Two paradoxes of bounded rationality

Abstract

My aim in this paper is to develop a unified solution to two paradoxes of bounded rationality. The first is the regress problem that incorporating cognitive bounds into models of rational decisionmaking generates a regress of higher-order decision problems. The second is the problem of rational irrationality: it sometimes seems rational for bounded agents to act irrationally on the basis of rational deliberation. I review two strategies which have been brought to bear on these problems: the way of weakening which responds by weakening rational norms, and the way of indirection which responds by letting the rationality of behavior be determined by the rationality of the deliberative processes which produced it. Then I propose and defend a third way to confront the paradoxes: the way of level separation.

1   Introduction

The study of bounded rationality presents new problems and paradoxes. My aim in this paper is to set out and study two of the best-known paradoxes of bounded rationality which I will term the problem of rational irrationality and the regress problem.

These paradoxes are often studied separately. However, many of the same normative ideas have been brought to bear on each of the paradoxes. As a result, I will study the paradoxes together. Studying the paradoxes together avoids duplication of normative labor, gives a broader view of the normative strategies underlying solutions to the paradoxes, and allows us to assess these strategies globally by their track records across the board. This discussion will point the way towards a new strategy which can be used to shed light on both paradoxes.

Here is the plan. Section 2 presents two paradoxes of bounded rationality.[1] Sections 3-4 review two normative strategies that have been used to resolve the paradoxes. Section

---

[1]Like many bounded rationality phenomena, these paradoxes arise both in the practical domain of

5 proposes a new strategy, and Section 6 shows how this strategy can be used to confront the paradoxes.

2   Two paradoxes of bounded rationality

2.1   The problem of rational irrationality

It is widely agreed that bounded agents should sometimes decide how to act using a toolbox of decisionmaking heuristics (Gigerenzer and Selten 2001; Gigerenzer and Gaissmaier 2011). Heuristics typically use only some of the information available to agents or make only some of the inferences warranted by that information.

There are at least three reasons why rational decisionmaking is often heuristic. First, in many cases there is an accuracy-effort tradeoff between the quality of decisions and the cost of producing them (Johnson and Payne 1985). Heuristics often strike a better balance than nonheuristic methods between decision quality and decision costs. Second, we sometimes lack the cognitive ability to carry out complex nonheuristic inferences, for example by fully calculating expected utilities or deducing complex theorems (Chickering 1996; Dagum and Luby 1993). We cannot be rationally required to use decisionmaking strategies that are beyond our abilities. Third, we face less-is-more effects: in some contexts simpler decision rules reliably outperform more complex decision rules on a variety of performance metrics (Geman et al. 1992; Gigerenzer and Brighton 2009).[23]

To illustrate the structure of heuristic decisionmaking, consider decisionmaking by

---

rational decisionmaking and in the theoretical domain of rational belief formation. For brevity, I present each paradox in the special case of rational decisionmaking. When possible, I give examples designed to illustrate how the same paradoxes arise in rational belief formation.

[2]A related phenomenon is that sometimes it may be better for agents to be ignorant of some facts than to know those facts (Schooler and Hertwig 2005; Hertwig and Engel 2021).

[3]We can also support the rationality of heuristic cognition by holding that rationality is *ecological*, or environment-relative (Berg 2003, 2014; Katsikopoulos et al. 2020; Todd and Gigerenzer 2012) and noting that heuristics often perform well in the environments to which they are adapted. Similar arguments may follow from adopting a naturalized perspective towards rationality (Bishop and Trout 2004), or from thinking about the performance of heuristics in unstable environments (Bookstaber and Langsam 1985).

satisficing (Selten 1998; Simon 1955). Imagine that you are selecting a hotel room for a conference. A satisficer first selects an aspiration level in one or more goods to guide her search. For example, the hotel room should cost less than $150, be within five miles of the conference, have at least three stars and offer continental breakfast. The satisficer then searches for options, in this case hotel rooms, one at a time until she finds an option which meets all of her aspirations. She then takes that option, in this case by booking the hotel room.

In many contexts satisficing is a rational decisionmaking heuristic, returning effective decisions at reasonable cost. The problem is that satisficing sometimes produces actions which look irrational.[4] To see the problem, note three facts about satisficing.[5] First, satisficing displays only a threshold sensitivity to quantity. Although hotel rooms are penalized for costing more than $150, satisficing does not discriminate between hotel rooms which cost $149, $100, or $1. Second, satisficing is insensitive to goods which are not reflected in the agent's aspirations. Our satisficer's decision is unaffected by the presence or absence of a hotel rewards program. Third, satisficing is noncompensatory: shortfalls along one dimension cannot be compensated by excesses along another. Our satisficer will not book a hotel without continental breakfast, even if the cost savings would be enough to buy the best breakfast in town.

Such is the price of cognitive frugality. These vulnerabilities do not mean that satisficing is irrational in this environment.[6] But they do mean that satisficing will sometimes produce actions which look irrational, even when by all appearance satisficing was a rational decisionmaking strategy and was applied in a rational manner. Because satisficing is noncompensatory, ignores relevant goods, and displays only a threshold sensitivity to

_____

[4]We will see below that this is not to be equated with the claim that satisficing itself is irrational in this environment. One reason for that might be that satisficing does well in areas such as speed, frugality, and resilience to environmental shocks.

[5]An anonymous referee notes that alternative satisficing rules may lack some of these features, and some optimizing rules may have them.

[6]Here and throughout I take an *ex ante* approach to the rationality of decision procedures as well as the rationality of their results.

quantity, satisficing can produce choices which go against the balance of an agent's normative reasons; do not maximize expected utility; and would be withdrawn by the agent on considered reflection. For example, our satisficer may pass up an excellent hotel room which costs $100, exceeds most of her aspirations and comes with a top-notch rewards program, but is six miles away, only to end up booking a barely-satisfactory hotel room at the top of her price range. On most ways of filling out the story, the satisficer has passed up the room she has most reason to book in favor of a room she has less reason to book, and which produces significantly less expected utility. If confronted with these facts, our satisficer would probably reverse her choice on considered reflection.[7]

This example raises the *problem of rational irrationality*. On the basis of apparently rational decisionmaking procedures, agents can take actions which look irrational.[8] Instances of the problem of rational irrationality abound. It has been argued that the poor are led to dramatically undersave and overborrow based on rational patterns of allocating scarce cognitive resources such as attention and computational bandwidth (Mani et al. 2013; Shah et al. 2012). And it is well known that even our most cherished heuristics can lead agents to take dominated actions (Birnbaum 2008), or ignore features of options such as their scope (Desvousges et al. 1992) and duration (Fredrickson and Kahneman 1993) which matter a great deal.[9] In face of these problems, judgment and decisionmaking researchers have generally retained their insistence that decisionmaking heuristics can be rational ways for bounded agents to make up their minds.[10] But how can heuristics be rational if they lead to irrational actions?

