

Satisficing, maximizing and bounded rationality

David Thorstad

Forthcoming in *Phil Quarterly*

Penultimate draft, please cite published version

Abstract

It is often argued that maximizing normative theories are inappropriate for bounded agents. I argue that leading motivations for theories of bounded rationality are substantially compatible with some traditional maximizing normative theories. This finding expands the audience for theories of bounded rationality by making room for advocates of maximizing normative theories, and has implications for knowledge-action principles, the relationship between descriptive and normative theorizing, and the definitions of central tenets of bounded rationality.

Keywords: Satisficing, maximizing, bounded rationality, knowledge and action, minimal rationality, rationality and cognitive science.

1 Introduction

Bounded agents satisfice. They do not maximize. The first sentence expresses a truism. The second is often treated as such by leading scientists:

Information-processing systems typically need to *satisfice* rather than optimize.
(Gigerenzer and Goldstein 1996, p. 651)

When we leave the thin theory of rationality, the link between rationality and optimality is completely broken. (Elster 1983, p. 18)

Human rationality [is] non-maximizing and non-optimizing, but rather bounded, heuristic and satisficing. (Viale 2021, p. 1)

Herbert Simon once jokingly threatened to sue people who used his term ‘bounded rationality’ to refer to any form of optimization (Gigerenzer 2004).

A similar view is expressed by many philosophers:

A ‘satisficer’ content with good enough ... does not seek to maximize (optimize) his expectations ... Such a person clearly fails to exemplify the maximizing and optimizing model of individual rationality. (Slote 1984, p. 143)

The standard model of practical rationality that we use in philosophy is that of expected utility maximisation. But there are both theoretical and experimental reasons to think that this is not the right model for [all] choices. (Weatherson 2024, p. 162)

Rather than seeking to optimize in each instance, Ecological Rationality also endorses procedures designed to do good enough (“satisficing”). (Schmidt 2019, p. 522)

My goal in this paper is to show that contemporary knowledge about bounded rationality is surprisingly compatible with a traditional philosophical form of maximization.

Here is the plan. Section 2 clarifies my project and the relevant notion of maximization. Section 3 discusses the philosophical and scientific importance of the compatibility between bounded rationality and maximization. Sections 4-8 discuss leading reasons why bounded rationality has been taken to be incompatible with maximization. Section 9 concludes.

2 Clarifications

Before beginning, I want to make three clarificatory remarks. First, it is important to distinguish the descriptive project of saying how bounded agents in fact cognize from the normative project of assessing their cognition (Bermúdez 2009). My project is normative, rather than descriptive. However, many bounded rationality theorists (Gigerenzer and Selten 2001b; Sturm 2012; Thorstad 2024) share with Bayesian rational analysts (Anderson 1990; Icard 2018) the assumption that human cognition is often fairly rational. If that is right, then normative theorizing can often provide a good defeasible starting point for

descriptive theorizing. Nevertheless, it remains important to clearly distinguish between normative and descriptive projects.

Second, for concreteness it will help to fix a particular maximizing normative standard. In this paper, I am interested in the prospects for a traditional form of maximizing consequentialism. Say that option X is maximal for S at t if no alternative available to S at t is better than X . Maximizing consequentialism holds:

(Maximizing Consequentialism) For all agents S , acts X and times t :

(Maximizing Criterion of Correctness) S may (must) X at t just in case X is maximal (uniquely maximal) for S at t .

(Information-Sensitive Value) The value of S 's X -ing at t is the expected value of X 's consequences, where the expectation in question is taken relative to evidential probabilities. That is, $V(X) = \sum_w Pr_E(w)V(X, w)$, where Pr_E are the evidential probabilities generated by S 's total evidence at t .¹

I suspect that the discussion in this paper can generalize to many, but perhaps not all leading maximizing views.²

Third, some elements of my project may be partly anticipated by earlier work. Many theorists including Herbert Simon have admitted that the line between maximizing and satisficing approaches is less sharp than we might suppose.³ Some have gone further and discussed the possibility of taking maximizing theories on board as normative standards, which is the proposal that I endorse.⁴ I think that this background helps to reduce the

¹For discussion of information-sensitivity see Kolodny and MacFarlane (2010). Readers who prefer objective or subjective readings of deontic modals are invited to explore the implications of these readings throughout the paper.

²A referee notes that this is a maximality rule in Sen's (1997) sense, which might be contrasted to other conceptions of maximization.

³For example, Simon (1978, p. 8) sets out "to show how large a plot of common ground is shared by optimizing and satisficing analyses." See also Pettit (1984).

⁴For example, Thomas Sturm (2012, pp. 77-8) notes that "formal rules, such as Bayes' theorem or

seeming heterodoxy of my approach, without eliminating the need for a new and detailed defense. We saw in Section 1 that some elements of my proposal are deeply controversial, and we will see in Section 3 why it matters. For this reason, it is important to present a careful defense of the view and to take time exploring its implications.

Summing up, my project is the normative project of saying how bounded agents should cognize. I concentrate on the prospects for a traditional form of maximizing consequentialism. I claim that reconciling maximizing consequentialism with bounded rationality is important. The next section says why this is so.

3 Why it matters

Why would it matter if maximizing normative theories were compatible with considerations of bounded rationality? This would matter because many philosophical (Section 3.1) and scientific projects (Section 3.2) are motivated or defined by the assumption that maximization is incompatible with bounded rationality. If this assumption is wrong, then leading programs need new motivations and definitions. I revisit these implications in Section 9.

3.1 Philosophical projects

3.1.1 Minimal rationality

One of the best-known philosophical approaches to bounded rationality is Christopher Cherniak's (1981; 1986) minimal rationality. Cherniak argues that agents cannot be subject to maximizing normative standards, and hence must be subject to weaker minimal standards.

Say that an act is *apparently appropriate* if it maximizes expected utility given an agent's

principles of optimization (such as maximizing expected utility) continue to play an important role even within attempts to construct fast and frugal heuristics, namely as a normative standard." However, Sturm also holds that such norms play different functional roles from traditional norms of bounded rationality.

beliefs and desires. Instead of requiring agents to take all apparently appropriate actions:

(Full Rationality) If *A* has a particular belief-desire set, *A* would undertake *all* and only actions that are apparently appropriate. (Cherniak 1986, p. 8)

minimal rationality requires agents to take some apparently appropriate actions:

(Minimal Rationality) If *A* has a particular belief-desire set, *A* would undertake some, but not necessarily all, of those actions that are apparently appropriate. (Cherniak 1986, p. 8)

Minimal Rationality is defended on the grounds that maximizing standards such as Full Rationality cannot be appropriate for bounded agents when compliance is intractable or impossible. Minimal rationality has been extended to support cognate approaches to zetetic epistemology (Hoek forthcoming), theories of philosophical naturalism (Kitcher 1992), and assessments of animal cognition (Dretske 2006).

If bounded rationality does not support rejecting maximizing normative standards, then new motivations are needed for minimal rationality and the projects it supports.

