

BRIDGING THE GAP BETWEEN THE NORMATIVE AND THE DESCRIPTIVE: BOUNDED EPISTEMIC RATIONALITY

Nastja Tomat*

University of Ljubljana
Ljubljana, Slovenia

DOI: 10.7906/indecs.22.1.6
Regular article

Received: 9 December 2023.
Accepted: 20 February 2024.

ABSTRACT

The aim of the article is to propose bounded epistemic rationality as a concept that blurs the divide between normative and descriptive approaches to the study of rationality. I illustrate the contrast between philosophy as a normative discipline and psychology as the empirical study of cognition and show that unattainable standards and the arbitration problem pose a challenge for normative theories of rationality. I then outline three possible types of relations between normative and descriptive theories of rationality, the third being the proposal for hybrid concepts, such as bounded epistemic rationality, that include both normative and descriptive elements. I continue by describing Herbert Simon's notion of bounded rationality and Gerd Gigerenzer's ecological rationality, and consider the role of bounded rationality in epistemology. I reflect on the relationship between norms of epistemic and bounded rationality and finally, drawing on the work of David Thorstad, I suggest some features that I believe should be included in an account of bounded epistemic rationality. I aim to show that an understanding of epistemic rationality that is compatible with bounded rationality can help to avoid overly strict, idealized, as-if theories of rationality, narrow the gap between the normative and the descriptive, and bring us closer to a comprehensive understanding of epistemic rationality.

KEY WORDS

epistemic rationality, bounded rationality, ecological rationality, normative, descriptive

CLASSIFICATION

APA: 2340, 2630

JEL: I12

*Corresponding author, *η*: nastja.tomat@ff.uni-lj.si; -; -

INTRODUCTION

What it means to be rational is one of the core questions of epistemology. Epistemic rationality concerns epistemic attitudes, states, and processes [1], with a focus on beliefs [2]. It is distinguished from other kinds of rationality in terms of pursuing a strictly epistemic goal [3], for example truth, knowledge, or understanding [4-6]. It stands at the intersection of epistemology and investigation of rationality, and as such is closely connected with core concepts in epistemology, especially justification [7]. Although epistemic rationality and epistemic justification are often used interchangeably, I will distinguish between them, as I agree with Foley's [7] claim that this allows us to investigate epistemic rationality in a similar way to rationality of other phenomena, for example decisions or actions, and frees it from the preoccupations of traditional epistemology, such as inquiries about the connection between rationality and knowledge.

It is often said that there is a division of labour between philosophy and psychology [8, 9]. Philosophy, including epistemology, is considered a normative discipline which is primarily concerned with how people should reason, inquire, form and update beliefs, and make judgments in order to be rational. Rationality is often defined as adherence to a particular normative system, such as following the rules of logic or probability theory. Stein calls such a view "a standard picture of rationality", and he writes that "to say that someone is rational is roughly to say that she reasons in the way she ought to" [10; p.17]. Rationality thus has a strong evaluative component and can be understood as a normative concept. On the other hand, empirical disciplines such as psychology study how reasoning, judgement, decision making, belief formation and updating actually occur and aim to provide processing of cognition and describe the underlying mechanisms. Although they often refrain from making evaluative judgments about rationality, they are still closely related to normative theories, as they usually employ a specific normative system (e.g., reasoning according to the rules of logic or probability theory) as a yardstick for evaluating the (ir)rationality of beliefs, judgments or decisions [11].

With normative philosophical theories on the one hand and descriptive, empirical research on the other, one is faced with the question of what is the appropriate relationship between the two approaches. Normative theories are often criticized for not considering empirical insights about the limitations of human abilities and consequentially proposing norms of rationality that are too demanding and unachievable for real human cognizers [12]. It seems that norms of rationality should be informed in some way by empirical data about human cognition, but evaluative judgments about how we should think and reason to be rational cannot be derived from descriptive, empirical premises without committing the "is-ought" fallacy [13]. This leaves us with a question of how to define epistemic rationality in a way that would consider what psychology and other disciplines can tell us about human cognition, but still keep the normative element and guidance for epistemically good reasoning without succumbing to the abovementioned fallacy.

The aim of this article is to propose bounded epistemic rationality as a hybrid concept that blurs the divide between the normative and the descriptive. Bounded rationality is a term first proposed by Herbert A. Simon in the 1950, and it emphasizes that the notion of rationality should consider the limitations of human cognition and the environment in which we operate. An account of bounded epistemic rationality recognises that we are bounded agents, proposes norms that are empirically grounded and attainable, and as such can provide good epistemic guidance. I begin by describing how the concept of rationality is intertwined with the notion of normativity and highlight two challenges for normative theories of rationality. I summarize some possible ways of relating normative and descriptive theories of rationality, and then move on to a description of Herbert A. Simon's concept of bounded rationality and Gerd Gigerenzer's

ecological rationality, which is based on Simon's work. In the last part, I reflect on the role of bounded rationality in epistemology, drawing on David Thorstad's work, and – relying on the abovementioned concepts – suggest some features that I believe are crucial for an account of bounded epistemic rationality that includes both normative and descriptive elements.

RATIONALITY AND NORMATIVITY

The connection of rationality and normativity is a much debated issue in epistemology [14, 15]. Normative theories of rationality involve some sense of “ought”, but Knauff and Spohn [1] argue that the terms “ought” and “norms” are ambiguous when they apply to rationality and point to the difference between empirical and genuine normativity. Empirical normativity consists of the norms that are established in a particular community, but cannot help us answer the question of genuine normativity; if there is a rule to stop at the red light when driving a car, this does not answer the question of whether we should really stop at the red light. What one really should believe and what should be the case cannot be determined by empirical observation, but by normative deliberation and accepting a particular normative stance [1]. Knauff and Spohn [1] emphasize that the goal of normative theorizing is not the search for normative truth but for normative agreement, and that normative theories of rationality are merely hypotheses about how one should achieve rational beliefs. In contrast to empirical theories, that are based on experimental or other empirical data, the starting points of normative theories are intuitions or assessments of various normative claims. In a similar way as empirical theories are meant to explain and predict data, normative theories are meant to justify these intuitions or assessments and arrive at a reflective equilibrium. However, using intuitions as a starting point for defining rationality has also been problematized [9].