---

[7]To be clear, there are other ways of filling out the story that might reverse these normative verdicts: for example, if we assume that the website is strategizing to make money against satisficers.

[8]A referee suggests that we restrict these normative claims to a range of well-behaved decision problems, such as Savage's (1954) small-world decision problems; problems involving individuals rather than groups; or those not involving Knightian uncertainty. While that is not my view, readers who hold such views are welcome to adopt that reading throughout this paper.

[9]A referee suggests another example: skepticism about expert advice may be a rational heuristic stance, although it will sometimes lead to irrational choices. A related phenomenon is also raised by Daniel Newark's work on the logic of absurdity (Newark 2018).

[10]See for example Gilovich and Griffin (2002, p. 3) and Tversky and Kahneman (1974, p. 1130) in the heuristics and biases tradition; Lieder and Griffiths (2020) in the Bayesian tradition; and Gigerenzer and Selten (2001) in the adaptive toolbox tradition.

Some theorists bite the bullet and insist that the problem of rational irrationality is no problem at all. Frank Knight held that:

> It is evident that the rational thing to do is to be irrational, where deliberation
> and estimation cost more than they are worth. (Knight 1921, p. 67).

On Knight's view, if acting rationally is more trouble than it is worth, then rationality requires us to act irrationally. This view carries an uncomfortable ring of paradox that most philosophers will want to avoid.

As a consequence, a good solution to the problem of rational irrationality should do three things. First, it should dispel the paradox by explaining why it cannot be rational to act irrationally, even as a result of rational decisionmaking. Second, it should explain why we are nonetheless tempted to classify these cases as instances of rational irrationality. And third, it should help us to see how the many near-cousins of Knight's position in the contemporary empirical literature may be in an important sense correct.

2.2   The regress problem

Suppose you are deciding how much money to save for retirement this year. We might model the base-level decision problem that you face by letting the options available to you be numbers of dollars saved and building a probabilistic model of the effects of each savings strategy. Call this problem $\mathcal{D}_0$.

Bounded rationality theorists protest that $\mathcal{D}_0$ is an incomplete description of the decision problem that you face. Investment decisions are the outcome of costly and ability-constrained processes of deliberation. Instead of immediately making an investment, you may first choose to deliberate about your investment strategy using one of these processes. So we should enrich our decision problem by incorporating new options representing decision procedures over monetary investments. Call this enriched problem $\mathcal{D}_1$.

But decision procedures are not produced in a vacuum any more than the decisions that result from them. Instead of deliberating straightaway, you may engage in metacognitive strategy selection (Lieder and Griffiths 2017; Marewski and Schooler 2011) aimed at picking a good decision procedure over investment strategies. So we need to enrich the decision problem to a new problem $\mathcal{D}_2$, with an enlarged option space that includes metacognitive processes of strategy selection.

And now we seem to have talked ourselves into a muddle. Metacognitive strategies can also be chosen as a result of higher-order deliberation, so it looks as though we will have to enrich the decision problem to a new problem $\mathcal{D}_3$ with an option space incorporating these meta-metacognitive procedures.[11] Proceeding in this way, we can generate an ascending hierarchy of decision problems which look to improve on their predecessors in exactly the same way at $\mathcal{D}_1$ improved on $\mathcal{D}_0$.

This line of reasoning has been taken to generate a family of *regress problems* for theorizing about bounded rationality.[12] To begin, it may be unclear if there is any complete description of the decision problem that bounded agents face. Many authors have resisted this way of taking the problem by holding that decision problems facing bounded agents can be described as a limit, for example the limit $\mathcal{D}_\omega$ of all the finite-level decision problems posed above, or even more speculatively the limit $\mathcal{D}_\alpha$ of decision problems below some larger transfinite ordinal α (Lipman 1991; Mongin and Walliser 1988; Vassilakis 1992).[13] But this raises a new threat: it may turn out that there is nothing which it is rational for bounded agents to do unless we can prove the existence of a solution to these limit problems which satisfies even minimal rationality constraints connecting the rationality of deliberation procedures to the rationality of their results.[14]

---

[11]To be clear, this is only a claim about the nature of the decision problem that agents face, not a claim about how agents must model that problem to themselves.

[12]For discussion of what the regress problem could be, see Lin (2014) and Smith (1991).

[13]We will see in the next section that there is some reason to be skeptical of these latter solutions.

[14]For example, Mongin (2000) finds that a solution exists only if the agent's deliberation procedure is a constant function, returning the same base-level action in each world state. Constant deliberation procedures are typically irrational because they are insensitive to evidence.

Another difficulty is that as we enrich our decision problems they come to look less and less like problems that bounded agents face. If we know anything about bounded rationality, it is that boundedly rational agents typically do not operate with rich problem descriptions full of higher-order facts about deliberative and metacognitive processes. One reason for this is that the cognitive cost of solving each problem $\mathcal{D}_i$ is a rapidly increasing function of i, and the cognitive benefits of increased complexity soon diminish.[15] As we iterate problem-hierarchies into the transfinite, it becomes increasingly less plausible that the decision problems described have anything much to do with the problems that bounded agents face.

The upshot of this discussion is that a solution to the regress problem should say three things. First, it should say precisely which decision problem bounded agents face. Is it $\mathcal{D}_0$, $\mathcal{D}_1$, $\mathcal{D}_\omega$, or something else? Second, it should show that, at least in typical cases, there exists a rational action which can be taken in such decision problems. And third, a solution should show how the type of decision problem described can plausibly be said to be a problem that bounded agents confront.

Now we have our marching orders. We want to see what bounded rationality theorists should say about the problem of rational irrationality and the regress problem. And we have three criteria on a good solution to each problem. In Sections 3-4, I review two strategies that are commonly brought to bear on these problems. Then in Sections 5-6, I propose a third way to confront the paradoxes.

3   The way of indirection

Our first solution begins with the problem of rational irrationality. This problem arises because seemingly rational decisionmaking procedures such as satisficing can produce irrational-looking actions, such as passing up an excellent hotel room in favor of a barely satisfactory room. If we want to avoid the Knightian conclusion that it can be rational

---

[15]Other standard arguments for the rationality of heuristic cognition, including those advanced in Section 2.1, apply here as well.

to act irrationally, a natural idea is to find some guarantee that the rational status of deliberation procedures cannot come apart from the rational status of the actions they produce.

