

# Ecological rationality without externalism

Forthcoming in *Philosophy of Science*

Penultimate draft, please cite published version

## Abstract

Theories of bounded rationality join process reliabilists in holding that rationality is ecological, or environment-relative. Most theories of ecological rationality, like most versions of reliabilism, have been externalist. In this paper, I develop a de-externalized account of ecological rationality. I show how the account retains many advantages of externalist accounts while avoiding key challenges. I conclude with an application to the psychology of poverty, focusing on the rationality of agents caught in poverty traps.

## 1 Introduction

Humans are bounded agents. Some of our bounds are internal. We have limited cognitive abilities and pay costs to exercise them. Other bounds are external. We find ourselves in environments not of our choosing, which structure the strategies available to us and the consequences of using them. Theories of bounded rationality have held on this basis that rationality is *ecological* or environment-relative.<sup>1</sup>

We will see in Section 4 that most theories of ecological rationality have been externalist. That is, they instruct agents to use strategies which in fact perform well in their actual environment, regardless of what their evidence suggests about the environment they occupy or the performance of strategies within this environment.

Although externalist approaches enjoy important advantages, we will see that they also face a range of challenges. As a result, this paper aims to develop a de-externalized theory of bounded rationality and to assess its performance against candidate externalist views.

Here is the plan. Section 2 works toward a precise statement of ecological rationality. Section 3 motivates ecological rationality. Section 4 motivates externalist approaches to ecological rationality. Section 5 raises three challenges for externalist approaches to ecological rationality. Section 6 develops an alternative information-sensitive approach. Section 7 argues that the information-sensitive account avoids the challenges facing externalist accounts. Section 8 argues that the information-sensitive account captures many

---

<sup>1</sup>See Morton (2017); Schmidt (2019) and Thorstad (2024) for philosophical perspectives on ecological rationality. See Gigerenzer and Selten (2001); Simon (1956) and Todd and Gigerenzer (2012) for scientific perspectives.

of the motivations for externalist accounts. Section 9 considers an application to poverty traps. Section 10 concludes.

## 2 Ecological rationality

### 2.1 Formulating ecological rationality

On a traditional non-ecological approach to rationality, rational behavior is relative to agents, acts and times. Rational permissions are determined by some normative standard  $\Psi$ , and rational requirements are unique rational permissions.<sup>2</sup>

**(Non-Ecological Rational Permission)** For all agents  $S$ , acts  $X$  and times  $t$ ,  $S$  is rationally permitted to  $X$  at  $t$  iff  $\Psi(S, X, t)$ .

**(Non-Ecological Rational Requirement)** For all agents  $S$ , acts  $X$  and times  $t$ ,  $S$  is rationally required to  $X$  at  $t$  iff  $X$  is the unique permissible act for  $S$  at  $t$ .

Specifying a normative standard  $\Psi$  yields more specific accounts of rationality. For example, on a maximizing, value-first approach we would have:

**(Non-Ecological Maximizing Value-First Account)** For all agents  $S$ , acts  $X$  and times  $t$ ,  $S$  is rationally permitted (required) to  $X$  at  $t$  iff there is no alternative  $X'$  to  $S$ 's  $X$ -ing at  $t$  with value greater than (greater than or equal to) the value of  $S$ 's  $X$ -ing at  $t$ .

Readers who prefer other approaches are invited to replace  $\Psi$  with their preferred normative standard throughout this paper.<sup>3</sup>

To say that rationality is ecological is to say that rational behavior also depends on the structure of an agent's environment  $\mathcal{E}$ . That is:

**(Ecological Rational Permission)** For all agents  $S$ , acts  $X$ , environments  $\mathcal{E}$  and times  $t$ ,  $S$  is rationally permitted to  $X$  at  $t$  in  $\mathcal{E}$  iff  $\Psi(S, X, \mathcal{E}, t)$ .

**(Ecological Rational Requirement)** For all agents  $S$ , acts  $X$ , environments  $\mathcal{E}$  and times  $t$ ,  $S$  is rationally required to  $X$  at  $t$  in  $\mathcal{E}$  iff  $X$  is the unique permissible act for  $S$  at  $t$  in  $\mathcal{E}$ .

For example, on an ecological, maximizing, value-first approach we would have:

---

<sup>2</sup>A referee notes that ecological rationality theorists are concerned with the evaluation of diachronic processes of heuristic deliberation (Simon 1976). Procedural rationality can be incorporated into the criteria in this section by evaluating the components of search, stopping and decisionmaking required at each time (Gigerenzer and Selten 2001b), or by modifying the criteria to incorporate diachronic processes as the fundamental object of rational evaluation.

<sup>3</sup>Some readers may be concerned that the rationality criteria in this section are too heavily decision-theoretic. One remedy might be replacing the illustrative value-first views with other views that do not explicitly reference the value of states. Further reformulations are possible, such as replacing states with other entities throughout this section.

**(Ecological Maximizing Value-First Account)** For all agents  $S$ , acts  $X$ , environments  $\mathcal{E}$  and times  $t$ ,  $S$  is rationally permitted (required) to  $X$  at  $t$  in  $\mathcal{E}$  iff there is no alternative  $X'$  to  $S$ 's  $X$ -ing at  $t$  in  $\mathcal{E}$  with value greater than (greater than or equal to) the value of  $S$ 's  $X$ -ing at  $t$  in  $\mathcal{E}$ .

As before, varying the normative standard  $\Psi$  yields other ecological accounts.

The Ecological Maximizing Value-First Account raises three questions.

**(Ontological Question)** What sort of thing is an environment?

**(Identification Question)** Which environment does any given agent occupy?

**(Axiological Question)** How is the value of acts in an environment determined?

The Ontological and Identification Questions recur for all normative standards  $\Psi$ . As we move away from the value-first account, the Axiological Question gives way to cognate questions, such as the question of how duties are determined in an environment. For concreteness, I focus on the questions raised by the Ecological Maximizing Value-First Account.

There are at least two natural ways to answer these questions: the State View (Section 2.2) and the Class View (Section 2.3).<sup>4</sup>

## 2.2 State view

One natural approach is to treat environments as world states. On an externalist reading, this State View holds:

**(Ontological Question - State View)** Environments are world states.

**(Identification Question - State View)** An agent's environment is the current state of the world she occupies.

**(Axiological Question - State View)** The value of an action is the value of the world state that results.

Letting  $X(w)$  be the world state resulting from act  $X$  in world  $w$  and letting  $V$  be a value assignment to states, the Ecological Maximizing Value-First Account then becomes:

**(Ecological Maximizing Value-First Account - State View)** For all agents  $S$ , acts  $X$ , worlds  $w$  and times  $t$ ,  $S$  is rationally permitted (required) to  $X$  at  $t$  in  $w$  iff there is no alternative  $X'$  to  $S$ 's  $X$ -ing at  $t$  in  $\mathcal{E}$  with  $V(X'(w)) > (\geq) V(X(w))$ .

---

<sup>4</sup>In offering these views, my intention is not to gloss the existing practice of ecological rationality theorizing but rather to cover a range of plausible externalist views. This increases the robustness of my argument against externalism in Section 5 and allows me to speak precisely in Section 5 about the various patterns of verdicts that externalists might endorse. The views are chosen to represent externalist positions which philosophical readers might endorse.

This is a traditional version of externalism familiar to many philosophers.

Some theorists may not be satisfied with the State View. The State View reduces environments to world states, rather than introducing a new unit of analysis. And we will see in Section 4 that many externalists would like to make room for uncertainty and anti-luck requirements. However, we will see in Section 5 that the State View struggles to do this.

## 2.3 Class view

An alternative familiar from the history of process reliabilism is to treat environments as sets of world states (Alston 1995; Goldman 1986; Weatherson 2024).

**(Ontological Question - Class View)** Environments are sets of world states.

On the identification question, the Class View has at least two natural accounts available. On an *Indistinguishability View*, environments are states which the agent cannot distinguish from reality.