Normative theories of rationality are faced with various challenges. One of the much-discussed criticisms is that they place unrealistic, impractical or even impossible demands on agents [16-18]. Normative theories of rational reasoning or decision making often require optimization – employing the best possible solution to a given class of problems. However, problems in the real world are often so complex that optimal solutions are computationally intractable even for hypothetical machines with infinite time and computational power. If we want to establish optimization as a norm of rationality, hardly any human in the real world could ever be rational [12]. Humans operating in the real world are far from ideal cognizers equipped with computational capacities and relevant information from the environment that would enable them to search for optimal solutions – regardless of whether we are talking about belief formation, reasoning, decision making or behaviour [19, 20].

Another challenge for normative theories is the question of what makes a particular normative system the right one. The idea that rational reasoning follows the rules of logic and probability seems intuitively plausible, but this alone cannot provide sufficient justification for it. Sometimes called “a problem of arbitration”, the issue consists of the lack of clear criteria for choosing one particular normative system over another [17]. An example is Wason selection task [21] in which the participants are presented with four cards that have a colour on one side and a number on the other. The fronts of the cards show the numbers 3 and 8 and the colours brown and red. Participants have to decide which cards to turn over to test the rule: “If there is an even number on one side of the card, the opposite side is red”. According to the rules of deductive logic, which were the predominant normative system for evaluating responses to the task, participants should turn over a card with the number 8 to test the modus ponens and a brown card to test the modus tollens. The majority of participants in different studies and different variations of the tasks answered that they had to turn over a card with the number 8 and the red colour, the latter being a logical fallacy (affirming the consequent). However, these answers are only wrong if they are interpreted as violating the rules of deductive logic, and

various authors have proposed alternative normative systems for evaluating the correctness of the answers. One example is Oaksford and Chater's rational analysis [22], according to which participants choose the cards that are expected to yield the greatest information gain, and such answers are considered rational. This shows that the evaluation of a particular type of answer on a reasoning task as correct or incorrect depends on which normative system is used as a benchmark. Normative theories must therefore address the question of how the preference for one normative system over another can be justified.

RELATIONSHIP BETWEEN NORMATIVE AND DESCRIPTIVE THEORIES OF RATIONALITY

The question about the relationship between the normative and the descriptive can be traced back to Hume and the well-known "is-ought" fallacy. According to the dominant interpretation, Hume [13] states in the *Treatise* that no evaluative conclusions can be drawn from descriptive premises [23]. Whenever we want to infer "ought" from "is", or, in other words, whenever we want to infer from what is the case to what ought to be the case, we are committing a logical fallacy. The same applies to inferences in the other direction: it is not valid to infer from "ought" to "is". The debate applies not only to ethics and moral reasoning, but is also highly relevant for research of reasoning and rationality. In order to avoid is-ought fallacy, some authors [24] suggest that theories of rationality should adopt another principle of normativity: ought implies can. According to this principle we can only require an agent to perform a certain action if she is able to do so. This means that epistemic norms we propose should be attainable, and if agents cannot execute a demanding cognitive operation or conduct a complex inquiry, they should not be expected to do so.

Due to human receptivity for norms, there are many defeasible connections between normative and descriptive theories of rationality, but Knauff and Spohn [1] are critical of the types of relations proposed in the literature so far and do not consider any of them satisfactory. I believe that we have three different possibilities. The first is to keep normative and descriptive theories completely separate – let philosophical disciplines, like epistemology, to determine how we should think and act in order to be rational, and leave psychologists to work on processing accounts of human reasoning, judgement, decision making and other cognitive processes. This is the position advocated by Elqayam and Evans [17]: To avoid is-ought fallacy, descriptive research should dispense with the notion of normativity and focus exclusively on providing processing accounts of cognition. They do not argue for excluding normativism from scientific research completely and they acknowledge its importance in various domains, but claim that psychological research on cognition would do better without normativism. However, other authors argue that normative perspective is indispensable in rationality research [1], but nevertheless urge to consider the option that normative and descriptive approaches are at least logically independent, since both is-ought and ought-is inferences are considered unacceptable [1].

The second option is to continue what we are already doing: we keep conducting empirical research on human cognition with some degree of normativism involved and keep engaging with philosophical theories of rationality, but this means that somewhere along the way we should address the arbitration problem and argue why empirical data on human cognition is not relevant to norms of rationality. At this point, it is worth considering the possibility that the term rationality has a different meaning in normative theories than in descriptive ones; an example of this would be Evans and Over's [25] rationality₁ and rationality₂. They define rationality₁ as a personal and instrumental kind of rationality that is evaluated in terms of achieving one's goals, while rationality₂ or impersonal rationality refers to following the rules of logic and probability in one's reasoning and decision making. In their words, rationality₁ is "thinking, speaking, reasoning, making a decision, or acting in a way that is generally reliable

and efficient for achieving one's goals", while rationality₂ is "thinking, speaking, reasoning, making a decision, or acting when one has a reason for what one does sanctioned by a normative theory" [25, p.8]. This means that an individual can at the same time be rational in one sense but not in the other. This distinction suggests that, in addition to views that equate rationality with reasoning in accordance with a normative system, another possibility is to understand rationality instrumentally, i.e. in terms of achieving one's behavioural goals, whatever they may be.