The *way of indirection* does this by making the normative status of actions derivative on the normative status of deliberation procedures. The most common way to do this is to say that an action is rational if it results from a rational deliberation procedure. Then apparent cases of rational irrationality will turn out to be rational actions resulting from rational decision procedures. If it is rational to make up your mind by satisficing, then it is rational to book the hotel room that results.

The way of indirection is the orthodox solution to the problem of rational irrationality among bounded rationality theorists. Because important bounds such as the costs of computation are felt most strongly at the level of decision procedures, Herbert Simon held that the fundamental turn in the study of bounded rationality is the turn to a notion of procedural rationality which reflects the rational importance of decision procedures. Simon understood the notion of procedural rationality as the way of indirection suggests: "Behavior is procedurally rational when it is the outcome of appropriate deliberation" (Simon 1976, p. 66). Most theorists in the bounded tradition have followed Simon in holding that apparently irrational behavior is rational when it results from appropriate deliberation.

Some authors have extended the way of indirection to solve the regress problem by allowing indirect normative assessment to cover higher levels of deliberation such as metacognition. We can treat options in the base-level decision problem $\mathcal{D}_0$ as nondeliberative actions such as booking a hotel room, and options which first appear in the $\alpha + 1$-st level $\mathcal{D}_{\alpha+1}$ as decision procedures over options at level α. This representation can be used to generate two types of indirect solutions to the regress problem.

The first solution class, *limit models*, proves that after a (typically uncountable) number of iterations the decision problems $\mathcal{D}_\alpha$ converge to a limit $\mathcal{D}^*$ which reflects all deliberative and nondeliberative options that agents face (Lipman 1991). The regress stops at this level.

Limit models are used to claim that rational agents should take the optimal deliberative or non-deliberative option o^* in $\mathcal{D}^*$. Applying the way of indirection, it is held that lower-level deliberative and nondeliberative options are rational just in case they are entailed by o^*. For example, if o^* is decisionmaking by satisficing then it is also rational to book whichever hotel room results from satisficing and irrational to engage in any metacognitive or higher-order deliberation on the matter.

The second solution class, *convergence theorems* identifies a level α and a base-level option o. It is shown that for any $\beta > \alpha$, the rational option in $\mathcal{D}_\beta$ eventually leads the agent to do o (Mongin 2000). On this basis, it can be held that o must be the eventual result of rational deliberation. By indirection, this means that o is rational.

How should we evaluate the way of indirection? There are many ways to push back here. One worry is that the way of indirection may count irrational-looking base-level actions as rational when they are produced by rational deliberation procedures. I press this case in Section 4. Another worry is that the way of indirection models agents as responding to long chains of decision problems or transfinite limits of these chains, which seems a substantial departure from the situation of bounded decisionmakers. I want to set these worries aside for now and focus on a challenge that gets to the heart of my disagreement with the way of indirection.

The challenge for the way of indirection is that it collapses normative structure. It takes rich examples full of many objects for normative assessment such as action, deliberation and metacognition, then forces us to assess each of these objects in the same way. This is admirably parsimonious, but sometimes it is too parsimonious to capture the normative data.

As an example, consider the Cognitive Reflection Test. The Cognitive Reflection Test asks questions such as the following:

> **(Bat and Ball Problem**) A bat and a ball cost $1.10. The bat costs $1.00 more
> than the ball. How much does the ball cost? (Frederick 2005).

The Bat and Ball Problem triggers an initial, frugal reasoning process based on attribute

substitution (Kahneman and Frederick 2002).[16] This process simplifies the problem statement by replacing '$1.00 more than' with '$1.00', and from there quickly reasoning to the conclusion that the ball costs 10 cents. Many highly intelligent people fail to override this initial reasoning process, outputting the incorrect judgment that the ball costs 10 cents.

Something irrational has happened here, but what does that irrationality consist in? Assume for simplicity that reasoning processes go no higher than $\mathcal{D}_2$, the level of metacognition. This is the highest level about which we have good descriptive data in most contexts, and the problems that I raise for the way of indirection will only get worse as we move to higher levels.

There is mounting evidence that the Bat and Ball Problem does not involve irrational metacognition. It was initially thought that the Bat and Ball Problem involves irrational metacognitive *monitoring* of deliberation: agents are not paying attention to the possibility that their strategies could be mistaken (Kahneman and Frederick 2002). But this is too hasty: across a variety of measures, agents show good evidence that they are aware of the risk of error.[17] So we do not yet have evidence for irrational metacognitive monitoring in the Bat and Ball Problem.

This might suggest that the Bat and Ball Problem involves irrational metacognitive *control* of deliberation. After all, if agents are aware that their strategies could be unreliable, then why do they not override those strategies in favor of other, more reliable strategies? But again, this argument is too quick. To see the point, it may help to redescribe the problem in terms familiar from signal detection theory (Jackson et al. 2016).

In problems such as the Bat and Ball Problem, we are confronted by an initially tempting

---

[16]The time course of heuristic and nonheuristic reasoning in these problems is a matter of controversy (Bago and De Neys 2017; Travers et al. 2016). To my knowledge, all of my arguments in this section can be reformulated to accommodate other leading views.

[17]For evidence of rational metacognitive monitoring, compare 'conflict' tasks such as the Bat and Ball Problem to 'no-conflict' versions of the same tasks in which no risky heuristic process tempts reasoners. In conflict tasks agents take longer to produce answers and confidence judgments (Bonner and Newell 2010; De Neys et al. 2013; Johnson et al. 2016); report lower metacognitive feelings of rightness and increased feelings of error (Gangemi et al. 2015; Thompson et al. 2011; Thompson and Johnson 2014); and suggest sensitivity to conflict in process measures such as gaze tracking and talk-out-loud procedures (De Neys and Glumicic 2008).

heuristic strategy: in this case, attribute substitution. The metacognitive control task is to decide whether to let heuristic deliberation go ahead or to override it in favor of explicit calculation. Rational metacognitive control strives to balance two risks. First, there is a risk of false negatives: failing to block unreliable heuristic strategies. False negatives drive up the error cost from false judgments and bad decisions. Second, there is a risk of false positives: intervening to block reliable heuristic strategies. False positives drive up the cost of reasoning because explicit calculation is expensive. Rational metacognitive control must balance these risks by accepting nontrivial risks of both false positives and false negatives. The only way to avoid false negatives altogether, that is to never allow unreliable heuristics to go ahead, would be to always override heuristic reasoning. But we saw in Section 2 that it is often rational for bounded agents to reason heuristically.

