



Does rationality demand higher-order certainty?

Mattias Skipper¹

Received: 21 February 2020 / Accepted: 27 July 2020 / Published online: 3 August 2020
© Springer Nature B.V. 2020

Abstract

Should you always be certain about what you should believe? In other words, does rationality demand *higher-order certainty*? First answer: Yes! Higher-order uncertainty can't be rational, since it breeds at least a mild form of *epistemic akrasia*. Second answer: No! Higher-order certainty can't be rational, since it licenses a dogmatic kind of insensitivity to *higher-order evidence*. Which answer wins out? The first, I argue. Once we get clearer about what higher-order certainty *is*, a view emerges on which higher-order certainty does not, in fact, license any kind of insensitivity to higher-order evidence. The view as I will describe it has plenty of intuitive appeal. But it is not without substantive commitments: it implies a strong form of internalism about epistemic rationality, and forces us to reconsider standard ways of thinking about the nature of evidential support. Yet, the view put forth promises a simple and elegant solution to a surprisingly difficult problem in our understanding of rational belief.

Keywords Higher-order evidence · Higher-order defeat · Higher-order uncertainty · Epistemic akrasia · Inter-level coherence

1 Introduction

There are many stars in the universe. How many, exactly? I'm not sure. It's not that I don't care. I'm just not sure. And I'm not irrational for that. There is nothing irrational in being uncertain about matters for which you lack decisive evidence. And I lack decisive evidence when it comes to the exact number of stars in the universe.

When you're less than fully certain about whether a given proposition is true or false, let's say that you have *first-order uncertainty* about that proposition. Here is an uncontroversial claim: first-order uncertainty can be rational. For instance, I'm rational in having first-order uncertainty about whether the number of stars in the universe is

✉ Mattias Skipper
filmsr@cas.au.dk

¹ Department of Philosophy, Aarhus University, Jens Chr. Skous Vej 7, 8000 Aarhus C, Denmark

odd or even. This is a familiar kind of uncertainty, and one that has received plenty of attention in the epistemological literature.

But first-order uncertainty is not the only kind of uncertainty you might have about a proposition. Sometimes you might be uncertain about not only whether a proposition is *true*, but also about how *confident* you should be that the proposition is true. I might, for example, be uncertain about how confident I should be that there is an odd number of stars in the universe. Is 60% the rational credence for me to have? Or should I rather be 70% confident?

When you're less than fully certain about what credence you should have in a given proposition, let's say that you have *higher-order uncertainty* about that proposition. This paper is about higher-order uncertainty. The central question to be addressed is this: can higher-order uncertainty be *rational*? It's a simple question to state. But the answer is surprisingly elusive. The recent "higher-order evidence" literature has brought out two different lines of argument that, although individually compelling, give opposing answers to our question. If I had to summarize the arguments in one sentence each, they would go as follows:

The Anti-Akrasia Argument: Higher-order uncertainty can't be rational, since it breeds at least a mild form of *epistemic akrasia*.

The Anti-Dogmatism Argument: Higher-order certainty can't be rational, since it licenses a dogmatic kind of insensitivity to *higher-order evidence*.

As stated, neither argument may sound all that compelling. This will change once the flesh is put on the bones (Sects. 2, 3). But just from these snippets, it should already be clear that at most one of the arguments can be sound: if we accept the Anti-Akrasia Argument, we must deny the Anti-Dogmatism Argument; if we accept the Anti-Dogmatism Argument, we must deny the Anti-Akrasia Argument. Something has to give, on pain of contradiction.¹ The question is: what?

At the highest level of generality, we face two options. The first option is to accept the Anti-Dogmatism Argument and deny the Anti-Akrasia Argument. This option has been explored in recent writings by Dorst (2019a, b). Very roughly, he argues that higher-order uncertainty need only commit us to certain *moderate* instances of epistemic akrasia, which we should be prepared to accept as rational. While I have some tentative misgivings about this proposal (to which I'll briefly return in Sect. 2), my aim in this paper is not to argue that the Anti-Akrasia Argument is ultimately sound. As such, I'll simply leave Dorst's proposal on the table for us to consider.

The second option is to accept the Anti-Akrasia Argument and deny the Anti-Dogmatism Argument. This is the option that I'd like to explore in this paper. My strategy will not be to argue, as some epistemologists have, that it is rational to be insensitive to higher-order evidence.² Rather, I want to show that when we get clearer

¹ I say "on pain of contradiction" although, strictly speaking, it is logically possible for both conclusions to be true at the same time. This would be the case if it were neither rational to be higher-order certain nor rational to be higher-order uncertain. Some epistemologists have expressed sympathy for views on which such "epistemic dilemmas" are possible—e.g., Christensen (2007b, 2010a) and Hughes (2019). However, it is usually assumed that they are not, and I'll proceed on this assumption here.

² E.g., Titelbaum (2015), Smithies (2019), Tal (2020) and Weatherson (2019).

about what higher-order certainty *is*, a view emerges on which higher-order certainty does *not*, in fact, license any kind of insensitivity to higher-order evidence (Sect. 4). The view as I will describe it has plenty of intuitive appeal. But it is not without substantive commitments: it implies a strong form of internalism about epistemic rationality, and forces us to reconsider standard ways of thinking about the nature of evidential support (Sect. 5). Yet, the view put forth promises a simple and elegant solution to a surprisingly difficult problem in our understanding of rational belief.

The project described in this paper is in certain respects still a work in progress. There are, as we shall see, various open questions and problems that deserve a more extensive treatment than I'm able to provide here. But even without these loose ends tied up, I think the central proposal of the paper is an important one to consider.

2 The Anti-Akrasia Argument

Although not the main focus of the paper, I want to begin by devoting a few more words to the Anti-Akrasia Argument. In doing so, my aim is not to decisively show that the argument can't be resisted. But I find it worthwhile explaining, at least in rough outline, why some epistemologists (myself included) have been compelled to think that rationality demands higher-order certainty.

The Anti-Akrasia Argument (as the name suggests) starts from the assumption that epistemic akrasia is irrational. To say that epistemic akrasia is irrational is to say, roughly, that certain “mismatches” between an agent's *first-order attitudes* and *higher-order attitudes* can never be rational. For example, if you believe that q despite believing that you *shouldn't* believe that q , then there is an intuitive sense in which your first-order belief (that q) fails to “line up” with your higher-order belief (that “I shouldn't believe that q ”). You somehow seem to exhibit a kind of incoherence, which is (at least superficially) similar to the kind of incoherence exhibited by traditional Moorean beliefs of the form “ q , but I don't believe q .” While it has proven difficult to pin down exactly wherein the incoherence lies, there seems to be a robust and widely held intuition to the effect that epistemic akrasia is indeed irrational.³ Let's call this the *Enkratic Intuition*.

Now, it turns out that there is a close—and, to my mind, surprising—connection between epistemic akrasia and higher-order uncertainty: if higher-order uncertainty can be rational, then so can epistemic akrasia. To get an initial feel for this connection, let's begin by considering an extreme case of higher-order uncertainty. Suppose that you should be (a) 80% confident of q , and (b) certain that you should be only 10% confident of q . Then, if you believe as you should, your first-order and higher-order attitudes will come radically apart. You will think to yourself: “I'm certainly way too confident!” That's blatantly akratic.

Less extreme instances of higher-order uncertainty tend to give rise to less extreme instances of epistemic akrasia. Suppose, for example, that you should be (a) 80% confident of q , (b) 90% confident that you should be 80% confident of q , and (c) 10%

³ See, e.g., Greco (2014) and Horowitz (2014). See also Smithies (2012) for a related discussion of Moore's Paradox.

confident that you should be 60% confident of q . Then, if you believe as you should, your first-order and higher-order attitudes will still come apart, but in a less radical way. You will think to yourself: “I’m certainly not *under*confident, but I may well be *over*confident!” That’s moderately akratic.

How general is this connection between epistemic akrasia and higher-order uncertainty? Does higher-order uncertainty *always* give rise to epistemic akrasia? It’s natural to think “no.” More specifically, it’s natural to think that higher-order uncertainty doesn’t give rise to epistemic akrasia as long as your first-order attitude properly “balances the different pulls” of your higher-order attitudes. Here is a simple example to illustrate the idea: suppose that you are uncertain about whether you should have a credence in q of 60% or 80%, and suppose that you place the same credence in either possibility. That is, you are 50–50 about whether the rational credence for you to have in q is 60% or 80%. Given this, it seems that a credence of 70% in q strikes the right kind of balance between your higher-order attitudes (since $.5 \times 60\% + .5 \times 80\% = 70\%$).

Generalizing this thought, we arrive at the following suggestion: as long as your credence in q equals your expectation of what your credence in q *should* be, you are not epistemically akratic. This is the familiar principle of “Rational Reflection:”⁴

Rational Reflection: $P(q) = \mathbb{E}_P(P(q))$,

where P is the rational credence function,⁵ and \mathbb{E}_P is the mathematical expectation calculated with respect to P .⁶

Many philosophers have taken Rational Reflection to be a natural precisification of the Enkratic Intuition.⁷ Here I will follow suit. That is, I will assume that an agent is epistemically non-akratic if, and only if, the agent complies with Rational Reflection. If we combine this assumption with the natural thought that Rational Reflection is compatible with higher-order uncertainty, we arrive at the conclusion that higher-order uncertainty need not give rise to epistemic akrasia.