If we prefer the second option, the question arises as to what to do when empirical findings about human cognition deviate from the norms of rationality. One possibility is to stick to the chosen normative standard and interpret the deviations as irrationality; another is to modify the normative standard in order to reduce the discrepancy (e.g. provide alternative explanations of the answers to the Wason selection task in terms of optimal data selection [22] instead of deductive logic or falsification); and the third is to take it as a starting point for criticising normative systems, which might lead us to the third possible relationship between the normative and the descriptive.

The third option is to propose some kind of hybrid concepts or models of rationality that contain both normative and descriptive elements [26]. Such theories should be informed by empirical research on human cognition, but should not dispose of the normative and evaluative questions of what is good reasoning, what is rationality, and how should we conduct inquiries. Gigerenzer and Sturm [8], for example, argue for such a naturalized account of rationality. They believe that it is – to some extent – possible to determine norms of rationality on the basis of empirical research, but not in all domains of reasoning; the normative-descriptive divide cannot be bridged everywhere. Their naturalized view of rationality is based on Gigerenzer's work on fast and frugal heuristics, which I will describe in more detail in the section on ecological rationality. According to them, rationality should be assessed in terms of fit between a strategy and the environment, and they aim to replace an instrumental understanding of rationality with an ecological one. Their naturalism about rationality is, in their words, "normative, ecological, and limited" [8, p. 245].

Schurz and Hertwig [9] also urge philosophers and psychologists to work together to find a new definition of rational cognition. They argue that rationality should not be measured against specific, "universal" benchmarks, such as consistency, coherence, deductive logic or Bayes' rule, but should be understood in terms of cognitive success in the real world. They propose a consequentialist approach to rational cognition, in which normative standards are justified by the success of their outcomes. They problematize the inevitably subjective nature of intuitions and argue that intuition-based account of rationality is prone to strong cognitive relativism. Instead of intuitions, they propose basing epistemic normativity on cognitive success, which is defined in terms of successful prediction. They claim that in their consequentialist account, empirical data gain normative weight and normative claims guide new empirical questions, and argue that such an approach could help to overcome the division of labour between philosophy and psychology.

BOUNDED RATIONALITY

One of the most prominent critics of normative views on rationality was a political scientist Herbert A. Simon, who introduced the concept of bounded rationality. He criticized the way in which traditional economic theory, e.g. expected utility theory [27] viewed agents: this "economic man" has information about all relevant aspects of the environment, has complete and stable preferences, and is equipped with a cognitive apparatus that enables him to perform complex calculations that determine which of the possible alternatives in a decision situation yields the highest expected utility. In real life, however, we usually act under uncertainty and

practically never know all possible alternatives, have only partial knowledge of the consequences and imperfect estimations of the values [28]. Simon urged that global rationality postulated by economic theories should be replaced by a concept of rationality that is compatible with both the computational capacities of the agents and the structure of the environment in which they operate. Due to the limits of human cognitive processing, particularly in terms of computational and predictive capacity, “actual human rationality can at best be extremely crude and simplified approximation to the kind of global rationality that is implied, for example, by game-theoretical models” [29, p.101]. He urged to replace the requirement of optimization – the search for solutions that would maximize the payoff – with the notion of satisficing, i.e. the search for solutions that are good enough, but not necessarily the best [29, 30]. He emphasized that to understand human rationality, we must take into account both the limitations of human cognition and the characteristics of the environment in which humans operate [30]. To illustrate his point, he used a metaphor of a scissors: “Human rational behaviour (and the rational behaviour of all physical symbol systems) is shaped by a scissors whose two blades are the structure of task environments and the computational capabilities of the actor” [31, p.7]. In order to understand human rationality, we have to take into account both blades. The answer to the question of how to behave rationally differs depending on whether we look at it from the perspective of global, economic theories of rationality or from the perspective of theories that consider the limitations of human cognition and environment [30]. Another important feature of bounded rationality is its procedural character. In contrast to substantive rationality, according to which a behaviour is rational if it helps us to achieve our goals, procedural rationality also considers the process; a behaviour is procedurally rational if it is the consequence of an appropriate process of deliberation [32].

It is worth noting that the term “bounded rationality” is nowadays used in different disciplines and not necessarily in exactly the same sense as Simon described it; according to Gigerenzer [18], some authors understand bounded rationality as optimization under constraints [33] or as irrationality. In the first case, bounded rationality is often used in economic models as a kind of optimization under constraints, in part because such models are mathematically much less complex than non-optimizing ones. In the second case, bounded rationality is interpreted as a deviation from rational choice theory, which is considered normatively correct, and thus equated with irrationality [18]. In the article, however, I intend to use the term bounded rationality in Simon’s terms.

ECOLOGICAL RATIONALITY

Simon’s bounded rationality has gained considerable recognition in various fields, particularly in psychology and economics, and has greatly influenced the way we think about rationality. One of the concepts widely used in psychology of reasoning and decision making that stems from Simon’s bounded rationality is ecological rationality, investigated by Gerd Gigerenzer & ABC research group. They argue that a particular strategy of reasoning, problem solving or decision making is ecologically rational to the extent that it is adapted to the structure of the task; (ir)rationality of a strategy should not be judged according to a priori normative criteria, but by its degree of fit with the environment [34, 35].

Gigerenzer has established a research programme of fast and frugal heuristics. As the name suggests, his focus is the investigation of heuristics: strategies that ignore part of the information and are not computationally demanding, which can lead to faster, frugal and more accurate judgments. Heuristics and the core capacities on which they rely (working memory, attention, object tracking, etc.) are part of what Gigerenzer calls the mind’s adaptive toolbox. The function of an organism’s adaptive toolbox is defined in evolutionary terms: to reach proximal goals, such as finding food, avoiding predators, finding a mate, etc. Adaptiveness is therefore an important component of Gigerenzer’s understanding of rationality [34, 35].