The Bat and Ball Problem induces a false negative: an unreliable process of attribute substitution that has been allowed to go ahead. But we have seen that the occasional existence of false negatives is a predictable consequence of rational metacognitive control, so that in itself is no evidence of irrational metacognitive control. More generally, many theorists have thought that there is a good deal of evidence that agents exercise rational metacognitive control in the Bat and Ball Problem.[18]

But then where is the irrationality in the Bat and Ball Problem? We would like to say that this problem involves irrational belief and deliberation on the basis of rational metacognitive strategy selection. Agents select an inappropriate form of deliberation, namely attribute substitution, and as a result form inappropriate beliefs, but they do so on the basis of rational metacognitive monitoring and control of deliberation. However, the indirect approach cannot say this. The indirect approach will have to say that there is

---

[18]At first two strands are worth noting. First, numbers in the Bat and Ball Problem are carefully chosen: when the numbers are changed, agents tend not to make any unreliable attribute substitutions (Frederick 2005). This suggests that their metacognitive control processes have a low rate of false negatives, producing few false negatives over a range of similar problems. Second, individual differences in performance on the Bat and Ball Problem and a range of similar problems are explicable in terms of three things: differences in metacognitive monitoring attitudes; differences in abilities to execute first-order reasoning processes; and different thresholds for the strength of error signal needed to override heuristic reasoning (Jackson et al. 2016). These are exactly the factors that normative theory predicts should be relevant to metacognitive control.

nothing irrational in our response to the Bat and Ball Problem if that response results from rational metacognition. And here it looks as though the indirect view errs by refusing to divorce the normative status of deliberation and its outcomes from the normative status of the metacognitive processes that produced them.

Exactly the same structure is found in decisionmaking problems. Sometimes agents act irrationally and deliberate irrationally about how to act, but do so on the basis of rational metacognition. This would happen, for example, if I asked you to make a small wager about the cost of the ball. Here the way of indirection says that irrational-seeming actions and deliberations are rational when they result from rational metacognition. This claim removes the possibility of paradox by preventing us from saying what we would like to say: that the agent has acted and deliberated irrationally, but metacognized rationally. The way of indirection collapses normative structure by forcing us to make the same assessment of all three objects.

4   The way of weakening

Our second solution begins with reflection on the relationship between rationality and maximization. Many philosophers hold that actions must maximize some important quantity such as choiceworthiness or value, otherwise they would not be rational. But in some applications, maximizing requirements can seem overly demanding.

Suppose you are at the grocery store choosing among a shelf full of artisanal vinegars. It is very easy for you to make a satisfactory choice. Just pick any affordable vinegar with the right flavor profile and a few reasonable indicators of quality. But if choosing rationally requires making the best possible choice, then it looks like you are doomed to spend a half-hour buying vinegar or else run a substantial risk of acting wrongly.

Many authors have taken such examples to suggest that rational norms for bounded agents need to be weakened in order to make them less demanding. The *way of weakening* confronts the paradoxes by developing weakened forms of traditional rational norms

which are appropriate for bounded agents.

The best-known philosophical exemplar of the way of weakening is Michael Slote's satisficing consequentialism (Slote 1984).[19] Consequentialists traditionally apply a maximizing standard of correctness on which actions are right or rational just in case they are best.[20] Satisficing consequentialism weakens the traditional consequentialist view by applying a satisficing standard of correctness on which actions are right just in case they are sufficiently good, even if not the best.

Satisficing consequentialism explains our intuitions about rational vinegar purchasing. It is right to buy any sufficiently good bottle of vinegar, on this view, because rightness demands only that we do what is good enough, not what is best.

Satisficing consequentialism will also reduce the problem of rational irrationality. Many cases used to motivate the problem of rational irrationality involve satisfactory action on the basis of rational deliberation. For example, our satisficing hotel-shopper books a merely satisfactory hotel room on the basis of rational heuristic deliberation. In these cases, satisficing consequentialism will say that agents both act and deliberate rationally because their actions and deliberations are both satisfactory. As a result, many of the cases used to motivate the problem of rational irrationality will not involve any type of irrationality.

Some authors hold that satisficing solves the regress problem. The regress problem gets off the ground because it is hard to see how agents can act optimally without engaging in higher-order deliberation. But it may be easier to imagine that agents could act satisfactorily well without further deliberation. Most famously, John Elster claims that the

---

[19]Another prominent exemplar is Christopher Cherniak's minimal rationality program (Cherniak 1981, 1986). As always, each position deserves its own separate treatment, although many of the remarks made here will generalize to other programs. It is important to recall that Cherniak's primary motivation was the predictive-explanatory project of finding the minimal conditions for the aptness of describing agents using intentional vocabulary. That may be a separate and compatible project (Bermúdez 2009).

[20]Here I will not be concerned with the difference between rightness and rationality as normative categories. One way to justify this omission would be to adopt a reason-responsiveness view on which rationality and rightness are both determined by the totality of relevant normative reasons. On many versions of this view, what it is right to do and what we are rationally required to do will be the same thing (Lord 2017).

right way out of the regress problem is to decide, at some level $\mathcal{D}_n$ to deliberate no further and act (Elster 1983). Elster thinks that the right way to take this action will be through satisficing, and that this will be a fully rational thing to do.

One worry for the weakening approach is that it opens the door to irrational forms of submaximization.[21] Slote considers the example of a home-seller who lists her home for a price that she considers satisfactory. Slote asks us to imagine that the home-seller could fetch a higher price with no added uncertainty or delay. Asked why she does not seek more money for her house, the home-seller replies that she is a satisficer: for her, a satisfactory price is good enough. Satisficing consequentialism permits her to set any satisfactory price, even if more money could be obtained with no additional cost.[22]

I think that any plausibility attaching to the satisficing consequentialist's verdict in this case rests on the difficulty of imagining what Slote asks us to imagine: a case in which raising the price would not increase the time-to-sale, decrease the likelihood of sale, or invite greedy relatives to ask for their share of the pie. If we set these assumptions aside, it is difficult to see how the satisficer can rationally turn down free money. Perhaps we are to imagine that the satisficer has no desire to spend more money on herself. But then she could take the money and donate it to charity. Or perhaps the satisficer doubts that she could spend the money more productively than the home-buyer. But then we no longer have a failure to maximize value. These difficulties have led many authors to think that satisficing consequentialism should be construed in some other way, for example as a point about virtuous dispositions rather than a revision to standards of correctness.[23]

Turn next to the problem of rational irrationality. This problem arises in many classic heuristics and biases tasks, where rational-seeming heuristics produce irrational-seeming

---

[21]On this point see Bradley (2006), Byron (1998), Mulgan (2001), and Pettit (1984) among others.