However, the natural thought turns out to be mistaken: Rational Reflection is not, in fact, compatible with higher-order uncertainty. In other words, if Rational Reflection is true, then rationality requires you to be certain of what credences you should have. As far as I know, this result was first recognized by Elga (2013), with a predecessor in an unpublished manuscript by Samet (1997). The formal details will not matter for our purposes, but it will be helpful to see an informal walk-through of the proof.⁸

⁴ So-called because of its affinity with van Fraassen’s (1984) “Reflection Principle,” which says (roughly) that you should defer to your own future credences.

⁵ Here and throughout, I’ll assume that P is an *ideally* rational credence function, not a *non-ideally* (or “boundedly”) rational credence function that is sensitive to human cognitive limitations. In doing so, I’m not making any substantive assumptions about what it takes for an agent’s credences to be ideally rational. In particular, I’m not assuming that P obeys the probability axioms.

⁶ As usual, the mathematical expectation of a discrete random variable, X , calculated with respect to a probability function, P , is defined as follows: $\mathbb{E}_P(X) = \sum_x P(X = x)x$.

⁷ E.g., Christensen (2010b), Sliwa and Horowitz (2015), and Dorst (2019a, b).

⁸ For those interested in the technical details, Dorst (2019a, b) provides an excellent formal characterization of the result. Elga’s own formulation of the proof goes as follows: let P' be any credence function that the agent thinks might be ideally rational. By Rational Reflection, then, $P(q \mid P' \text{ is ideal}) = P'(q)$. Letting q be the proposition “ P' is ideal,” we have: $P(P' \text{ is ideal} \mid P' \text{ is ideal}) = P'(P' \text{ is ideal})$. And since $P(P' \text{ is ideal} \mid P' \text{ is ideal}) = 1$, it follows that $P'(P' \text{ is ideal}) = 1$. That is, P' is certain that it itself is the ideally rational credence function.

Let P be the credence function you *should* have, and let Cr be the credence function you *do* have. In other words, P is the *rational* credence function in your situation, and Cr is your *actual* credence function. So you're rational just in case Cr is identical to P . Now suppose that you are higher-order uncertain: that is, you are less than fully certain that P is the rational credence function. Given this, two things must be true of you. First, your credence function conditional on P being the rational credence function should equal P : that is, if you believe as you should, $Cr(\cdot | P \text{ is rational}) = P(\cdot)$. This is a straightforward consequence of Rational Reflection.⁹ Second, your credence function should differ from your credence function conditional on P being the rational credence function. That is, if you believe as you should, $Cr(\cdot) \neq Cr(\cdot | P \text{ is rational})$. After all, you are (by hypothesis) not unconditionally certain that P is the rational credence function. But then Cr cannot be identical to P : that is, your actual credence function cannot be identical to the rational one.¹⁰ So, if you *should* have higher-order uncertainty, you *should* violate Rational Reflection. Contraposing: since you *shouldn't* violate Rational Reflection (on pain of epistemic akrasia), you *shouldn't* have higher-order uncertainty. That's the bare bones of the Anti-Akrasia Argument.

How might the argument be resisted? Some writers have argued that the Enkratic Intuition is radically misleading. According to them, even the most blatant violations of Rational Reflection can be rational.¹¹ Others have argued that the Enkratic Intuition, although ultimately mistaken, is nonetheless onto something. According to them, certain *moderate* violations of Rational Reflection can be rational, whereas *blatant* ones cannot.¹² On the face of it, neither response seems all that attractive: the former goes against a robust intuition of ours; the latter runs the risk of *ad hocness*.¹³ This is not to say that neither strategy can be made to work. Appearances can sometimes be deceptive. But I hope the foregoing remarks convey why someone might feel attracted to the view that higher-order uncertainty is irrational. That was the modest goal of this section, anyway. The more ambitious goal lies ahead: to argue that there is an independently plausible way of resisting the Anti-Dogmatism Argument, which allows us to maintain Rational Reflection (and hence allows us to vindicate the Enkratic Intuition). But first, let's take a closer look at what the Anti-Dogmatism Argument says.

⁹ It is perhaps easier to see this implication from Elga's own (equivalent) formulation of rational reflection, stated in terms of conditional probabilities instead of expectations: $P(q | P' \text{ is ideal}) = P'(q)$.

¹⁰ This is not, to be sure, to say that $Cr(\cdot | P \text{ is rational})$ and $Cr(\cdot)$ must assign different values to *every* proposition. The implication is just that $Cr(\cdot | P \text{ is rational})$ and $Cr(\cdot)$ must assign different values to *at least one* proposition.

¹¹ E.g., Lasonen-Aarnio (2015, 2020) and Williamson (2014).

¹² E.g., Dorst (2019b). I would include Elga (2013) in this camp as well, although, as Dorst (2019a, §3.2) shows, the "New Rational Reflection" principle allows for some rather blatant instances of epistemic akrasia.

¹³ I say "runs the risk" to leave open the possibility that there is, on closer inspection, a principled way to distinguish between those instances of epistemic akrasia that are too blatant to be rational, and those that are not. Dorst (2019b) has recently offered an interesting proposal along these lines, which, if correct, would allow us to deny Rational Reflection while maintaining that there is something importantly right about the Enkratic Intuition.

3 The Anti-Dogmatism Argument

The Anti-Dogmatism Argument starts from another intuitive idea, namely that evidence *about* what you should believe can affect *what you should believe*. In the common epistemological parlance: *higher-order evidence* can affect what you should believe.

To illustrate what this idea amounts to, let's consider a paradigmatic case:¹⁴

Miss Modest and the Self-Serving Bias: Miss Modest believes herself to be more intelligent than the average person. Her belief is, in fact, fully rational: the belief has been formed on the basis of an extensive body of evidence, and she has judged the bearing of that evidence in an impeccable manner. However, as she reads through the latest issue of *Psychology Today*, she learns a disturbing fact: it turns out that most people suffer from a pronounced *self-serving bias*, which causes them to overrate themselves on a number of desirable traits and abilities, including intelligence, social skills, musicality, etc.

How, if at all, should Miss Modest revise her belief state in light of this new information about the self-serving bias? Assuming (as we may) that she has good reason to trust the psychological studies reported in *Psychology Today*, there seems to be a robust and widely held intuition to the effect that she should lose confidence, at least to some extent, that she is more intelligent than the average person. Not all philosophers agree that this intuitive verdict is ultimately correct.¹⁵ But for those who do, the more general lesson is this: you should be prepared to revise your beliefs in light of evidence about whether those beliefs are rational. In short: you should take higher-order evidence seriously.

What, exactly, is wrong with *not* taking higher-order evidence seriously? Christensen (2010a, 2011) has offered what I think is a plausible initial diagnosis. Suppose, as we can imagine, that Miss Modest *didn't* lose confidence in her own superior intelligence upon receiving the higher-order evidence (that is, upon learning about the self-serving bias). Given this, she must effectively take the higher-order evidence to be misleading. After all, if the higher-order evidence *wasn't* misleading, she *should* lose confidence. But the higher-order evidence is (*ipso facto*) misleading only if Miss Modest doesn't suffer from the self-serving bias. So, in reacting as she does, Miss Modest effectively treats herself as an exception to the rule, as if she thinks to herself: "Most people suffer from the self-serving bias... but I don't!" Yet, this seems like a patently dogmatic way of conceiving of her own situation.¹⁶ After all, Miss Modest doesn't have any independent reason (as so we may suppose) to think that she is less

¹⁴ Similar cases can be found in many other places; e.g., Christensen (2010a), Elga (2005), Horowitz (2014), and Skipper (2019a, b).

¹⁵ E.g., Tal (2018, 2020), Titelbaum (2015), and Weatherson (2019).

¹⁶ Although not "dogmatic" in the traditional sense of the term, which is associated with the *dogmatism puzzle* (Kripke 2011). A detailed comparison would take us too far astray, but interested readers are referred to Christensen (2010a, 2011), who offers an illuminating discussion of how the present kind of dogmatism differs from the kind of dogmatism exhibited by someone who ignores a body of evidence merely on the grounds that it speaks against his or her prior opinions.

prone to suffer from the self-serving bias than other people.¹⁷ Thus, it seems that Miss Modest *should*, on pain of dogmatism, lose confidence in her own superior intelligence when learning about the self-serving bias. Let's call this the *Modesty Intuition*.

However, and this is the crux of the Anti-Dogmatism Argument, there is a seemingly compelling line of argument to the conclusion that someone who has no higher-order uncertainty *shouldn't* take higher-order evidence seriously. The most detailed formulation of this argument is due to Dorst (2019a, b), who embeds his argument in a formal, model-theoretic framework. For present purposes, however, the informal story will do the job.

Suppose that Miss Modest has no higher-order uncertainty before she learns about the self-serving bias. She is, in other words, certain that she is rational in believing herself to be more intelligent than the average person. Given this, how should she react to the higher-order evidence? She should, it seems, think to herself: "I'm certainly rational in believing as I do. And the higher-order evidence suggests that I'm *not* rational in believing as I do. So, the higher-order evidence must be misleading!" Put differently, if Miss Modest has no higher-order uncertainty to begin with, she would, it seems, be rational to be certain that the higher-order evidence is misleading. And if you are rationally certain that a body of evidence is misleading, whatever it is, you should presumably not take the evidence seriously. Contraposing: since Miss Modest *should* take the higher-order evidence seriously (on pain of dogmatism), she *shouldn't* start out higher-order certain. That's the bare bones of the Anti-Dogmatism Argument.