**(Identification Question - Indistinguishability View)** An agent's environment consists of the set of world states that she cannot distinguish from the current state.

The Indistinguishability View could be further elucidated by specifying the relevant notion of indistinguishability. For example, we might say that an agent's environment consists of the set of world-states that she does not know she is not in.

For a less epistemicized account, we could appeal to the similarity of environments rather than an agent's ability to distinguish them.

**(Identification Question - Relevant Similarity View)** An agent's environment consists of the set of world states relevantly similar to her own.

Again, the Relevant Similarity View could be further elaborated by specifying the notion of relevant similarity. For example, we might say that an agent's environment consists of the set of states that are sufficiently similar along specified relevant dimensions, or which could be produced by an underlying causal mechanism.

On the axiological question, any number of decision-theoretic criteria might be applied. However, a natural answer closely tied to the experimental methodology used by ecological rationality theorists would be to value acts by their average performance in an environment.

**(Axiological Question - Class View)** For all agents  $S$ , acts  $X$ , environments  $\mathcal{E}$  and times  $t$ , the value of  $S$ 's  $X$ -ing at  $t$  is  $\sum_{w \in \mathcal{E}} \frac{V(X(w))}{\|\mathcal{E}\|}$ .

On the Class View, the Ecological Maximizing Value-First Account becomes:

**(Ecological Maximizing Value-First Account - Class View)** For all agents  $S$ , acts  $X$ , environments  $\mathcal{E}$  and times  $t$ ,  $S$  is rationally permitted (required) to  $X$  at  $t$  in  $\mathcal{E}$  iff there is no alternative  $X'$  to  $S$ 's  $X$ -ing at  $t$  in  $\mathcal{E}$  with  $\sum_{w \in \mathcal{E}} \frac{V(X'(w))}{\|\mathcal{E}\|} > (\geq) \sum_{w \in \mathcal{E}} \frac{V(X(w))}{\|\mathcal{E}\|}$ .

We will consider an alternative probability-weighted version of the Class View in Section 6. However, we will see that this account is less naturally classified as an externalist view, and that it threatens to collapse into the State View.

So far, we have introduced an externalist account of ecological rationality (Section 2.1) and developed two versions of that account (Sections 2.2-2.3). The next order of business is to understand why theories of bounded rationality have typically been ecological (Section 3) and externalist (Section 4).

## 3 Why ecological rationality?

Why have bounded rationality theorists adopted ecological accounts of rationality? At least three arguments suggest themselves.

### 3.1 Performance varies with environments

Process reliabilists discovered that the reliability of belief-forming methods varies with environments (Goldman 1979). Processes which reliably produce true beliefs in one environment may not reliably produce true beliefs in another. For this reason, process reliabilists hold that belief-forming processes are not rational or irrational in themselves. Rather, they are rational in environments where they perform well and irrational elsewhere.

Beyond reliability, bounded rationality theorists are often concerned with many other features of belief-forming processes including speed (Heitz 2014), coherence (Staffel 2020; Thorstad 2025), effort (Johnson and Payne 1985), bias (Gigerenzer and Brighton 2009) and variance (Geman et al. 1992). The performance of belief-forming processes may vary along all of these dimensions across environments. For this reason, bounded rationality theorists hold that belief-forming processes are not rational or irrational in themselves, but only in an environment.

### 3.2 Anti-luck

Performance in any given case may be a matter of luck. For example, suppose that the answer to some question  $Q$  is  $p$ . Then the following belief-forming process is a reliable way of answering  $Q$ : immediately conclude deliberation with the conclusion that  $p$ .

For this reason, process reliabilists learned to focus not on the reliability of a process in a given case, which may be due to luck, but rather on the reliability of the process across an environment of similar problems. The method of immediately concluding that  $p$  may be a reliable way to answer  $Q$ , but it is not a reliable way to answer other questions in the environment. An ecological conception of rationality therefore correctly classifies this method as irrational because its performance in a single case was a mere matter of luck.

In the same way, bounded rationality theorists have stressed that complex belief-forming processes may succeed in a given problem by overfitting models to available data (Gigerenzer and Brighton 2009; Gigerenzer and Sturm 2012; Wheeler 2020). Many of the trends captured by overfitted models are not genuine regularities, but rather accidental features of a given problem. This means that complex, overfitted models may perform

poorly across an environment of similar problems as they continue to rely on spurious regularities to make predictions. In such cases, it is natural to say that the processes in question were not rational, since their performance on a given problem was due to luck. For this reason, bounded rationality theorists join process reliabilists in holding that rationality must be ecological, so that lucky applications of overfitted models will not be classified as rational (Rich 2016).

### 3.3 Vindictory epistemology

Bounded rationality theorists have a vindictory project. In the 1950s and 1960s, neoclassical economists held that human cognition is highly rational. In the 1970s and 1980s, the heuristics and biases program found that humans often fall short of neoclassical rationality standards, concluding on this basis that humans may not be as rational as previously supposed (Kahneman et al. 1982; Tversky and Kahneman 1974). Since the 1990s, the project of *vindictory epistemology* has aimed to vindicate the rationality of many seemingly irrational cognitions by showing how apparent irrationalities result from rational responses to normatively relevant factors that are not taken into account by neoclassical economic models (Bishop and Trout 2004; Icard ms; Thorstad 2024).

Some neglected normative factors are features of agents, such as limited cognitive abilities and computational costs (Icard ms). But other neglected factors are features of environments. For example, many experiments in the heuristics and biases program have the following structure.

**(Heuristics and Biases Case)** Heuristic  $H$  is known to perform well in many environments, but can be shown from the armchair to perform badly in environment  $\mathcal{E}$ . Agents are dragged into a laboratory and induced to use  $H$  in  $\mathcal{E}$ .

On an ecological picture, Heuristics and Biases Case does not cast any doubt on the rationality of  $H$ . After all, theories of ecological rationality instruct agents to use  $H$  in environments where  $H$  performs well, and those environments do not include  $\mathcal{E}$ . Furthermore, it is agreed by all parties that agents tend to use many heuristics primarily in environments to which they are well suited.<sup>5</sup> On an ecological approach, those findings imply that most uses of  $H$  are rational.<sup>6</sup>

Summing up, an ecological approach to rationality captures the insights that strategy performance varies with environments, that performance in a single case may be due to luck, and that vindicating the rationality of seemingly irrational cognitions often requires appeal to the structure of the agent's environment. In the next section, I ask why theories of ecological rationality are typically externalist.

---

<sup>5</sup>See (Gilovich and Griffin 2002, pp. 3-4) and Chapter 2 of Thorstad (2024) for discussion.

<sup>6</sup>We might be forced to conclude that isolated experimental uses of  $H$  are irrational, although there is plenty of room to push back. But even then, we would be able to vindicate a substantial part of human cognition by largely isolating conclusions about laboratory studies from conclusions about rationality in the wild.

## 4 Why externalism?

Many traditional accounts of ecological rationality are externalist. Externalist accounts make the rationality of strategies depend on the actual structure of the environment and the actual performance of strategies within the environment, rather than on the agent's information or beliefs about environments and strategy performance. For example:<sup>7</sup>

Ecological rationality is the study of under what conditions cognitive strategies are successful. (Gerd Gigerenzer, cited in (Mousavi and Kheirandish 2014, p. 1782))

An agent  $A$  should deliberate using those norms  $N$  that allow her to reliably achieve her ends  $E$ , given her cognitive capacities, in those contexts  $C$  in which she regularly finds herself. (Morton 2017, p. 554)

*Ecological rationality*: a person's decision is procedurally rational in an environment to the extent that, given her particular psychological makeup, the decision-making procedures she uses allow her to reliably achieve her ends in this type of environment. (Schmidt 2019, p. 521)

Here it is that actual fact of success or reliability in an environment that determines a strategy's ecological rationality, whether or not this fact is believed by an agent or supported by her total evidence.