3.1.2 Knowledge and action

Many philosophers defend principles linking knowledge to rational action (Fantl and McGrath 2009). For example, Brian Weatherson defends:

(Knowledge Allows Exclusion) If a chooser knows that a possibility does not obtain, then it is permissible to use a decision table where that possibility is excluded. (Weatherson 2024, p. 94)

A common objection to principles such as Knowledge Allows Exclusion is that they conflict with leading theories of expected utility maximization. For example, Maximizing Consequentialism conflicts with Knowledge Allows Exclusion if agents can know propositions with evidential probability below one.

Weatherson is aware of this objection and argues against maximizing normative theories, drawing extensively on considerations of bounded rationality such as cognitive costs and choice paralysis. If these considerations are compatible with Maximizing Consequentialism, then Weatherson's arguments will not go through.

3.1.3 Satisficing consequentialism

We will see in Section 5 that considerations of bounded rationality have been taken to support alternatives to maximizing consequentialism, such as Michael Slote's (1984; 2004) satisficing consequentialism:

(Satisficing Criterion of Correctness) *S* may (must) *X* at *t* just in case *X* is sufficiently good (the only sufficiently good option) for *S* at *t*.

More generally, consequentialist normative theories have often been rejected on the grounds that they are computationally intractable for bounded agents. For example, Rawls took the computational demands of consequentialism as a decisive objection to consequentialist normative theories (Rawls 1971; Simmons 2010). If Maximizing Consequentialism is compatible with considerations of bounded rationality, then these arguments do not work.

3.2 Scientific projects

Bounded rationality is an area in which philosophical research can usefully inform scientific research. Many scientific programs use submaximization in order to motivate bounded rationality (Section 3.2.1), define key concepts (Section 3.2.2) and resist Bayesian approaches (Section 3.2.3). If bounded rationality is not at odds with maximization, then revisions may be called for.

3.2.1 Motivating bounded rationality

Many of the best-known arguments for theories of bounded rationality begin with criticisms of maximizing approaches. John Elster's (1983) *general argument for satisficing* relies on the claim that maximizing approaches are vulnerable to infinite regress. Reinhard Selten (2001) motivates his aspiration adaptation approach to satisficing by arguing that maximization is infeasibly complex. And Herbert Simon holds that we must study bounded rationality because it studies "the behavior of human beings who satisfice because they have not the wits to maximize" (Simon 1957, p. xxiv). If bounded rationality is not at odds with maximizing approaches, then these motivations do not succeed.

3.2.2 Definitional challenges

The notion of sub-maximization is often used to define bounded rationality and related notions. For example, the editors of a recent anthology on the ecological approach to bounded rationality define it as asking a question which arises only once maximization is ruled out:

Ecological rationality's question: Given a problem in an environment, which strategies *should* humans rely on when optimization is not feasible? (Todd and Gigerenzer 2012, p. 489)

And Gerd Gigerenzer defines heuristics as submaximizing processes:

A heuristic is a mental process that ignores part of the available information and does not optimize. (Gigerenzer 2010, p. 529)

If the opposition between bounded rationality and maximization is broken, then these definitions require reformulation.

3.2.3 Bayesian approaches

There has been a resurgence of Bayesian approaches to bounded rationality (Dorst 2023; Icard ms; Wheeler 2020). Recent Bayesian paradigms such as cognitively bounded rational

analysis (Howes et al. 2009), boundedly rational analysis (Icard 2018) and resource-rational analysis (Lieder and Griffiths 2020) are used to construct detailed models of the demands facing bounded agents.

Despite this, leading scientists have roundly criticized Bayesian models as inappropriate for bounded agents because they assume that agents maximize objects such as utility functions (Elster 1983; Gigerenzer and Selten 2001b; Hertwig et al. 2022). For example, the editorial introduction to a leading anthology begins its discussion of what bounded rationality is not by clarifying that bounded rationality should not be confused with Bayesian models of constrained optimization (Gigerenzer and Selten 2001b, p. 4). Similar charges are sometimes leveled by philosophers (Pasnau 2013). If maximization cannot be ruled out as an approach to bounded rationality, then these objections to Bayesian theories of bounded rationality will not work.⁵

3.3 Taking stock

So far, we have seen that many philosophical (Section 3.1) and scientific projects (Section 3.2) are motivated or defined by the assumption that maximization is incompatible with bounded rationality. On what grounds is maximization taken to be incompatible with bounded rationality? Sections 4-8 discuss five leading arguments and suggest that they do not show maximizing normative theories to be substantially incompatible with considerations of bounded rationality.

4 Modest changes

At least two familiar complaints against maximizing theories can be handled through at most modest changes in the view.

⁵This is not to say that other objections, such as inadequate or overfitted modeling, may not apply.

4.1 All options are dominated

In some cases, there is no best option for agents to take. Consider:

(Everbetter Wine) An immortal possesses a bottle of wine which gets better every year. When should she drink it? (Pollock 1983)

(Finite Bliss) God asks you to name a number n . You will be rewarded with n years of bliss, followed by annihilation. Which number should you name? (Tucker 2016)

For any n , the option of drinking the wine in year n or asking for n years of bliss is suboptimal, because it would be better to opt for $n + 1$ years of bliss or wine maturation. But this means that no matter what the agent chooses, her choice will be worse than another option she could have selected. This means that maximizing consequentialism rules out all possible options in both cases. Does this pose a decisive objection to maximizing consequentialism (Elster 1983; Schmidtz 2004)?

We can react to such cases in two ways. On the one hand, some theorists have recommended holding our ground, retaining maximizing consequentialism and treating the cases as rational dilemmas in which no option is permissible (Sorensen 2006). This would require no change to maximizing consequentialism.

One way to defend this stalwart reaction is to note that the cases in question threaten not only maximizing consequentialism, but also much weaker principles. For example, consider:

(Very Weak Statewise Dominance) For any agent S and options o, o' available to S , if o is guaranteed to provide a strictly better outcome than o' in every state of the world, then o' is impermissible.

The intuition behind Very Weak Statewise Dominance is that it would be impermissible to choose o' because it would be better to choose o instead. Dominance principles are widely

regarded as among the most minimal and normatively compelling decision-theoretic principles, and Very Weak Statewise Dominance is one of the weakest dominance principles on offer. But note that all options in Finite Bliss and Everbetter Wine are dominated, so that if we do not treat these cases as rational dilemmas, we need not only to reject maximizing consequentialism, but also a great number of decision-theoretic principles which include or imply Very Weak Statewise Dominance. And that is a very costly thing to do.

On the other hand, we might take cases such as Finite Bliss and Everbetter Wine to show that maximizing consequentialism is inapplicable in cases where there is no best option. That would involve a gentle relaxation of maximizing consequentialism, but it would not be a devastating change, since cases like Everbetter Wine and Finite Bliss are rarely encountered in our daily lives. Nor would it be an unprincipled move, for when asked why we decided to relax maximizing consequentialism in exactly these cases, we would reply that we relaxed maximizing consequentialism because it cannot be coherently applied to these cases.⁶ At the same time, it is important to recall that relaxing maximizing consequentialism comes with a number of costs. In this case, we would also have to deny Very Weak Statewise Dominance, and that is something we may not be willing to do.