How might the argument be resisted? Some writers have argued that the Modesty Intuition is ultimately mistaken. According to them, Miss Modest shouldn't lose confidence upon learning about the self-serving bias.¹⁸ An alternative reaction (which I've encountered in conversation, though I haven't seen it defended in print) would be to say that Miss Modest can be rationally certain that the higher-order evidence is misleading *before*, but not *after*, having received the higher-order evidence. This amounts to denying the standard assumption that rational certainty is indefeasible.¹⁹ On the face of it, neither response seems all that attractive: the former goes against a robust intuition of ours; the latter doesn't seem to get to the heart of the matter.²⁰ Again, this is not to say that neither response can be made to work. The goal of this

¹⁷ What would count as a suitably "independent" reason here? This turns out to be a difficult question. To a first approximation, we can think of the "independence" requirement as ruling out considerations that are produced by the very kind of reasoning whose reliability is called into question by the relevant higher-order evidence. For a more careful treatment of the issue, see Christensen (2019).

¹⁸ Different versions of this view have been defended by Lasonen-Aarnio (2014), Tal (2018, 2020), Titelbaum (2015), and Weatherson (2019).

¹⁹ To say that rational certainty is indefeasible is to say, roughly, that someone who is rationally certain of a given proposition should remain certain of that proposition no matter what evidence is received in the future. In Bayesian parlance: if the prior probability of q is 1, the posterior probability of q is also 1, conditional on any new evidence. This "certainty preservation" principle is an immediate consequence of the ratio formula for conditional probabilities (cf. Joyce 2019).

²⁰ Why not? Briefly put, even if rational certainty is defeasible in *some* contexts—say, contexts involving memory loss—this does little to show that rational certainty is defeasible in the *present* context. Indeed, I see no independent reason to think that rational certainty is defeasible in contexts involving higher-order evidence. In any case, however, I will eventually offer a way to resist The Anti-Dogmatism, which does not rely on any controversial assumptions about the defeasibility of rational certainty.

section is just to convey why someone might feel attracted to the view that higher-order certainty cannot be rational.

Here, then, is the story so far. There are two individually compelling arguments, the Anti-Akrasia Argument and the Anti-Dogmatism Argument, which give opposing answers to the question of whether rationality demands higher-order certainty. What to make of this tension? I submit that there is an independently plausible way of resisting the Anti-Dogmatism Argument, which at the same time allows us to vindicate both the Enkratic Intuition and the Modesty Intuition. That is the more ambitious claim I promised to defend, and to which I now turn.

4 Higher-order certainty without dogmatism: a “Top-Down” View

In the introduction, I said that you are higher-order certain just in case you are “certain of what you should believe.” But what, exactly, are you certain *of* when you are certain of what you should believe? Assuming that it is your *evidence* that determines what you should believe, the answer comes to something like this: you are higher-order certain if, and only if, you are certain of what your total evidence supports. Or, when applied to a single proposition: you are higher-order certain about q if, and only if, you are certain about how strongly your total evidence supports q .

This way of characterizing higher-order (un)certainly obviously leaves open various further questions about what, exactly, you are certain of when you are “certain of what your total evidence supports.” That’s is a deliberate choice on my part, since I don’t want to make any unnecessary assumptions about the nature of evidence or the evidential support relation. The only assumption I want to make at this point is that it is your total evidence that solely determines what you should believe.

Or, rather, that is the only *substantive* assumption I want to make at this point. I will also be relying on a distinction between two different ways in which evidence can bear on propositions. Following Christensen (2016, 2019), let’s call it the “first-order” way and the “higher-order” way. We have already seen the distinction at play in connection with the case of Miss Modest and the Self-Serving Bias. The idea there was that the evidence about the self-serving bias counted as “higher-order evidence” for Miss Modest, because it indicated that her initial belief was *irrational*, not merely *false*. By comparison, if we instead suppose that Miss Modest learns that she has scored low on a reliable IQ test, this would count as “first-order evidence” for her, because it would not indicate that her belief in her own superior intelligence was *irrational*, but merely *false*.

More generally, we may characterize the distinction between first-order evidence and higher-order evidence as follows: a piece of evidence bears in the first-order way on a proposition just in case it bears on whether the proposition is true or false; and it bears in the higher-order way on a proposition just in case it bears on what credence it is rational for you to have in the proposition.²¹ Needless to say, this way of drawing the distinction between first-order evidence and higher-order evidence amounts to

²¹ For other, similar ways of characterizing the distinction between first-order evidence and higher-order evidence, see Christensen (2010a), Lasonen-Aarnio (2014), and Schoenfield (2015, 2016).

little more than an intuitive gloss. But since I won't be relying on any particular precisification of the distinction, I'll resist the temptation to go into further detail here. Let me, however, highlight two features of the distinction, which may help clear away some potential misunderstandings in what follows.

First, note that the same piece of evidence can be “first-order” relative to one proposition while being “higher-order” relative to another proposition. For example, the evidence about the self-serving bias is higher-order relative to the proposition “Miss Modest is more intelligent than the average person” while being first-order relative to the proposition “Miss Modest suffers from the self-serving bias.” As such, it would be a mistake to think of a piece of evidence as being first-order or higher-order *in and of itself*. Rather, we should think of it as being first-order or higher-order *relative to a proposition*. In this respect, the distinction between first-order evidence and higher-order evidence is similar to the distinction between confirming and disconfirming evidence. It would be a mistake to think of a piece of evidence as being intrinsically confirming or disconfirming. Rather, we should think of it as being confirming or disconfirming relative to a proposition. And just as a piece of evidence can confirm one proposition while disconfirming another, so a piece of evidence can bear in the first-order way on one proposition while bearing in the higher-order way on another.

Second, note that the same piece of evidence can be both first-order and higher-order relative to the *same* proposition. Most obviously, this can happen if we take the conjunction of a piece of first-order evidence and a piece of higher-order evidence relative to the same proposition, in which case the conjunction bears in both the first-order way and the higher-order way on that proposition. But there might be less trivial cases as well. Suppose, for example, that I learn that I'm likely to suffer from a highly general confirmation bias, which makes me bad at responding to all sorts of evidence, regardless of the subject matter. Arguably, this piece of evidence is both first-order and higher-order relative to the proposition “I'm bad at responding to all sorts of evidence.”²² As such, we shouldn't think of the distinction between first-order evidence and higher-order evidence as inducing a (proposition-relative) partition on bodies of evidence. Rather, we should think of it as capturing two different, but not mutually exclusive, ways in which evidence can bear on propositions.

Distinction in hand, we can pose a question that will be of key importance in what is to come: is higher-order certainty compatible with uncertainty about what your first-order evidence supports? In other words, can you be certain about what your *total* evidence supports, while being uncertain about what your *first-order* evidence supports? For reasons that will emerge, I think the answer is “yes.” But before we get too far into the weeds, let me explain how this positive answer can help us to resist the Anti-Dogmatism Argument.

In rough outline, the idea goes as follows. Suppose that you are uncertain about what your first-order evidence supports. Given this, you should take seriously any evidence you might later get about what your first-order evidence supports. After all, you cannot be certain that such evidence would be misleading. But this is just to say that you should take higher-order evidence seriously: you should be prepared to

²² Thanks to Jack Spencer for suggesting this example. See also Christensen (2019) for a similar kind of example.

revise your beliefs in light of evidence about what your first-order evidence supports. Thus, as long as you are uncertain about what your first-order evidence supports, you should be sensitive to higher-order evidence. This is so even if you have no uncertainty about what your *total* evidence supports—that is, even if you have no higher-order uncertainty.

To illustrate the idea, let's apply it to Miss Modest's situation. Let q be the proposition that "Miss Modest is more intelligent than the average person," and let's suppose that Miss Modest has no higher-order uncertainty about q prior to learning about the self-serving bias: that is, she is certain that her total evidence supports such-and-such a credence in q . But let's also suppose that she *does* have uncertainty about what her *first-order* evidence supports: that is, she is not certain that her first-order evidence supports such-and-such a credence in q . (This might seem like a puzzling set of assumptions. How could it be rational to be certain about what one's *total* evidence supports while being uncertain about what one's *first-order* evidence supports? I'll have more to say about this question in Sect. 4.2, but here is a provisional way of making sense of it. Imagine that Miss Modest has accrued a complex body of first-order evidence about her own and other people's intelligence, and that she, modest as she is, is doubtful that she has managed to judge the first-order evidence in a fair and unbiased way. Suppose also, however, that she realizes that all of her reasons for doubting her own ability to assess the first-order evidence, or any other body of evidence for that matter, is part of her total evidence. In light of this realization, she sees no reason to be uncertain about what *it*, the total evidence, supports. In other words, since she realizes that her assessment of what her total evidence supports already takes into account, or compensates for, all of her doubts about her own cognitive fallibility, she is left with no *residual* doubts about what she should, all-things-considered, believe.)

Given these background assumptions, how should Miss Modest react to the evidence about the self-serving bias? She should think to herself: "The psychological studies indicate that I'm likely to have overestimated how strongly my first-order evidence supports q . And I'm not certain that my judgment of the first-order evidence was correct in the first place. So, I cannot be certain that the psychological studies are misleading. I better take them seriously!" In other words, as long as Miss Modest is uncertain about what her first-order evidence supports, she should take the higher-order evidence seriously. This is so even if she has no higher-order uncertainty to begin with.

That's my basic proposal for how to resist the Anti-Dogmatism Argument. To sharpen the proposal, I will make a few assumptions about how we can describe Miss Modest's epistemic situation. First, I will assume that we can use a real-valued function, P , to encode information about what the total evidence supports. This will allow us to write things like " $P(q) = 80\%$ " to say that the total evidence supports a credence of 80% in q . Second, I will assume that we can use another real-valued function, P_F , to encode information about what the first-order evidence supports. This will allow us to write things like " $P_F(q) = 80\%$ " to say that the first-order evidence supports a credence of 80% in q . Finally, I will assume that the rational credences are those that are supported by the total evidence: that is, Miss Modest should have a credence of c in q if, and only if, $P(q) = c$. This is just an implementation of the assumption that your total evidence determines what you should believe.