It is not a coincidence that traditional accounts of ecological rationality tend to be externalist. There are at least three reasons why many have thought that ecological rationality should be given an externalist articulation.

### 4.1 Evolution

How did humans come to possess strategies that are well adapted to the problems that we face, as well as the ability to select appropriate strategies in many environments? Ecological rationality theorists hold that humans have access to an *adaptive toolbox* of heuristics (Gigerenzer and Selten 2001). Natural selection shaped the toolbox to contain strategies that are well adapted to our evolutionary environment and provided us with strategy selection mechanisms that tend to select high-performing strategies in our evolutionary environment. Moreover, through a bidirectional process of *niche selection* humans selected and shaped environments which co-evolved to be better-suited to the toolbox of strategies available to us (Marewski and Schooler 2011). The twin processes of natural selection and niche selection are taken to explain why human cognition was well suited to our evolutionary environment. To the extent that our current environment resembles our evolutionary environment, natural selection and niche selection are also proposed as explanations of why human cognition tends to perform well in our current environment.

---

<sup>7</sup>Note that although ecological rationality theorists often give externalist statements of ecological rationality, much of their work is friendly to de-externalized approaches. For example, ecological rationality theorists advance computational models of how agents learn to select heuristics in environments where they are likely to perform well (Rieskamp and Otto 2006). Such computational models have, if anything, more rationalizing force on the information-sensitive view adopted in this paper.

Ecological rationality theorists seek a theory of rationality which can explain these data. That is:

**(Evolutionary Explanation Desideratum)** A good theory of ecological rationality should capture leading evolutionary explanations of human cognitive success.

Externalists meet the Evolutionary Explanation Desideratum by telling exactly the story given above, without modification. However, many internalists cannot tell the same story. After all, until recently humans did not know about natural selection and niche selection. Even now, some people do not know these things, and precious few have considered and appreciated the support that natural selection and niche selection provide for the rationality of a good part of human cognition. If it turns out that on an internalist story, evolutionary rationalizing explanations work only for a select few individuals, then externalism may perform best on the Evolutionary Explanation Desideratum.

## 4.2 Representational limitations

Humans have limited capacities to form and respond to representations of the environment. Theories of ecological rationality aim to show how such agents may nonetheless cognize rationally. That is:

**(Representational Limitation Desideratum)** A good theory of ecological rationality should allow representationally-limited agents to cognize rationally.

In many cases, externalists seem better posited to capture the Representational Limitation Desideratum.

For example, baseball players using the *gaze heuristic* position themselves to catch a fly ball by running until the ball appears to move horizontally towards them (McLeod and Dienes 1996). The gaze heuristic is a reliable method given a number of complex facts about gravity, optics and mechanics which make the optical trajectory of a ball a reliable guide to its landing spot. However, although agents using the gaze heuristic represent the ball's position visual trajectory, they needn't represent the facts about gravity, optics and mechanics which make the ball's visual trajectory a reliable guide to landing spots.

Externalists will be untroubled by these facts. On their view, rationality requires responding appropriately to features of the environment, but does not require constructing and responding to internal representations of the environment.<sup>8</sup> Agents are not to be faulted for representational errors or omissions so long as they continue to make accurate judgments based on their mental models.

By contrast, internalists struggle to meet the Representational Limitation Desideratum in cases such as the gaze heuristic. On many internalist views, agents are not justified in forming intentions and beliefs based on a ball's optical trajectory unless they represent the facts inferentially linking the ball's trajectory to its landing zone. And it is not clear that representationally limited agents do, or should do this.

---

<sup>8</sup>This, they will hasten to add, dovetails nicely with a more general insistence within the bounded tradition that rationality is a matter of performance, not internal cognitive structure (Douven 2020; Gigerenzer 2019). On this view, internal features of an agent's cognition are valuable only to the degree that they conduce to good performance.

### 4.3 Explaining reliabilist disanalogies

We saw in Section 3 that there are deep analogies between ecological rationality and process reliabilism (Dusi 2024). Where process reliabilists instruct agents to use strategies that are reliable within an environment, ecological rationality theorists instruct agents to use strategies that balance reliability against other desiderata such as speed and effort within an environment. Ecological rationality, like process reliabilism, is motivated by the insights that performance varies across environments and that strategy performance must be assessed across an environment of problems to distinguish rationality from luck.

Given these analogies, we should be struck by the fact that most traditional articulations of reliabilism have been externalist.<sup>9</sup> If theories of ecological rationality depart from externalism, then they should explain what has changed to warrant their departure. That is:

**(Reliabilist Disanalogy Desideratum)** A good theory of ecological rationality should identify relevant disanalogies to explain divergences from process reliabilism.

In particular, non-externalist theories of ecological rationality should explain why it makes sense for theories of ecological rationality to break with reliabilists on the issue of externalism.

Together, the appeal to evolution (Section 4.1), representational limitations (Section 4.2) and reliabilist analogies (Section 4.3) go some way towards explaining why many traditional accounts of ecological rationality have been externalist. Nevertheless, we will see in the next section that externalist accounts of ecological rationality face at least three challenges.

## 5 Challenges

This section raises three challenges for externalist accounts of ecological rationality. We will see that externalist theories yield a number of verdicts that many ecological rationality theorists may not wish to accept. This will motivate the construction of an alternative approach which reverses the verdicts in question.

### 5.1 Information-sensitivity

Humans are informationally bounded agents. We have limited information about the world. Bounded rationality theorists are keen to stress the importance of informational bounds. For example, they hold that theories of rational judgment and decisionmaking should incorporate a detailed account of how agents should search for information before making a decision (Gigerenzer and Selten 2001b), and they frequently criticize neoclassical economic models when those models assume that agents possess perfect information about the world (Gigerenzer and Brighton 2009; Gigerenzer 2019). We might encapsulate this demand in a theoretical desideratum:

---

<sup>9</sup>Many versions of coherentist reliabilism (Goldberg 2012), evidentialist reliabilism (Comesaña 2010) and perhaps also virtue reliabilism (Greco 2010; Sosa 2015) soften some traditional externalist commitments.

**(Informational Limitation Desideratum)** A good theory of ecological rationality should account for the normative force of informational limitations.

Externalists recover some aspects of the Informational Limitation Desideratum, but not others.

Externalist accounts of ecological rationality can readily explain why it is important for rational processes of judgment and decisionmaking to gather and process information (Gigerenzer and Selten 2001b).<sup>10</sup> However, externalists struggle to capture the relevance of limited information to the choice between processes whose results are unknown.

On the State View, the value of a process is the value of its actual consequences in the actual world. If I believe, on the basis of excellent evidence, that a process will produce an accurate judgment or a good decision, but in fact the process will perform badly, then I ought to use a different process. And if I believe, on the basis of excellent evidence, that a process will produce an inaccurate judgment or a bad decision, but in fact the process will perform well, then I ought to use the process nonetheless. In this way, the State View does not make an agent's limited information about processes relevant to the rationality of processing decisions.

Some versions of the Class View give a more expansive role to informational bounds by letting informational bounds bear on the identity of an agent's environment. For example, the Indistinguishability View treats environments as worlds which the agent cannot distinguish from her own. This can be given an epistemic reading, on which for example environments are worlds which the agent does not know not to obtain. However, this version of the Indistinguishability View still limits the normative relevance of informational bounds in two ways. First, while information bears on the identification of an environment, it does not bear on the assessment of strategies within the environment. And second, information is relevant only if it constitutes knowledge. If an agent does not know which of two states the world is in, but has a moderate amount of evidence favoring one state over the other, that evidence will not be reflected in the identification of environments.<sup>11</sup>

Other versions of the Class View give a less expansive role to informational bounds. For example, the Relevant Similarity View treats environments as worlds which are relevantly similar to the agent's own. Here it is not evidence of similarity, but actual similarity that counts. If an agent has excellent evidence that it is warm outside but in fact it has suddenly become cold, then it is coldness rather than warmth which makes a world relevantly similar to the agent's own world.