Our intuition often tells us that more information and computation will always lead to better results than less information and computation. The idea that more information is always better and that one should always use as much evidence as possible is deeply rooted in our notions of rationality, Carnap's [36] principle of total evidence being just one example. In contrast to our intuitions, Gigerenzer's research has shown the so-called less-is-more effect. Under conditions of uncertainty, when not all alternatives, values and probability distributions are known, heuristics often provide more accurate predictions than more complex or optimizing strategies that use and weigh all available information. This means that heuristics are computationally less demanding, but not at the expense of accuracy. In other words, there is often no trade-off between frugality and accuracy [34, 35].

If we use a strategy that gives us more accurate predictions than other available strategies, we are ecologically rational. The question of fit between the strategy and the environment is, according to Gigerenzer, an empirical one, and one of the main goals of his research is to explain the characteristics of the environment that are relevant for determining the ecologically rational strategy: degree of uncertainty, number of possible alternatives, the size of a learning sample and so on [34, 35].

Ecological rationality is not defined by adherence to certain norms and emphasizes that no strategy is a priori rational. The idea of ecological rationality is therefore to suggest that problem-solving strategies should be evaluated in relation to the environment in which they are used, and that such an approach is a better starting point for addressing the normative questions than evaluating a strategy against a normative yardstick from the standard picture. Ecological rationality refers to the success of cognitive strategies, while success is defined in terms of accuracy and frugality. Ecological rationality is therefore intertwined with the notion of success and is instrumental in this sense, but still keeps the standard epistemic goals by emphasizing the importance of truth or accuracy [34, 35].

The normative claim of Gigerenzer's research programme is that in cases where heuristics lead to more accurate judgments than other strategies, people should use them – but this implies that it is crucial to recognise when to use a particular strategy. Classical logic may be a perfectly justified normative system, but it would hardly be rational to apply it in all situations of everyday reasoning. However, the fact that a particular rule should not be used in all situations does not mean that it is inadequate; no formal system such as logic or Bayes' theorem – or heuristics – can be the best solution in all situations. There is no general rule of reasoning and rationality that is a priori superior. As Gigerenzer and Sturm say, "mind and environment hang together, and a theory of rationality should respect and explicitly deal with both blades of Simon's scissors" [8, p.264].

BOUNDED RATIONALITY IN EPISTEMOLOGY

Sturm [37] finds it surprising that only a few philosophers have drawn on bounded rationality, since it is an empirically grounded approach which aims to encompass not only descriptive but also normative dimensions. Consequently, the philosophical aspects of bounded rationality have not been systematically explored. He explicates the goals of a systematic, comprehensive study of the philosophical dimensions of bounded rationality: to clarify and analyse in detail the assumptions of the concept of bounded rationality and its differences from non-bounded rationality; to assess in which domains of philosophy bounded rationality can be useful; and to investigate its theoretical and methodological foundations in various disciplines such as psychology, economics and political science [37].

Similarly, Thorstad [24] argues that bounded rationality is not as strongly represented in epistemology as it should be. He views it not as a theory but as a paradigm and describes it in

terms of five normative claims, the first being that bounds are important. In practical philosophy, it is generally accepted that physical limitations must be taken into account when we think about norms for rational action – for example, we cannot be required to lift an object that weighs a tonne. However, this is often not the case with rational beliefs, and norms of epistemic rationality frequently place unachievable demands on agents. Thorstad claims that bounds should also be considered when thinking about rational cognition: if we are unable to perform a complex cognitive operation, we should not be required to do so. The second claim is that we should focus not only on the final doxastic states, but also on the processes that led to them. This procedural outlook on rationality derives directly from Simon’s work and shifts the focus away from norms of belief to norms of inquiry. The third and fourth claims relate to ecological rationality and the use of heuristics: rationality is not bound only by internal cognitive factors, but also by the environment, and is therefore ecological. Moreover, the use of heuristics can be rational in many cases – either because non-heuristic processes are too demanding or even impossible or because heuristics provide more accurate predictions than other, more complex strategies. The final claim is that bounded rationality is consistent with the programme of vindicatory epistemology, which regards alleged violations of norms of rationality as the consequence of boundedly rational deliberation. We cannot comply with requirements such as coherence of beliefs because of our cognitive limitations, and we often inquire in the most rational way we can, considering our bounds [24].

NORMS OF BOUNDED EPISTEMIC RATIONALITY

If we want to understand epistemic rationality as bounded, we need to address the question of the relationship between traditional epistemic norms and norms of bounded rationality, which has already been raised in the literature [24, 37, 38]. Sturm [37] points out that we cannot simply map the norms of the standard picture into recommendations for rational reasoning in everyday life. One example is the rule of non-contradiction, which is elementary in classical logic, but is not easily translated into the requirement of a coherent set of beliefs. The belief sets we hold in real life are large and complex, and it seems impossible for us to check for possible contradiction every time we obtain new belief and adjust the set so it would satisfy the law of non-contradiction. According to other views, the norms of the standard picture and bounded rationality are not necessarily incompatible, quite the contrary. Gigerenzer and Sturm [8], for example, argue that their naturalized account of rationality does not undermine the standard picture. Sturm [38] also says that rules of formal logic are embedded in the formulation of Gigerenzer’s heuristics, such as the recognition heuristic, which presupposes the basic understanding of if-then conditional: “If one of two alternatives is recognized, infer that it has the higher value on the criterion” [39, p.130]. Bounded rationality is thus inevitably based on the concepts of logic [37].

