practice. This means that judgments of rationality should be rooted less in the formally optimal outcomes of decision situations, and more in the actual choice processes in the human mind itself. Simon called this a shift from *substantive* to *procedural* rationality.[26]

An example of a procedurally rational tool that agents may use instead of expected utility maximization is what Simon called satisficing. When an agent satisfices, she determines a level of satisfaction that she wants to reach, which is called her aspiration level. As soon as she finds an alternative that matches or exceeds this level, she chooses it, even if there might exist further, unknown alternatives that would strictly speaking be better. If prolonged search yields no results, the agent can lower her aspiration level to make a satisfactory outcome more likely. Conversely, if satisfaction criteria are met easily, she might raise her aspiration level and look for better alternatives. This model makes an important sacrifice, for it can fail to select optimal alternatives and instead opt for ones that are merely "good enough". Still, satisficing relies on steps that are entirely feasible within the bounds of reason, so that it is potentially more useful than expected utility theory.[27]

In the case of chess, Simon suggested, the agent's aspiration level might be decided based on each player's piece count, where better pieces are weighted more heavily. As soon as a certain move ensures that the agent's piece count is higher than the opponent's without creating obvious vulnerabilities, she chooses it, even if there might technically

[25] Ibid., 169-170.
[26] Simon, "Rationality," 12-14.
[27] Simon, "Theories," 167-169.

have been a better, but unknown move available. According to Simon, there is sufficient evidence that chess players not only use such techniques in real matches, but also have better results doing so compared to when they try to analyze all potential moves in search of an optimum. Regarding chess, therefore, satisficing is a model that might be both descriptively and prescriptively superior to expected utility theory.[28]

An important aspect of Simon's shift from substantive to procedural rationality is that it takes models of processes in the human mind as an object of research. This adds a psychological dimension to the study of decision making that was absent in Savage's framework. Notably, the consequences of this extend beyond the descriptive domain and affect prescriptive questions as well. After all, if subjective expected utility maximization is not deemed a good guide for rational decision making, the question arises why satisficing *is*. On its own, the fact that satisficing is descriptively more accurate than expected utility theory does not guarantee that it is also more rational. For this reason, Simon's procedural approach still requires some kind of rationality criterion. The next section explores Simon's approach to this issue based on his concept of relative efficiency.

3. Is it rational to satisfice?

In Simon's picture of satisficing, it requires neither superhuman computational skill, nor extended stretches of time, to help an agent decide on a course of action. Therefore, a prescription to satisfice can be operable in places where subjective expected utility theory is not. The question now arises why these prescriptions should be called rational. Critics of Simon have argued that satisficing does not challenge the classical picture of rationality at all, either because it saves time and energy in an optimal way, or because it optimizes overall life quality. Arguably, this means that satisficing is nothing but optimization on a deeper level. This section engages with these two arguments and shows that neither dispels Simon's concerns with optimization. It then presents Simon's attempt to posit an independent prescriptive account based on the concept of relative efficiency. In the end, however, it is argued that Simon does not explain *why* it is always rational to be efficient. This leaves a fundamental question for Simon's approach unanswered.

In response to Simon, some authors have suggested that his picture of rationality can be incorporated into subjective expected utility theory without too much trouble.

[28] Ibid.

Simon himself discussed an argument to this effect, which views satisficing as optimal search. In this picture, the rationally satisficing agent dynamically adjusts her aspiration levels to a point where the expected benefit of the chosen alternative precisely matches the expected marginal benefit of searching further per unit of time. If satisficing is rational, then, it is so purely because it strikes an optimal balance between the benefit of choosing now and the expected benefit of looking further. However, Simon rejected this picture because applying it requires a common measure of value for the alternatives to choose from and the cost of search, which is hard, if not impossible, to establish. Furthermore, it would require estimating the marginal productivity of search, which is exceedingly difficult under conditions of uncertainty. In other words, like classical subjective expected utility theory, this picture involves a level of computational complexity that makes it an inadequate guide for real-world decision making.[29]

A similar, but more philosophically oriented criticism of Simon's satisficing has been put forward by Byron, who argues that the only reason *why* satisficing might be rational is that it optimizes a general goal of making one's life go well. This argument starts from the notion that goals can be means to achieve further goals. For example, an agent may follow a particular diet for the sake of achieving the broader goal of being healthy. When she satisfices with respect to the lower-level goal, Byron claims, this can only be rational if it serves the broader goal: if an agent satisfices with respect to her diet, this can only be rational if it ultimately benefits her health. Now, while lower-level or "local" goals may relate to a multitude of broader goals, this chain must stop at the broadest level. According to Byron, the only goal that can reside at this level is the goal of making one's life go well, because agents pursue all other goals for the sake of it. Seeing as the goal at the broadest level cannot be constrained by any further goals, agents cannot rationally satisfice with respect to it. If local satisficing is rational at all, then, it is only so because it is a global form of optimizing.[30]

However, this picture of satisficing is hard to reconcile with Simon's conception of the boundedly rational agent. This is because it seems prohibitively difficult for real-world agents to construct the extensive ordering of goals that Byron presents. If an agent deliberates about following a diet, she might imagine that it will improve her health and fitness, while negatively impacting factors like her pleasure in eating, her general mood,

[29] Simon, "Theories," 170-171.
[30] Michael Byron, "Satisficing and Optimality," *Ethics* 109 (1998): 76-79.

and her feeling of social belonging. To determine the extent to which dieting maximizes her life's going well, she would have to decide on a common measure of value for all her goals and then form precise beliefs about how much her different goals contribute to her life quality respectively. What is more, seeing as her goals are presumably interlocked in an integrated whole, this procedure may require considering how her choice would affect her entire system of values. This seems excessively demanding in practice: while agents can arguably deliberate about how their goals and plans relate, they cannot possibly hope to tackle every decision by comparatively analyzing their entire value system. From a prescriptive standpoint, then, the injunction to optimize one's life is too complex to put into practice.

Given these complexities, satisficing is not easily folded into optimizing theories of rationality. Therefore, it is worthwhile to look at the way Simon addressed prescriptive questions himself. In this regard, the complication arises that Simon did not present a final or unique criterion of procedural rationality himself. Still, Till Grüne-Yanoff observes that a potential candidate for such a criterion in Simon's work is the concept of efficiency.[31] Indeed, efficiency seems to play an important prescriptive role for Simon. For example, he described its importance for procedural rationality as follows:

> "Procedural rationality is the rationality of a person for whom computation is the scarce resource – whose ability to adapt successfully to the situations in which he finds himself is determined by the efficiency of his decisionmaking and problem solving processes."[32]

When Simon discussed efficiency, he meant computational efficiency, which concerns the time and effort a system needs to solve a problem in terms of some basic operations. Here, Simon took the human mind to function similarly to a computer, capable of serial computation by recognizing symbols and patterns, storing, copying and comparing them and presenting some output. In order to find relatively efficient decision procedures, then, one needs to compare different procedures to each other in terms of the

[31] Till Grüne-Yanoff, "Bounded Rationality," *Philosophy Compass* 2, No. 3 (2007): 555-556. Grüne-Yanoff observes that in practice, the injunction to use efficient procedures requires deliberation about what decision procedure to use. For this deliberation to be rational, it needs to be efficient itself, and so forth, resulting in an infinite regress. While this is indeed an obstacle when the goal is to find the *most* efficient procedure, this problem dissolves when the goal is to find a procedure that is *relatively* efficient compared to optimization.
[32] Herbert A. Simon, "On How to Decide What to Do," *The Bell Journal of Economics* 9, No. 2 (1978): 504.

number and complexity of the operations they require to reach an outcome. Computational mathematics can be a tool for this comparison, which is why Simon called it a normative theory of procedural rationality. Importantly, this normative theory never culminates in an elegant demonstrative apparatus that selects some procedure as uniquely rational. This is because human decision environments and the strategies used to deal with them continually change throughout history. In other words, the rationality of procedures always depends on the context they are used in.[33]

In his early work *Administrative Behavior* (1945), Simon engaged with the criterion of efficiency in more detail, arguing that the question whether some process is efficient can only be answered in relation to some pre-defined value or goal. After all, it is this goal that the process is supposed to reach in an efficient way. Notably, Simon argued that agents can attach value to different ends for different reasons, and there is no *a priori* reason to assume that these different sources of values are commensurable or comparable. This means that Simon's approach has a strong instrumental dimension – Simon would later call reason "wholly instrumental"[34] –, where decision procedures are evaluated as appropriate *means* to achieve certain goals, without evaluating the ultimate rationality of these goals themselves.[35]

Given this evaluative framework, Simon argued that decision procedures can be judged based on two criteria: the quality of their outcomes measured in terms of given goals, and the efficiency by which they reach those outcomes. Notably, these two measures of success can conflict, since increased efficiency may lead to worse outcomes and decreased efficiency to better ones. Often, this conflict cannot be resolved, because it is not always possible to express the value of outcomes and the value of efficiency in terms of comparable units. In consequence, there is no way to fully balance the two measures of success to attain a perfect or absolute degree of efficiency, so that efficiency is only a relative concept.[36]

Given this framework, it is easy to see why satisficing is more procedurally rational than expected utility theory. Indeed, the process of setting an aspiration level and choosing the first alternative that matches or surpasses it costs much less time and computation

[33] Simon, "Substantive," in *Method and appraisal in economics*, ed. Spiro Latsis (Cambridge: Cambridge University Press, 1976), 132-135, 146.
[34] Simon, *Reason*, 7.
[35] Herbert A. Simon, *Administrative Behavior. A Study of Decision-Making Processes in Administrative Organizations. Fourth Edition* (New York: The Free Press, 1997), 73-74, 256.
[36] Ibid., 258.

than structuring one's preferences using Savage's seven postulates. In addition, it can lead to outcomes where expected utility theory cannot. However, an important problem of this approach is that it does not yet explain *why* it is rational to be efficient. This is especially true in cases where being efficient sacrifices decision outcomes. Within procedural rationality, this issue is hard to resolve, because the approach cannot rely on optimization criteria. Alternatively, it could be argued that it is simply irrational not to aim to be efficient, but this would impose a constraint on agent's ends that is hard to reconcile with Simon's instrumental approach. In sum, therefore, Simon's efficiency seems to be a relatively *ad hoc* criterion of rationality. While this does not invalidate bounded rationality as a prescriptive project altogether, it does call for further explication.