## 5.2 Verdicts in novel environments

Consider:

---

<sup>10</sup>Under many conditions, gathering and processing information cannot decrease the expected value of an agent's decisions (Blackwell 1953; Good 1966) or the expected accuracy of her beliefs (Maher 1990), and it often increases both quantities.

<sup>11</sup>Externalists could perhaps block this second complaint by adopting the Williamsonian identification of knowledge and evidence (Williamson 2000). But readers with sufficient externalist sympathies to identify knowledge and evidence are unlikely to be among those voicing any of the challenges in this section. Even then, the problem would remain for agents who know enough to make one state more likely to obtain, but not enough to know that the other state will not obtain.

**(Normal Foraging)** You are foraging in the forest for food, as you have done in the same forest for many years. You see wolf tracks to the west, which long experience has taught you to treat as a sign that predators lie in the direction of the tracks. You see rotten fruit to the east, which long experience has taught you to treat as a sign of abundant and potentially ripe fruit. On this basis, you head east in the direction of the fruit.

Many theorists would like to hold:

**(Normal Foraging Datum)** You behave rationally by heading east in Normal Foraging.

After all, you have excellent evidence in Normal Foraging that there will be fruit to the east and wolves to the west.

Now consider:

**(Abnormal Foraging)** While you slept last night, you were transported by a camera crew into a different forest, with your new sleeping place carefully constructed to resemble the old one. However, there are no wolves in this environment, as dead contestants are bad for ratings. Instead, all wolf tracks have been constructed by staff to point in the direction of fruit, and all rotten fruit has been littered to point in the direction of empty fields.

Many theorists would still like to hold:

**(Abnormal Foraging Datum)** You behave rationally by heading east in Abnormal Foraging.

After all, you see rotten fruit to the east, which long experience has led you to associate with food. And insofar as we are concerned to disentangle rationality from luck (Section 3.2), we should worry that heading west without evidence would make you very lucky to be successful.<sup>12</sup>

Externalists struggle to deliver the Normal and Abnormal Foraging Data together. On the State View, the Normal Foraging Datum holds so long as going east will in fact yield fruit. But the Abnormal Foraging Datum fails, since going east will lead you to an empty field.

Some versions of the Class View deliver the same verdicts as the State View does. For example, on the Relevant Similarity View, the environment consists of situations relevantly similar to your own. Rotten fruit reliably correlates with food in environments similar to the normal forest, giving the Normal Foraging Datum. But rotten fruit does not reliably correlate with food in environments similar to the abnormal forest, falsifying the Abnormal Foraging Datum.

Other versions of the Class View give different verdicts. For example, on the Indistinguishability View, it is rational to behave differently only if you are able to distinguish

---

<sup>12</sup>Many ecological rationality theorists will not be moved by the need to recover the Normal and Abnormal Rationality Data, though they will be moved by the Heuristics and Biases Rationality Datum. These theorists can productively read the discussion here as a set of cases aimed at revealing what is at issue in Heuristics and Biases Case and why we should accept the Heuristics and Biases Rationality Datum.

the new environment from the old. This may not be the case. But this does not give us what we wanted, for now the Indistinguishability View evaluates behavior even in the normal forest by considering performance across a wide range of indistinguishable forests including the abnormal forest. Across this broad environment, the strength of the correlation between tracks and predators is greatly diminished. As a result, both the Normal Foraging Datum and the Abnormal Foraging Datum are likely to fail.<sup>13</sup>

If this is right, then none of the externalist accounts from Section 2 uncontroversially deliver the Normal and Abnormal Foraging Data together. Readers who want to recover these data will need to look elsewhere.

### 5.3 Vindictory failures

Consider again:

**(Heuristics and Biases Case)** Heuristic  $H$  is known to perform well in many environments, but can be shown from the armchair to perform badly in environment  $\mathcal{E}$ . Agents are dragged into a laboratory and induced to use  $H$  in  $\mathcal{E}$ .

We saw in Section 3.3 that many ecological rationality theorists would like to hold:

**(Heuristics and Biases Rationality Datum)** Some agents in Heuristics and Biases Case cognize rationally.

For example, Heuristics and Biases Rationality Datum may be attractive for agents who are given misleading evidence that they are still in an environment where their strategies are reliable. However, externalists struggle to capture the Heuristics and Biases Rationality Datum without unwanted consequences.

The situation here largely parallels our discussion of the foraging data. As with our foraging agent, the State View and the Relevant Similarity View hold that your environment is now one in which  $H$  performs poorly, and hence it is irrational for you to use  $H$ . And as with our foraging agent, the Indistinguishability View can escape this consequence only by dramatically broadening the class of environments against which everyday behavior is assessed, so that everyday uses of  $H$  are assessed against a broad environment that includes the deliberately misleading laboratory experiments. As before, this will have the consequence of recommending reduced reliance on  $H$  not only in  $\mathcal{E}$ , but also in environments where  $H$  is often taken to be warranted.

## 6 Information-sensitivity

We saw in Section 5 that externalist accounts of ecological rationality face several challenges. This section develops an alternative account that avoids those challenges.

The key observation is that many of the arguments given for externalism in Section 4 are sharpest against traditional internalist views, on which rationality is relative to

---

<sup>13</sup>And if they do not, we can construct cases similar to Normal Foraging in which the Normal Foraging Datum fails.

an agent's beliefs about the world. It may then be difficult to explain how evolution rationalizes heuristic cognition (Section 4.1) and how representationally limited agents cognize rationally (Section 4.2).

It is increasingly recognized that there is a third option between belief- and fact-relativity. Namely, we can take an *information-sensitive* approach on which rationality is relative to the available evidence (Kolodny and MacFarlane 2010).<sup>14</sup> We will shortly see that an information-sensitive approach captures many motivations for externalist accounts (Section 8) while avoiding their challenges (Section 7).

We can generate information-sensitive analogs of the State and Class Views by retaining their answers to the ontological and identification questions while reformulating their axiological claims in an information-sensitive way.<sup>15</sup> For example, a modified State View continues to treat environments as world states and an agent's environment as the actual state of the world. However, the value of strategies in environments is no longer the value of their actual consequences, but rather their expected value given the evidential probabilities  $Pr_E$  generated by the agent's total evidence about her environment.

**(Axiological Question – Information-Sensitive State View)** For all agents  $S$ , acts  $X$ , environments  $\mathcal{E}$  and times  $t$ , the value of  $S$ 's  $X$ -ing at  $t$  in  $\mathcal{E}$  is  $\sum_w Pr_E(w)V(X(w))$ .

This is a traditional information-sensitive view.

A modified Class View continues to treat environments as sets of world states, and an agent's environment as the set of world states which the agent cannot distinguish from the current state (Indistinguishability View) or which are relevantly similar to the current state (Relevant Similarity View). A natural information-sensitive extension of the Class View's axiological commitments values acts not at their average performance within an environment, but rather by their probability-weighted performance within the environment. Then we recover:

**(Axiological Question – Information-Sensitive Class View)** For all agents  $S$ , acts  $X$ , environments  $\mathcal{E}$  and times  $t$ , the value of  $S$ 's  $X$ -ing at  $t$  in  $\mathcal{E}$  is  $\sum_{\mathcal{E}} Pr_E(\mathcal{E}) \sum_{w \in \mathcal{E}} Pr_E(w|\mathcal{E})V(X(w))$ .

By the law of total probability, this value is simply  $\sum_w Pr_E(w)V(X(w))$ , in agreement with the State View.