4.2 Incommensurability

Some theorists object that maximizing theories do not allow for the possibility of incommensurable options (Gigerenzer and Selten 2001; Klein 2001; Schmitz 1995). Here the problem cannot be with the Maximizing Criterion of Correctness:

(Maximizing Criterion of Correctness) *S* may (must) *X* at *t* just in case *X* is maximal (uniquely maximal) for *S* at *t*.

The Maximizing Criterion of Correctness allows for incommensurable options by allowing multiple options to be maximal. The Maximizing Criterion of Correctness would put pressure against incommensurability if it were not formulated in terms of maximality:

⁶Alan Hajek (2014) defends a similar position.

(Restrictive Maximizing Criterion of Correctness) S may (must) X at t just in case X is at least as good as (better than) all alternatives to X for S at t .

The Restrictive Maximizing Criterion of Correctness treats cases of incommensurability as rational dilemmas, since incommensurable maximal options will not be better than, or at least as good as one another. That may be unappealing. It is precisely to block this consequence that many theorists prefer the Maximizing Criterion of Correctness to the Restrictive Maximizing Criterion of Correctness.

We are not out of the woods yet, for incommensurability does conflict with my statement of information-sensitivity:

(Information-Sensitive Value) The value of S 's X -ing at t is the expected value of X 's consequences, where the expectation in question is taken relative to evidential probabilities. That is, $V(X) = \sum_w Pr_E(w)V(X, w)$, where Pr_E are the evidential probabilities generated by S 's total evidence at t .

However, incommensurability is not incompatible with the core insight of information-sensitivity:

(Value is Information-Sensitive) The value of S 's X -ing at t is assessed against S 's total evidence.

The core claim that Value is Information-Sensitive is compatible with incommensurability, but it will have to be given a different decision-theoretic articulation.

Our initial statement of Information-Sensitive Value built in a second claim, specifying the decision-theoretic relationship between the value of an action and the value of its consequences:

(Value is Expectational) The value of S 's X -ing at t is its expected value according to a relevant probability function Pr and utility function V .

Together with the claim that Value is Information-Sensitive, the claim that Value is Expectational yields Information-Sensitive Value by grounding the decision-theoretic claims made by Information-Sensitive Value.

The claim that Value is Expectational is incompatible with most views that allow value incommensurability, since the resulting decision theories cannot be given an expected value representation. Precisely for this reason, the claim that Value is Expectational must be replaced by an alternative decision theory supplied by the incommensurability theorist. But this is not particularly surprising, because to ask the maximizing theorist to make room for incommensurability just is to ask her to make room for alternatives to expected value theory. The maximizing theorist may not think that we need to do this, but she can make room by doing precisely as requested, removing the claim that Value is Expectational and retaining the rest of her theory.

In this section, we considered how maximizing consequentialists should react to cases in which there is no best option, because all options are dominated. We saw that they might plausibly treat the cases as rational dilemmas by appealing to a minimal dominance principle. Alternatively, consequentialists could treat these cases as a rare exception to maximizing normative theories. That would be disappointing, perhaps, but not a substantial change or the sort of thing to scuttle the broader theory. In the next section, I discuss a significant change that has been proposed to maximizing consequentialism. I argue that traditional philosophical ways of motivating this change are orthogonal to discussions of bounded rationality.

5 Orthogonal challenges

Suppose you are at the grocery store choosing among a shelf of artisanal vinegars. Making a satisfactory choice is easy: just choose any bottle that is in-budget and has a few reliable indicators of quality. But making an optimal choice is very hard. Are you really required to choose the best, or expectedly best bottle of vinegar?

Michael Slote (1984; 2004) takes the moral of such cases to be that the Maximizing Criterion of Correctness built into maximizing consequentialism should be replaced with a weaker satisficing criterion:

(Satisficing Criterion of Correctness) *S* may (must) *X* at *t* just in case *X* is sufficiently good (the only sufficiently good option) for *S* at *t*.

Slote took himself to be motivated by Simon's work on bounded rationality. But on reflection, it is not clear that Slote's proposal or the bulk of the ensuing philosophical debate (Byron 2004; Tucker 2016; Yetter Chappell 2019) turns strongly on considerations of bounded rationality.

Consider one of the cases that Slote (1984) uses to motivate his Satisficing Criterion of Correctness.

(Home-Seller) A home-seller sets the price of her home at a satisfactorily high number. When asked why she didn't ask for more, she replies that she is a satisficer: for her, enough is enough.

Importantly, Slote asks us to imagine that the home-seller is not concerned with the probability of making a sale, the time and effort involved in making a sale, or other traditional mainstays of bounded rationality. Slote's home-seller is not concerned that she could not get more money, or that it would be costly to do so. She is simply uninterested in receiving more money. The Satisficing Criterion of Correctness is meant to explain why it is permissible for Slote's home-seller to do this.

Bounded rationality theorists can respond to such cases in a variety of ways. Most directly, they might deny that it is rational for Slote's home-seller to turn down free money (Dreier 2004; Thorstad 2022). On this view, it is better to have more money than less, if nothing else because the extra money can be given to charity. But more to the point, they should note that Home-Seller is deliberately constructed to rule out most leading considerations of bounded rationality. Slote directly stipulates that the home-seller is not

concerned with time, effort, or other bounds. This suggests that Slote aimed to capture a motivation for satisficing distinct from bounded rationality. What else might Slote have been after?

Some readers have thought that Slote was making a point about the virtue of moderation: a moderate individual would have moderate desires, and hence would not want excessive amounts of money (Swanton 2004). Indeed, Slote (2001) himself later suggests this interpretation, linking moderation to a species of Stoic virtue. This claim would not pose any threat to the Maximizing Criterion of Correctness, which is a claim about rightness and not a claim about virtue.⁷

Others have taken Slote to be making a point about the demandingness of consequentialism: it is wrong to demand too much sacrifice of agents. For example, Slote holds that it is permissible for a hotelier to give a free room to a poor family, but not to give them the best room. Limiting the demandingness of consequentialism is a live challenge to maximizing consequentialism which runs deep through philosophical debates about satisficing (Yetter Chappell 2019). But it is not clear that this challenge turns on considerations of bounded rationality: it is fully possible to defend theories of bounded rationality on which morality is maximally demanding, or theories of unbounded rationality on which morality is less demanding, and indeed many theorists have done both.