Conclusion

In his work, Simon argued that expected utility theory can impose demands without operational significance for real-world agents, even in Savage's subjective version of the theory. For this reason, Simon put forward a more realistic decision procedure called satisficing, which he deemed rational because it selects satisfactory outcomes in a relatively efficient way. Against Simon, some authors have suggested that satisficing is only rational when it is optimal at a further level of analysis. However, this chapter has argued that these approaches do not resolve the complexities of decision making that Simon's procedural approach aimed to address. Therefore, procedural rationality needs an independent measure of rationality, part of which is the concept of relative efficiency.

Simon's vision of rationality has profound prescriptive implications. Indeed, procedural rationality does not concern the consistency of preferences or their ability to yield a maximum, but revolves around the success of decision procedures in attaining specified goals, whatever those goals may be, and doing so in a relatively efficient way. Compared to Savage's theory, it therefore represents a shift from a consistency-based to an instrumental account of rationality. In this chapter, it has been argued that a potential problem with this view is that it hinges on the concept of relative efficiency without real justification. Therefore, as a prescriptive project, Simon's bounded rationality seems to be relatively *ad hoc*. In light of this discussion, a challenge that faces bounded rationality is to provide a compelling prescriptive argument. The next chapter engages with one attempt to do this by deriving prescriptive requirements from empirical statements of fact, namely Gigerenzer's program of ecological rationality.

Chapter II. Gigerenzer's ecological rationality

Introduction

A contemporary attempt to bolster the prescriptive force of bounded rationality is Gigerenzer's program of ecological rationality, which tries to derive the norms of rationality from empirical research. It does this by studying cognitive decision algorithms, called fast-and-frugal heuristics, that are rooted in the evolved capacities of human agents. An example of such a heuristic is recognition, where an agent selects an alternative merely because she recognizes it, while not recognizing the other available alternatives. In the right environments, these heuristics can be more accurate, faster and require less information than optimization strategies. In these situations, Gigerenzer claims, the ecologically rational agent *ought* to use heuristics to make her decision. In this way, Gigerenzer uses the descriptive study of the human mind to establish prescriptions for rational behavior in the real world, thus bridging *is* and *ought*.[37]

This chapter presents Gigerenzer's program and assesses its prescriptive argument. To do this, the first chapter presents the problems that Gigerenzer observes in the rival approaches to bounded rationality that emerged after Simon's groundbreaking work, which did not aim to revise the norms of rationality in a prescriptive sense. It then presents the heuristics at the basis of Gigerenzer's own program, using the example of recognition. The second section focuses on Gigerenzer's prescriptive argument and shows how he attempts to derive *ought*-statements from *is*-statements based on a picture of rationality that, like Simon's, is instrumental in nature. One distinctive aspect of Gigerenzer's approach is that it advances accuracy, speed and frugality as the key measures of instrumental success. The third section engages critically with this position and argues that speed and frugality do not necessarily define instrumental success. In consequence, the program's prescriptive conclusions only follow when the goals of speed and frugality are already posited as sources of value. This means that Gigerenzer's ecological rationality, like Simon's procedural rationality, ultimately rests on the *ad hoc* assumption that it is rational to be efficient.

[37] Gigerenzer and Sturm, "How (far)?" 252, 260-262.

1. Fast-and-frugal heuristics

In his program of ecological rationality, Gigerenzer tries to elaborate Simon's original approach in a way that has consequences not only for the descriptive, but also for the prescriptive study of decision making. According to Gigerenzer himself, this sets his program apart from other bounded approaches to rationality, which often fail to reject "unbounded" ideals of rationality as correct standards of reasoning. To formulate a program with both descriptive and prescriptive implications, Gigerenzer studies heuristics, which are cognitive algorithms that human agents use to make decisions in the real world.[38] To give a clearer picture of Gigerenzer's program, this section first addresses the author's criticism of rival approaches to bounded rationality. Then, it introduces the heuristics at the basis of his own account using the example of recognition.

Like Simon, Gigerenzer is skeptical of theories that formulate prescriptions of rationality that transcend the cognitive capacities of human agents. For this reason, Gigerenzer criticizes the "axiomatic rationality" of Savage's subjective expected utility theory, although his focus extends beyond formal decision theory and is directed at models of deductive reasoning in general. This includes methods like multiple regression, Bayesian updating and even machine learning algorithms. What all these models have in common is that they present rationality as conformity to abstract axioms or categoric rules, whether it be the axioms of expected utility theory, the rules of probabilistic inference or the rules of logic. To denote the view of rationality that treats this full class of these coherence criteria as requirements of rationality, Gigerenzer uses the label "logical rationality".[39]

According to Gigerenzer, the axioms and rules of logical rationality pervade the rationality literature despite their unrealistic nature, even in presumably "bounded" approaches to the subject. One example of this is the optimization-under-constraints literature, which models cognitive limits and environmental uncertainty as mere constraints on optimization. This means that agents that decide under conditions of uncertainty are taken to weigh the cost of acquiring more information against the benefits of deciding based on the information they already possess. This makes the search for additional information part of an optimization process itself. As discussed in the previous

[38] Gerd Gigerenzer and Reinhard Selten, "Rethinking Rationality," in *Bounded Rationality: The Adaptive Toolbox*, eds. Gerd Gigerenzer and Reinhard Selten (Cambridge: The MIT Press, 2001): 3-6.
[39] Gerd Gigerenzer, "Axiomatic rationality and ecological rationality," *Synthese* 198 (2021): 3548-3552.

chapter, Simon already criticized a similar attempt to treat satisficing as an optimal search for alternatives, because it imposes high computational demands if it is possible at all. Likewise, Gigerenzer observes that determining the optimal level of information search can be prohibitively complex. Indeed, it can require even more information than straightforward optimization would in the same situation. Therefore, agents would have to possess extensive computational powers to put this approach into practice, so that it is not a bounded approach to rationality at all.[40]

Another example of a bounded account that relies on "unbounded" norms is the influential heuristics-and-biases program championed by Daniel Kahneman and Amos Tversky. Gigerenzer argues that authors in this tradition correctly observe that humans systematically violate the axioms of subjective expected utility theory. In part, this is because these axioms can be too complex to apply in the real world. However, the heuristics-and-biases program interprets these violations as "errors" and as proof of widespread irrationality. This means that they still define rationality in terms of Savage's axioms and the associated rules of logic and probabilistic inference. In consequence, the heuristics-and-biases literature presents unrealistic ideals as prescriptively valid, even though it accepts the fact that these ideals are not or cannot be matched by real humans.[41]

In contrast to these two approaches, Gigerenzer aims to use empirical research to revise the norms of rationality themselves. To do this, he presents human agents as evolved organisms that face decision problems in complex environments. In these environments, nothing guarantees the *a priori* compatibility or commensurability of the goals they pursue. This means that oftentimes, preferences cannot be structured in a way that yields expected utility functions with unique maxima, rendering expected utility maximization impossible. Rather, humans tend to make choices by simply "betting" on their environments based on previous experience. To this end, they use cognitive algorithms that exploit the structure of decision environments by deliberately filtering

[40] Gigerenzer and Selten, "Rethinking," 5.

[41] Gigerenzer, "Axiomatic," 3552. Notably, Gigerenzer's critique of Kahneman and Tversky evolved into a highly polemical debate with mutual accusations of misrepresentation. An important point by Kahneman and Tversky that deserves mention is that even if deviations from traditional norms are not considered irrational errors, they still require explanation. See Daniel Kahneman and Amos Tversky, "On the Reality of Cognitive Illusions," *Psychological Review* 103, No. 3 (1996): 591.

and ignoring information. These tools, called fast-and-frugal heuristics, are the descriptive basis of Gigerenzer's program.[42]

One of the simplest of these heuristics is so-called recognition.[43] When faced with more than one alternative, this heuristic selects one of the options for the sole reason that it is recognized while the others are not. Despite its simplicity, recognition can lead to good results in the real world. For example, Gigerenzer mentions an exploratory study by Daniel Goldstein where German and American students were asked whether Detroit or Milwaukee had a bigger population. Only 60% of the American students correctly answered "Detroit", while a surprising 90% of the German students did the same. One reason why German students did better than the American ones is that they were likelier to have heard of Detroit, but not of Milwaukee. They could use this limited knowledge to their advantage by inferring that the recognized city was bigger. After all, bigger cities are often better-known. Seeing as the American students knew both cities, they knew too much to rely on heuristic, thus making worse predictions.[44]

What sets heuristics like recognition apart from coherence criteria is that the former select alternatives based on only one discriminating reason. This challenges the common view that additional information and computation always result in better decisions. For this reason, it is perhaps not surprising that some heuristics can result in choice and inference patterns that violate even the most basic consistency criteria of subjective expected utility theory, like transitivity.[45] Moreover, unlike optimization, which is designed as a universal decision strategy, fast-and-frugal heuristics are never all-purpose. This means that no heuristic can be used across all potential decision situations. The recognition heuristic demonstrates this well, since it relies on partial ignorance: if an agent were to recognize all alternatives in some situation, the heuristic could not be used to discriminate at all. What is more, even where recognition is possible, it does not always perform well. In fact, it only works if the so-called recognition validity, which is the correlation between recognition and the target value, amounts to .5 or higher. If this is the case, relying on recognition works better than not relying on it would.[46]

---

[42] Gerd Gigerenzer, "The Adaptive Toolbox," in *Bounded Rationality: The Adaptive Toolbox*, eds. Gerd Gigerenzer and Reinhard Selten (Cambridge: The MIT Press, 2001), 40-42.
[43] For a list of ten well-studied heuristics, see Gigerenzer and Sturm, "How (far)?", 249-250.
[44] Gerd Gigerenzer and Daniel G. Goldstein, "The recognition heuristic: A decade of research," *Judgment and Decision Making* 6, No. 1 (2011): 100.
[45] Gerd Gigerenzer, *Adaptive Thinking: Rationality in the Real World* (Oxford: Oxford University Press, 2002), 194.
[46] Gigerenzer and Goldstein, "Recognition," 103.