[22]One might also worry that the way of weakening wrongly supposes that heuristics are inferior to nonheuristic strategies. Classic defenses of heuristic cognition hold that heuristic strategies often outperform nonheuristic strategies, once all normatively relevant factors are considered. If this is right, then there should be no need to weaken normative standards in order to rationalize heuristic cognition. Thanks to an anonymous referee for raising this point.

[23]On this point see Swanton (1993, 2004) and perhaps also Slote (2004).

judgments and decisions. One datum to explain in these cases is why participants, when confronted with their errors, often withdraw or repudiate them. If I tell you that you have borrowed too much money, booked the wrong hotel room, or judged a conjunction to be more probable than one of its conjuncts, you may retort that you did so as a result of rational deliberation. But you will certainly not say that your actions and beliefs are rational and hence do not need to be changed. However, insofar as satisficing consequentialism can make progress on the problem of rational irrationality in these tasks, it will do so by saying that the offending beliefs and actions are rational because they are satisfactorily good. This wrongly suggests that participants in many classic tasks should stick to their guns.

Another way of making this point is that many of the base-level actions used to motivate the problem of rational irrationality are themselves blameworthy. For example, return to the previously-cited claim that poverty traps such as overborrowing and undersaving can result from rational patterns of decisionmaking. Roughly, the idea is that agents have only so much time, computational bandwidth and attention to allocate to reasoning. The working poor are faced with an abundance of high-stakes, short-term challenges. As a result it can be rational for the poor to allocate most of their cognitive resources to evaluating the short-term consequences of their actions (Mani et al. 2013; Shah et al. 2012). However, this short-termist pattern of resource allocation produces substantial overborrowing, because short-term capital needs will be prioritized over long-term debt.

Suppose now that a mother borrows too much money because she is busy at work and occupied thinking about how to pay for food, educate her children, and keep the lights on. Tragically, she is unable to repay the loan and as a result her family is evicted. As much as we should stress the fact that she took the loan as a result of rational patterns of deliberation, we would also like to say that the mother is blameworthy for taking the loan. In particular, we know exactly who can blame her: her family. But if this is right, then she must have acted irrationally or else her actions would not be blameworthy.

15

On the regress problem, I confess that I cannot see how satisficing is supposed to help. If the point is that decisions at some level should be made without any higher-level deliberation, then we are in agreement. But this is precisely the point that a solution to the regress problem must argue for. The argument cannot be that in acting without deliberation we are guaranteed to act satisfactorily, because we have no such guarantee. And the point cannot be that decisionmaking by satisficing itself can be done without further deliberation. In its original heuristic sense, satisficing requires the specification of an aspiration level, the ordering of options during search, and procedures for adjusting aspirations as a result of new evidence. All of these can be and frequently are selected on the basis of higher-order deliberation (Caplin and Martin 2011). In asking how to conduct this deliberation, we take up a higher-order decision problem and by all appearances, the regress continues.

So far, we have reviewed and rejected two traditional solutions to the paradoxes: the way of weakening, and the way of indirection. The lessons learned from these discussions point us towards a third way to solve the paradoxes.

5   The way of level separation

In this section, I describe a third way of approaching the paradoxes, the *way of level separation*. The way of level separation consists of three normative claims. The first two of these claims deny the characteristic claims made by the ways of indirection and weakening, in order to avoid their pitfalls.

First, normative assessment should be *direct*, not indirect. The rationality of action, deliberation, metacognition, and higher-order deliberation is determined by separate applications of standards of correctness at each of these levels. For example, direct consequentialism would say that an action such as booking a hotel room is rational just in case the expected utility of booking that room is at least as great as the expected utility of booking any other room (Kagan 2000; Pettit and Smith 2000). And similarly, it is rational

16

to deliberate about booking a hotel room using the satisficing heuristic just in case the expected utility of deliberating by satisficing is at least as great as the expected utility of deliberating in any other way. These assessments can come apart, for example because the cognitive costs of deliberation bear on the expected utility of decision procedures but not on the expected utility of hotel bookings.

Direct normative assessment enforces a sharp level-separation between higher- and lower-order decision problems (Thorstad 2021). This means that unlike indirect normative theories, the way of level separation will not collapse normative structure. We can easily countenance cases in which rational metacognition leads to irrational first-level deliberation and irrational base-level actions. If rational metacognizers will not intervene in low-stakes decisionmaking without substantial evidence that they are deliberating irrationally, then rational metacognition will sometimes involve letting irrational deliberation continue. Because each of these normative claims applies at a separate level, there is no pressure to deny any of these claims by collapsing levels.

Second, standards of correctness should be *maximizing*, not satisficing.[24] It is rational to do what is most valuable or choiceworthy, and irrational to do anything less.[25] We can soften this claim in the usual ways.[26] For example, we could hold that goodness or even rationality is a scalar property and that satisfactory actions may be very good or even highly rational, even if not the best or rational full-stop. But by refusing to weaken the standards of rationality, we also avoid the unpalatable normative consequences of

[24]To be clear, this is a claim about standards of correctness, not about the types of processes that meet these standards. Defenders of heuristic cognition have rightly noted that explicitly calculating the noncognitive actions which would maximize expected utility is not itself a value-maximizing decision procedure. This point can be generalized to yield a preliminary defense of the rationality of heuristic cognition. Thanks to two anonymous referees for raising this point.

[25]For this to be plausible, we need to take a broad view of standards of correctness for inquiry such that not only the attitudes that result from inquiry, but also features such as decision costs and the structure of the social (Berg and Watanabe 2020) and natural environment are relevant. So for example, if a heuristic comes within a just noticeable difference of nonheuristic processes in the accuracy or quality of the resulting decision (Berg 2008) then that heuristic may be more rational than the nonheuristic process once other factors are included to tip the balance.

[26]We may also weaken this claim so that it does not apply in cases which it is impossible to maximize what matters (Pollock 1983; Mintoff 1997), or we may stand our ground and retain maximizing standards of correctness even in such cases (Sorensen 2006). Thanks to a referee for pushing me to address this point.

satisficing views. We can say why it is irrational for a home-seller to turn down free money, whatever the amount: it would be better to take the money in order to benefit herself or others.[27]

For the way of level separation to be acceptable as a theory of bounded rationality, we will also have to say something about the sense in which the way of level separation is a theory of procedural rationality. Here is something that Herbert Simon might easily have said about procedural rationality, and which I think he should have said.