I believe we have several possible types of relationship between the norms of the standard picture and bounded rationality:

- 1.) We can keep norms of epistemic rationality (e.g. forming true beliefs and avoiding false ones) and use the notion of bounded rationality as a framework for justifying the lowering of the threshold for epistemic rationality. This means that the norms should be empirically grounded – we should know what the limitations of our computational power and speed, working memory, attention, etc. are in order to propose feasible norms. In this case, we can apply ought-implies-can principle of normativity. In such a view, “bounded” in bounded epistemic rationality only serves as a justification for lowering the bar for traditional norms.
- 2.) We can use the framework of bounded rationality to put an emphasis on the procedural dimension of epistemic rationality, especially on the role of pragmatic factors. According to Simon [32], rationality of behaviour is judged not only by its final outcome, but also by the process that led to it. Shifting the focus from the final doxastic attitudes to the process is

already pursued by inquiry epistemology [40, 41], but bounded rationality might help us to consider how to put more emphasis on the costs of the process. For example, if a particular process is highly reliable (i.e. leads to a high ratio of true to false beliefs), but has a large computational cost, is very time consuming, and interferes with other important activities in life, it would hardly be considered rational – at least from the perspective of bounded rationality. If we do not evaluate the rationality of cognitive processes or inquiries only according to epistemic factors, such as reliability, but also consider the influence of pragmatic factors such as costs in terms of cognitive resources, effort and time, norms of rationality become more flexible and can be better applied to real, limited human agents. In a similar line, Bishop and Trout [42, 43] proposed a theory they call strategic reliabilism. They argue that epistemically good reasoning is connected to reliability, but one can still spend a lot of time reasoning about trivial matters, such as memorising the product descriptions in an online store or forming true beliefs about the number of bricks in a pavement. For this reason, they include pragmatic criteria for rational reasoning: it must be efficient in terms of time, energy and effort, and it must relate to topics that are relevant to the agent. According to them, “rational reasoning is reliable, cost-effective and focused on significant problems” [43, p.106]. Such an approach treats rationality as a function of factors that are usually considered purely epistemic (reliability) and factors that are considered purely practical (cost-effectiveness and relevance of issues to the agent). Bishop and Trout argue that strategic reliabilism has just the right mix of epistemic and practical elements to be used as a tool for ameliorative epistemologists who want to make recommendations for improving people’s reasoning. Because strategic reliabilism takes into account the limitations of human cognitive processing and the individual differences of reasoners, it is, as Bishop and Trout put it, an epistemological theory for real people. I believe that using the bounded rationality framework is consistent with Bishop and Trout’s approach and could help us to further explore the idea of incorporating practical factors into our procedural understanding of epistemic rationality.

- 3.) In notions of bounded and ecological rationality, the emphasis is on adaptiveness: rationality is defined in terms of fit between the strategy and the environment, there is no single set of rules of rationality that can be applied in all contexts, and heuristics are domain-specific. Perhaps it is possible to include adaptiveness as one of the norms of bounded epistemic rationality: for a process to be boundedly epistemically rational, it must lead to some epistemic goal and serve as an adaptive response to the environment. These two goals – the attainment of true beliefs and adaptiveness – often coincide, but there are many situations in which they come apart. If we have true beliefs about which food is poisonous and we therefore do not eat it, we have achieved both the epistemic and adaptive goals. However, we can think of many situations where one of the two goals may take precedence or even conflict with each other. If we devote enormous amounts of time and resources to gaining true beliefs about a particular trivial matter while neglecting other important activities, this will hardly be adaptive. We should know when to stop reasoning or inquiring, or, in other words, when the costs of reasoning outweigh its benefits. If we include adaptiveness as a norm, we must determine its relation to strictly epistemic goals. When we speak of bounded epistemic rationality, epistemic goals should always be present, but the question arises as to what role adaptiveness plays in evaluating the rationality of inquiry or belief.

BOUNDED EPISTEMIC RATIONALITY

It is not my aim to propose a comprehensive account of bounded epistemic rationality, but merely to describe some of the features that I think such an account should contain, drawing on the concepts and approaches described previously in the article.

I propose that bounded epistemic rationality means inquiring in a way that leads to good enough results: good enough accuracy, understanding, or prediction. In determining what is good enough – e.g., what is a belief that is sufficiently accurate or close to the truth – we must apply a form of ought-implies-can principle of normativity. Bounds that we need to consider are cognitive, practical, and environmental. The standard for how accurate beliefs we are expected to form, and in what way we should inquire, must be attainable by limited cognizers. In other words, we cannot be expected to have fully coherent belief systems about a complex issue; we cannot be expected to find an optimal solution to a complex decision problem; and we cannot be expected to perform perfect Bayesian conditionalization. Moreover, the norms of bounded epistemic rationality should take into account practical considerations. We are not bounded only by our cognitive abilities, but also by time. When we inquire, reason, and form beliefs, we do not do so in a bubble that isolates us from the practical considerations of our daily lives. Every day we have to decide (even if we do so automatically and implicitly) how much time and cognitive resources we want to devote to a particular task. Prioritising, allocating our resources, and knowing when to stop inquiring and move on to another task is a skill that bounded agents need to manage their daily lives. Bounded epistemic rationality should not require from us to reason and inquire in a way that interferes with other important activities in life [42]. Finally, the content of our beliefs and the process of our inquiry depend on the characteristics of the environment in which we operate: the access to information, the prevalence of misinformation and so on. If we inquire in an epistemically polluted environment, where we do not know which sources are reliable and there is a large amount of inaccurate or incomplete information, the formation of false beliefs should not be considered a sign or a consequence of irrationality – this view has already been put forward by Levy [44]. Another feature of our epistemic environment is uncertainty; we can hardly be expected to perform expected utility maximizations if we do not have a complete knowledge of alternatives and probability distributions. Bounded epistemic rationality must therefore apply to conditions of uncertainty, not just to conditions of risk, which are rare in everyday situations.