This brief discussion of recognition shows that under the right conditions, simple heuristic strategies can attain good results. In other words, complex problems may require simple solutions. To find these strategies, Gigerenzer studies how humans deal with complex decisions in the real world. However, when it is established that a certain heuristic is within the human cognitive apparatus, this raises the subsequent question under what conditions it works better than alternative strategies. After all, the success of a heuristic depends entirely on the context in which it is applied. This means that a sensible complement to the empirical study of heuristics is to determine the circumstances under which they work well. Gigerenzer takes this to be the prescriptive side of his program, which he calls ecological rationality and is discussed in what follows.

2. Ecological rationality

An important aspect of Gigerenzer's ecological rationality is that it evaluates decision strategies not in terms of intuitive axioms, but in terms of their real-world performance. To this end, Gigerenzer introduces measurable criteria like accuracy, speed and informational frugality, which allow for competitive empirical testing of decision strategies to determine their success in different environments. According to Gigerenzer, this information can be used to derive prescriptions: if an agent finds herself in an environment where a heuristic works relatively well, she *ought* to use it instead of an optimization strategy. He calls the view of rationality that underlies this prescription *ecological*, because it formulates prescriptions based partially on the agent's decision environment.[47] This section traces Gigerenzer's prescriptive argument, which is shown to rely on a particular instrumental conception of rationality that defines instrumental success itself in terms of accuracy, speed and informational frugality.

According to Gigerenzer, heuristics can be prescriptively better than the complex strategies associated with logical rationality. To defend this position, he essentially posits two arguments. The first argument is reminiscent of Simon and holds that the postulates of subjective expected utility theory are too demanding to have operational significance. This means that the theory does not pose practically useful prescriptions in most real-world situations, except in a few simplified settings.[48] While this argument does not yet show that heuristics are *better* than optimization, it does show that unbounded ideals of

[47] Gigerenzer, "Axiomatic," 3555.
[48] Ibid., 3548-3551.

rationality can be prescriptively weak due to their narrow domain of real-world applicability. Presumably, this opens the way for alternative approaches to rationality.

The second argument is stronger and relies on the observation that in the right circumstances, heuristics are superior to the strategies associated with traditional rationality in terms of accuracy, speed and informational frugality. In these cases, Gigerenzer contends, it would be *irrational* not to use heuristics. This is a different argument than is often put forward in the psychological literature, which holds that heuristic save time and energy while sacrificing some modicum of accuracy. According to this received view, an agent should use a heuristic only when she judges that the benefit in terms of saved effort outweighs the cost of sacrificing accuracy. Differently, Gigerenzer highlights that heuristics can capitalize on less-is-more effects, where a reduction in effort leads to more, not less, accuracy. In these cases, using heuristics is not just expedient, but a requirement that ecologically rational agents *ought* to follow.[49]

Clearly, then, ecological rationality is a program with strong prescriptive implications. On one place, Gigerenzer formulates these implications as follows:

> "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 (…). These results lead to a (…) prescription: If conditions E hold, then people *ought* to rely on one-reason heuristics rather than on linear models."[50]

This formulation reveals that like Simon's approach, Gigerenzer's understanding of rationality has a fundamentally instrumental dimension. This is because it derives prescriptive judgments from a decisions strategy's ability to reach a given end, in this case to make an accurate inference in a way that is fast and requires relatively little information. Indeed, Gigerenzer calls ecological rationality a *formalization* of instrumental rationality.[51]

At this point, a question that arises is what instrumental rationality means for Gigerenzer. Borrowing from Robert Nozick – who described instrumental rationality as "the efficient and effective achieving of goals"[52] – Gigerenzer calls it the "default" notion of

[49] Gigerenzer and Sturm, "How (far)?" 260-262.
[50] Gigerenzer, "Axiomatic," 3555.
[51] Ibid., 3547.
[52] Robert Nozick, *The Nature of Rationality* (Princeton: Princeton University Press, 1993), 133.

rationality. However, he does not elaborate what this notion requires from the agent. Instead, he observes rather generally that ecological rationality involves looking for "means suited to certain goals", and that these goals may range from satisfying desires to making correct predictions.[53] Elsewhere, he writes that in his program, "success" has "several aspects, depending on the problem", including accuracy, speed and frugality.[54] In a helpful analysis of ecological rationality, Gerhard Schurz observes that in the experiments of Gigerenzer and his colleagues, strategies are indeed deemed successful when agents make correct inferences or predictions, while success is deemed to decrease when cognitive costs are incurred in terms of information and time.[55] In other words, Gigerenzer treats the triplet of accuracy, speed and frugality not as arbitrary goals, but as measures of what it means to successfully reach a goal in the first place. This means that for Gigerenzer, these three criteria define instrumental success itself.

Notably, this picture of instrumental success extends beyond inferential problems. A good illustration of this is the so-called gaze heuristic. Agents use this heuristic to catch objects like baseballs that are on a downward trajectory from the air. When using the gaze heuristic, the agent fixates her gaze on the object, starts running, and adjusts her running speed to keep the angle of her gaze constant. Relying on this heuristic, people tend to have better success catching an object than if they try to estimate its trajectory and run to the expected landing point. While the gaze heuristic is thus a good tool to catch an object, it is expressly *not* fit for predicting its landing point.[56] This shows that for some fast-and-frugal heuristics, success is not captured by the notion of *inferential* accuracy, but rather revolves around effectively reaching a goal at all.

In the literature, Gigerenzer's treatment of success has been met with some criticism. For example, Patricia Rich observes that Gigerenzer and his colleagues tend to focus on problems with clear-cut answers, like predicting observable values. However, many decision problems do not have such straightforwardly "accurate" solutions, as evidenced by decades of philosophical discussion about topics like risky decisions and games. This means that accuracy is not always clearly defined, so that ecological

[53] Gigerenzer and Sturm, "How (far)," 245, 255.
[54] Gerd Gigerenzer, "Bounded and Rational," in *Philosophy: Foundations and Applications*, eds. Ansgar Beckermann et al. (Paderborn: Mentis, 2008), 245.
[55] Gerhard Schurz, "Cognitive success: instrumental justifications of normative systems of reasoning," *Frontiers in Psychology* 5, No 625 (2014): 5.
[56] Gigerenzer and Selten, "Rethinking," 6-7.

rationality is silent about what to do in an important class of cases.[57] Depending on the pervasiveness of this issue, however, it might not be of great concern. After all, ecological rationality is limited to those cases that allow for competitive testing between decision strategies. If in some situations the concept of accuracy is not defined, ecological rationality might simply not be in force there. While this reduces the program's prescriptive scope, this might not be a fundamental problem for an approach that abstains from positing all-purpose norms.

A second potential problem of Gigerenzer's view of success is that it raises the question what degree of speed, frugality and accuracy is rational. For instance, Klaus Fiedler et al. observe that in ecological rationality, the accuracy and, by implication, the rationality of a heuristic depend on how well it corresponds to the decision environment it is used in. Yet what level of correspondence should be deemed "successful" can only be determined arbitrarily: if a heuristic makes accurate inferences in 70% of relevant cases, the question is whether this is enough to speak of success. The problem is that there seems to be no principled reasons to fix this benchmark at any certain point.[58] However, this criticism fails to account for the comparative nature of Gigerenzer's approach. Indeed, one of the defining aspects of ecological rationality is that it does not need a hard benchmark of success. Rather, it competitively tests strategies to gauge their *relative* performance. In this respect, Gigerenzer writes that the goal of ecological rationality is to find "heuristics that are *better than* others and *good enough* to solve a problem in real time with real resources".[59] Notably, this is reminiscent of Simon's procedural rationality as discussed in Chapter 1, which was about the relative, not the optimal efficiency of decision strategies.