Perhaps the gap between bounded rationality and satisficing consequentialism is clearest on views such as Richard Yetter-Chappell's (2019) willpower satisficing. This holds, relative to an effort ceiling X , that:

(Willpower-Satisficing Consequentialism) An act A is permissible iff it produces no less utility than any alternative action the agent could perform at the cost of up to either (i) X willpower, or (ii) the willpower cost of A , whichever is greater. (Yetter Chappell 2019, p. 255)

⁷This claim would only threaten the Maximizing Criterion of Correctness on views which used claims about virtue to ground claims about right action (Hursthouse 2001). Even then the threat is not so clear: absent a strong unity of virtues thesis (Penner 1973), from the fact that an action would exhibit one virtue, such as moderation, it does not follow that a virtuous person would do it.

Whatever the fate of such views, we should not take them to be aimed at capturing the full demands of bounded rationality. Otherwise, we would be saddled with the implausible consequence that limitations on willpower are the only bounds that matter rationally, whereas limitations on computation, attention, memory and the like do not. This stands in stark contrast to traditional theories of bounded rationality.

In this section, we have seen that a traditional way of motivating a Satisficing Criterion of Correctness does pose a challenge to maximizing consequentialism, but that this challenge is largely orthogonal to debates about bounded rationality. It is best understood as a debate about something else, such as the demandingness of morality or the virtue of moderation. The next section considers how maximizing views make sense of the rationality of heuristic cognition.

6 Level separation

One of the most widely accepted claims about bounded rationality is:

(Heuristic Rationality) It is often rational for agents to make decisions using a toolbox of fast-and-frugal heuristics (Gigerenzer and Selten 2001; Gigerenzer and Gaissmaier 2011).

Heuristics do not involve explicitly calculating expected values. Heuristics often make decisions based on a small subset of the available evidence, or draw only some of the many inferences warranted by that evidence. But maximizing consequentialism seems to require agents to explicitly calculate the expected values of all available options before acting. If that is true, then maximizing consequentialism conflicts with Heuristic Rationality. How can maximizing consequentialism be squared with Heuristic Rationality?

6.1 Level separation

To see the way forward, consider a traditional complaint against maximizing consequentialism. Maximizing consequentialism seems to require agents to explicitly calculate the expected utility of every option they face. If this is right, then it raises two objections.

The first is the *computational demandingness objection* (Williams 1973a). Explicitly calculating expected utilities is often expensive, and sometimes entirely beyond an agent's ability. When calculating expected utilities is unduly expensive, maximizing consequentialism seems to demand cognitive wastefulness. And when agents are unable to calculate expected utilities, maximizing consequentialism threatens to violate the principle that ought implies can.

A related objection is the *alienation objection* (Stocker 1976; Williams 1973b). Agents who always make decisions by coldly calculating expected utilities will alienate those around them. Their families, friends and kin will feel unloved. Society at large will distrust them. For these reasons, it would often be better for agents to make decisions without calculating expected utilities, perhaps allowing themselves to be moved to act directly out of love. But if that is right, then consequentialism appears self-defeating. Maximizing consequentialism says to always do what is best, but what's best is not always to calculate expected utilities.

The traditional response to these worries is to posit a sharp *level separation* (Kagan 2000; Parfit 1984; Thorstad 2024) between two questions: the lower-order question of which nondeliberative acts an agent ought to take, and the higher-order question of how she ought to deliberate about which nondeliberative act to take. For example, it is one question whether I ought to buy coffee and another question entirely how I ought to deliberate about whether to buy coffee. We can apply maximizing consequentialism directly to both questions: I ought to buy coffee just in case it is expectedly best to buy coffee, and I ought to deliberate using the process which it would be expectedly best for me to use.

Level separation solves the computational demandingness objection by letting costs

and abilities bear directly on the question of how I ought to deliberate. If it is too costly to explicitly calculate the expected utility of buying coffee, then precisely for that reason I should use a simpler decision procedure. And if I cannot explicitly calculate the expected utility of buying coffee, then explicitly calculating expected utilities is not an option for me, and hence I cannot be required to do it.

Level separation solves the alienation objection by letting alienation bear on the rightness of decision procedures. If the decision procedure of coldly calculating expected utilities would breed hurt and distrust, then that is a strong reason not to use it. It might be better to decide by allowing myself to be moved directly out of love. In this case, maximizing consequentialism says that I should decide in this way.

One feature of level separation is that it allows for mismatch between the nondeliberative actions I ought to take and the outputs of the deliberative processes I ought to use. It might be that I ought to buy coffee, because it is expectedly best to buy coffee. It might also be that I ought to deliberate about whether to buy coffee using some heuristic *H*, because *H* is the expectedly best decision process given a full specification of its speed, cost, reliability and other factors. And it might be that although *H* is generally reliable, *H* will, alas, misfire in this case, and output the decision not to buy coffee.

This type of *level tension* has long been noted as a feature of the orthodox consequentialist view (Kagan 2000; Parfit 1984; Thorstad 2024). It may make some theorists uncomfortable, but it has been defended at length, and some theorists have gone further, claiming that the ability to make sense of level tension is a virtue of the consequentialist view (Driver 2001; Thorstad 2024). Because my interest in this paper is not to throw sparks on old debates, but rather to evaluate the novel contributions of bounded rationality, I want to take the orthodox consequentialist solution on board and ask how that solution, if sound, could make sense of Heuristic Rationality.

6.2 Applying level separation

A traditional objection to maximizing theories of bounded rationality is that they cannot accommodate Heuristic Rationality (Gigerenzer and Selten 2001b).⁸ However, level separation allows us to not only reconcile maximizing consequentialism with bounded rationality, but also to do this by capturing traditional arguments for Heuristic Rationality.

There are at least three traditional arguments for Heuristic Rationality. First, in many cases there is an *accuracy-effort tradeoff* between the quality of decisions and the effort of making them (Johnson and Payne 1985). Heuristics often perform well on the accuracy-effort tradeoff, returning high-quality decisions at low cognitive cost. Second, sometimes it is beyond an agent's ability to use complex nonheuristic methods. In this case, heuristics are the only game in town. Finally, in some contexts heuristics reliably outperform nonheuristic methods, even ignoring cognitive costs, due to *less-is-more effects* (Gigerenzer and Brighton 2009). In such cases, heuristics are better in every way than nonheuristic methods, returning better decisions at lower cost.

Maximizing consequentialism accommodates Heuristic Rationality by grounding all three of these traditional arguments. Maximizing consequentialism incorporates accuracy-effort tradeoffs because cognitive resources wasted on one inquiry could have been better spent on another. Maximizing consequentialism incorporates limited cognitive abilities, because ought implies can. If it is not within our ability to use some cognitive process, then maximizing consequentialism does not require us to do so. And maximizing consequentialism incorporates less-is-more effects, because in the presence of less-is-more effects, heuristics have better expected consequences of all types than nonheuristic methods.