Another feature of bounded epistemic rationality is that it is not defined by adherence to a priori normative criteria, but by the fit between the strategy and the environment. Bounded epistemic rationality is therefore ecological (in Gigerenzer's sense) and does not strive for a rigid, all-encompassing set of rules of rationality, but allows the question of which strategies give the best results in certain tasks to be investigated empirically. In this sense, bounded epistemic rationality is consequentialist – it cares about cognitive success in the sense of Schurz and Hertwig [9] – and as such it allows for different strategies, from complex reasoning to simple heuristics, to be rational.

Such an account of bounded epistemic rationality is compatible with moderate epistemic naturalism [16]: cognitive psychology and other disciplines can help us understand how human belief formation, inquiry, and other cognitive processes work and that can serve as a baseline for determining what can be expected from us. However, defining what is good enough – enough knowledge, enough understanding, close enough to the truth – is still a philosophical question. Bounded epistemic rationality does not collapse into adaptiveness or fitness, as it emphasizes the notion of epistemic goals and still asks the question of how to arrive at the most accurate beliefs possible and what constitutes good inquiry, so it encompasses descriptive, normative and ameliorative aspects. Such a view of epistemic rationality can encourage us to think about what Bishop and Trout [43] call epistemology for real people and help us move away from idealized theories to notions of rationality that are psychologically realistic and applicable to real human cognizers in the world in which they live. If we assume that one of the tasks of epistemology is to provide guidance on how to conduct better inquiries and achieve our epistemic goals, it must be based on non-ideal, empirically grounded models of human cognizers and the environment [42, 43, 45].

QUESTIONS FOR FUTURE RESEARCH

Numerous aspects of bounded epistemic rationality require further investigation. One of these is how to determine which characteristics of human cognition should be considered a factor in lowering the rationality threshold; in other words, which limitations matter. This question has already been raised by Carr [46], and Thorstad suggested that the bounds that should lower the bar are those imposed by the agents' cognitive architecture (e.g. limitations of working memory capacity, computational power, or processing speed), but not the processes implemented within it. However, he points out that his aim is not to provide a detailed account of which bounds matter, but to stimulate further discussion on the topic [24].

Another question relates to the practical limitations for inquiries that we face in everyday life. If a concept of bounded epistemic rationality states that we should prioritize and decide which tasks are worth our time and resources, should it not also specify how we do so? The problem with bounded epistemic rationality is that it does not provide an account of which problems are important enough to merit long and careful inquiry. Bishop and Trout [42] have already pointed out that an epistemological theory should determine which problems are relevant. In their theory of strategic reliabilism, they attempted to describe significant problems in terms of so-called objective reasons, and a similar account may be adopted in a theory of bounded epistemic rationality.

Another need for further research lies in the social aspects of bounded epistemic rationality. In this article I have focused on the cognitive limitations of individual agents, but belief formation, knowledge acquisition, and inquiry are of course strongly influenced by factors that have long been neglected by traditional analytic epistemology. The emergence of the field of social epistemology has emphasised the study of the epistemic properties of social interactions, social systems, and groups, for example the role of social institutions in knowledge production and dissemination and the epistemic effects of agents' social roles and identities [45, 47]. I believe that the nature of social institutions, social interaction between agents, and other epistemically relevant aspects of our social situatedness should be taken into account when we think about bounded epistemic rationality.

Another question is how to determine what is accurate enough, true enough or explained enough. The question of which strategy in which environment leads to more accurate predictions than others is an empirical one, but determining the threshold for acceptable accuracy is not. An account of bounded epistemic rationality should specify how to determine this threshold. For example, if we are going to act on the basis of a belief and the consequences of acting on the basis of an inaccurate belief are severe, we can set the threshold for accuracy higher; this has already been proposed under the notion of pragmatic encroachment [48]. The question of what is true enough, especially in scientific research that relies heavily on models and idealizations, has been thoroughly explored by Catherine Elgin [6, 49]. I believe that the connection between bounded epistemic rationality and the concept of "true enough" is worthy of further research.

Another important feature of bounded epistemic rationality is strategy selection – to be boundedly epistemically rational, cognizers should know when to use a particular strategy. This is a question of both descriptive and normative importance, as the empirical study of metacognition is crucial when determining what can be expected of cognizers and when we think about what is good or bad metacognition.

CONCLUSION

In the article, I aimed to show that combining epistemic and bounded rationality in a hybrid concept could be a way of narrowing the gap between descriptive and normative approaches

to the study of rationality. I pointed out that I distinguish epistemic justification from epistemic rationality, and defined the latter as a type of rationality aimed at pursuing epistemic goals. I described the division of labour between philosophy and psychology, with the former providing normative theories of rationality and the latter proposing descriptive, processing accounts of cognition based on empirical data. I described how investigations of rationality are intertwined with the notion of normativity and pointed out the unattainable norms and the arbitration problem as the challenges for normative theories. I then proposed three possible types of relationships between normative and descriptive theories of rationality, pointing out the importance of avoiding is-ought fallacy and contrasting it with an acceptable ought-implies-can principle of normativity. I then described Herbert Simon's concept of bounded rationality and Gigerenzer's notion of ecological rationality as examples of non-ideal accounts of rationality that aim to overcome the problems of normative theories mentioned above. I pointed out that the notions of bounded or ecological rationality are rarely used in the philosophical literature, and summarized David Thorstad's arguments for a greater emphasis on bounded rationality in epistemology. I reflected on the relationship between the norms of bounded rationality and the norms of epistemic rationality in traditional analytic epistemology. Finally, I outlined some features that I think should be included in the hybrid concept of bounded epistemic rationality: satisficing instead of optimizing; the consideration of cognitive, environmental, and practical limitations; its ecological nature; and its compatibility with moderate epistemic naturalism.