In summary, Gigerenzer argues that heuristics are prescriptively superior to optimization strategies when and because they are more accurate, faster and more frugal than the latter. In this picture, these three criteria are presented as measures of instrumental success itself. It has been shown that the conception of rationality that underlies this argument is not only instrumental, but also comparative. As a result, the

[57] Patricia Rich, "Axiomatic and ecological rationality: choosing costs and benefits," *Erasmus Journal for Philosophy and Economics* 9, No. 2 (2016): 103-105.
[58] Klaus Fiedler et al., "Heuristics and Biases," in *The Handbook of Rationality*, eds. Markus Knauff and Wolfgang Spohn (Cambridge: The MIT Press, 2021), 285.
[59] Gerd Gigerenzer and Peter M. Todd, "Ecological Rationality: The Normative Study of Heuristics", in *Ecological Rationality: Intelligence in the World*, eds. Peter M. Todd and Gerd Gigerenzer, (Oxford: Oxford Academic, 2012), online edition, https://doi-1org-1ka9foxbx03b6.erf.sbb.spk-berlin.de/10.1093/acprof:oso/9780195315448.003.0142.

prescriptive force of Gigerenzer's argument emerges not from intuitive axioms, but from the empirically observable performance of heuristics compared to optimization strategies. However, a question that has not been addresses is why instrumental success should be defined in terms of accuracy, speed and frugality in the first place. The next section assesses this assumption and argues that it is not as intuitive as Gigerenzer's treatment of heuristics suggests.

3. Defining instrumental success

As discussed in the previous section, a heuristic is deemed ecologically rational when it is relatively accurate, fast and frugal in a specific environment. Gigerenzer argues that this amounts to an inference from *is*-statements to *ought*-statements, presumably bridging two domains of reasoning that have been separate in the philosophical literature for centuries.[60] Yet on closer inspection, this claim requires much nuance. This is because Gigerenzer's conclusions about rationality follow only when specific prescriptive assumptions are assumed to hold, for example that speed and frugality are measures of instrumental success. As this section will argue, this assumption can be challenged based on plausible examples where instrumental rationality does not unequivocally favor speed and frugality. As a result, the prescriptive scope of ecological rationality is much more limited than Gigerenzer suggests.

To begin with, Gigerenzer claims that the historic philosophical distinction between *is* and *ought* has created a division of labor between psychology and philosophy. While the former studies empirical behavior without prescriptive conclusions, the latter reasons about the requirements of rationality without empirical evidence. Ecological rationality intends to overcome this divide by "naturalizing" rationality, meaning that it derives the norms of reasoning from empirical statements. While this does not mean that ecological rationality is a universal theory of normativity, Gigerenzer believes that it has prescriptive force in specific favorable environments. In consequence, empirical psychology can be both a descriptive and a prescriptive endeavor.[61]

A concrete example that illustrates this position well is the case of optimal asset allocation. Gigerenzer discusses the groundbreaking work of the economist Harry Markowitz, who demonstrated that the optimal rule to invest wealth among assets is the so-called mean-variance portfolio. This rule maximizes financial returns while minimizing

[60] Gigerenzer, "Axiomatic," 3547-3548, 3554.
[61] Gigerenzer and Sturm, "How (far)," 243-245.

risk based on the observed performance of different assets. Nevertheless, despite this rule's formal optimality, Gigerenzer observes that there exists a fast-and-frugal heuristic that outperforms it in most real-life cases, namely the 1/N rule. With this rule, the agent simply invests equally in all alternatives under consideration, which is faster and requires much less information. Yet despite its simplicity, this heuristic earns better results than the mean-variance portfolio in terms of many standard financial measures. In fact, it has been found that when there are 50 funds to invest in, Markowitz's rule consistently outperforms 1/N in terms of these measures only when it has access to 500 years of historical stock data and it is assumed that the same funds will exist in 2500 years after the investment. Since these conditions do not hold in realistic cases, 1/N has better real-world results.[62]

From the descriptively observed success of 1/N, it follows that investors *ought* to prefer it over Markowitz's rule, at least according to ecological rationality. After all, 1/N outperforms the optimal rule in terms of accuracy, speed and frugality, so that using the latter rule would make the investor not only inefficient, but by that very fact irrational. Notably, however, this strong conclusion follows only if a few prescriptive assumptions are taken to hold. First, it must be assumed that rationality is instrumental. Second, it must be assumed that the criteria of instrumental rationality are accuracy, speed and frugality. This means that Gigerenzer's argument relies on premises concerning the nature of rationality that are expressly prescriptive in nature. Accordingly, it is wrong to say that ecological rationality infers *ought*-statements from *is*-statements alone.

Importantly, this does not mean that ecological rationality must be rejected. Indeed, if it is accepted that rationality is instrumental and that accuracy, speed and frugality are measures of instrumental success, it follows that in the right environments, heuristics can be more rational than optimization strategies. In this picture, descriptive statements could legitimately be used to inform prescriptive conclusions, as long as there are no *ought*-statements derived *exclusively* from *is*-statements. However, a problem with this view is that Gigerenzer does not explain *why* his prescriptive assumptions are the right ones to begin with. The previous section has already mentioned that he deems instrumental rationality the "default" conception of rationality, although he does not

[62] Ibid., 245-7.

elaborate what this entails.[63] Similarly, the claim that accuracy, speed and frugality are the uniquely correct measures of rationality is not well-explained either.

Now, in defense of instrumental rationality, it could be argued that bounded approaches to rationality cannot do without it. This is because they aim to improve the decision making of goal-directed agents in the world. Accordingly, at least some normative value must be attached to ends. However, the subsequent assertion that instrumental success is to be defined in terms of accuracy, speed and frugality is not as easily justified. One reason for this is that it is conceivable that a goal-directed agent may not find speed and frugality very important at all. In these cases, the instrumentally rational agent has little reason to rely on fast-and-frugal heuristics.

A small thought experiment may clarify this intuition. Imagine that our agent of Chapter 1, Bea, has arrived at the bakery she wanted to go to and has the choice between two types of pastry, one with raisins and the other with an unknown exotic fruit. She knows she loves raisins more than almost anything, so it seems smart for her simply to rely on the recognition heuristic. However, all the other clients choose the second pastry, which sparks doubt in Bea's mind. As a result, she decides to ask around, and when she is repeatedly told that the unknown fruit tastes like apricot, which she dislikes, she chooses the raisin pastry, which indeed leaves her most satisfied than the other option would have. In this situation, Bea would clearly have spent less time and information relying purely on recognition. What is more, expending all this effort for no added benefit seems superfluous from an instrumental point of view. At the same time, it seems radical to call Bea instrumentally *irrational*, considering that her ends are reached all the same.

A similar point has been made by John Brunero in a discussion of the rationality of being ambivalent. In response to the idea that having conflicting attitudes towards something impedes effective decision making and is therefore instrumentally irrational, Brunero argues that agents are not instrumentally required to be as effective as possible. Indeed, while instrumental rationality requires that agents intend the *necessary* means to their ends, the same cannot be said about intending the *most effective* means to those ends. To illustrate this, Brunero considers the situation of a university professor who wants more students to attend his classes. While the professor might reach his goal most effectively by bribing his students with good grades, this does not mean he is rationally

[63] Ibid., 245.

required to use this method.[64] Similarly, a student following this class might save much time and effort by cheating instead of studying for an exam. However, this does not mean that the student is *required* to cheat, nor that she would be irrational *not* to cheat. In cases like these, speed and frugality simply seem to be bad measures of instrumental success.

Naturally, it might be objected that without being efficient, agents risk never reaching any of their goals at all. Therefore, any serious account of instrumental rationality should acknowledge the importance of efficiency-related concerns like speed and informational frugality in actual decision making. Yet while it is true that a complete lack of efficiency would leave the agent instrumentally unsuccessful, this does not mean that instrumental success must be positively defined in terms of efficiency criteria. Indeed, the above examples show that some goals can be reached satisfactorily even when the agent takes her time and uses more information. In other words, the fact that some procedures are slower or more information-heavy than others is not enough to make them instrumentally irrational. This suggests that speed and frugality should not be treated as parts of what it means to reach a goal successfully.

In consequence, Gigerenzer's ecological rationality relies on a prescriptive assumption whose truth is not obvious. Indeed, he seems to treat efficiency similarly to how Savage treated consistency, namely as a criterion that rational agents will intuitively and voluntarily follow. Just as there are examples where reasonable agents might sacrifice consistency for the sake of efficiency, however, there are examples where those agents would sacrifice efficiency for the sake of some other goal. In other words, efficiency is not as essential to instrumental rationality as Gigerenzer presents it to be. If the rationality of efficiency is given up, however, one of the main reasons for using heuristics disappears. After all, they owe their ecological rationality partially to their very efficiency. Therefore, without the assumption that efficiency is rational, heuristics seem to lose much of their prescriptive force.

Conclusion

This chapter has discussed Gigerenzer's program of ecological rationality, with a focus on its prescriptive aspects. The first section describes Gigerenzer's criticism of alternative approaches to bounded rationality, which fail to dispense with "unbounded" prescriptive

[64] John Brunero, "Ambivalence, Incoherence and Self-Governance," in *The Philosophy and Psychology of Ambivalence. Being of Two Minds*, eds. Berit Brogaard and Dimitria Electra Gatzia (London: Routledge, 2021), 32.

ideals. It also presents the fast-and-frugal heuristics at the center of Gigerenzer's own program, like the recognition heuristic. The second section engages with the prescriptive side of Gigerenzer's program, or ecological rationality proper. It is shown that Gigerenzer deems heuristics prescriptively superior to optimization strategies when and because they reach the agent's aims more accurately, faster and more frugally than the latter. This argument relies on an instrumental conception of rationality that presents accuracy, speed and frugality as measures of instrumental success itself. However, the third section argues that there exist plausible scenarios where agents can forego speed and frugality without being instrumentally irrational. If this argument is correct, instrumental success should not be positively defined in terms of these criteria.