This answer to the Axiological Question brings the Class View perilously close to collapsing into the State View. Officially, the Class View comes apart from the State View by treating environments as sets of world states and giving a specific account of which set of world states constitutes an agent's environment. However, in evaluating the rationality of acts, these details will wash away and we will be left with the State View's injunction

---

<sup>14</sup>Information-sensitive approaches have gained popularity for their ability to explain normative intuitions in a variety of cases (Kolodny and MacFarlane 2010), to account for embedding behavior (Cariani 2016), and to build on the success of information-sensitive approaches to epistemic modality (Yalcin 2007).

<sup>15</sup>Some answers to the identification question may require modification. For example, not all information-sensitive theories will be friendly to the Relevant Similarity View, on which environments are composed of states relevantly similar to the current state. We might then modify the Relevant Similarity View to treat environments as states which the agent's evidence suggests are relevantly similar to the current state. I omit these modifications for the sake of brevity.

to maximize the expected value of world states given evidential probabilities. Neither the specifics of environment-individuation, nor even the bare insistence that environments be distinguished from world-states will make a difference at the level of axiology.

If we wanted to avoid this consequence, we might continue to value acts at their mean performance within the agent's expected environment, as in Section 2. That is:

**(Axiological Question – Information-Sensitive Class View with Environmental Means)** For all agents  $S$ , acts  $X$ , environments  $\mathcal{E}$  and times  $t$ , the value of  $S$ 's  $X$ -ing at  $t$  in  $\mathcal{E}$  is  $\Sigma_{\mathcal{E}} Pr_E(\mathcal{E}) \Sigma_{w \in \mathcal{E}} \frac{V(X(w))}{\|\mathcal{E}\|}$ .

This version of the Class View is information-sensitive across but not within environments.<sup>16</sup> For this reason, some advocates of information-sensitive views will not be comfortable with this answer to the axiological question, but it may be interesting to explore the consequences of this view in the discussion below.

In the next two sections, I argue that this information-sensitive reformulation of the State and Class Views avoids the challenges for externalism posed in Section 5 while capturing a good part of the motivations for externalism given in Section 4.

## 7 Addressing challenges

### 7.1 Informational bounds

Our first challenge to externalism drew on informational bounds (Section 5.1).

**(Informational Limitation Desideratum)** A good theory of ecological rationality should account for the normative force of informational limitations.

While externalists can hold that rational cognitive processes gather and assess information, they cannot make informational limitations bear directly on the normative status of processes themselves. If all evidence suggests that some process will perform well in an agent's environment but in fact the process will perform poorly, then on externalist views the agent should not use the process.

The information-sensitive view meets the Informational Limitation Desideratum by building evidential probabilities directly into the value of processes. If all evidence suggests that some process will perform well, then it enjoys high expected value on the agent's evidence and should therefore be used.

### 7.2 Verdicts in novel environments

Our second challenge to externalism was that it produces counter-intuitive verdicts in novel environments (Section 5.2). Externalism struggles to recover the Normal Foraging Datum that you behave rationally by heading east in normal foraging cases together with the Abnormal Foraging Datum that you behave rationally by heading east in abnormal foraging cases.

---

<sup>16</sup>That is, it makes evidential probabilities relevant to the selection of environments against which performance is assessed, but not to the scenarios within an environment against which performance is assessed.

All versions of the information-sensitive view avoid this consequence. In Normal Foraging, all evidence suggests that food lies east, so the evidentially-expected value of heading east is high. In Abnormal Foraging, your evidence has not changed, so neither has rational behavior.<sup>17</sup>

### 7.3 Vindictory failures

Our final challenge to externalism concerned:

**(Heuristics and Biases Case)** Heuristic  $H$  is known to perform well in many environments, but can be shown from the armchair to perform badly in environment  $\mathcal{E}$ . Agents are dragged into a laboratory and induced to use  $H$  in  $\mathcal{E}$ .

Vindictory epistemologists would like to hold:

**(Heuristics and Biases Rationality Datum)** Some agents in Heuristics and Biases Case cognize rationally.

However, we saw in Section 5.3 that externalists struggle to hold this.

All versions of the information-sensitive view recover the Heuristics and Biases Rationality Datum. For example, suppose that an agent is provided with misleading evidence that she is in an environment where  $H$  is reliable. Then all versions of the information-sensitive view rationalize the use of  $H$ . Although the agent is in fact in  $\mathcal{E}$ , where  $H$  is unreliable, this does not impugn the rationality of using  $H$  when the evidence suggests that the agent is not in  $\mathcal{E}$ .

So far, we have seen that an information-sensitive account avoids the challenges raised to externalist accounts in Section 5. Below, I argue that the account also captures many of the motivations given for externalist accounts in Section 4.

## 8 Capturing motivations

### 8.1 Evolution

A first motivation for externalism was capturing the relevance of evolutionary explanations such as natural selection and niche selection (Section 4.1).

**(Evolutionary Explanation Desideratum)** A good theory of ecological rationality should capture leading evolutionary explanations of human cognitive success.

We need to distinguish two readings of the Evolutionary Explanation Desideratum.

The first is a causal reading:

---

<sup>17</sup>Some externalists think that your evidence has changed in Abnormal Foraging (Williamson 2000). However, advocates of the information-sensitive view are unlikely to favor externalist accounts of evidence.

**(Evolutionary Explanation Desideratum - Causal Reading)** A good theory of ecological rationality should capture leading evolutionary explanations of how agents came to possess and be justified in using strategies which lead to cognitive success.

Information-sensitive views capture the Causal Reading without a hitch. They treat natural selection and niche selection as causal stories about how agents came to possess strategies that perform well in many of the environments that we encounter. By using these strategies successfully throughout our lives, humans accumulate evidence which justifies their use.

A second reading of the Evolutionary Explanation Desideratum is justificatory rather than causal:

**(Evolutionary Explanation Desideratum - Justificatory Reading)** A good theory of ecological rationality should allow evolutionary histories to directly rationalize heuristic strategies.

For example, theorists in the ecological rationality tradition sometimes appear to suggest that strategies are rational because they were shaped through evolution and niche selection to perform well in our environment.<sup>18</sup>

Information-sensitive accounts reject the Justificatory Reading of the Evolutionary Explanation Desideratum. In this, they join many philosophers (Fodor and Piattelli-Palmarini 2010) and some ecological rationality theorists (Boudry et al. 2015) in holding that evolution lacks direct normative force.

An advantage of rejecting the Justificatory Reading is that we can accommodate changes to the modern environment. Critics of ecological rationality often charge that heuristics which were well adapted to our evolutionary environments no longer perform well in modern urban environments (Tooby and Cosmides 1994; Stanovich and West 2000). Without exaggerating the prevalence of this phenomenon, we would like to agree with critics that heuristics which we have good evidence to suspect will not perform well in our modern environment are no longer rational. An information-sensitive view directly delivers this verdict by rejecting the Justificatory Reading of the Evolutionary Explanation Desideratum.

## 8.2 Representational limitations

A second motivation for externalism was capturing:

**(Representational Limitation Desideratum)** A good theory of ecological rationality should allow representationally-limited agents to cognize rationally.

We saw in Section 4.2 that belief-relative accounts struggle to capture the Representational Limitation Desideratum. For example, agents following the gaze heuristic may count as irrational unless they represent complex facts about optics and mechanics which rationalize their beliefs. But if rationality is evidence-relative, no such consequence follows. Agents have evidence about the performance of the gaze heuristic. Agents respond rationally to their evidence by continuing to employ the gaze heuristic, regardless of their knowledge of optics.

---

<sup>18</sup>See Boudry et al. (2015) for discussion.

### 8.3 Explaining reliabilist disanalogies

A final motivation for externalism was accommodating:

**(Reliabilist Disanalogy Desideratum)** A good theory of ecological rationality should identify relevant disanalogies to explain divergences from process reliabilism.

In particular, we need to explain divergences from externalism. Externalist theories of ecological rationality have an advantage here, because there are no divergences from externalism to explain.

The most natural explanation of why it might make sense for process reliabilists, but not ecological rationality theorists to be externalists is that externalism has been claimed to enjoy a number of philosophical advantages that largely cross-cut discussions of bounded rationality. These motivations and the ensuing debates can be largely settled off-stage.

