Higher-Order Evidence and the Normativity of Logic

Mattias Skipper

Penultimate draft. To appear in *Epistemic Dilemmas: New Arguments, New Angles*, S. Stapleford, K. McCain, and M. Steup (eds.), Routledge.

Abstract: Many theories of rational belief give a special place to logic. They say that an ideally rational agent would never be uncertain about logical facts. In short: they say that ideal rationality requires *logical omniscience*. Here I argue against the view that ideal rationality requires logical omniscience on the grounds that the requirement of logical omniscience can come into conflict with the requirement to proportion one's beliefs to the evidence. I proceed in two steps. First, I rehearse an influential line of argument from the "higher-order evidence" debate, which purports to show that it would be dogmatic, even for a cognitively infallible agent, to refuse to revise her beliefs about logical matters in response to evidence indicating that those beliefs are irrational. Second, I defend this "anti-dogmatism" argument against two responses put forth by Declan Smithies and David Christensen. Against Smithies' response, I argue that it leads to irrational self-ascriptions of epistemic luck, and that it obscures the distinction between propositional and doxastic justification. Against Christensen's response, I argue that it clashes with one of two attractive deontic principles, and that it is extensionally inadequate. Taken together, these criticisms will suggest that the connection between logic and rationality cannot be what it is standardly taken to be—ideal rationality does not require logical omniscience.

Keywords: Higher-order evidence, logical omniscience, normativity of logic, rational ideals, epistemic dilemmas, epistemic rationality

1. Introduction

When an agent is certain about all logical facts, we can say that the agent is *logically omniscient*. Such an agent never has any doubts about what logical relations hold between what propositions: what propositions entail what propositions, what propositions are tautological or contradictory, and so on.

Would a *rational* agent be logically omniscient? The answer arguably depends on just what kind of agent we have in mind. But when epistemologists consider the matter,

they often abstract away from human cognitive limitations, and ask: would an *ideally rational agent*—someone who doesn't suffer from any cognitive limitations—be logically omniscient?

Many theories of rational belief say yes. They see logical uncertainty as a sure sign of irrationality: someone who is logically non-omniscient could never be perfectly rational. The precise nature of this commitment depends, among other things, on the operative notion of "belief." Those who model belief as a binary, all-or-nothing, kind of attitude have traditionally claimed that ideal rationality requires *deductive cogency*: that is, beliefs which are mutually consistent and closed under logical consequence.[1] Those who model belief as a graded kind of attitude, one that comes in degrees, have traditionally claimed that ideal rationality requires *probabilistic coherence*: that is, degrees of belief (or "credences") which obey the standard axioms of probability theory.[2]

What has led philosophers to think that ideal rationality requires logical omniscience? Part of the motivation comes from sophisticated arguments involving, e.g., Dutch books (Ramsey 1926) and accuracy-dominance principles (Joyce 1998).[3] But perhaps the most basic line of motivation has its roots in a long-standing tradition of viewing logic as a *normative* discipline.[4] On this way of conceiving of logic, the most fundamental question of the discipline is "How should I reason?" (or some variant thereof: "What constitutes good reasoning?," "How would a rational agent reason?," etc.). If this is the proper subject matter of logical theorizing, it would seem to leave little room for doubt about whether logic is normative for reasoning. And, given the close connection between reasoning and belief, it then becomes natural to think that logic would be normative for belief as well, and that—under the right kinds of idealized circumstances—a rational agent would exemplify a certain kind of logical perfection akin to deductive cogency or probabilistic coherence.

---

[1] Two influential examples come from standard epistemic logic in the tradition of Hintikka (1962) and the AGM theory of belief revision (Alchourrón et al. 1985).

[2] Here the best-known example comes from orthodox Bayesian epistemology—see Titelbaum (forthcoming) for background.

[3] See Pettigrew (2020) for background on Dutch book arguments. For background on accuracy-first epistemology, see Greaves (2013) and Pettigrew (2016).