But how can heuristics be rational, if they sometimes output decisions that do not maximize expected utility? Here the consequentialist relies on level separation. It can be rational to use heuristic decision processes, even though those processes sometimes lead to suboptimal actions. In such cases, maximizing theories say that the agent used a

⁸Indeed, in one traditional usage, satisficing denotes a particular decisionmaking heuristic (Gigerenzer and Gaissmaier 2011; Selten 1998).

rational decision procedure, but was led to take an irrational action as a result. We saw in Section 6.1 that such cases involve a familiar type of level separation which has been defended at length by consequentialists. It may be that some bounded rationality theorists do not like to admit level tension into normative theorizing, but unless we are provided with a novel argument against level tension, it is not clear that bounded rationality adds anything new to the debate.

Level separation allows us to make sense of a range of related challenges to maximizing consequentialism. For example, if maximizing consequentialism requires agents to explicitly calculate the expected values of each decision process, then maximizing consequentialism is often demanding and sometimes perhaps physically impossible (Elster 1983; Cherniak 1986; Wheeler 2026). But this is not what maximizing consequentialism says. Maximizing consequentialism posits another level separation between the second-order question of which deliberative process agents should use, and the third-order or metacognitive question of how they should decide which deliberative process to use (Thorstad 2022). It is quite right that the metacognitive process of explicitly calculating expected utilities of each lower-order decision process is often more trouble than it is worth.⁹ Precisely for this reason, maximizing consequentialism says that agents should often use simpler metacognitive decision processes.¹⁰

In this section, we have seen how a traditional level separation between the rationality of processes and the rationality of their outputs may recover three traditional arguments for the rationality of heuristic cognition and answer concerns about the complexity of consequentialist metacognition. In the next section, I discuss the argument that maximizing standards are inappropriate under Knightian uncertainty.

⁹Maximizing consequentialism is friendly towards an ecological account on which simpler metacognitive processes succeed because they are well-adapted to human cognitive environments (Marewski and Schooler 2011).

¹⁰For example, many authors think that agents can often make adequate metacognitive choices by allowing themselves to be directly guided by a special class of noetic feelings (Proust 2013) without any explicit calculation at all. When this is a frugal and reliable strategy, maximizing consequentialism may recommend it.

7 Knightian uncertainty

One of the best known arguments against maximization for bounded agents is that under conditions of deep or ‘Knightian’ uncertainty, maximizing standards are inappropriate (Gigerenzer and Selten 2001b; Klein 2001). Knightian uncertainty is a difficult issue to broach, for two reasons. First, many things have been meant by Knightian uncertainty, so it is important to keep track of what we are talking about. And second, this is an area of theorizing in which it is often difficult to move established positions through argument. Nevertheless, in this section I explore two ways of filling out the notion of Knightian uncertainty that could underlie an objection to maximization. I argue that plausible ways of developing each conception of Knightian uncertainty leave substantial room for bounded agents to be subject to maximizing normative standards.

7.1 Unawareness

Many bounded rationality theorists understand Knightian uncertainty to involve unawareness of possible states that the world could be in (Artinger et al. 2014; Gigerenzer and Sturm 2012). They hold that maximization is impossible or inappropriate under conditions of unawareness. But what is the problem with maximizing under conditions of unawareness?

Sometimes the objection is simply that heuristics are uniquely appropriate under conditions of unawareness (Mousavi and Gigerenzer 2017). But this is no objection to maximizing consequentialism. We saw in Section 6 that maximizing consequentialism is fully compatible with the rationality of heuristic cognition.

Perhaps the objection is that agents who maximize expected utility according to less-aware probability and utility functions will often fail to maximize expected utility relative to a fully-aware probability and utility function. But at least from the *ex ante* perspective, this is a feature of the view, not a bug. Bounded agents are ignorant in two ways: they lack information about how the world is, and also about how the world could

be.¹¹ Maximizing consequentialism incorporates agents' limited information about how the world is by calculating expected utilities relative to an agent's evidence about the world. Similarly, the natural extension of maximizing consequentialism incorporates agents' limited information about how the world could be by calculating expected utilities relative to the world states that they are aware of.¹² Theories of bounded rationality would be wrongly insensitive to an agent's limited information about how the world is if they did not restrict attention to expected utilities. And likewise, *ex ante* theories of bounded rationality would be wrongly insensitive to an agent's limited information about how the world could be if they did not restrict attention to states of which agents are relevantly aware.

Sometimes the objection is made that under conditions of unawareness, the notions of probability and utility become ill-defined (Artinger et al. 2014; Mousavi and Gigerenzer 2014; Gigerenzer and Selten 2001b). But this objection requires separate treatment. The problem cannot be that probability and utility functions are formally undefinable under conditions of unawareness. Recent work in Bayesian decision theory has shown how probability and utility functions can be defined relative to a restricted state of awareness and how they might change with growing awareness (Bradley 2017; Karni and Vierø 2013). Furthermore, this objection is driven by many concerns beyond unawareness. So it seems most prudent to group this objection with a broader view of Knightian uncertainty that involves the ill-definedness of probabilities or utilities.

¹¹There are some technical notions of information on which unawareness is not a type of informational lack. It is not my aim to argue over terminology: those who hold this view should replace 'information about' with 'awareness of' in relevant parts of the subsequent discussion.

¹²This is not quite true. For example, on a popular 'catchall' model, agents also assign probabilities and utilities to a 'catchall' state capturing all possibilities they are not aware of. But this does not change the general point that maximizing consequentialism, as I have formulated it, captures a broadly *ex ante* perspective in which agents calculate expected values based on limited information and awareness.

7.2 Undefinedness

Some bounded rationality theorists hold that under Knightian uncertainty, probabilities or utilities become undefined (Artinger et al. 2014; Mousavi and Gigerenzer 2014; Gigerenzer and Selten 2001b; Schmidt 2019). To motivate this view, consider:

(Jellyfish) A stranger approaches you and begins pulling objects from a sack. He takes out a tube of toothpaste, a jellyfish, and then a smaller tube of toothpaste. How likely is it that the stranger will take out another tube of toothpaste?
(Elga 2010)

It may seem compelling to think that there is no fact about how likely it is that the stranger will take out another tube of toothpaste. Given the depth of our uncertainty about what the stranger will do, it might appear that the relevant probabilities are so unclear that they should be treated as ill-defined.

Three observations may serve to limit the class of cases on which this view gets a take. First, note that the ill-definedness of probabilities and utilities poses a radical challenge to a great number of important philosophical notions. For example, the notion of evidential probability used by maximizing consequentialism is also a mainstay of traditional epistemology. And the notion of utility used by maximizing consequentialism is the same notion that makes axiological theorizing possible. To deny that evidential probabilities or utilities are defined in some case is thereby to deny that important philosophical projects get a take on that case. While this is a consequence we might be willing to tolerate in some rare cases, it is not a consequence to be taken lightly.

Second, some ways of treating probabilities and utilities as ill-defined may retain a substantial core of maximizing consequentialism. For example, we might take an imprecise probabilistic approach on which probability functions are replaced with sets of probability functions (Troffaes 2007; Wheeler 2022). Many of the strongest decision rules for imprecise probabilities retain analogs of the requirement to maximize expected utility. These include Γ -maximin, on which agents maximize worst-case expected utility across

their set of probability functions, and Γ -maximax, on which agents maximize best-case expected utility. Some weaker rules retain a maximizing flavor, such as E -admissibility, on which agents take options maximizing expected utility on at least one probability function.

