Nastja Tomat
Univerza v Ljubljani

Bounded Epistemic Rationality and Norms of Inquiry

1 Introduction

Every day we conduct inquiries about various topics, from trivial to important ones. In recent years, the question of how we should inquire has gained increased attention in epistemology, and this shift in focus from the study of epistemic norms or norms of belief to the study of zetetic norms or norms of inquiry has been called the zetetic turn (Friedman, forthcoming). While epistemology has often relied on idealized models of epistemic agents and their cognitive capacities (McKenna, 2023, 2), it seems clear that real humans are bounded in various ways and that they cannot inquire in the same way as ideal agents.

The aim of the paper is to examine the norms of rational inquiry for non-ideal, psychologically realistic epistemic agents. Following Herbert Simon's work on bounded rationality, I argue that epistemic rationality should be understood as bounded and that we should focus on its procedural aspect and the norms of rational inquiry, rather than on the rationality of final doxastic states. I aim to show that an interdisciplinary, empirically informed, non-ideal approach to the study of norms of inquiry can provide good epistemic guidance for ordinary agents and contribute to the development of an "epistemology for real people" (Bishop and Trout, 2004, 2016, 106).

The plan for the paper is as follows. First, I will explain the concepts of non-ideal epistemology and non-ideal rationality, focusing on Gigerenzer's ecological rationality. I will then briefly describe Simon's notion of bounded rationality and the concept of bounded epistemic rationality, which I proposed and elaborated on in a previous work (Tomat, 2024a, 2024b). I will then provide an overview of selected norms of inquiry proposed in the literature. Finally, I will suggest some features that non-ideal norms of inquiry for bounded agents should include. I aim to show that collaboration between philosophy and empirical disciplines is crucial if we are to develop zetetic norms that are applicable to real human agents and could help them improve their inquiries.

DOI:10.4312/ars.19.1.51-66

2 Non-ideal epistemology and non-ideal epistemic rationality

There is ample empirical evidence from disciplines such as psychology and cognitive science that we are bounded agents with limited cognitive abilities, e.g. computational and predictive power, memory and attention. Although this has often not been reflected in the norms for rational belief and rational inquiry, there are several authors who acknowledge the relevance of our cognitive limitations to epistemological questions and work on non-ideal conceptualizations of epistemic rationality (Begby, 2021; Hughes, 2024; Morton, 2012; Thorstad, 2024a, 2024b; Siscoe, 2022; Singer, 2023; Smithies, 2015; Staffel, 2019).

The need to incorporate empirical evidence about various kinds of limitations of real, ordinary epistemic agents into epistemological theorizing has recently been articulated in terms of a meta-epistemological approach of non-ideal epistemology (McKenna, 2023). While ideal epistemology is concerned with how perfectly rational and cognitively unlimited agents would think, and often proposes norms that require logical omniscience, consistency between beliefs, immediate updating of beliefs through Bayesian conditionalization, etc. (Carr, 2022, 1132), non-ideal epistemology aims to rely on psychologically realistic conceptions of epistemic agents and explores what norms are attainable for them and would help them improve their epistemic position in real-world situations[1] (McKenna, 2023, 10).

Authors working on epistemic rationality[2] through the lens of non-ideal epistemology fall into two categories (Hughes, 2024, 75). In the first category, there are non-ideal Bayesians who are concerned with norms that approximate those of standard Bayesianism (Staffel, 2019); in the second category, there are authors such as Gigerenzer (2000, 2008) who are concerned with ecological rationality. Ecological rationality is a concept that rejects defining rationality as adherence to *a priori* defined norms, such as following the rules of logic, probability or decision theory, but understands rationality as a fit with the environment. A cognitive strategy for judgment, belief formation, problem-solving is ecologically rational to the extent that it is adapted to the structure of the task. If we choose a strategy that leads to more accurate predictions than other possible strategies at a given task, we are ecologically rational. A large part of Gigerenzer's research programme is devoted to the empirical study of fast and frugal

1 Note that the line between ideal and non-ideal epistemology can be drawn in multiple ways (McKenna, 2023, 12), and that they are not separate categories, but should be thought of as a continuum (McKenna, 2023, 25). There are various types of idealizations that can be used in epistemological theorizing: about the psychology of epistemic agents, interactions between them, epistemic environment, and social institutions (McKenna, 2023, 5). In the paper, I focus on the idealizations of cognitive capacities and epistemic environment and leave other types of idealizations aside.