A large part of normative theory asks questions about *substantive rationality*: the rationality of base-level actions such as borrowing money. But it is also important to ask questions about *procedural rationality*: the rationality of decisionmaking procedures which produce base-level actions, or the rationality of higher-order decisionmaking procedures which guide lower-level decisionmaking procedures.[28] This is important because key cognitive bounds such as the cost of deliberation are felt most strongly at the procedural level.

Traditional normative theory is *outcome-focused*, asking questions primarily about substantive rationality and assuming as a default stance that these substantive questions will settle important normative questions in other areas. By contrast, the bounded tradition holds that normative theory should be *process-focused*, primarily asking and applying questions about procedural rationality.[29]

This is not merely a change in emphasis. Consider once again the woman who over-

---

[27]Some readers may wish to weaken this view in order to make it compatible with pluralistic approaches to rationality (Berg 2003). A plurality of normative standards may benefit groups and individuals, for example by increasing resilience to environmental shocks (Berg and Watanabe 2020) and increasing the speed of scientific discoveries. Some theses in the neighborhood of pluralism, such as the need to make room for conflicting or incommensurable goals, may already be captured by the present approach. Others may want to distinguish my approach from theories of structural or axiomatic rationality (Berg 2014), or a number of other competing approaches. On this view, the story told in this paper would hold true of one central and important species of rational assessment, and may have implications for others. Thanks to an anonymous referee for pressing me to address these points.

[28]The epistemological analogue of the procedural turn is the *zetetic turn* in the epistemology of inquiry. See Friedman (2020) and Thorstad (2021).

[29]Note that a process-focused approach to normative theorizing pronounces only on the questions normative theorists should focus on, and not yet on the answers we should give to those questions.

borrows as a result of rational patterns of attention allocation and as a result is evicted from her home. Now suppose you are a hard-line luck egalitarian. You think it is unacceptable for people to be deprived of homes as an outcome of anything less than their own irrational choices. But you think it is no injustice for people to wind up homeless as a result of their own choices.

Here you could take either an outcome-focused or a process-focused interpretation of irrational choice. On an outcome-focused interpretation, the woman has made an irrational choice, namely borrowing too much money. As a result, she has not been treated unfairly if she is left homeless.[30] On a process-focused interpretation, the woman has made a rational choice, namely allocating her attention as best she can to the many cognitive challenges that she faces. As a result, she will be treated unfairly unless she is housed. The third claim made by the way of level separation is that normative assessment should be process-focused. In most contexts, facts about procedural rationality deserve at least as much emphasis as facts about substantive rationality.

Summing up, the way of level separation advocates direct rather than indirect forms of normative assessment; maximizing rather than satisficing standards of correctness; and process-focused rather than outcome-focused normative theorizing. What light can the way of level separation shed on the paradoxes?

6 Solving two paradoxes of bounded rationality

6.1 The problem of rational irrationality

The problem of rational irrationality is that it sometimes seems rational for bounded agents to act irrationally on the basis of rational deliberation. We saw in Section 2 that a good solution to the problem of rational irrationality should do three things. First, it should dispel the problem by showing why it cannot be rational to act irrationally. Second,

---

[30]This is a version of the harshness objection (Voigt 2007) to luck egalitarianism. One of the upshots of my discussion is that a process-focused interpretation eases, but does not eliminate the harshness objection.

a good solution should explain why we were tempted to say that it can be rational to act irrationally. And third, a good solution should explain how neighboring statements made in the empirical literature can be true in a nonparadoxical sense.

On the first point, the way of level separation dissolves the problem of rational irrationality by applying standards of correctness directly at each level. Actions are rational if they meet these standards and irrational otherwise. The fact that an agent deliberated rationally cannot change the fact that she acted irrationally as a result.

On the second point, the way of level separation explains why we are tempted to talk of rational irrationality by emphasizing that normative assessment should be process focused. In an intellectual climate where most normative theory is outcome-focused, the claim that a bounded agent has acted irrationally sometimes obscures an important normative fact: that she did so as a result of rational deliberation. In our zeal to emphasize the fact that agents have deliberated rationally, we are tempted to go too far and claim that the actions which result from this deliberation are also rational. That is an understandable temptation which level separation trains us to resist.

On the third point, researchers in the heuristics and biases tradition sometimes say in one breath that typical heuristics are rational, but also that these heuristics violate normative standards.[31] The introduction to a leading anthology combines these claims in near-Knightian fashion.

> Although ... heuristics are distinguished from normative reasoning processes by biased judgments, the heuristics themselves are sensible estimation procedures that are by no measure 'irrational'. (Gilovich and Griffin 2002, p. 3).

In what sense can heuristic decision procedures be rational if they also violate normative

---

[31]Other traditions may be less prone to making such claims. Thanks to an anonymous referee for pressing me to clarify this point.

standards?[32] Level separation provides the natural answer.[33] The heuristics themselves are rational decision procedures. On occasion, the decisions that they produce violate rational norms on action. But this fact does not impugn the rationality of the heuristics themselves, because it does not undermine any of the grounds given for the rationality of heuristic decisionmaking in Section 2.1.

6.2   The regress problem

The regress problem is that bounded agents seem to be confronted with an infinite regress of higher-order decision problems. We saw in Section 2 that a good solution to the regress problem should do three things. First, it should say which of these decision problems a bounded agent actually faces. Second, it should reassure us that these problems will typically have at least one rational solution. And third, it should tell us how the type of decision problem described is something that a bounded agent could face.

On the first point, level separation holds that bounded agents are confronted with multiple decision problems: how to act, decide, metacognize and the like. More precisely, we face any of these problems in which there are options available to us in our capacity as rational agents. On many views, this means that we face only the first few levels of decision problems, either because we lack the capacity to act in response to higher-order considerations or else because we can only do so in a sense that is too subpersonal and insufficiently volitional or reason-responsive to count as an act of rational agency. But some Bayesians may think that we agentially confront each problem in an infinite hierarchy, and the way of level separation does not rule this view out in principle.

On the second point, level separation ensures that the multiplicity of decision problems

---

[32]One way to interpret this passage would be as evidence that axiomatic standards of rationality, which bounded agents often violate, are missing an important part of the story about human rationality.