Some of the weaker decision rules for imprecise probabilities move further away from maximizing consequentialism. These include the maximality rule, on which agents avoid options that are dominated in the sense that some other option has higher expected utility on all admissible probability functions. They also include interval dominance, on which agents avoid options whose best-case expected utility lies beneath the worst-case expected utility of another option. Those attracted to maximizing consequentialism will land closest to their original target by staying away from such rules.

At the same time, it should be noted that many decision rules for imprecise probabilities face objections. For example, many allow failures of Good's Theorem in which agents can rationally pay to avoid cost-free evidence (Wheeler 2020). And imprecise probabilities are vulnerable to dilation, in which learning experiences cause them to widen rather than narrow their uncertainties (Pedersen and Wheeler 2014). These phenomena may lead some maximizing theorists to hold their ground and retain precise probabilities.

Finally, in specifying the scope of Knightian uncertainty, the objection from undefinedness faces a dilemma. To illustrate the problem, consider a famous taxonomy of decision situations due to Luce and Raiffa (Luce and Raiffa 1957, p. 13). On this taxonomy, *decision-making under certainty* occurs when the consequences of all acts are known; *decisionmaking under risk* occurs when some consequences are unknown but their probabilities are known; and *decisionmaking under uncertainty* occurs when outcome probabilities are "completely unknown, or not even meaningful". Where, in this taxonomy, should we put Knightian uncertainty?

Taking the first horn of the dilemma involves giving a broad conception of Knightian uncertainty. For example, we could say that Knightian uncertainty involves anything deeper than decisionmaking under risk. This means that even when probabilities are

based on broad and robust evidence, unless every one of the relevant probabilities meets the demanding standards to constitute knowledge, we face Knightian uncertainty.¹³ But once we take such a broad conception of Knightian uncertainty, the idea that probabilities or utilities become undefined under Knightian uncertainty loses its appeal. Most philosophers are willing to accept that evidential probabilities are well-defined in cases where an agent's evidence base is broad and robust. If the objection from Knightian uncertainty involves denying this claim, then we are well within our rights to ask for further argument.

Taking the second horn of the dilemma involves giving a narrow conception of Knightian uncertainty. At the extreme, we could identify Knightian uncertainty with Luce and Raiffa's conditions of decisionmaking under uncertainty, under which probabilities are completely unknown, or not even meaningful. Or, broadening a little, we might take Knightian uncertainty to obtain in conditions of especially deep uncertainty such as Jellyfish, and also in problems such as long-run climate modeling (Lempert et al. 1996). Because these conceptions of Knightian uncertainty are highly restricted, they lend more plausibility to the idea that probabilities and utilities could be ill-defined under Knightian uncertainty. But precisely because these conceptions of Knightian uncertainty are so narrow, they leave open a great number of cases in which bounded agents do not face Knightian uncertainty. And in those cases, there will be no objection to the applicability of maximizing normative standards.

Summing up, in this section we considered the objection that maximizing is inappropriate under conditions of Knightian uncertainty. We looked at two ways in which the notion of Knightian uncertainty can be understood: as unawareness, or as ill-definedness of probabilities and utilities. We saw that in both cases, there may not be decisive reason to relax maximizing normative standards in a large number of interesting cases of rational decisionmaking for bounded agents. As a result, current appeals to Knightian uncertainty cannot sustain the idea that maximizing is always, or even usually inappropriate

¹³Even if we follow Moss (2018) in thinking that credences can constitute knowledge, we needn't think that all decision-relevant credences typically do constitute knowledge.

for bounded agents.

8 Constructed decision problems

It is sometimes held that maximizing normative theories may be appropriate when agents have fully constructed decision problems in a way that makes maximizing standards applicable and desirable. However, the objection continues, bounded agents have not always mentally constructed decision problems to a degree that makes maximizing standards applicable or desirable (Simon 1955). If this is right, then maximizing normative theories are true in the thin sense of holding for suitably-constructed decision problems, but not the thick sense of holding for all decision problems.

In particular, Maximizing Consequentialism helps itself to at least four things: a probability function, a utility function, a state space and a set of options. Maximizing Consequentialists must give accounts of each item that make Maximizing Consequentialism applicable and normatively attractive in the decision problems faced by bounded agents. Thinking through the space of possible accounts will help to illuminate the views most friendly to Maximizing Consequentialism and other maximizing normative theories.

Beginning with probabilities, there are three leading accounts of the decision-theoretically relevant probability function Pr . For subjectivists, Pr is the agent's own credence function. For objectivists, Pr is the truth-indicator function: $Pr(w) = 1$ if w is actual and 0 otherwise. For information-sensitive theorists, Pr_E is the evidential probability function determined by the agent's total evidence.

Subjectivism may not do, since bounded agents may not have settled on probabilistic credences or may not trust those they have settled on. But information-sensitive and objective theories pose no such obstacles. This paper adopts an information-sensitive account on which Pr is the evidential probability function determined by the agent's total evidence. Evidential support is standardly conceived as a two-part relationship between propositions and bodies of evidence, so that evidential probabilities exist independently

of an agent's construction of their decision problem.

Moving beyond probabilities, terminology differs, but the theories closest to the subjectivist's emphasis on mental entities explicitly constructed by decisionmakers will tend to be least friendly towards maximizing normative theories. Turning to utilities, there are two leading accounts of the decision-theoretically relevant utilities u . On the first, u is the agent's own utility function. On the second, u is the true moral value function. The first theory poses the same problems that arose with subjective probability functions. Bounded agents may not assign utilities to all states, and may not trust the utilities they have assigned. The second theory does not pose either problem. For this reason, Maximizing Consequentialism takes u to be the true moral value function.

Turning to options, opinions are less settled, but the broad lesson is the same. Theories of options which move further from the subject's own conceptualization of them will be more friendly towards maximizing normative theories. For example, consider two accounts proposed by Brian Hedden (Hedden 2012). On the first, options are belief-relative:

A set of propositions is a set of options iff it is a maximal set of mutually exclusive propositions, each of which is such that the agent believes she has the ability to bring it about. (Hedden 2012, p. 350)

Maximizing Consequentialism produces strange results on this theory of options for bounded agents. For example, Maximizing Consequentialism would say that agents who have only identified one option should take that option, no matter what it is.

By contrast, on Hedden's favored account, options are decisions:

A set of propositions is a set of options for agent S at time t iff it is a maximal set of mutually exclusive propositions of the form S decides at t to ϕ , each of which S is able to bring about. (Hedden 2012, p. 352)

This account does not raise similar problems for Maximizing Consequentialism. On this account, Maximizing Consequentialism would say that agents who have only identified

one option should not take that option unless it is the best decision they could take. If it is not, they should keep looking for better options.