2 By epistemic rationality I mean rationality that is directed towards reaching cognitive or epistemic goals, such as truth, knowledge, accuracy, or understanding (Foley, 1987; Knauff and Spohn, 2021).

heuristics: rules of thumb that use only part of the available information and are computationally less demanding, but sometimes lead to more accurate judgments than more complex strategies. An ecological conception of rationality is non-ideal, as it rejects the traditional view of epistemic rationality as adherence to requirements of coherence, deductive closure, logic, probability, etc.; instead, it relies on data about human cognitive processing and empirically investigates how different strategies work in real-world settings (Gigerenzer, 2000, 2008).

2.1 Bounded rationality in epistemology

An important concept that aims to overcome some of the issues of idealized theories of rationality is Herbert Simon's bounded rationality. The core idea of his theory is that we must replace highly idealized conceptions of rationality with a notion of rationality that is compatible with the limitations of the human cognitive system. Simon claimed that humans do not strive for the best possible solution, but merely for a solution that is good enough; in other words, humans are satisficers, not optimizers, and optimization should not be the requirement of rationality. He emphasized that rationality should be understood through the interaction between the strategy and the environment. He used the metaphor of scissors, where one blade represents the structure of the task environment and the other blade represents the computational capabilities of the actors, and claimed that one must examine both blades and the interaction between them to understand rationality. He also pushed for a procedural view of rationality, viewing behaviour as procedurally rational when it results from appropriate deliberation (Simon, 1955, 1956, 1976, 1990, 1992).

Although Simon's theory was primarily applied to rational decision-making and behaviour and would therefore fall within the realm of practical philosophy, I believe that bounded rationality can be fruitfully applied to research in the field of epistemology. Several philosophers have related bounded rationality to various epistemological issues; a comprehensive overview is beyond the scope of this paper, but see Dusi, 2024; Greco, 2023; Petracca and Grayot, 2023; Sturm, 2019; and Viale, 2020. Some authors rely, for example, on Simon's notion of satisficing. Talbot (2024, 56) argues that consequentialist epistemic norms should require satisficing instead of maximizing. Given that epistemic goods are gradable – we can be more or less rational, closer or further from the truth, have more or less understanding – epistemic practices will often involve a range of strategies: the epistemically best option (e.g., a method of inquiry that gets us as close to the truth as possible); a range of suboptimal options; and at least one option that is much worse than the rest. While maximizing norms only distinguish between optimal alternatives – e.g. the best way to form a belief – and suboptimal alternatives, satisficing norms can distinguish between optimal, suboptimal and worst

strategies. Talbot (2024, 53) argues that the difference between worst and suboptimal strategies, which is captured in satisficing norms but not in maximizing norms, is much more relevant for epistemically good practices than the difference between optimal and suboptimal strategies. There are many cases in which the difference between the optimal and the second-best method of inquiry leads to such a small difference in outcome (e.g., in the degree of understanding) that it is negligible; in contrast, the difference between suboptimal strategies and the worst strategy, which leads to missing the target altogether, is very important. Epistemic norms should therefore be satisficing and not maximizing.[3]

The author who has provided probably the most comprehensive work on bounded rationality in epistemology to date is David Thorstad (2021, 2024a, 2024b, forthcoming). He sees bounded rationality as a paradigm and lists five normative claims that define it. First, bounds should be relevant to our understanding of rational cognition; second, we should focus on the procedural aspect of rationality, which means that we should set aside normative questions about beliefs and focus on normative questions about inquiry; third and fourth, we should allow that heuristics are rational and consider rationality to be ecological; and fifth, we should consider that violations of traditional rationality requirements are not cases of irrationality but cases of the most rational deliberation possible given our limitations.

On the basis of Thorstad's work, I have proposed a concept of bounded epistemic rationality, which I have unpacked in earlier work (Tomat, 2024a, 2024b). Briefly summarized, I have argued that bounded epistemic rationality is directed toward various epistemic goals such as truth, knowledge, understanding, or accuracy. In the sense of Simon's satisficing, it does not require optimal solutions, but only sufficiently good ones. A person is boundedly epistemically rational if they inquire in a way that leads to sufficiently good results – good enough understanding, good enough accuracy or true enough beliefs.[4] Bounded epistemic rationality applies the ought-implies-can principle of normativity: rational requirements for inquiry can only be those that can in principle be achieved by bounded epistemic agents. Bounded epistemic rationality is ecological (in Gigerenzer's sense), meaning that the rationality of a particular type of inquiry is defined by its success in the real world, not by its adherence to *a priori* normative criteria.