A definition of epistemic rationality that is compatible with bounded rationality might help us avoid overly strict, idealized, *as-if* theories of rationality and bring us closer to understanding rationality through the interaction of our cognitive faculties and the structure of the environment. If we allow that rationality of a particular strategy should be understood as the degree of fit between the strategy and the environment, and assume that this is an empirical question, then bounded epistemic rationality blurs the divide between the normative and the descriptive. In this way, we can construct a theory that is not detached from the real world, but can provide good epistemic advice and guidelines for epistemically good inquiring and belief formation. Linking bounded and epistemic rationality can therefore provide a more comprehensive understanding of rationality that is applicable to real human subjects in the world in which they live.

REFERENCES

- [1] Spohn, W. and Knauff, M.: *The Handbook of Rationality*. MIT Press, Cambridge, 2021,
- [2] Audi, R.: *Theoretical Rationality: Its Sources, Structure, and Scope*. In: Rawling, P. And Mele, A.R., eds.: *The Oxford Handbook of Rationality*. Oxford University Press, New York, pp.17-44, 2004, <http://dx.doi.org/10.1093/oxfordhb/9780195145397.003.0002>,
- [3] Foley, R.: *The Theory of Epistemic Rationality*. Harvard University Press, 1987, <http://dx.doi.org/10.4159/harvard.9780674334236>,
- [4] David, M.: *Truth as the Epistemic Goal*. In: Steup, M., ed.: *Knowledge, Truth, and Duty*, Oxford University Press, New York, pp.151-169, 2001, <http://dx.doi.org/10.1093/0195128923.003.0010>,
- [5] Kvanvig, J.L.: *Truth is Not the Primary Epistemic Goal*. In: Steup, M. and Sosa, E., eds.: *Contemporary Debates in Epistemology*, Wiley-Blackwell, Malden, pp.285-295, 2005,
- [6] Elgin, C.Z.: *True Enough*. The MIT Press, Cambridge, 2017, <http://dx.doi.org/10.7551/mitpress/9780262036535.001.0001>,

- [7] Foley, R. *Epistemic Rationality*.
In: Bernecker, S. and Pritchard, D., eds.: *The Routledge Companion to Epistemology*, Routledge, New York, pp.37-45, 2011,
- [8] Gigerenzer, G. and Sturm, T.: *How (far) can Rationality be Naturalized?*
Synthese **187**(1), 243-268, 2012,
<http://dx.doi.org/10.1007/s11229-011-0030-6>,
- [9] Arkes, H.R.; Gigerenzer, G. and Hertwig, R.: *How Bad Is Incoherence?*
Decision **3**(1), 20-39, 2016,
<http://dx.doi.org/10.1037/dec0000043>,
- [10] Stein, E.: *Without Good Reason: The Rationality Debate in Philosophy and Cognitive Science*.
Clarendon Press, Oxford, 1996,
- [11] Kahneman, D.; Tversky, A. and Slovic, P.: *Judgment Under Uncertainty: Heuristics and Biases*.
Cambridge University Press, New York, 1982,
<http://dx.doi.org/10.1017/CBO9780511809477>,
- [12] Gigerenzer, G.: *Rationality for Mortals: How People Cope with Uncertainty*.
Oxford University Press, New York, 2008,
<http://dx.doi.org/10.1093/oso/9780195328981.001.0001>,
- [13] Hume, D.: *A Treatise of Human Nature*.
Clarendon Press, Oxford, 1739,
- [14] Wedgwood, R.: *The Nature of Normativity*.
Oxford University Press, Oxford, 2007,
<http://dx.doi.org/10.1093/acprof:oso/9780199251315.001.0001>,
- [15] Broome, J.: *Normativity, Rationality and Reasoning: Selected Essays*.
Oxford University Press, Oxford, 2021,
<http://dx.doi.org/10.1093/oso/9780198824848.001.0001>,
- [16] Goldman, A.: *Epistemology and Cognition*.
Harvard University Press, Massachusetts, 1986,
- [17] Elqayam, S. and Evans, J.S.B.T.: *Subtracting Ought from Is: Descriptivism versus Normativism in the Study of Human Thinking*.
Behavioral and Brain Sciences **34**(5), 233-248, 2011,
<http://dx.doi.org/10.1017/S0140525X1100001X>,
- [18] Gigerenzer, G. *What is Bounded Rationality?*
In: Viale, R., ed.: *Routledge Handbook of Bounded Rationality*, Routledge, New York, pp.55-71, 2021,
<http://dx.doi.org/10.4324/9781315658353-2>,
- [19] Gigerenzer, G.; Hertwig, R. and Pachur, T.: *Heuristics: The Foundations of Adaptive Behavior*.
Oxford University Press, New York, 2011,
<http://dx.doi.org/10.1093/acprof:oso/9780199744282.001.0001>,
- [20] Rich, P.; Blokpoel, M.; de Haan, R. and van Rooij, I.: *How Intractability Spans the Cognitive and Evolutionary Levels of Explanation*.
Topics in Cognitive Science **12**(4), 1382-1402, 2020,
<http://dx.doi.org/10.1111/tops.12506>,
- [21] Wason, P.C.: *On the Failure to Eliminate Hypotheses in a Conceptual Task*.
Quarterly Journal of Experimental Psychology **12**(3), 129-140, 1960,
<http://dx.doi.org/10.1080/17470216008416717>,
- [22] Oaksford, M. and Chater, N.: *A Rational Analysis of the Selection Task as Optimal Data Selection*.
Psychological Review **101**(4), 608-631, 1994,
<http://dx.doi.org/10.1037/0033-295X.101.4.608>,
- [23] Cohon, R.: *Hume's Moral Philosophy*.
<https://plato.stanford.edu/entries/hume-moral>, accessed 6th December, 2023,