Importantly, this discussion does not entail that it is never rational to be efficient. What is more, it does not follow that ecological rationality is without prescriptive value altogether. Rather, this chapter argues that simply positing speed and frugality as criteria of rationality is not the same as "naturalizing" rationality. A more rewarding strategy for ecological rationality might not be to deny the prescriptive assumptions it is built on, but to keep them to a minimum. In this spirit, the next chapter tries to construct a prescriptive argument for efficiency that respects the division between the descriptive and the prescriptive domains. For this, it relies on Morton's ecological theory of rationality, which presents the norms of rationality as stable habits that feature at the background of human deliberation.

Chapter III. An habitual approach to ecological rationality

Introduction

The previous chapter has argued that speed and frugality should not be taken to define instrumental success. Importantly, this does not mean that being efficient is incompatible with instrumental rationality altogether. One need only imagine an agent who lets her deliberation be paralyzed by complex decision strategies to realize that saving time and computational effort can be of keen instrumental value. For this reason, it seems that a sound instrumental argument can be constructed for efficient procedures like heuristics after all. This chapter argues that one way of doing this is to understand efficiency as a rational habit that conveys general instrumental benefit to cognitively bounded agents. This argument is based on Morton's ecological theory of rationality, which treats the requirements of rationality as habitual modes of deliberation rather than universally valid norms.[65] Notably, this approach furnishes a prescriptive argument for using heuristics that does not attempt to overcome the divide between *is* and *ought*.

To construct this argument, the first section gives a general overview of Morton's ecological theory of rationality, in which she presents the norms of rationality as habits. The second section argues that Morton's habitual approach can explain why it can be rational to be efficient for agents that face cognitive scarcity. To make this point, the section relies on Morton's own analysis of the reasoning habits exhibited by agents facing material scarcity through poverty. The third section addresses an important challenge this habitual approach to efficiency faces, namely that it can be cognitively taxing to use heuristics on a regular basis, just like it can be taxing to use optimization strategies. This has implications for bounded rationality as a prescriptive project, for it suggests that increasing the descriptive adequacy of a decision model does not automatically lead to better prescriptions. To resolve this issue, the requirements of bounded rationality may have to be formulated in a more modest way than often suggested.

[65] Morton, "Reasoning," 546.

1. Morton's ecological theory

Like Gigerenzer, Morton proposes an account of rationality that is ecological in nature. This means that prescriptions are formulated relative to the cognitive capacities of agents and the environments they make their decisions in. Differently to Gigerenzer, however, Morton does not focus on heuristics, but on universal norms of rationality. She conceives of these norms as habits that can function at the background of practical reasoning, namely as stable modes of deliberation. According to Morton, if a certain reasoning habit reliably helps an agent achieve her goals in the situations she generally finds herself in, that habit is instrumentally justified as a norm of rationality.[66] This section briefly presents the view and the picture of deliberation it relies on.

Morton observes that in philosophy, rationality is often discussed in terms of universal principles that hold unconditionally. Examples of this are the Instrumental Principle, which she formulates as "intend the means believed to be necessary for your ends" and the Consistency Principle "if you intend to F, do not intend anything you believe to be incompatible with F-ing". However, the problem with these principles is that for all norms, the literature presents examples where they can lead the agent astray. A famous example is Niko Kolodny's critique of consistency, in which he suggests that the requirement to be consistent offers little guidance as to how to resolve inconsistencies in practice. For example, when an agent has two incompatible beliefs and only one of them is correct, she can resolve the inconsistency by giving up either one of the beliefs. By giving up the correct belief, however, she would achieve consistency for the wrong reasons. What is more, it is possible that the belief that two beliefs are incompatible is itself faulty. The Consistency Principle is therefore a bad guide for deliberation, and similar examples seem to exist for other norms of rationality.[67]

Despite these examples, Morton argues that norms can still have prescriptive authority in a global sense. She situates this position in the tradition of Simon's bounded rationality and presents her approach as a philosophical counterpart to Gigerenzer's ecological program.[68] Accordingly, she relies on at least three basic insights reminiscent of the works of Simon and Gigerenzer. First, she observes that human agents are cognitively limited beings that reason in changing real-world environments. Second, she

[66] Jennifer M. Morton and Sarah K. Paul, "The Norms of Practical Reasoning," in *The Routledge Handbook of Practical Reason*, eds. Ruth Chang and Kurt Sylvan (London: Routledge, 2020), 541-542, 548-549.
[67] Ibid., 545
[68] Morton, "Reasoning," 544, 554.

argues that prescriptive approaches to rationality should factor in the time and effort it takes to deliberate in the real world. Third, she insists that the norms of rationality must be evaluated based on their results in the environments they are used in, not based on their intuitive validity. As a result, she criticizes philosophical approaches that justify norms in universal terms. At the same time, she argues that these norms need not be rejected altogether. Rather, she treats them as general rules that may or may not improve the decision making of the agent depending on the situations she reliably find herself in. This means that the norms of rationality can themselves be evaluated based on their ecological merit.[69]

This treatment of norms stems from a particular interpretation of how they function in deliberation. Morton argues that agents can only deliberate about their decisions when they do so according to some norm or set of norms. After all, without a norm to license the transition from premises to a conclusion, practical reasoning would not result in outcomes and would accordingly get "stuck". Furthermore, while agents can evaluate the norms they use to deliberate, they can only do so when they are not actively using them at the same time. To explain how this works, Morton compares the norms of deliberation to the rules of a game: before beginning, players can discuss and evaluate the game's rules extensively, but to license actual moves they must accept at least some of them without immediate revision. Similarly, when agents deliberate about what to do in a decision situation, the norms they use to do so function at the background of the deliberative process itself. Only when the agent takes a step back from the decision at hand and starts deliberating about deliberation do norms become the focus of her attention. In consequence, the norms of reasoning enjoy a certain level of stability.[70]

According to Morton, this makes the norms of reasoning much like habits. This means that they should be evaluated not based on their universal validity, nor their success in some individual decision situation, but on how well they serve the agent's *global* interests. In this picture, when habitually following some norm serves the agent's general interest better than not following it would, the norm is rationally sanctioned. Presumably, this overcomes the problem mentioned earlier, namely that for every purported norm, there seem to exist cases where the norm leads the agent astray. Morton embraces this possibility and argues that a norm can still be assessed globally even if it

[69] Morton and Paul, "Norms," 541-2, 550-1.
[70] Morton, "Reasoning," 546.

sometimes leads to bad individual decisions. After all, the norm's net effect may still be beneficial. It is important to note that this does not make norms immune from revision altogether. For instance, an agent who refuses to change her reasoning habits on the long term even if they start undermining her interests should be considered irrational at some point. Such judgments of irrationality, however, can only be based on how beneficial the norm is for the agent in practice.[71]

In line with Simon and Gigerenzer, Morton understands the agents' interests in instrumental terms. This means that the rationality of norms is based on how well they contribute to reaching an agent's *ends*. At first sight, this seems to imply a circularity. After all, if the rationality of a norm is based on whether it is an appropriate means to reliably achieve an agent's ends, rationality seems to be treated exclusively in terms of the Instrumental Principle mentioned earlier. However, this principle is itself a particular norm of rationality that requires justification. To address this risk of circularity, Morton analyzes Broome's thought experiment of a "quirky world". In this world, the contingent circumstances are such that means-ends deliberation consistently leads agents away from achieving their ends. In fact, when they settle on some end, agents happen to achieve it better by relying purely on instinct. In the quirky world, then, agents achieve their ends better by *not* reasoning along means-ends lines and should give up the Instrumental Principle altogether. According to Morton, Broome's analysis of the quirky world is correct to the extent that agents in this world should not deliberate using the Instrumental Principle. However, this does not mean that we must stop *evaluating* their deliberation in instrumental terms. Indeed, by giving up the Instrumental Principle as a *guide* for deliberation, the agent achieves her ends after all. In other words, we can still *assess* the agent's surrender to instinct in terms of how well her ends are reached.[72]

What Morton's analysis of the quirky world shows is that the assumption that ends have normative value does not automatically declare the Instrumental Principle to be a universal norm of rational deliberation. While this does not explain why norms should be assessed instrumentally in the first place, this need not be a problem for bounded approaches to rationality. This is because these approaches seek to improve agents'

[71] Ibid., 553-556.
[72] Jennifer M. Morton, "Toward an Ecological Theory of the Norms of Practical Deliberation," *European Journal of Philosophy* 19, No. 4 (2010), 568-569.

decision making by helping them achieve their ends in the world. Therefore, ends must be assumed to have value from the outset, which is an outspokenly prescriptive assumption.

With these considerations in mind, Morton formulates what she calls her Ecological Theory of Rationality as follows:

> "An agent A should deliberate using those norms N that allow her to reliably achieve her ends E, given her cognitive capacities, in those contexts C in which she regularly finds herself."[73]

Elsewhere, Morton surmises that many "classical" norms of rationality, like taking the necessary means to one's ends, being consistent or sticking to one's resolutions and commitments, are likely to be of general benefit in this ecological way. This means that many traditional philosophical accounts of rationality are instrumentally justified in "normal" lives. Nevertheless, since no norm is posited as universally valid, Morton admits that this is something that can be challenged empirically based on the circumstances of the agent.[74]

Now, even though Morton's terminology resembles Gigerenzer's ecological program, there are at least two important differences between the two approaches. Firstly, Gigerenzer focuses on heuristics, while Morton focuses on habits. Accordingly, Gigerenzer tends to discuss the environmental conditions that hold in individual decision situations, while Morton stresses the environments in which agents reliably and regularly find themselves. Secondly, Gigerenzer's ecological rationality gives special importance to accurate, fast and frugal decisions strategies, while Morton's theory is open to any mode of deliberation whatsoever as long as it generally contributes to achieving the agent's ends. Therefore, Morton's theory does not define instrumental success in terms of speed and frugality, rendering it more inclusive toward non-heuristic standards of rationality than Gigerenzer's project. However, these differences need not lead to conflict. This is because Morton's theory may explain *why* it is rational for cognitively bounded agents to deliberate using fast-and-frugal techniques. The following section explores this line of reasoning in more detail.