Process reliabilists can quickly avoid many forms of external world skepticism. For example, I may be justified in believing that I have two hands because I can see them, and my vision is reliable under normal conditions (Sosa 1994). More generally, reliabilism avoids many common struggles in accounting for non-inferential justification, including but not limited to the justification of perceptual beliefs (Goldman 1986). Process reliabilism allows cognitively unsophisticated agents to hold justified beliefs, even if they lack reflective access to those beliefs, so long as the beliefs were reliably formed (Bergmann 2006). Reliabilism incorporates the view that causal histories matter to justification (Goldman 1979). And reliabilism is sometimes held to provide a naturalistic epistemology tied to the truth-aim of belief (Goldman 1994).

Most of these justifications are largely orthogonal to considerations of bounded rationality.<sup>19</sup> If that is right, then the fact that most (though not all) reliabilists have been externalists is not a decisive reason for theories of ecological rationality to be externalist, despite strong analogies between reliabilism and ecological rationality. Most of what drives reliabilists towards externalism turns on philosophical considerations that are primarily orthogonal to debates about bounded rationality.

## 9 An application: Poverty traps

One way to see what is at stake in debates about ecological rationality and why an information-sensitive view might be attractive is to think about concrete applications of ecological rationality to our practical lives. This section considers a case study from the psychology of poverty, focusing on the rationality of agents caught in poverty traps. I show how a leading externalist account due to Jennifer Morton (2017) goes a long way towards rationalizing agents' behavior, and how existing behaviors can be rationalized by adopting an information-sensitive account of ecological rationality.

---

<sup>19</sup>Some, such as allowing unsophisticated agents to hold justified beliefs, will be relevant to bounded rationality. But an information-sensitive account enjoys this advantage too: it instructs cognitively limited agents to believe what their evidence supports, but does not require them to do so on the basis of complex internal representations or metacognitive reasoning processes. The information-sensitive account can, for example, agree that everyday social inferences made at the water cooler are justified if they are supported by, and based on the available evidence, even if most of us could not give a good account of how we make such inferences.

## 9.1 What must be explained

Here is an unfortunate fact. Material poverty correlates strongly with a constellation of short-termist patterns of thought and action. The poor tend to overborrow (Shah et al. 2012), undersave (Bernheim et al. 2015), and underinvest in education (Blanden and Gregg 2004). These behaviors contribute to poverty traps, perpetuating challenges such as debt and limited education that keep people locked in poverty.

At first glance, these findings appear to suggest that the poor are irrationally short-termist. Worse, they seem to suggest that many people remain poor because of irrational short-termism. This may in turn seem to imply that the poor are partly to blame for their plight. And at a policy level, it may seem to suggest that the poor are owed less recompense because their plight is partly due to their own irrationality.

Many scientists and policymakers have not wanted to draw these conclusions. How can they be avoided? Poverty is a complex and multifaceted phenomenon with no simple explanation. But at least two promising types of rationalizing explanations for the link between material poverty and short-termism have gained favor in recent years.<sup>20</sup>

## 9.2 Cognitive scarcity

First, material poverty leads to cognitive scarcity (Mani et al. 2013; Shah et al. 2012; Shafir and Mullainathan 2013). On the demand side, poverty creates large numbers of immediately pressing challenges with large short-term consequences, each of which requires considerable thought and attention to solve.<sup>21</sup> For the financially well-endowed, all of these problems are easily soluble, but for the poor, they are difficult problems requiring constant attention. On the supply side, poverty exacerbates problems such as sleep deprivation (Patel et al. 2010), stress (Cohen et al. 2006), malnutrition (Gailliot et al. 2007), noise pollution (Dean 2024), serious illness (Banerjee and Duflo 2007), and inadequate education (Blanden and Gregg 2004) which leave agents less equipped to devote cognitive resources to important problems.

Many theorists have suggested that a rational response to cognitive scarcity under material poverty involves preferentially allocating cognitive resources towards short-term problems and short-term consequences (Haushofer and Fehr 2014; Morton 2017; Thorstad forthcoming). In a slogan:

**(Cognitive Scarcity Explanation)** Many seemingly irrational cognitions under material poverty are rational responses to cognitive scarcity.

Before planning for college or attending an annual physical examination, it is a good idea to figure out how to keep your lights on. This may not be irrational, because any other pattern of cognitive response would be worse still, but it does predictably lead to many of the short-termist patterns of thought and action surveyed above. For example, the poor may be more likely to take payday loans at high interest rates because they are not always able to spare substantial thought and attention for the long-term consequences of payday loans (Shafir and Mullainathan 2013).

---

<sup>20</sup>For pushback see Haushofer and Salicath (2023) and de Bruijn and Antonides (2022).

<sup>21</sup>These include challenges such as deciding how to pay bills and which bills to pay, how to arrange child-care while working a full-time job, how to keep an adequate amount of nutritious food on the table, and how to arrange transportation to work and important meetings.

### 9.3 Changed prospects

A second cause of short-termism is that material poverty leads to meaningful changes in agents' prospects (Pepper and Nettle 2017). In a slogan:

**(Changed Prospect Explanation)** Many seemingly irrational cognitions under material poverty are rational responses to changed life prospects.

The poor face heightened morality rates (Fiscella and Franks 1997) and rates of serious illness and disability (Banerjee and Duflo 2007). It can make less sense to save for the future or invest in your future health if you expect this future to be short or of lower quality than the present (Pepper and Nettle 2014a,b). Poverty also leads to increased uncertainty about whether future gains will be realized and accompanying concerns about loss of realized gains through family needs (Jakiela and Ozier 2016), theft (Dupas and Robinson 2013) and other causes. Increased uncertainty about future gains is standardly taken to be an excellent reason for increased short-termism (Gollier 2002). Finally, poverty leads to diminished control over future prospects through reduced agency (Pepper and Nettle 2017) and limited access to credit markets (Santos and Barrett 2011). Seemingly short-termist borrowing behaviors may then sometimes result from a lack of viable alternatives (Bhutta et al. 2015), or from the use of borrowing to smooth consumption over periods of fluctuating income (Dobridge 2018).

### 9.4 Capturing the rationalizing explanations

The Cognitive Scarcity and Changed Prospect Explanations are available to theorists of many stripes. For example, we saw in Section 4 that Jennifer Morton adopts an externalist account of ecological rationality:

An agent  $A$  should deliberate using those norms  $N$  that allow her to reliably achieve her ends  $E$ , given her cognitive capacities, in those contexts  $C$  in which she regularly finds herself. (Morton 2017, p. 554)

Morton uses this externalist account of ecological rationality to recover the Cognitive Scarcity Explanation. Under conditions of material poverty and the accompanying cognitive scarcity, Morton argues, agents reliably achieve their ends through cognitive short-termism. Similar arguments could be made to capture the Changed Prospect Explanation.

However, agents often retain short-termist patterns of thought and action learned during periods of poverty for some time after their situations have improved (Cohen et al. 2010; Griskevicius et al. 2013; Grohmann et al. 2015). For example, faced with windfall gains, they often spend large portions of the windfall rather than saving or investing the windfall in ways that would cement their newfound economic status (Altındağ and O'Connell 2023; Hankins et al. 2011; Karlan et al. 2019). By way of illustration, Dean Karlan and colleagues (2019) studied market vendors in Chennai, India, who financed their daily inventory through single-day loans from a moneylender at a nominalized annual interest rate exceeding 1,700%. Karlan and colleagues gave the vendors a lump sum of money sufficient to pay their daily debts. Months later, the vendors were almost as indebted as before.

What we would like to say is that it can be rational for agents to take time to adapt to their changed life circumstances.<sup>22</sup> In a slogan:

**(Rational Short-Termist Delay)** Agents exiting poverty may rationally continue to use short-termist strategies for some time after those strategies have ceased to be beneficial.