For reasons to which I shall return in Sect. 5, I will not assume that P and P_F are standard probability functions, which are obtained from a shared urprior by conditionalizing on the total evidence and first-order evidence, respectively. All I assume at this point is that P and P_F are *precise* and *unique*: that is, they each return a single real number between zero and one, for any given proposition. Neither assumption is uncontroversial.²³ But they are, I think, harmless for present purposes, and they will help to simplify the exposition.

Next, let's apply our modeling assumptions to the story about Miss Modest. We begin by making two stipulations about her evidential situation prior to learning about the self-serving bias. First, her total evidence supports a credence of c in q : that is, $P(q) = c$. Second, her first-order evidence supports a credence of c_F in q : that is, $P_F(q) = c_F$ (where c_F may or may not equal c). Given this, how should Miss Modest think about what her own total evidence and first-order evidence supports, respectively? Recall that, as I have told the story, Miss Modest has no higher-order uncertainty about q to begin with. That is, she is certain that her total evidence supports a credence of c in q : $P(P(q) = c) = 1$. At the same time, she *does* have uncertainty about what the first-order evidence supports. In particular, she is not certain that the first-order evidence supports a credence of c_F in q : that is, $P(P_F(q) = c_F) < 1$. In fact, she *expects* the first-order evidence to support a credence of c in q : that is, $\mathbb{E}_P(P_F(q)) = c$. That's why she adopts a credence of c in q .

What, then, happens when Miss Modest learns about the self-serving bias? She should think to herself: "I'm likely to have overestimated how strongly the first-order evidence supports q . I better lower my credence in q !" In other words, upon realizing that her expectation for the value of $P_F(q)$ is most likely too high, she should lower this expectation, and hence lower her credence in q accordingly. The result: although Miss Modest has no higher-order uncertainty about q to begin with, she nevertheless ends up revising her credence in q in light of the higher-order evidence, because she is uncertain about what the first-order evidence supports.

I submit that this is how we should understand how higher-order evidence can influence what an agent should believe. I don't have a knock-down argument for it. The story above strikes me as an intuitively plausible way of describing why Miss Modest should lose confidence upon learning about the self-serving bias. Perhaps you'll agree. But in any case, the story will in large part earn its keep by the work it can do in helping us to resolve the tension between the Anti-Akrasia Argument and the Anti-Dogmatism Argument. The question that remains, then, is whether our intuitive story about Miss Modest can be turned into a credible theory of how your first-order and higher-order credences should hang together. The rest of this section aims to do just that.

4.1 Top-Down Guidance

The foregoing considerations suggest that there is a close connection between, on the one hand, what you should believe, and, on the other hand, what you should expect

²³ For discussions of Precision, see White (2009) and Elga (2010). For discussions of Uniqueness, see Schoenfield (2013) and Schultheis (2018).

your first-order evidence to support. More specifically, the story about Miss Modest suggests that you should have the credence that you should expect to be supported by your first-order evidence. Formally:

$$\text{Top-Down Guidance: } P(q) = \mathbb{E}_P(P_F(q))$$

I offer Top-Down Guidance as a precisification of the idea that Miss Modest should lower her credence in q upon learning about the self-serving bias, because she should lower her expectation for how strongly q is supported by her first-order evidence.

To see Top-Down Guidance in action, let's put some numbers on the table. Suppose that Miss Modest should initially (that is, before learning about the self-serving bias) be 50–50 about whether the first-order evidence supports a credence in q of 80% or 90%: that is, $P(P_F(q) = 80\%) = P(P_F(q) = 90\%) = .5$. Given this, she should expect the first-order evidence to support a credence of 85% in q : $\mathbb{E}_P(P_F(q)) = (.5 \times 80\%) + (.5 \times 90\%) = 85\%$. So, according to Top-Down Guidance, Miss Modest should have a credence of 85% in q prior to learning about the self-serving bias.

Now let's suppose that Miss Modest, *after* having received the evidence about the self-serving bias, should be 50–50 about whether her first-order evidence supports a credence in q of 20% or 30%: that is, $P(P_F(q) = 20\%) = P(P_F(q) = 30\%) = .5$. Given this, she should now expect the first-order evidence to support a much lower credence in q : $\mathbb{E}_P(P_F(q)) = (.5 \times 20\%) + (.5 \times 30\%) = 25\%$. So, in this toy example, Miss Modest should respond to the higher-order evidence by lowering her credence in q from 85 to 25%.

Of course, higher-order evidence will not always have the effect of *lowering* the rational credence. Sometimes, it will have the effect of *increasing* the rational credence instead. For example, we can imagine that Miss Modest goes on to learn that she is likely to be one of the few people who doesn't suffer from the self-serving bias. Given this, she should increase her expectation for the value of $P_F(q)$, and hence increase her credence in q accordingly.

That's how Top-Down Guidance *works*. But why should we *accept* it? Above I presented Top-Down Guidance as a natural precisification of the intuitive idea that Miss Modest should lower her credence upon receiving the higher-order evidence, because she should lower her expectation of what the first-order evidence supports. This will remain the central motivation of the principle. However, I'd like to highlight a few characteristics of Top-Down Guidance that may help shore up some additional motivation.

First, Top-Down Guidance offers a vindication of the familiar idea that you should treat your evidence as a "guide to the truth." This idea, albeit notoriously vague, has often been taken as a basic desideratum for any theory of rational belief.²⁴ Yet, it isn't immediately clear what it could mean to treat your higher-order evidence as a guide to the truth. After all, higher-order evidence doesn't bear on the *truth* of the believed proposition in any straightforward way, but rather bears on the *rationality* of the belief itself. How, then, could it make sense to say that you should treat your higher-order evidence as a guide to the truth? When seen through the lens of Top-Down Guidance, the idea is this: you should treat your first-order evidence as a guide to the truth about

²⁴ See, e.g., Kelly (2014).

q , and you should treat your higher-order evidence as a guide to the truth about how your first-order evidence bears on q . In short: you should treat your evidence as a guide to the truth in a “top-down” manner.

Second, if we test Top-Down Guidance against extreme cases, we get what looks like plausible results. For example, if you are rationally certain about what your first-order evidence supports, then Top-Down Guidance implies that you should respond dogmatically to any higher-order evidence you might later get, since no such evidence could make it rational for you to revise your expectation of what the first-order evidence supports.²⁵ This seems like the right advice. After all, if you are rationally certain that your first-order evidence supports such-and-such a credence, then you can be certain that any higher-order evidence that suggests otherwise is misleading. Needless to say, we will rarely (if ever) find ourselves in a situation where we can be absolutely about what our first-order evidence supports, since we typically (if not always) have at least some reason to doubt our own cognitive abilities. But if you *do* find yourself in such a situation, it seems that you *should* take a dogmatic stance towards any further higher-order evidence you might get.

Finally, Top-Down Guidance does not face the kinds of “bootstrapping” worries that have been raised against various other inter-level principles, which take the form of *narrow-scope* norms.²⁶ The worry stems from a notorious feature of narrow-scope norms, namely that they allow for what is sometimes called *detaching*. To illustrate this property, let’s consider the following narrow-scope norm: “If you believe that you should ϕ , then you should intend to ϕ .”²⁷ As stated, this norm implies that *merely believing* that you should ϕ suffices to make it the case that you *should* intend to ϕ . However, as many philosophers have pointed out, this often seems like the wrong advice. In many cases, you should *drop the belief* rather than *form the intention*. Or so the criticism goes.²⁸

It might be tempting to think that Top-Down Guidance falls prey to a similar kind of worry. After all, doesn’t Top-Down Guidance say that *merely expecting* that your first-order evidence supports such-and-such a credence suffices to make it the case that you *should* have such-and-such a credence? No. Top-Down Guidance says something importantly different, namely that *rationally expecting* that your first-order evidence supports such-and-such a credence suffices to make it the case that you *should* have such-and-such a credence. To make this difference vivid, consider the following wide-scope norm: “You should see to it that you intend to ϕ if you believe that you should ϕ .” This norm does not allow us to infer from the fact that you *do* believe that you should ϕ that you should intend to ϕ . It only allows us to infer from the fact that you *should* believe that you should ϕ that you should intend to ϕ .²⁹ The same goes for

²⁵ Assuming that rational certainty is indefeasible; cf. our discussion in Sect. 3.

²⁶ Examples of such norms include “Calibration” (White 2009), “Guess Calibration” (Sliwa and Horowitz 2015), and “Judgment Calibration” (Schoenfield 2015). See also Kelly’s (2010) bootstrapping worry about the Equal Weight View of peer disagreement, as defended by Christensen (2007a) and Elga (2007).

²⁷ Discussions of this kind of norm go back as far as to Aristotle’s ethics, but have recently seen a revival due to the work of Broome (1999, 2007), Kolodny (2005), and others.

²⁸ For further discussion of this general issue, see Broome (1999), Kolodny (2005), and Kiesewetter (2017).

²⁹ Assuming that ‘should’ distributes over the material conditional (as in standard deontic logic): if you should ‘ ψ , if ϕ ’, then you should ψ , if you should ϕ .

Top-Down Guidance: it does not allow us to infer from the fact that you *do* expect your first-order evidence to support such-and-such a credence that you should have such-and-such a credence. It only allows us to infer from the fact that you *should* expect your first-order evidence to support such-and-such a credence that you should have such-and-such a credence.