3 Norms of inquiry

Thorstad (forthcoming) argues that zetetic epistemology should be seen as a study of bounded rationality. One of the arguments for this claim is that our cognitive

3 See also Pils (2022) for a satisficing account of epistemic justification.

4 Here I refer to Catherine Elgin's notion of true enough (2004, 2017).

limitations affect the process of inquiry more than the resulting attitudes, and if we want to study the effects of human cognitive limitations on our rationality, we should study the inquiry and not the final attitudes. Thorstad claims that a zetetic turn in epistemology brings us closer to developing theories of epistemic rationality and other epistemic phenomena that are "humanly adequate" (Thorstad, forthcoming).

We can think of inquiry as a series of actions accompanied by specific kinds of attitudes, that occur over a limited interval of time, with the aim of figuring something out or settling a question (Friedman, 2017a). Inquiry consists of various phases. First, an inquirer becomes curious and begins to wonder about a question, then begins to actively investigate it, and she finally settles the question (Friedman, forthcoming). Traditional epistemic norms focus only on a small part of the whole process of inquiry, specifically on the last part when one is already at the stage of settling a question, and they tell us whether a person has enough evidence to form a belief, whether a particular belief can be considered knowledge, and so on. Friedman (forthcoming) argues that epistemology should provide norms that govern the whole process of inquiry.

One of the norms proposed by Friedman (2017a, 2017b) is that a person who has an interrogative attitude towards a question Q at time t must suspend her judgment about Q at time t. In other words, one should not inquire about something that one already believes. This norm is a kind of requirement of coherence, as it prohibits inquirers from having a certain combination of attitudes: believing that p and inquiring about p (Fredman, 2017a, 8). Once we enter the stage of active investigation, we need norms of inquiry that tell us how to proceed in order to settle the question. This suggests that in addition to the negative norms of inquiry that prescribe negative requirements – not to know, to believe, to be certain and to inquire at the same time – it is crucial to propose positive norms of inquiry that prescribe what one ought to do when inquiring. One of these is the zetetic instrumental principle, which states that "if one wants to figure out Q, then one ought to take the necessary means to figuring out Q" (Friedman, 2020, 503). Furthermore, Willard-Kyle (2023, 7) proposes a knowledge norm that says that one ought inquire into Q only if one knows that Q has a true answer, since it seems that something is epistemically wrong to inquire about the question if one does not know that a true answer exists. A related norm is the Meno norm, which states that "one ought: inquire into Q only if one can recognize a correct answer to Q" (Haziza, 2023, 6), reflecting the intuition that it is wrong to inquire about a question if one cannot recognize the correct answer to it.

It seems, however, that we need more specific norms of inquiry – ones that could provide concrete guidance in particular situations in which ordinary inquirers find themselves. How should we gather and evaluate evidence? What judgments and conclusions should we make based on the available data? Should we select strategies of inquiry based only on epistemic criteria such as reliability, or could we also consider

non-epistemic criteria such as the time required? Such questions suggest that zetetic norms should evaluate and rank the possible strategies of inquiry according to certain criteria and prescribe to agents which strategy they should use in a given situation. What should also be acknowledged is that we cannot pursue all questions at the same time and that we have considerable limitations on what can be done simultaneously. We need to allocate our resources and zetetic norms should tell us how to do this. Should we prioritize the questions that are most relevant to us right now, the questions we are most curious about, the questions we are most likely to get answers to (Friedman, forthcoming)?

4 Norms of inquiry for boundedly rational agents

Even if we adopt a non-ideal approach to studying epistemic rationality, incorporating the insights of bounded rationality into norms of inquiry is not an easy task. This brings us to the main question of this paper: how should boundedly epistemically rational agents inquire? First, we should explain what we mean by bounded agents. The bounds I think should be recognized in the norms of rational inquiry are cognitive, environmental, and practical. In all three aspects, norms should employ ought-implies-can principle of normativity: they should require that inquiry is conducted in a way that is achievable for ordinary human inquirers. In terms of our cognitive bounds, this means that norms should consider empirical data provided by disciplines such as cognitive and social psychology and cognitive science about the mechanisms of human cognitive processing and its limitations: judgment and decision-making, reasoning, hypothesis testing, inference and problem-solving, attention, memory, computational and predictive abilities, and so on. Such data should put a constraint on norms of inquiry: we should not be required to allocate attention on ten different tasks simultaneously, to perform a complex statistical analysis to determine which brand of yogurt in the store is the optimal choice given our preferences and budget, or to immediately update our beliefs in light of new evidence.[5]

Furthermore, norms of inquiry should recognize that we are limited not only by our cognition, but also by resources such as time. Inquiring is just one of the activities in our daily lives, and in many situations, rational inquiry – or the pursuit of epistemic goals – simply cannot not be our priority. Boundedly epistemically rational agents