[4] This conception of logic goes back at least to Frege (1893/1903).

Now, this way of conceiving of logic has obviously not gone unchallenged. Some philosophers have argued that there is no rational pressure to believe the logical consequences of what you believe, if those consequences are of no interest to you;[5] others have taken the Preface Paradox (and similar paradoxes) to show that rational belief is not closed under conjunction;[6] yet others have argued that logical pluralism (the view that there is more than one correct logic) clashes with the view that logic is normative;[7] and some have argued on more general grounds that requirements of formal coherence (including logical requirements) are not normative at all.[8]

My topic here is not these familiar concerns. Rather, I will argue against the view that ideal rationality requires logical omniscience on the grounds that the requirement of logical omniscience can come into conflict with the widely (though not universally[9]) accepted requirement to proportion one's beliefs to the evidence. I will proceed in two steps. First, I will rehearse an influential line of argument from the "higher-order evidence" debate, which purports to show that it would be dogmatic, even for a cognitively infallible agent, to refuse to revise her beliefs about logical matter in response to evidence about the rational status of those beliefs (§2). Second, I will defend this "anti-dogmatism" argument against two responses put forth by Declan Smithies (2015; 2019) and David Christensen (2007; 2010; 2016). Against Smithies' response, I will argue that it leads to irrational self-ascriptions of epistemic luck, and that it obscures the distinction between propositional and doxastic justification (§3.1). Against Christensen's response, I will argue that it clashes with one of two attractive deontic principles, and that it is extensionally inadequate (§3.2). Taken together, these criticisms will suggest that the connection between logic and rationality cannot be what it is standardly taken to be—ideal rationality does not require logical omniscience.

---

[5] See Harman (1986) and more recently Friedman (2018). Attempts to meet Harman's challenge can be found in MacFarlane (ms), Field (2009), and Steinberger (2019b; 2019c).

[6] See, e.g., Foley (1992) and Christensen (2004, ch. 3). For an attempt to avoid this result, see Leitgeb (2014).

[7] See Steinberger (2019a).

[8] See Kolodny (2005). For a response to Kolodny's challenge, see Kiesewetter (2017).

[9] For some influential defenses of this sort of "evidentialist" requirement, see Conee and Feldman (1985; 2004) and Shah (2006). For arguments against it, see Littlejohn (2018) and Rinard (2015), among others.

2. The conflict: from higher-order defeat to logical uncertainty

There is a simple and, I think, compelling line of argument, which has featured prominently in recent debates about higher-order evidence, and which can be used to put pressure on the view that ideal rationality requires logical omniscience. Below I present what I take to be the main thrust of the argument, leaving out various bells and whistles that won't matter for present purposes.

We begin by considering a paradigmatic case of higher-order defeat, which I've adapted from Christensen (2007):[10]

> **Dr. Smart and the Logic Drug:** Dr. Smart has written down a proof showing that *L* is a logical truth. The proof is impeccable; and not by accident. Dr. Smart is a brilliant logician. However, she now learns that she has been slipped an imperceptible, but powerful drug, which is likely to have impaired her ability to engage in logical reasoning while she was carrying out the proof.

The question we will be interested in is: How, if at all, should Dr. Smart revise her credence in *L* upon learning that she was drugged while carrying out the proof?

To tackle this question, it will help to first unpack the case a bit further. There are two salient factors that might be thought to have a bearing on what credence in *L* Dr. Smart should end up with. On the one hand, there is the fact that she has written down a correct and competently produced proof of *L*. On the other hand, there is the fact that she has received strong evidence indicating that she was cognitively impaired while carrying out the proof. This latter piece of evidence is an example of what has come to be known as *higher-order evidence*—so called because the evidence doesn't seem to bear straightforwardly on the "object-level" question of whether *L* is *true*, but rather seems to bear on the "meta-level" question of what credence is *rational* for Dr. Smart to have in *L*.[11]

---

[10] Similar cases have been discussed in many other places, including Coates (2012), Dorst (2019), Horowitz (2014), Lasonen-Aarnio (2014), Schoenfield (2015; 2018), Skipper (2019; forthcoming), Whiting (2020), and Worsnip (2018).