So far, so good. But problems crop up.

4.2 Top-Down Fixation

Recall that the view I've set out to explore is supposed to deliver the result that rationality demands higher-order certainty. That is what will allow us to resist the Anti-Dogmatism Argument while at the same time accepting the Anti-Akrasia Argument.

However, Top-Down Guidance does not by itself rule out rational higher-order uncertainty. The reason for this is that Top-Down Guidance is compatible with rational uncertainty about what you should expect your first-order evidence to support.³⁰ And if you are rationally uncertain about what you should expect your first-order evidence to support, Top-Down Guidance says that you should be uncertain about what you should believe. Formally put: since Top-Down Guidance permits that $P(\mathbb{E}_P(P_F(q)) = x) < 1$, it also permits that $P(P(q) = x) < 1$. As such, Top-Down Guidance does not by itself provide us with the necessary means to resist the Anti-Dogmatism Argument while accepting the Anti-Akrasia Argument.

What to do about it? The source of the problem, as I have said, is that Top-Down Guidance is compatible with rational uncertainty about what you should expect your first-order evidence to support. Thus, in order to solve the problem, we need to find a way of ruling out rational uncertainty about what you should expect your first-order evidence to support. In other words, we must find a way of establishing the following principle:

Top-Down Fixation: $[\mathbb{E}_P(P_F(q)) = x] \rightarrow P[\mathbb{E}_P(P_F(q)) = x] = 1$

According to Top-Down Fixation, it is never rational to be uncertain about what you should expect your first-order evidence to support: that is, if you should in fact expect your first-order evidence to support such-and-such a credence in q , then you should be certain of that fact.

When combined with Top-Down Guidance, Top-Down Fixation gets us the result that rationality demands higher-order certainty. According to Top-Down Fixation, you should be certain of the value of $\mathbb{E}_P(P_F(q))$, whatever it is. And according to Top-Down Guidance, the value of $\mathbb{E}_P(P_F(q))$ *just is* the credence that you should have in q . So, Top-Down Guidance and Top-Down Fixation jointly entail that you should be certain of what credence you should have in q , which is just to say that you should have no higher-order uncertainty about q .

This also means that we have a way of meeting our dual aim of resisting the Anti-Dogmatism Argument while accepting the Anti-Akrasia Argument. We have already

³⁰ Notice the difference between (a) having uncertainty about what your first-order evidence *in fact* supports, and (b) having uncertainty about what you should *expect* your first-order evidence to support. Top-Down Guidance allows for both kinds of uncertainty, but it is only the latter kind of uncertainty that creates trouble for Top-Down Guidance.

seen (in Sect. 4.1) how Top-Down Guidance by itself allows us to resist the Anti-Dogmatism Argument. And with the addition of Top-Down Fixation, we now have a way of accepting the Anti-Akrasia Argument. How so? The thing to notice here is that anyone who is rationally higher-order certain about q will satisfy an instance of the following conditional: $[P(P(q) = c) = 1] \rightarrow [P(q) = c]$. And this conditional is just a special case of Rational Reflection. So, anyone who is higher-order certain will automatically satisfy Rational Reflection, and hence be epistemically non-akratic.

All of this just goes to show that Top-Down Fixation is what we *need*. But why should we *accept* it? Once again, the key motivation is that Top-Down Fixation features in an intuitively plausible story about why we should be sensitive to higher-order evidence; a story that also allows us to resolve a deep tension between the Anti-Akrasia Argument and the Anti-Dogmatism Argument. This might not seem like much of a motivation. Absent any independent reasons in its favor, isn't Top-Down Fixation just an ad hoc fix?

I don't think so. In the wider scheme of things, we can see Top-Down Fixation and Top-Down Guidance as forming part of a broadly "internalist" view of rational belief, according to which facts about what you should believe are determined solely by factors that are internally accessible to the agent. To bring out this connection, notice that Top-Down Guidance and Top-Down Fixation jointly entail the following principle:

Access Internalism: $[P(q) = x] \rightarrow P[P(q) = x] = 1$

According to this principle, agents are always in a position to "access" the requirements of rationality, whatever they are. More precisely, if you should in fact have such-and-such a credence in q , then you should be certain of that fact. This is, of course, just to say that rationality demands higher-order certainty. So, whatever reasons we have to think that rationality demands higher-order certainty carry over as motivation for Top-Down Fixation and Top-Down Guidance.

Now, I realize that this line of motivation is not independent of the issue at hand. After all, the reasons I've offered for thinking that rationality demands higher-order certainty stem, in large part, from the need to resolve the tension between the Anti-Akrasia Argument and the Anti-Dogmatism Argument. But at the very least, I take these considerations to show that Top-Down Fixation, far from being an ad hoc fix, dovetails quite nicely with the overall picture that I've set out to explore in this paper.

There are, to be sure, general worries that one might have about Access Internalism and other internalist principles in epistemology. For example, those who are attracted to the kinds of "anti-luminosity" considerations that have motivated Timothy Williamson's externalist program in epistemology will be skeptical of the idea that we have privileged access to certain facts about what we should believe.³¹ However, note that someone who rejects Access Internalism on such general externalist grounds will already have decided his or her mind against the view that rationality demands higher-order certainty. So, to avoid trivializing the question under discussion, I will proceed on the assumption that it hasn't already been settled by whatever reasons one might have to pick sides in the internalism/externalism debate about epistemic rationality.

³¹ Williamson (2000, 2014, 2019). See also Lasonen-Aarnio (2014, 2019), Worsnip (2018), and Salow (2019).

I would, however, like to address two more specific worries about Top-Down Fixation and Top-Down Guidance that have come to my attention.

The first worry is that Top-Down Fixation might seem to reinstate the very kind of dogmatic insensitivity to higher-order evidence that fueled the Anti-Dogmatism Argument in the first place. After all, doesn't Top-Down Fixation effectively require you to be certain of what your *higher-order* evidence supports, even if it allows you to be uncertain of what your *first-order* evidence supports? And if you are certain of what your higher-order evidence supports, shouldn't you take a dogmatic stance towards any *further* higher-order evidence you might later get about what your *initial* higher-order evidence supports?

Let's straightaway concede that the answer to the latter question is "yes:" if you are certain of what your higher-order evidence supports, you shouldn't be sensitive to any further higher-order evidence that you might get about what your initial higher-order evidence supports. But the answer to the former question is "no:" Top-Down Fixation doesn't require you to be certain of what your higher-order evidence supports. All Top-Down Fixation says is that you should be certain of what your *total* evidence supports. And this leaves plenty of room for uncertainty about what your higher-order evidence supports.

How so? The picture I have in mind (which I previously hinted at in Sect. 4) goes as follows: let's suppose that Miss Modest, after having learned about the self-enhancement bias, remains uncertain about whether she has judged the higher-order evidence correctly. This only seems reasonable given the kind of fallible being that she is. However, let's also suppose that she recognizes that her *total* evidence encompasses all of her grounds for doubting her own ability to judge various bodies of evidence, and that she therefore sees no reason to be uncertain about what *it*, the total evidence, supports. In other words, suppose that, upon recognizing that her overall assessment of what her total evidence supports already takes into account all of her uncertainty about her own cognitive fallibility, she is left with no residual doubts about what she should, all-things-considered, believe. (In a moment, I'll consider some challenges to this picture, but let's first see what follows from it.)

Now, let's suppose that Miss Modest goes on to learn that most people have a tendency to *overcompensate* when revising their credences in light of evidence about their own biases. That is, she learns that people tend to revise their credences *too much* in light of such information. How should Miss Modest respond to this further piece of higher-order evidence? She should think to herself: "I'm likely to have reduced my confidence too much in light of the evidence about the self-enhancement bias. I better undo some of that reduction!" In other words, as long as Miss Modest is uncertain about whether she has taken the evidence about the self-serving bias correctly into account, she should take seriously any further higher-order evidence that she might get about what the evidence about the self-serving bias supports. That is why Top-Down Fixation does not, in fact, reinstate any kind of dogmatic insensitivity to higher-order evidence.

The response above might seem a bit too quick. Haven't we merely push the problem back one step? The worry I have in mind here goes as follows: the reason why Miss Modest should be sensitive to the *third-order* evidence (i.e., the evidence about people's tendency to overcompensate when revising their credences in light of information

about their own biases) is that she is uncertain about what the *second-order* evidence (i.e., the evidence about the self-serving bias) supports. But, presumably, she should also be sensitive to any *fourth-order* evidence that she might later get about what the *third-order* evidence supports, in which case she must also be uncertain about what the *third-order* evidence supports. Continuing this line of thought, we seem forced to conclude that Miss Modest should, on pain of dogmatism, have higher-order uncertainty “all the way up:” that is, she should have uncertainty about what her *n*th-order evidence supports, for all values of *n*. Doesn’t this mean that she must be uncertain about what her *total* evidence supports?

What I say in response to this worry will have to be somewhat tentative. My contention is that, although Miss Modest should have level-uncertainty “all the way up,” she shouldn’t be uncertain about what her total evidence supports. Why not? Because there something special about your total evidence—something that sets it apart from any subset of your evidence—namely that it encompasses all of the evidence you possess about your own ability to judge various bodies of evidence. As such, when you assess what your total evidence supports, you should take into account all of the reasons you might have to doubt your own ability to assess various parts of your evidence. In particular, a complete assessment of your total evidence should take into account any uncertainty that you might rationally have about what your *n*th-order evidence supports, for any *n*. Once you have done so, however—that is, once you have taken into account all of your doubts about your own cognitive fallibility—you should have no *residual* doubts about what you should, all-things-considered, believe. Put differently, once you have “calibrated” your credence in light of your uncertainty about what your 2nd-order evidence, 3rd-order evidence, 4th-order evidence, etc.... supports, you shouldn’t have any residual uncertainty about whether that credence is rational. If you *did*, your calibration process would remain incomplete. That’s the tentative proposal.