[73] Morton, "Reasoning," 554.
[74] Morton and Paul, "Norms," 548-549.

2. Efficiency under scarcity

There exists a potential conflict between Morton's ecological theory and ecological rationality as formulated by Gigerenzer. This is because Morton uses her theory to justify norms like consistency, while Gigerenzer focuses on heuristics that he defines in opposition to all-purpose norms. This section argues that the two approaches might nevertheless be compatible, based on Morton's analysis of rational deliberation under material scarcity. According to Morton, it is rational for agents facing poverty to allocate their resources in a way that prioritizes efficient short-term reasoning, even if this is at the cost of important long-term aims. While this argument relies on questionable assumptions about the cognitive capacities of agents facing poverty, the argument can safely be applied to agents facing the cognitive and temporal scarcity that characterize human decision making. Therefore, Gigerenzer's fast-and-frugal reasoning may be conceptualized as an ecologically rational norm, or habit, of deliberation. Incidentally, this treatment of rationality also resonates with Simon's procedural rationality, which is predicated on the computational scarcity that human agents face.

In Morton's ecological theory, habits serve as a descriptive basis. Indeed, Morton brings up multiple real-world examples to demonstrate that the norms of reasoning can be observed to function as habits. For example, she cites a 2013 study by Sendhil Mullainathan and Eldar Shafir, where a group of street vendors in India was given money to release them from a borrowing trap. After a year, the researchers observed, the group of vendors was in as much debt as their colleagues who had not gotten financial support. This is because they prioritized unexpected and urgent short-term problems, like paying for wedding gifts, over long-term projects like maintaining a healthy level of capital. The researchers concluded that despite the money influx, the vendors were stuck in a "scarcity mindset", which made them address pressing needs to the detriment of long-term concerns. According to Morton, this shows that the vendors' reasoning habit persisted even when their general decision context was altered.[75]

When this habit is prescriptively assessed, however, it seems instrumentally irrational. After all, the vendor's short-term focus necessarily undermines their long-term aims, even ones that might be more important by their own light. However, Morton disagrees with this diagnosis based on ecological considerations. This is because the

[75] Morton, "Reasoning," 548-549.

vendors' mode of reasoning might have been well-adapted to the structural situation of scarcity they were in before the cash infusion. Under scarcity, Morton observes, agents face urgent needs that require efficient solutions, leaving little cognitive room for long-term planning. In general, therefore, agents that face scarcity must restrict their time and energy to short-term ends alone to obtain instrumental success at all. Hence, given their previous poverty, the vendors had good reason to be in a scarcity mindset. Their reasoning *habit* was therefore well-grounded, even if it led to worse individual decisions in some instances.[76]

A sensible objection to this view seems that the vendors' short-term focus may have been justified before they received financial support, but stopped being so as soon as they received the money. From that moment on, the vendors had the financial means to focus on long-term projects, so they were in a material position to form a new deliberative habit. For this reason, their persistent habit of prioritizing the short term can rightfully be called irrational. However, Morton argues that this objection mistakenly assumes that agents have the time, cognitive resources, and cognitive flexibility to change their deliberative habits on a whim. What is more, it assumes that agents can reliably distinguish cases where short-term reasoning is required from ones where long-term reasoning would be better. However, agents facing poverty do not have the cognitive leeway to perform these deliberative tasks. Therefore, the vendors' persistent scarcity mindset should not be deemed irrational, at least not on the short term.[77]

Notably, this approach can be challenged on empirical grounds. For example, Eric Sampson argues that it relies on the dubious claim that resource-scarce agents are cognitively unable to switch their modes of reasoning from short-term to long-term. However, he observes that agents often display such switching, for example when they forego activities with high short-term gains, like stealing, out of concern for their long-term damage. Indeed, even under conditions of resource scarcity, agents can see clearly that some behavior is destructive over time without too much cognitive effort. This means that they have the cognitive capacity to switch their modes of deliberation after all, so that they can reasonably be expected to sometimes prioritize long-term aims. For this reason,

[76] Ibid., 549-550, 553-554.
[77] Ibid., 554-556.

it is right to call agents irrational when they undermine long-term aims for the sake of less important short-term goals, even if these agents face poverty.[78]

While this objection is strong, it does not mean that Morton's argument lacks prescriptive bite altogether. This is because Morton's considerations can also be applied to a different type of scarcity, namely the cognitive scarcity that human agents face throughout their lives. Indeed, human decision making is characterized by permanent scarcity, whether it be scarcity of time, attention, information, or computation. While this is certainly different from material poverty, it poses similar requirements on the deliberating agent. After all, if the cognitively bounded agent wants to make satisfactory decisions at all, she cannot spend all her time and energy in pursuit of optima. Indeed, she would be well-advised to limit the use of resource-intensive methods like expected utility theory, and to sometimes use efficient procedures like heuristics instead.

Notably, this way of talking about efficiency does not close the door for unbounded ideals of rationality altogether. In some cases, agents may still have *some* instrumental success by optimizing: even the person who approaches decisions using subjective expected utility theory might achieve some of her ends in the world, albeit after much time and effort. For this reason, instrumental success cannot be *defined* in terms of efficiency. Nevertheless, it seems plausible that the hopelessly inefficient agent makes much fewer and less satisfactory decisions than her efficient counterpart. In other words, efficient reasoning can be deemed instrumentally rational based on the way humans relate to their environment as cognitively bounded agents. This picture of efficiency avoids Sampson's objection to Morton's position, because using procedures like subjective expected utility theory can be so costly that it undermines the agent's instrumental interests altogether. This makes it different from switching from a short-term to a long-term focus, which is arguably not as cognitively challenging.

At this point, however, the question might arise to what extent this picture of rationality is truly empirical. When presenting her approach, Morton writes little on how to test the relative success of norms. This poses a stark contrast to Gigerenzer's project, which is rife with carefully devised models to competitively test heuristics against other decision strategies. In fact, empirical testability is so fundamental to Gigerenzer's program that he only considers those heuristics that can be modeled to be part of ecological

[78] Eric Sampson, "Do the Standards of Rationality Depend on Resource Context?" *Acta Analytica* 38 (2023): 331-332.

rationality in the first place.[79] It is unclear whether Morton's approach can be conceptualized with the same precision. After all, it is hard to establish what makes one norm "better" than another over longer stretches of time, whether it be the quantity of the ends achieved, their quality or something else. This renders Morton's normative criteria much vaguer than Gigerenzer's criteria of accuracy, speed and frugality.

However, a systematic literature review by Hal Arkes and co-authored by Gigerenzer and Ralph Hertwig might be of some help regarding this issue. In the review, the authors look for evidence that it is costly to violate standard coherence norms, like having transitive preferences or following standard statistical rules. They find little such evidence, and cautiously suggest that there exists little reason to assume that coherence norm violations have negative results in the real lives of agents. In this way, the authors use empirical studies about reasoning in specific contexts to make inferences about the role of norms in the lives of agents in a more global sense.[80] While this does not address the methodological issue of testing Morton's norms, which would be beyond the scope of this thesis, it shows that in principle, it is within the bounds of empirical research to speak of norms in relatively global terms, albeit with some caution.

In an important way, therefore, Morton's approach seems compatible with Gigerenzer's program. Indeed, if being efficient is understood as a habit of using heuristics, it might be possible to justify this type of reasoning as an ecologically rational norm of deliberation. This approach is more parsimonious than Gigerenzer's prescriptive argument, for it does not posit that speed and frugality are measures of instrumental success itself. Rather, speed and frugality derive their rationality from the instrumental benefit they have for the cognitively bounded agent. Notably, this approach is reminiscent of one of Simon's claims mentioned in Chapter 1 of this thesis, namely that bounded rationality is intended for agents for whom computation is the scarce resource. In the habitual approach presented here, it is precisely this scarcity that makes it rational to be efficient. Nevertheless, a problem arises with this view, namely that it is hard to formulate Gigerenzer's highly flexible and context-sensitive heuristics as a stable and persistent habit. The next section discusses this problem and suggests that in the end, it might necessitate a certain prescriptive modesty in ecological rationality.

[79] Gigerenzer, "Axiomatic," 3555.
[80] Hal R. Arkes et al., "How Bad is Incoherence?," *Decision* 3, No. 1 (2016): 22-26.

3. Can we form heuristic habits?

So far, it has been argued that it may be rational for cognitively bounded agents to be in a habit of using heuristics. From a prescriptive viewpoint, however, this view faces a problem. This is because heuristics are never all-purpose, so that their nature seems to preclude the formation of stable habits. Indeed, different environments call for different heuristics, so that using them always requires careful assessment of the individual decision at hand. For instance, no matter how successful the stock market can be predicted using the 1/N rule mentioned in Chapter 2, an agent cannot fall back on this heuristic when trying to catch a ball. Yet determining which heuristic is appropriate in a specific case can itself be a costly process. As a result, a norm enjoining heuristic reasoning might be too cognitively demanding to be ecologically rational. This section explores the seriousness of this problem and concludes that Morton's framework can succeed in giving a normative grounding to Gigerenzer's ecological rationality only if its prescriptions are formulated in a more modest way.