Moving finally to states, decision theorists often historically assumed an unrestricted state space Ω that contains all possible world-states, individuated as finely as possible. This poses few problems for Maximizing Consequentialism, but may also not do justice to bounded agents' limited awareness of states. More recently, we saw in Section 7.1 that theories of awareness-restricted decisionmaking have shown how a variety of maximizing normative theories including Maximizing Consequentialism can be applied to spaces that contain only some possible world-states, or that individuate states more coarsely. This account is likewise friendly towards Maximizing Consequentialism.

One reason why these distinctions matter is that they help us to see how maximizing normative theories can be defended without imposing coherence requirements on attitudes such as belief and subjective utility, which are unpopular among bounded rationality theorists (Thorstad 2024, 2025). If the decision-relevant probability and utility functions are the agent's own beliefs and desires, then maximizing normative theories may be ill-defined unless beliefs and desires satisfy familiar coherence requirements. By contrast, if the decision-relevant probability and utility functions are evidential probabilities and true moral values, then applying maximizing normative theories does not presuppose that the agent's beliefs and desires have any particular structure.

Summing up, maximizing normative theories may be inapplicable or normatively unattractive on accounts of decision-relevant probabilities, utilities, options and states which stay close to the subjectivist emphasis on mental entities explicitly constructed by decisionmakers. Moving towards other leading theories makes room for maximizing normative theories and helps to distance maximizing normative theories from coherence requirements.

9 Conclusion

This paper argued that the rationality of satisficing can be squared with a substantial core of traditional maximizing normative theories. By way of illustration, we considered the extent to which traditional arguments for satisficing could be reconciled with maximizing consequentialism.

(Maximizing Consequentialism) For all agents S , acts X and times t :

(Maximizing Criterion of Correctness) S may (must) X at t just in case X is maximal (uniquely maximal) for S at t .

(Information-Sensitive Value) The value of S 's X -ing at t is the expected value of X 's consequences, where the expectation in question is taken relative to evidential probabilities. That is, $V(X) = \sum_w Pr_E(w)V(X, w)$, where Pr_E are the evidential probabilities generated by S 's total evidence at t .

We saw in Section 4 that some objections, based on cases in which all options are dominated or some values are incommensurable, occasion at most modest and well-motivated revisions to Maximizing Consequentialism. We saw in Section 5 that traditional philosophical arguments for a satisficing criterion of correctness draw largely on concerns such as a desire to limit the demands of morality or to elucidate the virtue of moderation. These concerns, while philosophically important, do not turn essentially on considerations of bounded rationality, and as such can be assessed separately.

We saw in Section 6 that maximizing theories can incorporate the rationality of decisionmaking heuristics such as satisficing by invoking a traditional level separation between the rationality of decisionmaking processes and the rationality of the intentions or actions they produce. We saw in Section 7 that on leading conceptions of Knightian uncertainty as unawareness or undefinedness, maximizing criteria may be defended except possibly in the deepest cases of uncertainty. As with cases in which all options

are dominated (Section 4.1), such cases are rare and the rational reaction to such cases is disputed, but precisely because they are rare we could accept them as limitations on the scope of maximizing theories without substantial loss. And we saw in Section 8 how maximizing normative theories can be defended when agents have not fully constructed decision problems, by distancing decision-relevant items such as probabilities, utilities, states and options from the agent's conceptualization of them.

This discussion suggests that the truth of maximizing theories, such as Maximizing Consequentialism, may be entirely or largely compatible with considerations of bounded rationality thought to rule them out. This conclusion expands the audience for theories of bounded rationality by showing how central features of bounded rationality can be incorporated into a wider range of views. It also suggests that future work may productively revisit the role that non-maximization plays in motivating scientific and philosophical views about bounded rationality (Section 3). Clarifying the relationship between maximization and bounded rationality can help to clarify the extent to which considerations of bounded rationality must be supplemented by additional scientific and philosophical theses in order to motivate leading views.

References

- Anderson, John. 1990. *The adaptive character of thought*. Lawrence Erlbaum Associates.
- Artinger, Florian, Petersen, Malte, Gigerenzer, Gerd, and Weibler, Jürgen. 2014. "Heuristics as adaptive decision strategies in management." *Journal of Organizational Behavior* 36:533–52.
- Bermúdez, José. 2009. *Decision theory and rationality*. Oxford University Press.
- Bradley, Richard. 2017. *Decision theory with a human face*. Cambridge University Press.
- Byron, Michael (ed.). 2004. *Satisficing and maximizing: Moral theorists on practical reason*. Cambridge University Press.

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

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

Dorst, Kevin. 2023. "Rational polarization." *Philosophical Review* 132:355–458.

Dreier, James. 2004. "Why ethical satisficing makes sense and rational satisficing doesn't." In Michael Byron (ed.), *Satisficing and maximizing: Moral theorists on practical reason*, 131–154. Cambridge University Press.

Dretske, Fred. 2006. "Minimal rationality." In Susan Hurley and Matthew Nudds (eds.), *Rational animals?*, 107–117. Oxford University Press.

Driver, Julia. 2001. *Uneasy Virtue*. Cambridge University Press.

Elga, Adam. 2010. "Subjective probabilities should be sharp." *Philosophers' Imprint* 10:1–11.

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

Fantl, Jeremy and McGrath, Matthew. 2009. *Knowledge in an uncertain world*. Oxford University Press.

Gigerenzer, Gerd. 2004. "Striking a blow for sanity in theories of rationality." In Mie Augier and James March (eds.), *Models of a man: Essays in memory of Herbert A. Simon*, 389–409. MIT Press.

—. 2010. "Moral satisficing: Rethinking moral behavior as bounded rationality." *Topics in Cognitive Science* 2:528–554.

Gigerenzer, Gerd and Brighton, Henry. 2009. "Homo heuristics: Why biased minds make better inferences." *Topics in Cognitive Science* 1:107–43. doi:10.1111/j.1756-8765.2008.01006.x.

Gigerenzer, Gerd and Gaissmaier, Wolfgang. 2011. "Heuristic decision making." *Annual Review of Psychology* 62:451–82.

- Gigerenzer, Gerd and Goldstein, Daniel. 1996. "Reasoning the fast and frugal way: Models of bounded rationality." *Psychological Review* 103:650–9.
- 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.
- Hajek, Alan. 2014. "Unexpected expectations." *Mind* 123:533–67.
- Hedden, Brian. 2012. "Options and the subjective ought." *Philosophical Studies* 158:343–360.
- Hertwig, Ralph, Leuker, Christina, Pachur, Thorsten, Spillopoulos, Leonidas, and Pleskac, Timothy. 2022. "Studies in ecological rationality." *Topics in Cognitive Science* 14:467–91.
- Hoek, Daniel. forthcoming. "Minimal rationality and the web of questions." In Dirk Kindermann, Peter van Elswyk, and Andy Egan (eds.), *Unstructured content*. Oxford University Press.
- Howes, Andrew, Lewis, Richard, and Vera, Alonso. 2009. "Rational adaptation under task and processing constraints: Implications for testing theories of cognition and action." *Psychological Review* 116:717–51.
- Hursthouse, Rosalind. 2001. *On virtue ethics*. Oxford University Press.
- Icard, Thomas. 2018. "Bayes, bounds, and rational analysis." *Philosophy of Science* 85:79–101.
- . ms. *Resource rationality*.