5 When incorporating our bounds into norms of inquiry, we must address the question of which bounds legitimately lower the bar for rationality (Carr, 2022, 524). I find Thorstad's (2024b, 8) suggestion quite convincing: drawing on the distinction between a fixed cognitive architecture and representations or processes in cognitive science (Langley et al., 2009), he proposes that only features of cognitive architecture, but not the processes or representations realized in it, are of importance for our attributions of rationality. For example, we are not irrational if we cannot store ten pieces of information in our working memory at the same time when we inquire, but we are irrational if we fill our working memory with information that is irrelevant to the current inquiry.

should devote their cognitive and other resources to questions that are – or should be – important or relevant to them in some way. Providing an account of relevance is a difficult task. Bishop and Trout (2004) have, in their theory of strategic reliabilism, proposed an account of significant problems, and claimed that the main challenge for such an account is that it must allow for interpersonal differences in what we consider significant but at the same time avoid "anything goes" subjectivism (Bishop and Trout, 2004, 95). They proposed a reason-based account of significance, claiming that significance of a problem for S is "a function of the weight of the objective reasons S has for devoting resources to solving that problem" (Bishop and Trout, 2004, 95). Accepting this claim means accepting the existence of objective reasons for action, including inquiry. They can be based on moral, prudential or social role obligations and are in general conducive to human well-being.[6] In line with this position, we could accept that there are objective reasons for inquiry and that boundedly rational inquirers are obliged to prioritize inquiry about topics for which they have objective reasons. It seems that an account of epistemic relevance must strike a balance that allows agents to devote some resources to inquiring about topics that are personally relevant to them, such as learning astronomy as a hobby; to inquiring about topics that are relevant because of their professional or social obligations, such as attending seminars on new diagnostic tools if they are a doctor; and to inquiring about topics that are important to society, such as researching the programmes of political parties before an election.

Since the norms of inquiry for bounded agents rely on empirical data about human cognition, an account of relevant problems should incorporate insights about how people determine which problems are relevant to them in their daily lives. There is an ongoing debate in cognitive science about how people decide which aspects of their environment are important. Szollosi and Newell (2020, 1), for example, suggest that we should move away from the assumption that people represent their environment in a fixed way and instead examine how people actually determine which feature of the environment is important enough to represent. This perspective recognizes people's ability to "flexibly construct, modify, and replace the representations of the decision problems they face" (Szollosi and Newell, 2020, 1). Similarly, Strle (2016, 85) argues that decision-making should be understood as a sense-making activity and must be studied from the perspective of the decision-maker. Decision-making should therefore not only be examined from an "objective" third-person perspective, but must also include first-person data. Approaches that collect data on decision-making and other cognitive processes from a first-person perspective, such as empirical phenomenology, can help us gain insights into how people determine which situation to inquire about or decide in the first place, how they determine which features of situations are

6 For a detailed discussion, see Bishop and Trout (2004).

important and what are relevant alternatives, and what the subjective experience of such processes looks like. Such approaches to the study of relevance are consistent with bounded and ecological rationality, which aims to move away from *a priori*, fixed normative criteria for what is right and what is a mistake, what is rational and what is not. If we move away from the assumption that there is a statistical model of the environment that one should apply in order to maximize gains – in other words, that there is the best possible representation of an environment – we can explore the variety of ways in which an environment can be represented and study the factors that influence that representation, from inherent psychological characteristics such as categorization ability, cultural factors, motivation and goals, background knowledge, and so on (Szollosi and Newell, 2020, 8). Although such empirical data on *how* people determine relevance cannot, of course, alone answer the normative question of which problems *should* be epistemically relevant, they should, like other empirical data on cognitive processing, be included in an account of rational inquiry for bounded agents.

Suppose we have an account for determining the relevance of the issues and can identify the problems that a boundedly rational agent should start inquiring about. When we enter the stage of active inquiry, are we required to inquire in the most rational way possible – the way that gets us as close as possible to our epistemic goal given our cognitive, practical, and environmental limitations? I argue that this is not the case. Following Talbot (2024) and Pils (2022), I argue that zetetic norms for bounded agents should be of a satisficing and not a maximizing nature. They should allow inquiry in a way that leads to good enough results: good enough understanding, true enough beliefs, or accurate enough predictions. But how can the threshold for "good enough" be determined? Again, both epistemic and non-epistemic factors play a role. For example, what is "good enough" can be determined by practical considerations – if an inaccurate prediction or incomplete understanding of a particular problem leads to severe negative consequences, the threshold should be higher.[7] At the same time, we do not want to set the bar for "good enough" too low. Good enough solutions still need to meet some minimal epistemic criteria, but in terms of satisficing, there is a whole range of possible good enough solutions – from those who barely meet these criteria to the best possible solution. Let us take understanding as an example. We can say that some epistemic criteria must be met for something to be classified as understanding: a person must – to a certain degree – grasp explanatory and conceptual relationships between objects (Grimm, 2021; Kvanvig, 2018). It seems that understanding is gradable: one person's understanding of an issue can be deeper and more comprehensive than