It takes a good deal of evidence to paint an adequate picture of how an agent's circumstances have changed, and how those changed circumstances reduce the performance of cognitive and behavioral strategies that previous experience led agents to trust. While we would not like to rationalize arbitrarily large or arbitrarily long failures to adapt, it seems harsh to say that agents behave irrationally if they do not immediately and fully adapt their behaviors to their new environments.

Externalists may struggle to recover Rational Short-Termist Delay. In a changed environment, the Cognitive Scarcity Explanation and Changed Prospect Explanations now fail because resources are no longer scarce and prospects are no longer changed. On most externalist approaches, agents are required to immediately adapt to these changes.

By contrast, information-sensitive approaches recover Rational Short-Termist Delay. An information-sensitive approach predicts that agents should gradually adapt their strategies to new environments as they acquire evidence about environments and strategy performance. In this way, an information-sensitive approach helps us to avoid victim-blaming and policy recrimination not only against those who are led to short-termism by current material poverty, but also against those who are led to short-termism by past experience of material poverty.

## 10 Conclusion

This paper investigated the prospects for a de-externalized theory of ecological rationality. After stating (Section 2) and motivating (Section 3) ecological rationality, we saw that there are *prima facie* compelling reasons why existing theories of ecological rationality are often externalist (Section 4). Nevertheless, we saw that externalist approaches face challenges (Section 5) which motivate an alternative information-sensitive approach (Section 6). We saw that the information-sensitive approach avoids the challenges facing externalist accounts (Section 7) while capturing many of their motivations (Section 8). We also saw that the information-sensitive approach gives an expanded rationalizing treatment of agents caught in poverty traps (Section 9).

It may be productive for future work to explore the prospects for information-sensitive approaches to further contentions of the ecological rationality approach, such as the claims that rationality is procedural (Simon 1976; Thorstad forthcoming) and that heuristic cognition is often rational under deep uncertainty (Brighton and Gigerenzer 2012). Another fruitful project would be investigating alternative information-sensitive approaches, such as those incorporating evidence that agents irrationally failed to gather (Goldberg 2016, 2017). Finally, it would be productive to explore the vindictory implications of the information-sensitive approach beyond the applications to poverty traps and heuristics and biases cases.

---

<sup>22</sup>Relatedly, see Podgorski (2016a,b,c, 2017) and Na'aman (2021a,b) on rational delay.

## References

Alston, William. 1995. "How to think about reliability." *Philosophical Topics* 23:1–29.

Altındağ, Onur and O'Connell, Stephen. 2023. "The short-lived effect of unconditional cash transfers to refugees." *Journal of Development Economics* 169:102942.

Banerjee, Abhijit and Duflo, Esther. 2007. "The economic lives of the poor." *Journal of Economic Perspectives* 21:141–67.

Bergmann, Michael. 2006. *Justification without awareness: A defense of epistemic externalism*. Oxford University Press.

Bernheim, B. Douglas, Ray, Debray, and Yeltekin, Sevin. 2015. "Poverty and self-control." *Econometrica* 83:1877–1911.

Bhutta, Neil, Skiba, Paige Marta, and Tobacman, Jeremy. 2015. "Payday loan choices and consequences." *Journal of Money, Credit and Banking* 47:223–59.

Bishop, Michael and Trout, J.D. 2004. *Epistemology and the psychology of human judgment*. Oxford University Press.

Blackwell, David. 1953. "Equivalent comparisons of experiments." *The Annals of Mathematical Statistics* 24:265–72.

Blanden, Jo and Gregg, Paul. 2004. "Family income and educational attainment: A review of approaches and evidence for Britain." *Oxford Review of Economic Policy* 20:245–63.

Boudry, Maarten, Vlerick, Michael, and McKay, Ryan. 2015. "Can evolution get us off the hook? Evaluating the ecological defense of human rationality." *Consciousness and Cognition* 33:524–35.

Brighton, Henry and Gigerenzer, Gerd. 2012. "How heuristics handle uncertainty." In Peter Todd and Gerd Gigerenzer (eds.), *Ecological rationality: Intelligence in the world*, 33–60. Oxford University Press.

Cariani, Fabrizio. 2016. "Deontic modals and probabilities: One theory to rule them all?" In Nate Charlow and Matthew Chrisman (eds.), *Deontic modality*, 11–46. Oxford University Press.

Cohen, Sheldon, Doyle, William, and Baum, Andrew. 2006. "Socioeconomic status is associated with stress hormones." *Psychosomatic Medicine* 68:414–20.

Cohen, Sheldon, Janicki-Deverts, Denis, Chen, Edith, and Matthews, Karen. 2010. "Childhood socioeconomic status and adult health." *Annals of the New York Academy of Sciences* 1186:37–55.

Comesaña, Juan. 2010. "Evidentialist reliabilism." *Noûs* 571–600.

de Bruijn, Ernst-Jan and Antonides, Gerrit. 2022. "Poverty and economic decision making: A review of scarcity theory." *Theory and Decision* 92:5–37.

Dean, Joshua. 2024. "Noise, cognitive function, and worker productivity." *American Economic Journal: Applied Economics* 16:332–60.

Dobridge, Christine. 2018. "High-cost credit and consumption smoothing." *Journal of Money, Credit and Banking* 50:407–33.

Douven, Igor. 2020. "The ecological rationality of explanatory reasoning." *Studies in History and Philosophy of Science Part A* 79:1–14.

Dupas, Pascaline and Robinson, Jonathan. 2013. "Why don't the poor save more? Evidence from health savings experiments." *American Economic Review* 103:1138–71.

Dusi, Giovanni. 2024. "Reliabilist epistemology meets bounded rationality." *Synthese* 115.

Fiscella, Kevin and Franks, Peter. 1997. "Poverty or income inequality as predictor of mortality: Longitudinal cohort study." *British Medical Journal* 314:1724–7.

Fodor, Jerry and Piattelli-Palmarini, Massimo. 2010. *What Darwin got wrong*. Farrar, Straus and Giroux.

Gailliot, Matthew, Baumeister, Roy C., Dewall, Nathan, Maner, Jon, Plant, E. Ashby, Tice, Dianne, Brewer, Lauren, and Scheichel, Brandon. 2007. "Self-control relies on glucose as a limited energy source: Willpower is more than a metaphor." *Journal of Personality and Social Psychology* 92:325–36.

Geman, Stuart, Bienenstock, Elie, and Doursat, René. 1992. "Neural networks and the bias/variance dilemma." *Neural Computation* 4:1–58.

Gigerenzer, Gerd. 2019. "Axiomatic rationality and ecological rationality." *Synthese* 194:3547–64.

Gigerenzer, Gerd and Brighton, Henry. 2009. "Homo heuristicus: Why biased minds make better inferences." *Topics in Cognitive Science* 1:107–43.

Gigerenzer, Gerd and Selten, Reinhard (eds.). 2001. *Bounded rationality: The adaptive toolbox*. MIT Press.

Gigerenzer, Gerd and Selten, Reinhard. 2001b. "Rethinking rationality." In Gerd Gigerenzer and Reinhard Selten (eds.), *Bounded rationality: The adaptive toolbox*, 1–12. MIT Press.

Gigerenzer, Gerd and Sturm, Thomas. 2012. "How (far) can rationality be naturalized?" *Synthese* 187:243–68.

Gilovich, Thomas and Griffin, Dale. 2002. "Heuristics and biases: Then and now." In Thomas Gilovich, Dale Griffin, and Daniel Kahneman (eds.), *Heuristics and biases: The psychology of intuitive judgment*, 1–18. Cambridge University Press.

Goldberg, Sanford. 2012. "A reliabilist foundationalist coherentism." *Erkenntnis* 77:187–96.

—. 2016. "On the epistemic significance of evidence you should have had." *Episteme* 13:449–70.

—. 2017. "Should have known." *Synthese* 194:2863–2894.