[11] For a more detailed characterization of the distinction between first-order and higher-order evidence, see Christensen (2010) and Skipper (forthcoming).

There is also a third factor that might be thought to have a bearing on what credence in *L* Dr. Smart should end up with, namely, whether she is cognitively fallible or not. I have deliberately left this matter open in the description of the case above. Ultimately, we are interested in how Dr. Smart should respond to the higher-order evidence given that she is a cognitively infallible agent (in particular, given that she never makes any mistakes in logical reasoning). But it will be instructive to begin by supposing that she is a fallible agent like you and me.

So, assume that Dr. Smart is a cognitively fallible agent. How should she respond to the higher-order evidence? According to a very natural line of reasoning, she should *lower* her credence in *L*, at least to some extent. David Christensen has nicely articulated this line of reasoning in a series of papers.[12] He asks us to consider what it would look like if Dr. Smart *didn't* lower her credence in *L* upon learning that she was drugged while carrying out the proof. As he points out, it would look as if she took the higher-order evidence to be misleading. After all, the higher-order evidence is non-misleading just in case Dr. Smart was cognitively impaired while carrying out the proof, in which case she most likely made an error somewhere in the reasoning process. Yet, and this is the crucial point, it would seem patently dogmatic (at least in some sense of "dogmatic"[13]) if Dr. Smart were to simply assume the higher-order evidence to be misleading. After all, she has no independent reason (or so we may assume) to think herself immune to the relevant drug. Thus, we seem led to the conclusion that, on pain of dogmatism, Dr. Smart *should* lower her credence in *L* upon learning that she was drugged.

Now, some epistemologists have tried to resist this "anti-dogmatism" argument by arguing that, despite appearances to the contrary, higher-order evidence in fact never makes any difference to what one should believe (whether about logical or non-logical matters).[14] I won't discuss such arguments here; not because they are not

[12] Christensen (2007; 2010; 2011). See also Schechter (2013) for a similar line of reasoning, though he doesn't frame his discussion with an explicit focus on ideal epistemic agents.

[13] I add this qualification because the type of dogmatism at play here is importantly different from the type of dogmatism traditionally associated with the Kripke/Harman "Dogmatism Puzzle" (Kripke 2011; Harman 1973, pp. 147-49). For a comparison of these two kinds of dogmatism, see Christensen (2010; 2011).

[14] See, e.g., Tal (forthcoming), Titelbaum (2015), and Weatherson (2019, ch. 8).

worthy of serious consideration (they are), but because my goal here is not to argue on general grounds that higher-order evidence can make a difference to what one should believe. Rather, I want to argue that those who already accept that higher-order evidence can make a difference to what one should believe should also be willing to accept that ideal rationality does not require logical omniscience.

Bearing this in mind, let us now instead suppose that Dr. Smart is a cognitively infallible agent. In particular, let us suppose that she never makes any mistakes in deductive reasoning. Does this make a difference to whether Dr. Smart should lower her credence in *L* upon learning that she was drugged? It seems not. After all, Dr. Smart still has strong (albeit misleading) reason to think that she was cognitively impaired while carrying out the proof. So, if she *didn't* lower her credence in *L*, it would still look as if she took the higher-order evidence to be misleading, which (absent independent reasons to think herself immune to the effects of the drug) would still seem dogmatic.