As the arguments of Simon, Gigerenzer and Morton show, bounded rationality relies on the assumption that prescriptions can benefit from accurate descriptions. In case of ecological rationality, however, critics have argued that this move from description to prescription is far from automatic. For example, in a recent analysis, Ivan Moscati observes that heuristic decision models sometimes present choice processes in steps that require great cognitive effort. Therefore, even if heuristic models are within the bounds of human cognition, this does not guarantee that they are easy to put into practice.[81] Similarly, D. Wade Hands observes that the prescriptions of ecological rationality are based on the cognitive capacities of agents and therefore *not impossible*. However, this is hardly a sufficient condition for accepting it as a theory of rationality.[82]

A similar problem exists in case of heuristic selection. Surprisingly, this topic has received little attention from a prescriptive viewpoint, despite being a key concern in the descriptive literature on ecological rationality. Indeed, there exist a variety of hypotheses that explain why humans are able to select appropriate heuristics in the right situations, relying on mechanisms like evolutionary pressure and reinforcement learning.[83] However,

[81] Ivan Moscati, "Behavioural and heuristic models are as-if models too – and that's ok," *Economics and Philosophy*, First View, (2023): 16-17.

[82] D. Wade Hands, "Normative ecological rationality: normative rationality in the fast-and-frugal heuristics program," *Journal of Economic Methodology* 21, No. 4 (2014): 407.

[83] Andreas Ortmann and Leonidas Spiliopoulos, "Ecological rationality and economics: where the Twain shall meet," *Synthese* 201, No. 135 (2022): 13-14.

these descriptive mechanisms do not solve the prescriptive problem. After all, if the mind selects a heuristic automatically, there would be no need for deliberation, yet a prescriptive theory aims to counsel people how to resolve decisions through deliberation. This means that as a prescriptive project, ecological rationality must offer some guidelines as to how to select the appropriate heuristic in a specific decision instance.

To illustrate that this can be difficult, it is helpful to return to Bea. Chapter 1 presents her deliberation about going to a bakery, while Chapter 2 presents her deliberation about which out of two pastries to buy in that bakery. What makes her decision hard in the first situation is that she does not know whether the bakery is open, while in the second she only recognizes one of her two alternatives, while all other clients choose the unrecognized kind of pastry. If Bea wants to approach both decisions using heuristics, she first faces the issue of what heuristics she should rely on. For example, if most of her peers visit the bakery every morning with much satisfaction, she could use the imitate-the-majority heuristic to decide to do the same. However, when she faces the subsequent choice between two pastries, she realizes that she could use either imitate-the-majority or recognition, both of which recommend a different course of action. To make an ecologically rational decision, Bea must select the heuristics that yields the best result in her current environment. This might involve finding out the value of relevant environmental cues like the recognition validity, but this would require much time and effort. Alternatively, she could simply make a lucky guess and hope the heuristic she chooses turns out to be the best choice, but this would hardly amount to deliberation. In other words, Bea is stuck between two heuristics, unable to determine what course of action would be best for her within a reasonable time limit.

This example shows that even though heuristics can be very simple, this is not enough to result in a helpful guide for deliberation. Indeed, an increase in descriptive adequacy does not by itself guarantee a better prescriptive account. This problem becomes even bigger when an agent is expected to use a decision strategy not once, but on an habitual level. If Bea wants to tackle not only these cases, but more of her daily decisions using heuristics, she will have to spend a considerable amount of time deliberating about what heuristics are appropriate in each situation. As a result, heuristics may in practice be as paralyzing for Bea's decision making as the "unbounded" theories criticized by Simon and Gigerenzer.

For an ecological approach to rationality to have both descriptive *and* prescriptive value, the requirements of rationality therefore need to be better attuned to the practical needs of the agent. To get an idea of how this might be possible, it is worthwhile to take another look at the norms of rationality as proposed by Morton. In her theory, Morton attempts to rationalize divergent norms ranging from consistency to efficiency under scarcity. However, she formulates these norms in a very practicable way. For example, when she discusses consistency, she does not mean structuring preferences in a way that yields a maximizable expected utility function. Rather, she talks of avoiding blatantly incompatible intentions, partially to avoid becoming completely ineffective in reaching one's ends, and partially to remain interpretable to others and coordinate with them.[84] Therefore, this interpretation of consistency seems much more modest than Savage's axiomatic system.

Admittedly, even such a modest norm of consistency, when taken to the extreme, imposes draconian requirements. Avoiding contradictory beliefs in a strict sense might require subjecting all one's beliefs to rigorous logical analysis, which is a difficult and time-consuming task if it is possible at all. For example, Arkes et al. observe that in a community of 150 members where a person holds 20 beliefs about each person, checking for the consistency of these beliefs would require making almost 4.5 million binary comparisons.[85] However, this strict interpretation of consistency clearly risks making the agent completely ineffective. Therefore, obedience to a strict norm of consistency would not be instrumentally useful. Differently, a more modest version of the norm might focus at avoiding incompatibilities at places where they obstruct the agent's decision making, which will likely be of instrumental benefit. In this case, rationality might require following a modest, not a strict norm of consistency.

This modest treatment of consistency might offer a hint as to how an habitual approach to efficiency is possible as well. Strictly following Gigerenzer's program might sometimes require time-consuming analysis of available environmental cues to establish what heuristic works well, and then to use the heuristic in the appropriate way. However, a more modest approach seems possible. If agents can be taught to recognize certain situations they regularly find themselves in as ones where heuristics are of value, their decision making can be improved relying on these heuristics. For example, if it turns out

[84] Morton and Paul, "Norms," 550.
[85] Arkes et al., "How Bad?," 30.

that Bea is generally better off if she relies on the imitate-the-majority heuristic when purchasing food, she could use this heuristic to structure her deliberation in the relevant cases. While the heuristic may sometimes lead her astray, this is not enough to make it irrational, for what counts is how well it serves her interests in a global sense. In other words, this approach might justify a particular heuristic in a relevant class of environments purely based on the general instrumental benefit it conveys.

Notably, this modest treatment of efficiency requires no fundamental changes to the descriptive practice of ecological rationality. After all, the empirical study of which heuristics work well in which environments can offer invaluable clues about how agents can improve their real-world decision making. However, analogously to subjective expected utility theory, the descriptive study of heuristics might have to be understood as finding out how an *ideal* agent would behave, and not how a real agent *can* behave. This means that empirical findings must still be translated into practicable recommendations for human agents. In this interpretation, the prescriptive aim of ecological rationality is no longer to "naturalize" the norms of rationality. Rather, it is to formulate operational guidelines that can help agents make better decisions in the environments they regularly face. While this aim is admittedly less ambitious, it is much more perceptive of the practical issues the cognitively bounded agent faces in the real world.

Conclusion

This chapter has presented Morton's ecological theory of rationality and argues that it can address some of the prescriptive challenges faced by Gigerenzer. Specifically, Morton's theory can explain why it can be ecologically rational to use heuristics without defining instrumental success in terms of speed and frugality. This is because habitually using heuristics can help cognitively bounded agents achieve their ends relatively well. To construct this argument, the first section presents Morton's ecological theory of rationality. In this theory, an agent ought to use those norms that help her reliably achieve her ends given the contexts she regularly finds herself in. The second section discusses Morton's analysis of efficient reasoning under material scarcity. While Morton's argument rests on claims that can be challenged empirically, these concerns disappear when the argument is applied to the cognitive scarcity that human agents face. Accordingly, it may be rational for cognitively bounded agents to habitually prioritize efficient over resource-intensive decision strategies, as long as this serves their aims well.

The third section sketches some implications of this approach to ecological rationality. Most importantly, it is argued that it may be cognitively taxing and time-consuming to be a perfect heuristic reasoner on an habitual level. Therefore, everyday agents who need to allocate their time and energy efficiently cannot be expected to always select the correct heuristic. It has been argued that this is not a problem, for the requirements of ecological rationality can be formulated in more modest terms. If agents can be trained to recognize certain classes of environments as ones where a specific heuristic is useful, they can rely on heuristics in a habitual sense after all. Therefore, the descriptive study of heuristics can still yield valuable prescriptions that improve the agent's real-world decision making.

Taking stock

While all the main authors discussed in this thesis seek to connect the descriptive and prescriptive sides of rational decision making, they each take a distinct approach. Chapter 1 has shown how Simon broke the ground for bounded rationality by opposing theories like Leonard Savage's subjective expected utility theory, which ignored computational complexity. To address this issue, Simon proposed a new, procedural approach to rationality that relied on a decision procedure called satisficing, which makes efficient use of time and energy to reach satisfactory outcomes. However, the chapter has argued that Simon's approach still lacked a prescriptive argument *why* the rational agent ought to be relatively efficient, especially in cases where it sacrifices decision outcomes.

In Chapter 2, it has been shown how Gigerenzer seeks to construct a stronger prescriptive argument for bounded rationality than Simon could provide. To do this, Gigerenzer relies on the observation that some cognitive decision algorithms, called fast-and-frugal heuristics, can be both more efficient *and* more accurate than optimization strategies. Therefore, Gigerenzer concludes, people *ought* to use heuristics when they can. This conclusion relies on the joint assumptions that rationality is instrumental and that speed and frugality are part of what it means to have instrumental success. However, while being fast and frugal is doubtless instrumentally useful in many cases, it has been argued that this does not mean that people always *ought* to go for the fastest and most frugal decision strategy open to them. This means that instrumental success should not be defined in terms of these two properties, which weakens Gigerenzer's prescriptive claim.