7 This is compatible with the thesis of pragmatic encroachment (Fantl and Mcgrath, 2007); but here pragmatic considerations do not influence how we think about epistemic desiderata such as knowledge or justification, but are used to refine an account of what is good enough solution. Pragmatic factors can help us determine whether partial achievement of an epistemic goal is epistemically good enough for an inquirer with specific professional and social obligations in specific circumstances.

another's. This means that we have a range of solutions: from an understanding that barely meets the criteria for good enough to the most comprehensive understanding possible. Once we have established what minimum epistemic criteria must be met, pragmatic and other considerations become relevant: the threshold for good enough understanding for a layperson who is simply curious about a topic and has very little time for inquiring about it is different from the threshold for a scientist who is expected to contribute to policy making with their expertise.

Let us now turn to the environmental limitations. Following Simon, I want to emphasize that norms of inquiry should be sensitive not only to internal but also to external factors that constrain us; in particular, when we think of epistemic rationality, we should think not only of our cognitive limitations but also of the epistemic environment in which we inquire. We could have perfect deductive skills, memory, attention, etc., but would still fail to achieve our epistemic goals if the information available to us were false. Levy (2021) introduced the concept of polluted epistemic environments – environments with a high prevalence of misinformation in which it is difficult to recognize trustworthy sources of information. Levy claims that in such environments, failure to achieve our epistemic goals – for example, forming a true belief about anthropogenic climate change – should not be attributed to the inquirers' lack of rationality or virtue, but to features of the environment. Of course, agents can conduct more or less rational inquiries regardless of the characteristics of the environment; however, it is important to recognize that in certain environments, the majority of agents would not achieve their epistemic goals even if they were to inquire in accordance with satisficing zetetic norms. What does this mean for our ascriptions of rationality? I would say that rational inquirers in such an environment should suspend their judgment until it is possible for them to reliably identify trustworthy sources of information.[8] In practice, this does not mean that inquirers should learn to reason better, inquire more responsibly, etc., but that we should improve our epistemic environments (Levy, 2021, 130).

This leads me to my next point about the importance of the epistemic environment. Bounded epistemic rationality is ecological in Gigerenzer's (2000, 2008) sense and assumes that a strategy of inquiry is rational to the extent that it is adapted to the environment. The rationality of inquiry is therefore not defined *a priori* on the basis of a rigid set of normative criteria, but by the fit with the task. In different environments and for different tasks, different strategies lead to our epistemic goals. In a polluted epistemic environment, it is rational to suspend judgment; in some tasks, it is rational to rely on heuristic thinking – for example, when we have limited time, information, and cognitive resources, and when the consequences of a potentially

8 This is in line with Huges' (2024) reliabilist account of non-ideal epistemic rationality.

false belief are not too severe, or when heuristics lead to more accurate judgments than other, more cognitively demanding strategies; in some cases, it is rational to form a belief on the basis of a scientific consensus on the matter, and so on. However, the ecological nature of bounded epistemic rationality does not suggest that every kind of inquiry that leads us to a good enough achievement of our epistemic goals is rational. For example, if a person engaged in wishful thinking and motivated reasoning and conducted inquiries that, in most cases, happened to lead to epistemic goals, we would not say that she is rational. We want our inquiries to satisfy some epistemic conditions, such as reliability and responsiveness to evidence, in order to be considered boundedly rational.[9]

I believe that the greatest challenge for a theory of rational inquiry for bounded agents is to propose norms that can serve as a practical guidance for ordinary inquirers. Norms that are prevalent in the literature on inquiry epistemology, such as the ignorance norm and the zetetic instrumental principle, are crucial for developing theoretical accounts of rational and responsible inquiry and mapping out a field of zetetic normativity. However, telling an agent she should take the necessary means to figuring out Q will not be very helpful to her in her daily inquiries. Future research should identify more fine-grained norms that could guide agents on how they should inquire in specific situations to achieve their epistemic goals. I believe that bounded and ecological rationality can provide a good framework for such research, as they recognize various types of limitations that constrain us, rely on empirical data provided by cognitive psychology and cognitive science, and allow rationality to be determined by the fit between strategy and environment. The research programme on non-ideal norms of inquiry must be interdisciplinary and partly empirical, since the question of which strategies lead to epistemic goals in which situations is an empirical one – but its main components remain philosophical reflections on zetetic and epistemic normativity and debates about epistemic goals and values. Norms of inquiry, functioning as epistemic advice, should tell us, for example, how to proceed when confronted with contradictory evidence; how to determine what constitutes a good piece of evidence and when we have enough evidence to (at least temporarily) terminate our inquiry; how to recognize our cognitive biases and avoid them when they prevent us from achieving our epistemic goals; how to inquire when it comes to peer disagreement; when and to what extent we can rely on expert testimony and when it is permissible to "do our own research", and so on. This requires collaboration between philosophical research and the empirical disciplines that study human cognition, including third- and