As stated, the proposal is clearly incomplete. It relies on a highly non-trivial “Convergence Conjecture,” as we might call it, which says (roughly) that the process of calibrating your expectation of $P_F(q)$ in light of your higher-order uncertainty at various levels converges to a limit whose value is certain given your total evidence. A bit more precisely, it says that the iterated process of calculating the expected value of $P_F(q)$ in light of your rational uncertainty about what your 2nd-order evidence, 3rd-order evidence, 4th-order evidence, etc.... supports satisfies the following two properties: (a) the process converges to a value for the expectation of $P_F(q)$ in the infinite limit, and (b) your total evidence leaves no room for rational uncertainty about what the limit value for the expectation of $P_F(q)$ is. This Convergence Conjecture, if true, ensures that it can never be rational to be uncertain about what one should expect one’s first-order evidence to support, though it can be rational to be uncertain about what one’s *n*th-order evidence supports, for any *n*. Consequently, the Conjecture allows us to rely on Top-Down Fixation without undermining our ability to use Top-Down Guidance to resist the Anti-Dogmatism Argument. As such, the Conjecture isn’t just a technical curiosity, but in fact plays a central role in the proposed resolution of the tension between the Anti-Akrasia Argument and the Anti-Dogmatism Argument.

Yet, I don’t know of a way to defend the Convergence Conjecture on independent grounds. I suspect that the Conjecture can be defended on broadly internalist grounds.

But for now, the Conjecture must earn its keep through the role it plays in the overall proposal that I've tried to defend. If it turns out that the proposal must ultimately be rejected because the Convergence Conjecture cannot be upheld, this would be an interesting result in its own right.

The second worry I want to address starts from the observation that Top-Down Fixation, when combined with Top-Down Guidance, has a curious implication. The implication is that the actual value of $P_F(q)$ does not, in and of itself, make any difference to the value of $P(q)$. In other words, the fact of the matter as to what the first-order evidence supports does not, in and of itself, make any difference to what the total evidence supports. If it *did*, this would mean that it could be rational to have higher-order uncertainty (since we're assuming that it can be rational to be uncertain of what one's first-order evidence supports). However, and this is where the worry lies, it seems odd to say that what the first-order evidence in fact supports does not, in and of itself, make a difference to what the total evidence supports. Why else pay attention to what the first-order evidence supports in the first place?³²

Before I respond to this worry, let me try to spell out the worry in a little more detail. Suppose that you are rationally certain that the true value of $P_F(q)$ is either 30% or 80%, and suppose that you distribute (and rationally so) your credence evenly among these two possibilities: that is, $P(P_F(q) = 30\%) = 50\%$ and $P(P_F(q) = 80\%) = 50\%$. Now, suppose that the value of $P(q)$ depends on the value of $P_F(q)$ in the following way: if $P_F(q) = 30\%$, then $P(q) = 10\%$; and if $P_F(q) = 80\%$, then $P(q) = 90\%$. Given this, you should be uncertain about what the value of $P(q)$ is: you should be 50–50 about whether the value of $P(q)$ is 10% or 90%. After all, you are 50–50 about whether the value of $P_F(q)$ is 30% or 90%. But Top-Down Fixation and Top-Down Guidance jointly entail that it cannot be rational to have higher-order uncertainty. So, on pain of contradiction, the value of $P(q)$ cannot depend on the value of $P_F(q)$ after all. And yet this seems like an odd conclusion. Or so the worry goes.

While I think this worry is onto something interesting and important, I do not in fact think that it constitutes a *reductio* against the conjunction of Top-Down Fixation and Top-Down Guidance. Rather, I think it teaches us a surprising lesson about the normative role of first-order evidence. To appreciate what this role consists in, it will be worth reminding ourselves of how many epistemologists have traditionally conceived of the connection between *truth* and *rationality*. On the picture I have in mind, the actual truth-value of a proposition does not, in and of itself, make a difference to what it is rational to believe about that proposition. In other words, if we change the truth-value of a proposition without changing anything else about the agent's epistemic situation, the agent shouldn't change his or her credence. This does not render it mysterious why we should pay attention to the truth-value of a proposition in the first place. It just shows that the connection between truth and rationality is somewhat indirect: the truth-value of a proposition only has an indirect kind of influence on what it is rational to believe about that proposition—one that is mediated by the general truth-conduciveness of evidence.

I suggest that we think of first-order evidence as bearing a similar kind of connection to rationality. Just as the truth-value of q does not, in and of itself, make any difference

³² Thanks to an anonymous reviewer for raising this worry.

to the value of $P(q)$, so the value of $P_F(q)$ does not, in and of itself, make any difference to the value of $P(q)$. In other words, if we change the value of $P_F(q)$ without changing anything else about the agent's epistemic situation, the agent's credence shouldn't change. This does not render it mysterious why we should pay attention to the first-order evidence in the first place. It just shows that the connection between first-order evidence and rationality is somewhat indirect: what the first-order evidence supports only has an indirect kind of influence on what credence one should have in q —one that is mediated by the general truth-conduciveness of higher-order evidence.³³

This concludes the main positive part of the paper. Before we proceed, a brief summary is in order. I have proposed a view—let's call it *The "Top-Down" View*—which consists of two main principles: Top-Down Guidance and Top-Down Fixation. Together, these principles imply Access Internalism: the thesis that rationality demands higher-order certainty. There are two main attractions of the Top-Down View: (a) it offers an independently plausible explanation of why we should be prepared to revise our beliefs in response to higher-order evidence; and (b) it allows us to resolve a deep tension between the Anti-Akrasia Argument and the Anti-Dogmatism Argument. As announced in the introduction, however, the Top-Down View also comes with substantive commitments. In particular, it puts pressure on standard ways of thinking about rational belief within a Bayesian framework. That is the topic of our final section.

5 The Bayesian and the Top-Down Theorist

So far, I have imposed very few constraints on the functions P and P_F , aside from Top-Down Guidance and Top-Down Fixation. In particular, I haven't assumed that they obey the probability calculus (Probabilism), nor have I assumed that they are obtained from a shared urprior by conditionalizing on the total evidence and first-order evidence, respectively (Conditionalization). The reason for this is a substantive one: we run into problems once we try to integrate the Top-Down View with a standard Bayesian picture of rational belief. I'm not sure about how best to resolve this conflict. But I want to at least explain how the conflict arises, and gesture at some ways of responding to it.

First, let's recall one of the central features of the Top-Down View, namely that it can be rational to be uncertain about what your first-order evidence supports, while being certain about what your total evidence supports. This is what allowed us to resist the Anti-Dogmatism Argument. However, if we add Conditionalization to the mix, we can show that, at least under certain circumstances, it cannot in fact be rational to

³³ I also take these considerations to answer a related worry that is sometimes raised against broadly "conciliatory" views of higher-order evidence; see, e.g., Kelly (2005, 2010) for an influential development of this worry. The worry starts from an intuitive, albeit somewhat vague, idea, namely that you should "respect *all* the evidence" (cf. Sliwa and Horowitz 2015). In particular, you should respect both your first-order evidence and your higher-order evidence. Otherwise, you effectively ignore or "throw away" part of your evidence, which seems patently irrational. But how, exactly, should you respect both the first-order evidence and the higher-order evidence? I say: by treating both your first-order evidence and your higher-order evidence as "truth-guides." More specifically, you should treat your first-order evidence as a guide to the truth, and you should treat your higher-order evidence as a guide to the truth about what your first-order evidence supports.

be uncertain about what your first-order evidence supports while being certain about what your total evidence supports.

To see why, consider again Miss Modest's situation after she has learned about the self-enhancement bias. Let E_F be her first-order evidence, let E_H be her higher-order evidence, and let E be her total evidence. Now, suppose that there was a time when Miss Modest's total evidence was exhausted by E_F . That is, suppose that her current first-order evidence used to be her total evidence. At that time, she should then have the credences that are licensed by P_F . In particular, she should be certain about how strongly E_F supports q : that is, $P_F(P_F(q) = c_F) = 1$ (where c_F is the true value of $P_F(q)$). This follows from the conjunction of Top-Down Guidance and Top-Down Fixation. Given Conditionalization, this certainty is preserved in light of any new evidence, including E_H : that is, $P_F(P_F(q) = c_F | E_H) = 1$. Furthermore, Conditionalization implies that $P(P_F(q) = c_F) = P_F(P_F(q) = c_F | E_H)$. Thus, we get the result that Miss Modest should now be certain about what her first-order evidence supports: that is, $P(P_F(q) = c_F) = 1$. But this contradicts our initial assumption.

Does this result force the top-down theorist to reject Conditionalization? Not straightforwardly. There is at least one natural way we might try to block the result without rejecting Conditionalization, namely by denying that E_F could ever have exhausted Miss Modest's total evidence. On the face of it, this might seem blatantly ad hoc. But, on closer inspection, it is not all that easy to think of cases where an agent has no higher-order evidence whatsoever. This may be passed off as a mere shortcoming of our imagination (due, perhaps, to a lack of experience with real-world situations of the relevant kind). But I think the point runs deeper than that. After all, even if you *did* find yourself without any evidence about the reliability of your own cognitive faculties, it is far from clear, as Christensen (2019, §4) rightly points out, what you should think about your own reliability in that situation.