Chapter 3 has tried to construct a more modest prescriptive argument for heuristics. The chapter builds on Morton's ecological theory of rationality, which locates rationality in the norms agents use to deliberate on an habitual level. In this framework, efficient reasoning may be construed as an instrumentally rational habit for cognitively bounded agents. As a result, it is rational for humans to be in a habit of using heuristics. In this picture, however, it might be prohibitively difficult for real-world humans to be perfect heuristic reasoners. This is because selecting the right heuristic for each individual decision may be a costly endeavor itself. Therefore, individuals may sometimes be rationally permitted to violate the strict prescriptions of ecological rationality.

These observations help answer the main question of this thesis, namely how prescriptive arguments are possible in bounded rationality. First, all the arguments presented show that prescription must come from somewhere. Indeed, Simon, Gigerenzer

and Morton all accept the premise that reason is instrumental, which serves as a prerequisite for any further prescriptive claims. However, the difference between the authors is whether they rely on any additional prescriptive assumptions. It has been argued that both Simon and Gigerenzer treat efficiency-related criteria as independent sources of value, while Morton's approach can justify efficiency more parsimoniously in terms of its instrumental benefit alone. This speaks in favor of the latter approach, because it avoids the problem of treating efficiency as a measure of instrumental success.

Second, the writings of Simon, Gigerenzer and Morton all show the need for descriptive realism. Simon carefully demonstrated that without it, even a situation as transparent as chess does not allow for adequate resolution. Similarly, Gigerenzer highlights how important it is to treat the agent as an evolved human organism rather than an idealized abstraction, and Morton criticized traditional norms of rationality for being detached from actual reasoning contexts. The most important insight that emerges from these analyses is that any prescriptive approach to rationality should be sensitive to the contexts in which it is applied, at pain of losing operational significance altogether.

Third, however, empirical rigor is not enough to formulate good prescriptions. It is true that the empirically informed models of the human mind presented by Simon and Gigerenzer seem cognitively more realistic than Savage's system of axioms. However, the discussion of Morton's habitual approach to rationality shows that from the perspective of the agent, descriptively accurate models might be as difficult to put into practice as axiomatic postulates. This problem is clearest in the case of Gigerenzer's heuristics: the agent who wants to tackle every decision by using the most successful heuristic may be forced to deliberate so much that it undermines her instrumental success. As a result, even bounded rationality needs to be careful when moving from description to prescription.

All in all, this suggests that as a prescriptive project, bounded rationality should embrace its instrumental nature. This also means that it should abstain from tenuous prescriptive claims, like that efficiency is always rational or that heuristic models bridge *is* and *ought*. Nevertheless, one way of still giving efficient strategies like heuristics an important place is to treat them as habits of reasoning that serve the general interests of cognitively bounded agents. While this approach avoids the automatic progression from description to prescription present in Gigerenzer's program, it does this purely based on what is of practical use to the real-world agent. And it is precisely this perspective that seems most important in bounded rationality.

Bibliography

ARKES, HAL R., GERD GIGERENZER AND RALPH HERTWIG. "How Bad Is Incoherence?" Decision 3, No. 1 (2016): 20-39.

BELL, DAVID E., HOWARD RAIFFA AND AMOS TVERSKY. "Descriptive, Normative, and Prescriptive Interactions in Decision Making." In *Decision Making. Descriptive, Normative, and Prescriptive Interactions*, edited by David E. Bell, Howard Raiffa and Amos Tversky, 9-30. Cambridge: Cambridge University Press, 1988.

BRUNERO, JOHN. "Ambivalence, Incoherence and Self-Governance." In *The Philosophy and Psychology of Ambivalence. Being of Two Minds*, edited by Berit Brogaard and Dimitria Electra Gatzia, 25-48. London: Routledge, 2021.

BYRON, MICHAEL. "Satisficing and Optimality." *Ethics* 109 (1998): 67-93.

FIEDLER, KLAUS, JOHANNES PRAGER and LINDA MCCAUGHEY. "Heuristics and Biases." In *The Handbook of Rationality*, edited by Markus Knauff and Wolfgang Spohn, 273-293. Cambridge: The MIT Press, 2021.

GIGERENZER, GERD and REINHARD SELTEN. "Rethinking Rationality." In *Bounded Rationality: The Adaptive Toolbox*, edited by Gerd Gigerenzer and Reinhard Selten, 1-12. Cambridge: The MIT Press, 2001.

GIGERENZER, GERD. "The Adaptive Toolbox." In *Bounded Rationality: The Adaptive Toolbox*, edited by Gerd Gigerenzer and Reinhard Selten, 37-50. Cambridge: The MIT Press, 2001.

GIGERENZER, GERD. *Adaptive Thinking. Rationality in the Real World*. Oxford: Oxford University Press, 2002.

GIGERENZER, GERD. "Bounded and Rational." In *Philosophy: Foundations and Applications*, edited by Ansgar Beckermann, Holm Tetens and Sven Walter, 233-257. Paderborn: Mentis, 2008.

GIGERENZER, GERD AND DANIEL G. GOLDSTEIN. "The recognition heuristic: A decade of research." *Judgment and Decision Making* 6, No. 1 (2011): 100-121.

GIGERENZER, GERD and THOMAS STURM. "How (far) can rationality be naturalized?" *Synthese* 187 (2012): 243-268.

GIGERENZER, GERD AND PETER M. TODD. "Ecological Rationality: The Normative Study of Heuristics." In *Ecological Rationality: Intelligence in the World*, edited by Peter M. Todd and Gerd Gigerenzer, online edition. Oxford: Oxford Academic, 2012.

https://doi-1org-1ka9foxbx03b6.erf.sbb.spk-berlin.de/10.1093/acprof:oso/9780195315448.003.0142.

GIGERENZER, GERD. "Axiomatic rationality and ecological rationality." *Synthese* 198 (2021): 3547-3564.

GRÜNE-YANOFF, TILL. "Bounded Rationality." *Philosophy Compass* 2, No. 3 (2007): 534-563.

HANDS, D. WADE. "Normative ecological rationality: normative rationality in the fast-and-frugal heuristics program." *Journal of Economic Methodology* 21, No. 4 (2014): 396-410.

HUME, DAVID. *A Treatise of Human Nature Volume 1*, edited by David Fate Norton and Mary J. Norton. Oxford: Clarendon Press.

KAHNEMAN, DANIEL AND AMOS TVERSKY. "On the Reality of Cognitive Illusions." *Psychological Review* 103, No. 3 (1996): 582-591.

MORTON, JENNIFER M. "Toward an Ecological Theory of the Norms of Practical Deliberation." *European Journal of Philosophy* 19, No. 4 (2010), 561-584.

MORTON, JENNIFER M. "Reasoning Under Scarcity." *Australasian Journal of Philosophy* 95, No. 3 (2017), 543-559.

MORTON, JENNIFER M. AND SARAH K. PAUL. "The Norms of Practical Reasoning." In *The Routledge Handbook of Practical Reason*, edited by Ruth Chang and Kurt Sylvan, 541-552. London: Routledge, 2020.

MOSCATI, IVAN. *The History and Methodology of Expected Utility*. Cambridge: Cambridge University Press, 2023.

MOSCATI, IVAN. "Behavioural and heuristic models are as-if models too – and that's ok." *Economics and Philosophy*, First View (2023): 1-31.

NOZICK, ROBERT. *The Nature of Rationality*. Princeton: Princeton University Press, 1993.

OKASHA, SAMIR. "On the Interpretation of Decision Theory." *Economics and Philosophy* 32, No. 3 (2016): 409-433.

ORTMANN, ANDREAS AND LEONIDAS SPILIOPOULOS. "Ecological rationality and economics: where the Twain shall meet." *Synthese* 201, No. 135 (2022): 1-30.

RICH, PATRICIA. "Axiomatic and ecological rationality: choosing costs and benefits." *Erasmus Journal for Philosophy and Economics* 9, No. 2 (2016): 90-122.

SAMPSON, ERIC. "Do the Standards of Rationality Depend on Resource Context?" *Acta Analytica* 38 (2023): 323-333.

SAVAGE, LEONARD J. *The Foundations of Statistics. Second Revised Edition*. New York: Dover Publications, 1972.

SCHURZ, GERHARD. "Cognitive success: instrumental justifications of normative systems of reasoning." *Frontiers in Psychology* 5, No. 625 (2014): 1-16.

SIMON, HERBERT A. *Models of Bounded Rationality: Empirically Grounded Economic Reason. Volume 3*. Cambridge: The MIT Press, 1997.

SIMON, HERBERT A. "Theories of Bounded Rationality." In *Decision and Organization: A Volume in Honor of Jacob Marschak*, edited by C. B. McGuire and Roy Radner, 161-176. Amsterdam: North-Holland Publishing Company, 1972.

SIMON, HERBERT A. "On How to Decide What to Do." *The Bell Journal of Economics* 9, No. 2 (1978): 494-507.

SIMON, HERBERT A. "Rationality as Process and as Product of Thought." *The American Economic Review* 68, No. 2 (1978): 1-16.

SIMON, HERBERT A. "From substantive to procedural rationality." In *Method and appraisal in economics*, edited by Spiro Latsis, 129-148. Cambridge: Cambridge University Press, 1976.

SIMON, HERBERT A. *Administrative Behavior. A Study of Decision-Making Processes in Administrative Organizations. Fourth Edition*. New York: The Free Press, 1997.

SIMON, HERBERT A. *Reason in Human Affairs*. Stanford: Stanford University Press, 1983.

THOMA, JOHANNA. "Decision Theory." In *The Open Handbook of Formal Epistemology*, edited by Richard Pettigrew and Jonathan Weisberg, 57-106. The PhilPapers Foundation, 2019.