Goldman, Alvin. 1979. "What is justified belief?" In George Pappas (ed.), *Justification and knowledge*, 1–23. D. Reidel.

—. 1986. *Epistemology and cognition*. Harvard University Press.

—. 1994. "Naturalistic epistemology and reliabilism." *Midwest Studies in Philosophy* 19:301–20.

Gollier, Christian. 2002. "Discounting an uncertain future." *Journal of Public Economics* 85:149–66.

Good, I.J. 1966. "On the principle of total evidence." *British Journal for the Philosophy of Science* 17:319–321.

Greco, John. 2010. *Achieving knowledge: A virtue-theoretic account of epistemic normativity*. Cambridge University Press.

Griskevicius, Vladas et al. 2013. "When the economy falters, do people spend or save? Responses to resource scarcity depend on childhood environments." *Psychological Science* 24:197–205.

Grohmann, Antonia, Kouwenberg, Roy, and Menkhoff, Lukas. 2015. "Childhood roots of financial literacy." *Journal of Economic Psychology* 51:114–33.

Hankins, Scott, Hoekstra, Mark, and Skiba, Paige Marta. 2011. "The ticket to easy street? The financial consequences of winning the lottery." *Review of Economics and Statistics* 93:961–9.

Haushofer, Johannes and Fehr, Ernst. 2014. "On the psychology of poverty." *Science* 344:862–7.

Haushofer, Johannes and Salicath, Daniel. 2023. "The psychology of poverty: Where do we stand?" *Social Philosophy and Policy* 40:150–84.

Heitz, Richard. 2014. "The speed-accuracy tradeoff: History, physiology, methodology, and behavior." *Frontiers in Neuroscience* 8:1–19.

Icard, Thomas. ms. *Resource rationality*.

Jakiela, Pamela and Ozier, Owen. 2016. "Does Africa need a rotten kin theorem? Experimental evidence from village economies." *Review of Economic Studies* 83:231–68.

Johnson, Eric and Payne, John. 1985. "Effort and accuracy in choice." *Management Science* 31:395–414.

Kahneman, Daniel, Slovic, Paul, and Tversky, Amos (eds.). 1982. *Judgment under uncertainty: Heuristics and biases*. Cambridge University Press.

Karlan, Dean, Mullainathan, Sendhil, and Roth, Benjamin. 2019. "Debt traps? Market vendors and moneylender debt in India and the Philippines." *American Economic Review: Insights* 1:27–42.

Kolodny, Niko and MacFarlane, John. 2010. "Ifs and oughts." *Journal of Philosophy* 107:115–43.

Maher, Patrick. 1990. "Why scientists gather evidence." *British Journal for the Philosophy of Science* 41:103–119.

Mani, Anandi, Mullainathan, Sendhil, Shafir, Eldar, and Zhao, Jiaying. 2013. "Poverty impedes cognitive function." *Science* 341:976–80.

Marewski, Julian and Schooler, Lael. 2011. "Cognitive niches: An ecological model of strategy selection." *Psychological Review* 118:393–437.

McLeod, Peter and Dienes, Zoltan. 1996. "Do fielders know where to go to catch the ball or only how to get there?" *Journal of Experimental Psychology: Human Perception and Performance* 22:531–543.

Morton, Jennifer. 2017. "Reasoning under scarcity." *Australasian Journal of Philosophy* 95:543–59.

Mousavi, Shabnam and Kheirandish, Reza. 2014. "Behind and beyond a shared definition of ecological rationality: A functional view of heuristics." *Journal of Business Research* 67:1780–5.

Na'aman, Oded. 2021a. "Emotions and process rationality." *Australasian Journal of Philosophy* 99:531–46.

—. 2021b. "The rationality of emotional change: Toward a process view." *Noûs* 55:245–69.

Patel, Nirav, Grandner, Michael, Xie, Dawei, Branas, Charles, and Gooneratne, Nalaka. 2010. "'Sleep disparity' in the population: Poor sleep quality is strongly associated with poverty and ethnicity." *BMC Public Health* 10:e1–11.

Pepper, Gillian and Nettle, Daniel. 2014a. "Out of control mortality matters: The effect of perceived uncontrollable mortality risk on a health-related decision." *PeerJ* 26:e459.

—. 2014b. "Socioeconomic disparities in health behavior: An evolutionary perspective." In Mhairi Gibson and David Lawson (eds.), *Applied evolutionary anthropology*, 225–43. Springer.

—. 2017. "The behavioral constellation of deprivation: Causes and consequences." *Behavioral and Brain Sciences* 40:e314.

Podgorski, Abelard. 2016a. "Dynamic conservatism." *Ergo* 3:349–76.

—. 2016b. "Dynamic permissivism." *Philosophical Studies* 173:1923–39.

—. 2016c. "A reply to the synchronist." *Mind* 125:859–71.

—. 2017. "Rational delay." *Philosophers' Imprint* 17:1–19.

Rich, Patricia. 2016. "Axiomatic and ecological rationality: Choosing costs and benefits." *Erasmus Journal for Philosophy and Economics* 9:90–122.

Rieskamp, Jörg and Otto, Philipp. 2006. "SSL: a theory of how people learn to select strategies." *Journal of Experimental Psychology: General* 135:207–36.

Santos, Paulo and Barrett, Christopher. 2011. "Persistent poverty and informal credit." *Journal of Development Economics* 96:337–47.

Schmidt, Andreas. 2019. "Getting real on rationality – behavioral science, nudging, and public policy." *Ethics* 129:511–543.

Shafir, Eldar and Mullainathan, Sendhil. 2013. *Scarcity: Why having too little means so much.* Times Books.

Shah, Anuj, Mullainathan, Sendhil, and Shafir, Eldar. 2012. "Some consequences of having too little." *Science* 338:682–5.

Simon, Herbert. 1956. "Rational choice and the structure of the environment." *Psychological Review* 63:129–138.

—. 1976. "From substantive to procedural rationality." In T.J. Kastelein, S.K. Kuipers, W.A. Nijenhuls, and R.G. Wagenaar (eds.), *25 years of economic theory*, 65–86. Springer.

Sosa, Ernest. 1994. "Philosophical scepticism and epistemic circularity." *Proceedings of the Aristotelian Society, Supplement* 263–90.

—. 2015. *Judgment and agency*. Oxford University Press.

Staffel, Julia. 2020. *Unsettled thoughts: A theory of degrees of rationality*. Oxford University Press.

Stanovich, Keith and West, Richard. 2000. "Individual differences in reasoning: Implications for the rationality debate?" *Behavioral and Brain Sciences* 23:645–65.

Thorstad, David. 2024. *Inquiry under bounds*. Oxford University Press.

—. 2025. "The complexity-coherence tradeoff in cognition." *Mind* 134:422–57.

—. forthcoming. "The zetetic turn and the procedural turn." *Journal of Philosophy* forthcoming.

Todd, Peter and Gigerenzer, Gerd. 2012. "What is ecological rationality?" In Peter Todd and Gerd Gigerenzer (eds.), *Ecological rationality: Intelligence in the world*, 3–30. Oxford University Press.

Tooby, John and Cosmides, Leda. 1994. "The biological foundations of culture." In *The adapted mind: Evolutionary psychology and the generation of culture*, 85–116. Oxford University Press.

Tversky, Amos and Kahneman, Daniel. 1974. "Judgment under uncertainty: Heuristics and biases." *Science* 185:1124–31.

Weatherson, Brian. 2024. *Knowledge: A human interest story*. OpenBook Publishers.

Wheeler, Gregory. 2020. "Less is more for Bayesians, too." In Riccardo Viale (ed.), *Routledge handbook on bounded rationality*, 471–83. Routledge.

Williamson, Timothy. 2000. *Knowledge and its limits*. Oxford University Press.

Yalcin, Seth. 2007. "Epistemic modals." *Mind* 116:983–1026.