Suppose, for example, that you have no track-record whatsoever about your own ability engage in logical reasoning (nor, indeed, any other relevant evidence). How reliable should you take yourself to be at solving logic problems? It's hard to say. The question brings us right up against the "Problem of the Priors," which lies at the heart of the ongoing dispute between "subjective" and "objective" Bayesians. It is notoriously difficult to say anything much about what an agent's urpriors should look like. In particular, it is difficult to say what agents should think about the reliability of their own cognitive faculties in the absence of any evidence. But whatever else we might want to say about this matter, it seems plausible that one should not take oneself to be *infallible* at judging evidential relations. Contingent facts about one's own cognitive faculties are presumably not the kinds of things that one can be rationally certain of on purely a priori grounds.³⁴

Now, merely constraining Miss Modest's urpriors is clearly not enough to block the undesired result. After all, we have said nothing to rule out the possibility that E_F was once Miss Modest's total evidence, in which case it still follows from the conjunction of Access Internalism and Conditionalization that $P(P_F(q) = c_F) = 1$. However, we can block the result by adding a further assumption to the effect that, when you acquire

³⁴ This is not, to be sure, to deny that we have a kind of default "entitlement" to trust our own cognitive faculties (Dretske 2000; Pryor 2000; Wright 2004). The claim is just that this sort of default justification, if we have it, does not entitle us to be certain that our cognitive faculties are infallible.

a body of evidence and form an opinion on the basis of having judged that evidence, you also learn *that* you have formed an opinion on the basis of having judged that evidence.³⁵ This strikes me as an intuitively attractive assumption. I will not try to defend it.³⁶ But if it holds, it follows that Miss Modest’s total evidence could never have been exhausted by E_F . At the very least, her evidence would also include the fact *that* she has judged E_F . And when combined with the constraint on Miss Modest’s urpriors just introduced, this means that she should not become certain about what E_F supports after having acquired E_F . If so, the conjunction of Access Internalism and Conditionalization no longer has the undesired result that $P(P_F(q) = c_F) = 1$.

Of course, there is something a bit unseemly about making limiting assumptions about which bodies of evidence it is possible for agents to have. Would we do better to reject Conditionalization instead? I suspect that many formal epistemologists, and many philosophers more generally, will be hesitant to go this way. However, there are independent reasons to think that *anyone* who wants to maintain that higher-order evidence is normatively significant must ultimately reject Conditionalization, at least as the norm is standardly conceived.

To see why, let’s begin by considering a case adapted from Christensen (2010a, §4):

Science on Drugs: Dr. Scientia is going to perform a scientific experiment tomorrow, Monday, to test a hypothesis, H . The experiment has two possible outcomes: A and B. Dr. Scientia, being a brilliant scientist, is well aware of how the outcomes bear on H : A confirms H , whereas B disconfirms H . She is also aware, however, that she is going to be affected by a reason-disrupting drug tomorrow, which will make her unable to recognize whether the experimental outcome confirms H or not.

Following Christensen, let’s ask two questions about Dr. Scientia. First question: On Sunday, how confident should Dr. Scientia be of H on the supposition that the experiment delivers outcome A, and that she is going to be drugged on Monday? Intuitively, “very confident!” After all, the fact that she is going to be drugged some time in the future seems utterly irrelevant to the question of whether outcome A supports H . Second question: On Monday, given that the experiment in fact delivers outcome A, how confident should Dr. Scientia’s *then* be in H ? Intuitively, “not very confident at all!” After all, she is well aware in the given situation that she is affected by a drug that makes her unable to recognize how the experimental outcome bears on H . Taken together, these two verdicts put pressure on Conditionalization, since they suggest that Dr. Scientia’s posterior credence in H should not match her prior conditional credence in H .

I say “put pressure on” because, as Christensen (2010a, p. 201) also notes, it might be possible to rescue Conditionalization by understanding Science on Drugs as a case of

³⁵ Something like this principle is what Schoenfield (2015, p. 439) calls “Inclusion.”

³⁶ Although it might be worth emphasizing what the assumption is *not*: the assumption is not that, whenever you form an opinion on the basis of having judged some evidence, the fact that you have judged the evidence will be at the forefront of your mind. Nor is the assumption that you will be aware of *how* you went about judging the evidence. The assumption is merely that, when you form an opinion on the basis of having judged some evidence, the fact that you have judged the evidence will be part of your evidence.

self-locating evidence. Here is the idea in rough outline: Rather than conditionalizing on “I’m drugged on Monday,” Dr. Scientia should conditionalize on the self-locating proposition “I’m drugged now.” Given this, our first question from above becomes: On Sunday, how confident should Dr. Scientia be of H on the supposition that the experiment will deliver outcome A , and that she is drugged now? The intuitive answer to *this* question is “not very confident at all!” That’s good news for Conditionalization, since it removes the mismatch between Dr. Scientia’s posterior credence in H and her prior conditional credence in H .

However, it remains to be seen whether this “self-locating evidence” strategy holds up to scrutiny. In recent work, Schoenfield (2018) has cast doubt on the project.³⁷ But in any case, the problems for Conditionalization do not stop here.

Consider another case adapted from Christensen (2010a):

Logic on Drugs: Dr. Logos has just finished a proof that T is a tautology. The proof is impeccable. And not by accident. Dr. Logos is a brilliant logician. However, he now discovers that he has been slipped an imperceptible, but highly powerful, drug that is likely to have impaired his ability to engage in logical reasoning.

On the assumption that T is in fact a tautology, it is trivially the case that T is entailed by Dr. Logos’ total evidence. Thus, if P were obtained by conditionalizing on Dr. Logos’ total evidence, Dr. Logos would be rational to be certain of T : that is, $P(T) = 1$. However, anyone who thinks that higher-order evidence is normatively significant presumably wants to maintain that Dr. Logos *shouldn’t* be certain of T after he has learned that he has been drugged. Thus, Logic on Drugs seems to put pressure on Conditionalization.

For the same reason, Logic on Drugs also puts pressure on Probabilism. You comply with Probabilism only if you are certain of all tautologies. So, if Dr. Logos should not be certain of T , he should not comply with Probabilism.

What does the Top-Down View say about cases like Logic on Drugs? Let’s concede that the first-order evidence indeed supports T to the maximum degree: $P_F(T) = 1$. Even so, the top-down theorist can maintain that Dr. Logos should not be certain that the first-order evidence supports T to the maximum degree: $P(P_F(T) = 1) < 1$. From this it follows that Dr. Logos should *expect* the first-order evidence to support T to some degree below one: $\mathbb{E}_P(P_F(T)) < 1$. Thus, by Top-Down Guidance, we get the result that Dr. Logos should be less than fully certain of T : $P(T) < 1$.³⁸

³⁷ In brief, she argues that, if we treat higher-order evidence as a kind of self-locating evidence, then (given certain background assumptions) the update procedure that maximizes expected accuracy says that we should *not* be sensitive to higher-order evidence.

³⁸ An anonymous reviewer raises the worry that, in the absence of probabilism, Elga’s result discussed in Sect. 2 no longer goes through, thereby undermining the Anti-Akrasia Argument, which was a crucial part of the motivation behind the Top-Down View. However, as can be seen from the proof stated in footnote 8, Elga’s result in fact doesn’t presuppose probabilism, but only something much weaker, namely that the conditional probability of a proposition given itself is 1. This is what allows us to say that $P(P'$ is ideal | P' is ideal) = 1, which, when combined with Rational Reflection, implies that P' (P' is ideal) = 1. Thus, the move to non-probabilism does not undermine the motivation behind the Top-Down View.

All of this goes to show, at least in a preliminary way, that the Top-Down View does not sit well with orthodox Bayesianism. But it also looks like this might in fact be a feature of the view, not a bug.

6 Conclusion

I began this paper by describing two individually compelling arguments, which give opposing answers to the question of whether rationality demands higher-order certainty. On the one hand, the Anti-Akrasia Argument would have us think that rationality demands higher-order certainty, since rational higher-order uncertainty leads to violations of Rational Reflection, and hence runs afoul of the Enkratic Intuition. On the other hand, the Anti-Dogmatism Argument would have us think that rationality does *not* demand higher-order certainty, since rational higher-order certainty leads to a dogmatic kind of insensitivity to higher-order evidence, and hence runs afoul of the Modesty Intuition.

My claim has been that there is an independently plausible way of resolving this tension, which allows us to vindicate both the Enkratic Intuition and the Modesty Intuition. To make good on this claim, I have presented a view of the normative significance of higher-order evidence—I called it the “Top-Down View”—which consists of two core principles. The first principle, *Top-Down Guidance*, says that your credence in any given proposition should match your rational expectation of how strongly your first-order evidence supports that proposition. The second principle, *Top-Down Fixation*, says that you should never be uncertain about what you should expect your first-order evidence to support. Together, these two principles allow us to resist the Anti-Dogmatism Argument while accepting the Anti-Akrasia Argument: they allow us to resist the Anti-Dogmatism Argument by making higher-order certainty compatible with uncertainty about what your first-order evidence supports; and they allow us to accept the Anti-Akrasia Argument by validating Rational Reflection (in a trivial way). The upshot: the Top-Down View offers a new (and so far the only, to my knowledge) way of vindicating both the Modesty Intuition and the Enkratic Intuition. And it does so in a way that is grounded in an independently plausible story about why we should be prepared to revise our beliefs in response to higher-order evidence in the first place.