9 We would not require a boundedly epistemically rational person to gather all possible evidence before forming a belief. Heuristics, for example, ignore some of the available information and still lead to correct judgments in certain tasks (Gigerenzer, 2000, 2008). We could, for example, demand responsiveness only to a subset of evidence.

first-person approaches. As Strle (2016) points out, our accounts of cognitive processes are only partial if they fail to include experienced, subjective aspects of cognition. Gathering first-person data about how people determine that they do or do not have sufficient evidence for a conclusion; how they recognize that they are faced with conflicting evidence; how they decide to adjust their beliefs in light of new evidence or peer disagreement and so on, is an important part of understanding cognitive processes. Another interesting aspect of our cognition is metacognitive experiences such as curiosity, wonder, the desire to find something out, feeling of rightness, etc., which accompany people's inquiries and correlate with the start, unfolding, and the termination of inquiry. The phenomenological study of the experiential nature of these phenomena can help us gain a better understanding of human cognition that can inform norms of rational inquiry.

Norms of inquiry for bounded agents thus contain both normative and descriptive elements, blurring the line between "is" and "ought". They rely on empirical data about human cognition and the epistemic environment, but still raise normative, philosophical questions about what constitutes rational and responsible inquiry. Such an approach to the study of zetetic normativity is partially naturalized and requires – to some extent – bridging the gap between the normative and the descriptive. It does not replace normative questions with descriptive ones, but proposes a kind of limited normative naturalism (Gigerenzer and Sturm, 2012) that relies on empirical data to investigate which strategies of inquiry are most conducive to epistemic goals in particular circumstances. Such an approach is compatible with a consequentialist account of rational cognition, as proposed by Schurz and Hertwig (2018), who call on philosophers and psychologists to work together on a new, more comprehensive understanding of rationality that bridges the traditional gap between empirical studies of cognition and philosophical investigations of rationality.

There are debates in the literature about whether epistemology should be ameliorative or not (Ballantyne, 2019; Bishop and Trout, 2004; Hughes, 2021; Roberts and Wood, 2007). I think that a part of epistemology should be – without aiming to dismiss the importance of or replace traditional epistemological pursuits like identifying conditions for knowledge and justification. However, as Staffel (2019, 160) points out, a normative or evaluative theory cannot be simply translated into an ameliorative one. If we take, for example, a norm on evidence gathering saying something along the lines of "one should rely on information gained from reliable sources", this norm does not tell us *how* to do that. We cannot simply advise inquirers they select a trustworthy source of information, but we must provide some procedure or instruction on how to do so. Even if we have an account of which way of inquiry is rational in a specific environment, this will be no use to agents if they cannot recognize and select the strategy that our theory proposes. Many agents that inquire in a polluted environment do not

know that they are in a polluted environment; oftentimes we cannot judge if it is better to use a heuristic or a complex statistical method; and we frequently have poor insights into our cognitive shortcomings and overestimate our abilities. If theory of bounded epistemic rationality aims to be ameliorative and provide real-life, practical guidance for inquiry, it needs to unpack the metacognitive aspects of strategy selection and figure out if and how they can be trained – which, once again, calls for collaboration between philosophy, psychology, and cognitive science.

References

Ballantyne, N., *Knowing our limits*, New York 2019.

Begby, E., *Prejudice: A study in non-ideal epistemology*, Oxford 2021.

Bishop, M. et al., *Epistemology and the psychology of human judgment*, New York 2004.

Bishop, M. et al., Epistemology for (real) people. In: *A companion to applied philosophy* (ed. Lippert-Rasmussen, K., Brownlee, K., and Coady, D.), Chichester 2016, pp. 103–19.

Carr, J. R., Why ideal epistemology?, *Mind* 131 (524), 2022, pp. 1131–1162.

Dusi, G., Reliabilist epistemology meets bounded rationality, *Synthese* 203 (4), 2024, pp. 1–21.

Elgin, C., True enough, *Philosophical Issues* 14 (1), 2004, pp. 113–131.

Elgin, C., *True enough*. Cambridge 2017.