- [24] Thorstad, D.: *Why Bounded Rationality (in Epistemology)?*
 Philosophy and Phenomenological Research, 2023,
<http://dx.doi.org/10.1111/phpr.12978>,
- [25] Evans, J. and Over, D.E.: *Rationality and Reasoning*.
 Psychology Press, London, 1996,
- [26] Stroßner, C.: *Hybrid Frameworks of Reasoning: Normative-Descriptive Interplay*.
 Proceedings of the 45th Annual Conference of the Cognitive Science Society **45**, 1153-1160, 2023,
- [27] von Neumann, O. and Morgenstern, J.: *Theory of Games and Economic Behavior*.
 Princeton University Press, Princeton, 1944,
- [28] Simon, H.A.: *Administrative Behavior: A Study of Decision-Making Processes in Administrative Organizations*.
 Macmillan Publishers Limited, New York, 1947,
- [29] Simon, H.A.: *A Behavioral Model of Rational Choice*.
 The Quarterly Journal of Economics **69**(1), 99-118, 1955,
<http://dx.doi.org/10.2307/1884852>,
- [30] Simon, H.A.: *Rational Choice and the Structure of the Environment*.
 Psychological Review **63**(2), 129-138, 1956,
<http://dx.doi.org/10.1037/h0042769>,
- [31] Simon, H.A.: *Invariants of Human Behavior*.
 Annual Review of Psychology **41**(1), 1-20, 1990,
<http://dx.doi.org/10.1146/annurev.ps.41.020190.000245>,
- [32] Simon, H.A.: *From Substantive to Procedural Rationality*.
 In: Kastelein, T.J.; Kuipers, S.K.; Nijenhuis, W.A. and Wagenaar, G., eds.: *25 Years of Economic Theory*. Springer, Boston, pp.65-86, 1976,
http://dx.doi.org/10.1007/978-1-4613-4367-7_6,
- [33] Sargent, T.J.: *Bounded Rationality in Macroeconomics*.
 Oxford University Press, Oxford, 1993,
<http://dx.doi.org/10.1093/oso/9780198288640.001.0001>,
- [34] Gigerenzer, G. and Selten, R., eds.: *Bounded Rationality: The Adaptive Toolbox*.
 MIT Press, Cambridge, 2001,
<http://dx.doi.org/10.7551/mitpress/1654.001.0001>,
- [35] Gigerenzer, G.: *Why Heuristics Work*.
 Perspectives on Psychological Science **3**(1), 20-29, 2008,
<http://dx.doi.org/10.1111/j.1745-6916.2008.00058.x>,
- [36] Carnap, R.: *On the Application of Inductive Logic*.
 Philosophy and Phenomenological Research **8**(1), 133-148, 1947,
<http://dx.doi.org/10.2307/2102920>,
- [37] Sturm, T.: *Towards Critical Naturalism about Bounded Rationality*.
 In: Viale, R., ed.: *Routledge Handbook of Bounded Rationality*, Routledge, New York, pp.73-90, 2021,
<http://dx.doi.org/10.4324/9781315658353-4>,
- [38] Sturm, T.: *Formal versus Bounded Norms in the Psychology of Rationality: Toward a Multilevel Analysis of Their Relationship*.
 Philosophy of the Social Sciences **49**(3), 190-209, 2019,
<http://dx.doi.org/10.1177/0048393119842786>,
- [39] Gigerenzer, G. and Brighton, H.: *Homo Heuristicus: Why Biased Minds Make Better Inferences*.
 Topics in Cognitive Science **1**(1), 107-143, 2009,
<http://dx.doi.org/10.1111/j.1756-8765.2008.01006.x>,
- [40] Friedman, J.: *The Epistemic and the Zetetic*.
 Philosophical Review **129**(4), 501-536, 2020,
<http://dx.doi.org/10.1215/00318108-8540918>,

- [41] Thorstad, D.: *Inquiry and the Epistemic*.
Philosophical Studies **178**(9), 2913-2928, 2021,
<http://dx.doi.org/10.1007/s11098-020-01592-y>,
- [42] Bishop, M.A. and Trout, J.D.: *Epistemology and the Psychology of Human Judgment*.
Oxford University Press, New York, 2005,
<http://dx.doi.org/10.1093/0195162293.001.0001>,
- [43] Bishop, M. and Trout, J.D.: *Epistemology for (Real) People*.
In: Lippert-Rasmussen, K.; Brownlee, K. and Coady, D., eds.: *A Companion to Applied Philosophy*. Wiley-Blackwell, Malden, pp.103-120, 2016,
<http://dx.doi.org/10.1002/9781118869109.ch8>,
- [44] Levy, N.: *Bad Beliefs: Why They Happen to Good People*.
Oxford University Press, Oxford, 2021,
<http://dx.doi.org/10.1093/oso/9780192895325.001.0001>,
- [45] McKenna, R.: *Non-Ideal Epistemology*.
Oxford University Press, New York, 2023,
<http://dx.doi.org/10.1093/oso/9780192888822.001.0001>,
- [46] Carr, J.R.: *Why Ideal Epistemology?*
Mind **131**(524), 1131-1162, 2022,
<http://dx.doi.org/10.1093/mind/fzab023>,
- [47] Goldman, A.I. *Why Social Epistemology is Real Epistemology*.
In: Haddock, A.; Millar, A. and Pritchard, D., eds.: *Social Epistemology*, Oxford University Press,
New York, pp.1-29, 2010,
<http://dx.doi.org/10.1093/acprof:oso/9780199577477.003.0001>,
- [48] Fantl, J. and McGrath, M.: *On Pragmatic Encroachment in Epistemology*.
Philosophy and Phenomenological Research **75**(3), 558-589, 2007,
<http://dx.doi.org/10.1111/j.1933-1592.2007.00093.x>,
- [49] Elgin, C.Z.: *True Enough*.
Philosophical Issues **14**(1), 113-131, 2004,
<http://dx.doi.org/10.1111/j.1533-6077.2004.00023.x>.