- Johnson, Eric and Payne, John. 1985. "Effort and accuracy in choice." *Management Science* 31:395–414. doi:10.1287/mnsc.31.4.395.
- 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.
- Karni, Edi and Vierø, Marie-Louise. 2013. "'Reverse Bayesianism': A choice-based theory of growing awareness." *American Economic Review* 103:2790–2810.
- Kitcher, Philip. 1992. "The naturalists return." *Philosophical Review* 101:53–114.
- Klein, Gary. 2001. "The fiction of optimization." In Gerd Gigerenzer and Reinhard Selten (eds.), *Bounded rationality: The adaptive toolbox*, 103–21. MIT Press.
- Kolodny, Niko and MacFarlane, John. 2010. "Ifs and oughts." *Journal of Philosophy* 107:115–43. doi:10.5840/jphil2010107310.
- Lempert, Robert, Schlesinger, Michael, and Bankes, Steve. 1996. "When we don't know the costs or the benefits: Adaptive strategies for abating climate change." *Climatic Change* 33:235–74.
- Lieder, Falk and Griffiths, Thomas. 2020. "Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources." *Behavioral and Brain Sciences* 43:E1.
- Luce, R. Duncan and Raiffa, Howard. 1957. *Games and decisions: Introduction and critical survey*. Wiley.
- Marewski, Julian and Schooler, Lael. 2011. "Cognitive niches: An ecological model of strategy selection." *Psychological Review* 118:393–437. doi:10.1037/a0024143.
- Moss, Sarah. 2018. *Probabilistic knowledge*. Oxford University Press.
- Mousavi, Shabnam and Gigerenzer, Gerd. 2014. "Risk, uncertainty, and heuristics." *Journal of Business Research* 67:1671–1678.

- . 2017. "Heuristics are tools for uncertainty." *Homo Oeconomicus* 34:361–79.
- Parfit, Derek. 1984. *Reasons and persons*. Oxford University Press.
- Pasnau, Robert. 2013. "Epistemology idealized." *Mind* 122:987–1021.
- Pedersen, Arthur Paul and Wheeler, Gregory. 2014. "Demystifying dilation." *Erkenntnis* 79:1305–42.
- Penner, Terry. 1973. "The unity of virtue." *Journal of Philosophy* 82:35–68.
- Pettit, Philip. 1984. "Satisficing consequentialism." *Proceedings of the Aristotelian Society, Supplement* 58:165–176.
- Pollock, John. 1983. "How do you maximize expectation value?" *Noûs* 17:409–21.
- Proust, Joëlle. 2013. *The philosophy of metacognition*. Oxford University Press.
- Rawls, John. 1971. *A theory of justice*. Belknap press.
- Schmidt, Andreas. 2019. "Getting real on rationality – behavioral science, nudging, and public policy." *Ethics* 129:511–543. doi:10.1086/702970.
- Schmidtz, David. 1995. *Rational choice and moral agency*. Princeton University Press.
- . 2004. "Satisficing as a humanly rational strategy." In Michael Byron (ed.), *Satisficing and maximizing: Moral theorists on practical reason*, 30–58. Cambridge University Press.
- Selten, Reinhard. 1998. "Aspiration adaptation theory." *Journal of Mathematical Psychology* 42:191–214.
- . 2001. "What is bounded rationality?" In Gerd Gigerenzer and Reinhard Selten (eds.), *Bounded rationality*, 13–37. MIT Press.
- Sen, Amartya. 1997. "Maximization and the act of choice." *Econometrica* 65:745–79.
- Simmons, A. John. 2010. "Ideal and nonideal theory." *Philosophy and Public Affairs* 38:5–36.

- Simon, Herbert. 1955. "A behavioral model of rational choice." *Quarterly Journal of Economics* 69:99–118.
- . 1957. *Administrative behavior: a study of decision-making processes in administrative organization*. Macmillan.
- . 1978. "Rationality as process and as product of thought." *American Economic Review* 68:1–19.
- Slote, Michael. 1984. "Satisficing consequentialism." *Proceedings of the Aristotelian Society, Supplement* 58:139–63.
- . 2001. "Moderation and satisficing." In Elijah Millgram (ed.), *Varieties of practical reasoning*, 221–36. MIT Press.
- . 2004. "Two views of satisficing." In Michael Byron (ed.), *Satisficing and maximizing: Moral theorists on practical reason*, 14–29. Cambridge University Press.
- Sorensen, Roy. 2006. "Originless sin: Rational dilemmas for satisficers." *Philosophical Quarterly* 56:213–23.
- Stocker, Michael. 1976. "The schizophrenia of modern ethical theories." *Journal of Philosophy* 73:453–66.
- Sturm, Thomas. 2012. "The 'rationality wars' in psychology: Where they are and where they could go." *Inquiry* 55:66–81. doi:10.1007/s11229-011-0030-6.
- Swanton, Christine. 2004. "Satisficing and perfectionism in virtue ethics." In Michael Byron (ed.), *Satisficing and maximizing: Moral theorists on practical reason*, 176–189. Cambridge University Press.
- Thorstad, David. 2022. "Two paradoxes of bounded rationality." *Philosophers' Imprint* 22:1–16.
- . 2024. *Inquiry under bounds*. Oxford University Press.

- . 2025. “The complexity-coherence tradeoff in cognition.” *Mind* 134:422–57. doi:10.1093/mind/fzaf015.
- 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. doi:10.1093/acprof:oso/9780195315448.003.0011.
- Troffaes, Matthias C.M. 2007. “Decision making under uncertainty using imprecise probabilities.” *International Journal of Approximate Reasoning* 45:17–29.
- Tucker, Chris. 2016. “Satisficing and motivated submaximization (in the philosophy of religion).” *Philosophy and Phenomenological Research* 93:127–43.
- Viale, Riccardo. 2021. “Why bounded rationality?” In Riccardo Viale (ed.), *Routledge handbook of bounded rationality*, 1–54. Routledge.
- 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.
- . 2022. “A gentle approach to imprecise probability.” In Thomas Augustin, Fabio Cozman, and Gregory Wheeler (eds.), *Reflections on the foundations of probability and statistics*, 37–67. Springer.
- . 2026. “All rationality is bounded.” *Synthese* 207:246.
- Williams, Bernard. 1973a. “A critique of utilitarianism.” In J.J.C. Smart and Bernard Williams (eds.), *Utilitarianism: For and against*, 75–150. Cambridge University Press.
- . 1973b. *Problems of the self*. Cambridge University Press.
- Yetter Chappell, Richard. 2019. “Willpower satisficing.” *Noûs* 53:251–65.