[33]Recent Bayesian approaches to heuristic cognition suggest that maximizing standards of correctness may be integral to this answer as well. Bayesians increasingly argue that Bayesian normative principles, which are naturally seen as encoding maximizing standards of correctness, can be seen as high-level normative descriptions of computational problems, and that the most rational solutions to these problems are often fast-and-frugal heuristics (Sandborn and Chater 2016; Lieder and Griffiths 2020). But some defenders of fast-and-frugal heuristics may disagree with this interpretation (Gigerenzer and Selten 2001b).

does nothing to interfere with the existence of rational solutions to each problem. Of course, a lack of solutions may arise due to standard sources of rational dilemmas. And the rational solution to some problems may be to do nothing at all, for example because any type of twelfth-order deliberation would be too costly and unreliable to be worth engaging in. But each decision problem that agents face is a perfectly well-posed decision problem and there is no need to find an equilibrium or limiting solution to all of these problems at once in order to reassure ourselves that the problems are individually soluble.

On the third point, level separation helps us to see how the decision problems we face are problems that bounded agents could confront. As I have said, bounded agents may only face the first few levels of decision problems because we are incapable of agentially confronting higher-order decision problems. This fact does not mean that the decision problems we face are ill-formed or lead to regress. And as we confront these decision problems, the way of level separation does not demand that agents represent the problems to ourselves in any complex way. For example, an increasingly popular view of metacognition holds that we solve $\mathcal{D}_2$ largely by attending to a class of *noetic feelings* produced by our experience of deliberation, rather than by forming and analyzing metacognitive beliefs about our own deliberative processes. It is not especially difficult to see how bounded agents could be capable of forming and responding to such feelings, and a good deal of work has been done to show how humans (Arango-Muñoz and Michaelian 2014; Proust 2013) and even perhaps nonhumans (Hampton 2009) can use these feelings to metacognize effectively.

7   Conclusion

This paper considered two paradoxes of bounded rationality: the problem of rational irrationality (Section 2.1) and the regress problem (Section 2.2). I looked at two classes of solutions to the paradoxes: the way of indirection, which lets the outcomes of inquiry inherit the rational status of the processes that produced them (Section 3); and the way of

weakening, which weakens traditional norms to make them more achievable for bounded agents (Section 4). I used this discussion to develop a third solution to the paradoxes: the way of level separation (Sections 5-6). The way of level separation combines direct normative assessment with a maximizing standard of correctness, as well as a process-focused approach to normative theorizing.

This discussion raises important questions for further discussion. For example, how does the way of level separation interact with axiomatic, structural (Worsnip forthcoming), or coherence-based (Staffel 2020) theories of rationality? And what can we learn about recent appeals to level separation in the epistemology of inquiry (Thorstad 2021, forthcoming) by thinking about level separation in the theory of bounded rationality? But perhaps the most important lesson is methodological: empirically-grounded philosophical reflection on fundamental normative concepts may serve as an important source of new theories of bounded rationality, as well as a fertile source of novel solutions to normative puzzles.

References

Arango-Muñoz, Santiago and Michaelian, Kourken. 2014. "Epistemic feelings and epistemic emotions." *Philosophical Inquiries* 2:97–122.

Bago, Bence and De Neys, Wim. 2017. "Fast logic? Examining the time course assumption of dual process theory." *Cognition* 158:607–617.

Berg, Nathan. 2003. "Normative behavioral economics." *Journal of Socio-Economics* 32:411–27.

—. 2008. "Rational ignoring with unbounded cognitive capacity." *Journal of Economic Psychology* 29:792–809.

—. 2014. "The consistency and ecological rationality approaches to normative bounded rationality." *Journal of Economic Methodology* 21:375–95.

Berg, Nathan and Watanabe, Yuki. 2020. "Conservation of behavioral diversity: On nudging, paternalism-induced monoculture, and the social value of heterogeneous beliefs and behavior." *Mind and Society* 19:103–20.

Bermúdez, José. 2009. *Decision theory and rationality*. Oxford University Press.

Birnbaum, Michael H. 2008. "New paradoxes of risky decision making." *Psychological Review* 115:463–501.

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

Bonner, Carissa and Newell, Ben. 2010. "In conflict with ourselves? An investigation of heuristic and analytic processes in decision making." *Memory and Cognition* 38:186–196.

Bookstaber, Rick and Langsam, Joseph. 1985. "On the optimality of coarse behavior rules." *Journal of Theoretical Biology* 116:161–93.

Bradley, Ben. 2006. "Against satisficing consequentialism." *Utilitas* 18:97–108.

Byron, Michael. 1998. "Satisficing and optimality." *Ethics* 109:67–93.

Caplin, Mark Dean and Martin, Daniel. 2011. "Search and satisficing." *American Economic Review* 101:2899–2922.

Cherniak, Christopher. 1981. "Minimal rationality." *Mind* 90:161–183.

—. 1986. *Minimal rationality*. MIT Press.

Chickering, David Maxwell. 1996. "Learning Bayesian networks is NP-complete." In Doug Fisher and Hans-J. Lenz (eds.), *Learning from data*, 121–130. Springer.

Dagum, Paul and Luby, Michael. 1993. "Approximating probabilistic inference in Bayesian belief networks is NP-hard." *Artificial Intelligence* 60:141–153.

De Neys, Wim and Glumicic, Tamara. 2008. "Conflict monitoring in dual process theories of thinking." *Cognition* 106:1248–1299.

De Neys, Wim, Rossi, Sandrine, and Houdé, Olivier. 2013. "Bats, balls, and substitution sensitivity: Cognitive misers are no happy fools." *Psychonomic Bulletin and Review* 20:269–273.

Desvousges, William, Johnson, Reed, Dunford, Richard, Boyle, Kevin, Hudson, Sara, and Wilson, Nicole. 1992. *Measuring nonuse damages using contingent valuation: An experimental evaluation of accuracy*. Research Triangle Institute.

Elster, Jon. 1983. *Sour grapes: Studies in the subversion of rationality*. Cambridge University Press.

Frederick, Shane. 2005. "Cognitive reflection and decision making." *Journal of Economic Perspectives* 19:25–42.

Fredrickson, Barbara L and Kahneman, Daniel. 1993. "Duration neglect in retrospective evaluations of affective episodes." *Journal of Personality and Social Psychology* 65:45–55.

Friedman, Jane. 2020. "The epistemic and the zetetic." *Philosophical Review* 129:501–536.

Gangemi, Amelia, Bourgeois-Gironde, Sacha, and Mancini, Francesco. 2015. "Feelings of error in reasoning - in search of a phenomenon." *Thinking and Reasoning* 21:383–393.

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

Gigerenzer, Gerd and Brighton, Henry. 2009. "Homo heuristicus: Why biased minds make better inferences." *Topics in Cognitive Science* 1:107–143.