Where to go from here? There are at least two paths to be explored.

The first is toward *applications*: what are the consequences of the Top-Down View for other related debates about, e.g., “evidence of evidence” principles (Feldman 2007; Williamson 2019), “unmarked clock” puzzles (Elga 2013; Williamson 2014), and “logical omniscience” requirements (Christensen 2007b; Smithies 2015)? The view put forth will, no doubt, deliver broadly conciliationist answers to these questions. But the finer details await our exploration.

The second path is toward *modelling*: it is natural to hope that the Top-Down View can be embedded in a precise, model-theoretic framework that will allow us to study its consequences in a more rigorous way. We already know how to characterize principles such as Rational Reflection and Access Internalism in terms of conditions on probability frames.³⁹ However, the standard machinery of probabilistic modal logic

³⁹ See, e.g., Williamson (2000, 2014, 2019), Dorst (2019a, b), and Lasonen-Aarnio (2015).

is not straightforwardly applicable to principles like Top-Down Guidance and Top-Down Fixation, for at least two reasons. First, we cannot simply encode information about an agent's first-order evidence *relative to a proposition* in the same way that we would normally encode information about the agent's total evidence *simpliciter*. This is due to the proposition-relative nature of the distinction between first-order evidence and higher-order evidence described in Sect. 4. Second, we cannot simply define the functions P_F and P in terms of a shared urprior, which is conditionalized on the first-order evidence and total evidence, respectively. This is due to the conflicts with orthodox Bayesianism described in Sect. 5. If we want to study the Top-Down View from a model-theoretic perspective, we somehow need to overcome these challenges. I see no principled reason to think that this cannot be done. But even without a fully-fledged model theory on the table, it seems to me that the Top-Down View is an important view to consider.

Acknowledgements Earlier versions of this paper were presented at Cologne University, Aarhus University, and the “MATTI” seminar at MIT. I'd like to thank the audiences on those occasions for helpful comments and questions. A special thanks to David Christensen, Kevin Dorst, Mikkel Gerken, Jack Spencer, Eyal Tal, two anonymous reviewers at *The Philosophical Review*, and two anonymous reviewers at *Synthese* for feedback that greatly improved the paper.

References

- Broome, J. (1999). Normative requirements. *Ratio*, 12, 398–419.
- Broome, J. (2007). Wide scope or narrow scope? *Mind*, 116, 359–370.
- Christensen, D. (2007a). Epistemology of disagreement: The good news. *The Philosophical Review*, 116, 187–217.
- Christensen, D. (2007b). Does Murphy's law apply in epistemology? *Oxford Studies in Epistemology*, 2, 3–31.
- Christensen, D. (2010a). Higher-order evidence. *Philosophy and Phenomenological Research*, 81, 185–215.
- Christensen, D. (2010b). Rational reflection. *Philosophical Perspectives*, 24, 121–140.
- Christensen, D. (2011). Disagreement, question-begging, and epistemic self-criticism. *Philosophers' Imprint*.
- Christensen, D. (2016). Disagreement, drugs, etc.: From accuracy to Akrasia. *Episteme*, 13, 397–422.
- Christensen, D. (2019). Formulating independence. In M. Skipper & A. Steglich-Petersen (Eds.), *Higher-order evidence: New essays*. Oxford: Oxford University Press.
- Dorst, K. (2019a). Higher-Order Uncertainty. In M. Skipper & A. Steglich-Petersen (Eds.), *Higher-order evidence: New essays*. Oxford: Oxford University Press.
- Dorst, K. (2019b). Evidence: A guide for the uncertain. *Philosophy and Phenomenological Research*, 100, 586–632.
- Dretske, F. (2000). Entitlement: Epistemic rights without epistemic duties? *Philosophy and Phenomenological Research*, 60, 591–606.
- Elga, A. (2005). On overrating oneself... and knowing it. *Philosophical Studies*, 123, 115–124.
- Elga, A. (2007). Reflection and disagreement. *Noûs*, 41(3), 478–502.
- Elga, A. (2010). Subjective probabilities should be sharp. *Philosophers' Imprint*, 10.
- Elga, A. (2013). The puzzle of the unmarked clock and the new rational reflection principle. *Philosophical Studies*, 164(1), 127–139.
- Feldman, R. (2007). Reasonable religious disagreements. In L. Antony (Ed.), *Philosophers without god: Meditations on atheism and the secular life* (pp. 194–214). Oxford: Oxford University Press.
- Greco, D. (2014). A puzzle about epistemic Akrasia. *Philosophical Studies*, 167, 201–219.
- Horowitz, S. (2014). Epistemic Akrasia. *Noûs*, 48, 718–744.
- Hughes, N. (2019). Dilemmic epistemology. *Synthese*, 196, 4059–4090.

- Joyce, J. (2019). Bayes' theorem. In E. Zalta (ed.), *The stanford encyclopedia of philosophy*. <https://plato.stanford.edu/archives/spr2019/entries/bayes-theorem/>.
- Kelly, T. (2005). The epistemic significance of disagreement. In Gendler, T., & Hawthorne, J. (Eds.), *Oxford studies in epistemology* (vol. 1, pp. 167–196).
- Kelly, T. (2010). Peer disagreement and higher-order evidence. In A. Goldman & D. Whitcomb (Eds.), *Social epistemology: Essential readings* (pp. 183–217). Oxford: Oxford University Press.
- Kelly, T. (2014). Evidence. In Zalta, E. (Ed.) *The Stanford encyclopedia of philosophy*.
- Kiesewetter, B. (2017). *The normativity of rationality*. Oxford: Oxford University Press.
- Kolodny, N. (2005). Why be rational? *Mind*, 455, 509–563.
- Kripke, S. (2011). Two paradoxes of knowledge. *Philosophical troubles: Collected papers I* (pp. 27–51). New York: Oxford University Press.
- Lasonen-Aarnio, M. (2014). Higher-order evidence and the limits of defeat. *Philosophy and Phenomenological Research*, 88, 314–345.
- Lasonen-Aarnio, M. (2015). New rational reflection and internalism about rationality. In T. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology*. Oxford: Oxford University Press.
- Lasonen-Aarnio, M. (2019). Defeat as Evincibility. In M. Skipper & A. Steglich-Petersen (Eds.), *Higher-order evidence: New essays*. Oxford: Oxford University Press.
- Lasonen-Aarnio, M. (2020). Enkrasia or evidentialism? Learning to love mismatch. *Philosophical Studies*, 177, 597–632.
- Pryor, J. (2000). The skeptic and the dogmatist. *Noûs*, 34, 517–549.
- Salow, B. (2019). Elusive externalism. *Mind*, 128, 397–427.
- Samet, D. (1997). On the triviality of high-order probabilistic beliefs. Unpublished manuscript. <https://ideas.repec.org/p/wpa/wuwpaga/9705001.html>.
- Schoenfield, M. (2013). Permission to believe: Why permissivism is true and what it tells us about irrelevant influences on belief. *Noûs*, 47(1), 193–218.
- Schoenfield, M. (2015). A dilemma for calibrationism. *Philosophy and Phenomenological Research*, 91(2), 425–455.
- Schoenfield, M. (2018). An accuracy based approach to higher order evidence. *Philosophy and Phenomenological Research*, 96, 690–715.
- Schultheis, G. (2018). Living on the edge: Against epistemic permissivism. *Mind*, 507, 863–879.
- Skipper, M. (2019a). Higher-order defeat and the impossibility of self-misleading evidence. In M. Skipper & A. Steglich-Petersen (Eds.), *Higher-order evidence: New essays*. Oxford: Oxford University Press.
- Skipper, M. (2019b). Reconciling Enkrasia and higher-order defeat. *Erkenntnis*, 84, 1369–1386.
- Sliwa, P., & Horowitz, S. (2015). Respecting all the evidence. *Philosophical Studies*, 172, 2835–2858.
- Smithies, D. (2012). Moore's paradox and the accessibility of justification. *Philosophy and Phenomenological Research*, 85, 273–300.
- Smithies, D. (2015). Ideal rationality and logical omniscience. *Synthese*, 192(9), 2769–2793.
- Smithies, D. (2019). *The epistemic role of consciousness*. Oxford: Oxford University Press.
- Tal, E. (2018). Self-intimation, infallibility, and higher-order evidence. *Erkenntnis*, 85, 665–672.
- Tal, E. (2020). Is higher-order evidence evidence? *Philosophical Studies*.
- Titelbaum, M. (2015). Rationality's fixed point (or: in defense of right reason). In Gendler, T., & Hawthorne, J. (Eds.), *Oxford studies in epistemology* (Vol. 5, pp. 253–294).
- van Fraassen, B. (1984). Belief and the will. *The Journal of Philosophy*, 81, 235–256.
- Weatherston, B. (2019). *Normative externalism*. Oxford: Oxford University Press.
- White, R. (2009). Evidential symmetry and mushy credence. In Gendler, T., & Hawthorne, J. (Eds.), *Oxford studies in epistemology* (pp. 161–186).
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.
- Williamson, T. (2014). Very improbable knowing. *Erkenntnis*, 79(5), 971–999.
- Williamson, T. (2019). Evidence of evidence in epistemic logic. In M. Skipper & A. Steglich-Petersen (Eds.), *Higher-order evidence: New essays*. Oxford: Oxford University Press.
- Worsnip, A. (2018). The conflict of evidence and coherence. *Philosophy and Phenomenological Research*, 96, 3–44.
- Wright, C. (2004). On epistemic entitlement: Warrant for nothing (and foundations for free?). *Proceedings of the Aristotelian Society*, 78, 167–212.