Fantl, J., McGrath, M., On pragmatic encroachment in epistemology, *Philosophy and Phenomenological Research* 75 83), 2007, pp. 558–289.

Foley, R. *The theory of epistemic rationality*. Cambridge 1987.

Friedman, J., Inquiry and belief, *Noûs* 53 (2), 2017a, pp. 296-315

Friedman, J., Why suspend judging?, *Noûs* 51 (2), 2017b, pp. 302–326.

Friedman, J., The epistemic and the zetetic, *Philosophical Review* 129 (4), 2020, pp. 501–536.

Friedman, J. Zetetic epistemology. In: *Towards an expansive epistemology: Norms, action, and the social sphere* (ed. Reed, B. and Flowerree, A. K.), forthcoming.

Gigerenzer, G., *Adaptive thinking: Rationality in the real world*, New York 2000.

Gigerenzer, G., *Rationality for mortals: How people cope with uncertainty*, New York 2008.

Gigerenzer, G., et al., How (far) can rationality be naturalized? *Synthese* 187, 2012, pp. 243–268.

Greco, D., *Idealization in epistemology: A modest modeling approach*, New York 2023.

Grimm, S. Understanding, *Stanford Encyclopedia of Philosophy*, 2021, https://plato.stanford.edu/entries/understanding/, accessed 20. 2. 2025

Haziza, E., Norms of inquiry, *Philosophy Compass* 18 (12), 2023, e12952.

Hughes, N., Epistemology without guidance, *Philosophical Studies* 179 (1), 2021, pp. 163–196.

Hughes, N., Non-ideal epistemic rationality, *Philosophical Issues* 34 (1), 2024, pp. 72–95.

Knauff et al. *The handbook of rationality*. Cambridge 2021.

Kvanvig, J., Knowledge, Understanding, and Reasons for Belief. In: *The Oxford Handbook of Reasons and Normativity* (ed. Starr, D.), New York 2018, pp. 685–705.

Langley, P. et al., Cognitive architectures: Research issues and challenges, *Cognitive Systems Research* 10 (2), 2009, pp. 141–160.

Levy, N., *Bad beliefs: Why they happen to good people*, Oxford 2021.

McKenna, R., *Non-ideal epistemology*, New York 2023.

Morton, A., *Bounded thinking: Intellectual virtues for limited agents*, Oxford 2012.

Petracca, E. et al., How can embodied cognition naturalize bounded rationality?, *Synthese* 201 (4), 2023, pp. 1–28.

Pils, R., A satisficing theory of epistemic justification, *Canadian Journal of Philosophy* 52 (4), 2022, pp. 450–467.

Roberts, R. et al., *Intellectual virtues: An essay in regulative epistemology*, Oxford 2007.

Schurz, G., et al., Cognitive success: A consequentialist account of rationality in cognition. *Topics in cognitive science* 11, 2019, pp 7–36.

Simon, H. A., Rational choice and the structure of the environment, *Psychological Review* 63 (2), 1956, pp. 129–138.

Simon, H. A., A behavioral model of rational choice, *The Quarterly Journal of Economics* 69 (1), 1955, pp. 99–118.

Simon, H. A., From substantive to procedural rationality. In: *25 Years of Economic Theory: Retrospect and Prospect* (eds. T. J. Kastelein, T. J., Kuipers, K., Nijenhuis, W. A., and Wagenaar, G. R.) Boston 1976, pp. 65–86.

Simon, H. A., Invariants of human behavior, *Annual Review of Psychology* 41, 1990, pp. 1–19.

Simon, H. A., What is an explanation of behavior?, *Psychological Science* 3 (3), 1992, pp. 150–161.

Singer, D. J., *Right belief and true Belief*, New York 2023.

Siscoe, R. W., Real and ideal rationality, *Philosophical Studies* 179 (3), 2022, pp. 879–910.

Smithies, D., Ideal rationality and logical omniscience, *Synthese* 192 (9), 2015, pp. 2769–2793.

Staffel, J., *Unsettled thoughts: A theory of degrees of rationality*, New York 2019.

Strle, T. Embodied, enacted and experienced decision-making. *Phainomena* 25(98/99), 2016, pp. 83–107.

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), 2019, pp. 190–209.

Szollosi, A., et al., People as intuitive scientists: reconsidering statistical explanations of decision making. *Trends in Cognitive Science* 24(12), 2020, pp. 1008–1018.

Talbot, B., *The End of epistemology as we know it*, New York 2024.

Thorstad, D., *Inquiry under bounds*, Oxford 2024a

Thorstad, D., Inquiry and the epistemic, *Philosophical Studies* 178 (9), 2021, pp. 2913–2928.