Gigerenzer, Gerd and Gaissmaier, Wolfgang. 2011. "Heuristic decision making." *Annual Review of Psychology* 62:451–482.

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.

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.

Hampton, Robert R. 2009. "Multiple demonstrations of metacognition in nonhumans: Converging evidence or multiple mechanisms?" *Comparative Cognition and Behavior Reviews* 4:17–28.

Hertwig, Ralph and Engel, Christoph (eds.). 2021. *Deliberate ignorance*. MIT Press.

Jackson, Simon, Kleitman, Sabina, Howie, Pauline, and Stankov, Lazar. 2016. "Cognitive abilities, monitoring confidence, and control thresholds explain individual differences in heuristics and biases." *Frontiers in Psychology* 7:e1559.

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

Johnson, Eric, Tubau, Elisabet, and De Neys, Wim. 2016. "The doubting system 1: Evidence for automatic substitution sensitivity,." *Acta Psychologica* 164:56–64.

Kagan, Shelly. 2000. "Evaluative focal points." In Brad Hooker, Elinor Mason, and Dale E. Miller (eds.), *Morality, rules and consequences*, 134–55. Edinburgh University Press.

Kahneman, Daniel and Frederick, Shane. 2002. "Representativeness revisited: Attribute substitution in intuitive judgment." In Thomas Gilovich, Dale Griffin, and Daniel Kahneman (eds.), *Heuristics and biases: The psychology of intuitive judgment*, 48–81. Cambridge University Press.

Katsikopoulos, Konstantinos, Şimşek, Özgür, Buckmann, Marcus, and Gigerenzer, Gerd. 2020. *Classification in the wild: The art and science of transparent decision making*. MIT Press.

Knight, Frank. 1921. *Risk, uncertainty and profit*. University of Chicago Press.

Lieder, Falk and Griffiths, Thomas. 2017. "Strategy selection as rational metareasoning." *Psychological Review* 124:762–794.

—. 2020. "Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources." *Behavioral and Brain Sciences* 43:E1.

Lin, Hanti. 2014. "On the regress problem of deciding how to decide." *Synthese* 191:661–670.

Lipman, Barton. 1991. "How to decide how to decide how to ...: Modeling limited rationality." *Econometrica* 59:1105–1125.

Lord, Errol. 2017. "What you're rationally required to do and what you ought to do (are the same thing!)." *Mind* 126:1109–1154.

Mani, Anandi, Mullainathan, Sendhil, Shafir, Eldar, and Zhao, Jiaying. 2013. "Poverty impedes cognitive function." *Science* 341:976–980.

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

Mintoff, Joe. 1997. "Slote on rational dilemmas and rational supererogation." *Erkenntnis* 46:111–26.

Mongin, Philippe. 2000. "Does optimization imply rationality?" *Synthese* 124:73–111.

Mongin, Philippe and Walliser, Bernard. 1988. "Infinite regressions in the optimizing theory of decision." In Bertrand Munier (ed.), *Risk, decision and rationality*, 435–457. Springer.

Mulgan, Tim. 2001. "How satisficers get away with murder." *International Journal of Philosophical Studies* 9:41–46.

Newark, Daniel. 2018. "Leadership and the logic of absurdity." *Academy of Management Review* 43:198–216.

Pettit, Philip. 1984. "Satisficing consequentialism." *Proceedings of the Aristotelian Society, Supplement* 58:165–176.

Pettit, Philip and Smith, Michael. 2000. "Global consequentialism." In Brad Hooker, Elinor Mason, and Dale E. Miller (eds.), *Morality, rules and consequences*, 121–133. Edinburgh University Press.

Pollock, John. 1983. "How do you maximize expectation value?" *Noûs* 17:409–21.

Proust, Joëlle. 2013. *The philosophy of metacognition*. Oxford University Press.

Sandborn, Adam and Chater, Nick. 2016. "Bayesian brains without probabilities." *Trends in Cognitive Sciences* 20:883–893.

Savage, Leonard. 1954. *The foundations of statistics*. Wiley.

Schooler, Lael J. and Hertwig, Ralph. 2005. "How forgetting aids heuristic inference." *Psychological Review* 112:610–628.

Selten, Reinhard. 1998. "Aspiration adaptation theory." *Journal of Mathematical Psychology* 42:191–214.

Shah, Anuj, Mullainathan, Sendhil, and Shafir, Eldar. 2012. "Some consequences of having too little." *Science* 338:682–685.

Simon, Herbert. 1955. "A behavioral model of rational choice." *Quarterly Journal of Economics* 69:99–118.

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

Slote, Michael. 1984. "Satisficing consequentialism." *Proceedings of the Aristotelian Society, Supplement* 58:139–163.

—. 2004. "Two views of satisficing." In Michael Byron (ed.), *Satisficing and maximizing: Moral theorists on practical reason*, 14–29. Cambridge University Press.

Smith, Holly. 1991. "Deciding how to decide: Is there a regress problem?" In Michael Bachrach and S.L. Hurley (eds.), *Foundations of decision theory*, 194–217. Blackwell.

Sorensen, Roy. 2006. "Originless sin: Rational dilemmas for satisficers." *Philosophical Quarterly* 56:213–23.

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

Swanton, Christine. 1993. "Satisficing and virtue." *Journal of Philosophy* 90:33–48.

—. 2004. "Satisficing and perfectionism in virtue ethics." In Michael Byron (ed.), *Satisficing and maximizing: Moral theorists on practical reason*, 176–189. Cambridge University Press.

Thompson, Valerie and Johnson, Stephen. 2014. "Conflict, metacognition, and analytic thinking." *Thinking and Reasoning* 20:215–244.

Thompson, Valerie, Turner, Jamie, and Pennycook, Gordon. 2011. "Intuition, reason, and metacognition." *Cognitive Psychology* 63:107–140.

Thorstad, David. 2021. "Inquiry and the epistemic." *Philosophical Studies* 178:2913–28.

—. forthcoming. "The accuracy-coherence tradeoff in cognition." *British Journal for the Philosophy of Science* 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.

Travers, Eoin, Rolison, Jonathan, and Feeney, Aidan. 2016. "The time course of conflict on the Cognitive Reflection Test." *Cognition* 150:109–118.

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

Vassilakis, Spyros. 1992. "Some economic applications of Scott domains." *Mathematical Social Science* 24:173–208.

Voigt, Kristin. 2007. "The harshness objection: Is luck egalitarianism too harsh on the victims of option luck?" *Ethical Theory and Moral Practice* 10:389–407.

Worsnip, Alex. forthcoming. "Making space for the normativity of coherence." *Noûs* forthcoming.