Thorstad, D., Why bounded rationality (in epistemology)?, *Philosophy and Phenomenological Research* 108 (2), 2024b, pp. 396–413.

Thorstad, D., The zetetic turn and the procedural turn, *Journal of Philosophy*, forthcoming.

Tomat, N., Bridging the gap betw the normative and the descriptive: Bounded epistemic rationality, *INDECS* 22 (1), 2024a, pp. 107–121.

Tomat, N., Cilji omejene epistemske racionalnosti, *Analiza* 28 (2), 2024b, pp. 263–284.

Viale, R., *Routledge handbook of bounded rationality*, New York 2020

Willard-Kyle, C., The knowledge norm for inquiry, *Journal of Philosophy* 120 (11), 2023, pp. 615–640.

Bounded Epistemic Rationality and Norms of Inquiry

Keywords: epistemology, norms of inquiry, epistemic rationality, bounded rationality, ecological rationality, non-ideal rationality

Epistemology has recently taken a zetetic turn: a shift in emphasis from the study of norms of belief to norms of inquiry. In this paper, I argue that a non-ideal approach to zetetic normativity can provide norms of inquiry for bounded, psychologically realistic agents. I introduce the concepts of non-ideal epistemology and bounded epistemic rationality and provide a brief overview of inquiry epistemology. I continue with proposing some features that should be included in the norms of inquiry for bounded agents. I argue that such norms should recognize our cognitive, environmental, and practical limitations. They should direct bounded agents to inquire about relevant problems, and they should be satisficing in nature, meaning that they allow for good enough reaching of epistemic goals. I argue that rational inquiry is not determined by *a priori* normative criteria, but is defined as the fit of the strategy with the environment. Inquiry epistemology for bounded agents should provide concrete, fine-grained zetetic norms that are achievable for ordinary inquirers and can serve as epistemic

guidance. The study of norms of inquiry for bounded agents explores both normative and descriptive parts of inquiry, is ameliorative, partly empirical, and requires collaboration between philosophy and empirical disciplines such as psychology and cognitive science.

Omejena epistemska racionalnost in norme raziskovanja

Ključne besede: epistemologija, norme raziskovanja, epistemska racionalnost, omejena racionalnost, ekološka racionalnost, neidealna racionalnost

V epistemologiji se je v zadnjih letih zgodil zetetični obrat: preusmeritev pozornosti s preučevanja norm prepričanja proti preučevanju norm raziskovanja. V članku zagovarjam, da lahko z neidealnim pristopom k obravnavi zetetične normativnosti oblikujemo norme raziskovanja za psihološko realistične agente z omejenimi spoznavnimi kapacitetami. V prvem delu članka predstavim neidealno epistemologijo, koncept omejene epistemske racionalnosti in epistemologijo raziskovanja. V nadaljevanju predlagam, kako naj bi oblikovali norme raziskovanja za omejene agente. Zagovarjam, da morajo takšne norme upoštevati naše spoznavne, okoljske in praktične omejitve. Agente usmerjajo k zanje relevantnim problemom in ne zahtevajo optimizacije, temveč le dovolj dobre rešitve – dopuščajo torej dovolj dobro doseganje epistemskih ciljev. To, kaj je racionalno raziskovanje, ni določeno na podlagi apriornih normativnih kriterijev, temveč prek ujemanja med spoznavno strategijo in okoljem, kar pomeni, da je normativnost racionalnosti ekološka. Epistemologija raziskovanja za omejene agente mora ponuditi konkretne, specifične norme, ki so dosegljive za običajne spoznavalce in ki se lahko uporabljajo kot epistemsko vodilo. Preučevanje norm raziskovanja za omejene agente je ameliorativen in delno empiričen projekt, ki kliče po sodelovanju med filozofijo in empiričnimi disciplinami, kot je psihologija.

About the author

Nastja Tomat holds a master's degree in psychology and is working as a Young Researcher at the Department of Philosophy, Faculty of Arts, University of Ljubljana, Slovenia. Her primary research focus is epistemology. In her dissertation, she investigates epistemic rationality and epistemic normativity, the relationship between normative and descriptive aspects of rationality, and the connection between philosophy and empirical science.

Email: Nastja.tomat@ff.uni-lj.si

O avtorici

Nastja Tomat je magistrica psihologije in mlada raziskovalka na Oddelku za filozofijo Filozofske fakultete Univerze v Ljubljani. Njeno primarno raziskovalno področje je epistemologija. V svoji disertaciji raziskuje epistemsko racionalnost in epistemsko normativnost, odnos med normativnimi in deskriptivnimi vidiki racionalnosti ter povezavo med filozofijo in empirično znanostjo.

E-naslov: Nastja.tomat@ff.uni-lj.si