Of course, if Dr. Smart was rationally *certain* of her own cognitive infallibility before receiving the higher-order evidence, it presumably *would* be rational for her to take the higher-order evidence to be misleading, and so refuse to lower her credence in *L*. But as Christensen (2007, p. 14) has rightly pointed out, it is hard to see why we should think that ideally rational agents must invariably be certain of their own cognitive infallibility. After all, questions about the reliability of one's own cognitive faculties are presumably empirical in nature, and we don't normally consider it to be anything less than fully rational to be uncertain about empirical matters. On the contrary, it would seem irrational to be certain of one's own cognitive infallibility in the face of strong evidence to the contrary, even if the evidence happens to be misleading.

The upshot of these preliminary considerations is that the requirement of logical omniscience can come into conflict with the requirement to proportion one's beliefs to the evidence, even if we restrict our attention to cognitively infallible agents. By itself, this obviously doesn't challenge the view that ideal rationality requires logical omniscience any more than it challenges the view that ideal rationality requires one to proportion one's beliefs to the evidence. But if we want to retain the latter requirement

(which I'll simply take for granted here), the conflict presents a challenge for the view that ideal rationality requires logical omniscience.

3. Two responses

How might this challenge be met? In what follows, I will consider two proposals that have been offered in the literature. I don't think either proposal is ultimately convincing, but they fail for what I think are interesting and illuminating reasons.

3.1. Would an ideally rational agent be akratic?

The first response proceeds by conceding that *non-ideal* agents should be prepared to revise their logical beliefs in response to higher-order evidence, but denies that the same goes for *ideal* agents. A version of this response has been put forward by Declan Smithies in recent writings.[15] He argues that when we get clear about *why* cognitively fallible agents should revise their logical beliefs in response to higher-order evidence, it becomes apparent that the same story does not apply to cognitively infallible agents.

To lay out the proposal in more detail, consider a non-ideal agent who finds himself in a similar situation to that of Dr. Smart: that is, he has written down a correct and competently produced proof of *L*, but then learns that he was drugged while carrying out the proof. According to Smithies, such a non-ideal agent should think to himself: "My proof of *L* is most likely flawed, and *L* may well be false." In other words, the agent should lose confidence not only in the correctness of the *proof* of *L*, but also in the *truth* of *L*. So far, this dovetails nicely with the anti-dogmatism argument outlined in the previous section.

But *why*, exactly, should our non-ideal agent lower his credence in *L* in response to the higher-order evidence? According to Smithies, the answer is this: because the higher-order evidence acts as a *disabler*, which prevents the agent from converting his propositional justification into doxastic justification (Smithies 2015, §7). In other words, on Smithies' view, the higher-order evidence works as a *doxastic defeater*, not a *propositional defeater*.

[15] See Smithies (2015) and (2019, chs. 9-10).

The background story here is supposed to be that all agents, irrespective of their cognitive abilities, have the strongest possible level of propositional justification for every logical fact. However, non-ideal agents are sometimes prevented from turning this propositional justification into doxastic justification, because they are not, as Smithies puts it, "sufficiently sensitive to the logical facts" (Smithies 2015, p. 2786). By contrast, ideal agents are never prevented from turning their propositional justification into doxastic justification, because they are perfectly sensitive to the logical facts. In particular, Dr. Smart is not prevented from being certain of *L* in a way that is doxastically justified, because she is perfectly sensitive to the fact that *L* is a logical truth. As Smithies explains:

> The most probable hypothesis given [Dr. Smart's] total evidence is perhaps that she made an error in reasoning that leads to a true and certain conclusion. In that case, rationality requires that she doubt the quality of her reasoning, but since she is certain that her conclusion is true, there is no rational pressure to change her view. (Smithies 2015, p. 2789)

Put differently, when Dr. Smart learns that she has been drugged, she should think to herself: "My proof of *L* is most likely flawed, but *L* is no doubt true." Note that this is not to say that Dr. Smart should be certain of her own cognitive infallibility. Rather, the claim is that Dr. Smart should be certain of *L* *despite* considering it very likely that she was cognitively impaired while carrying out the proof.

This also reveals a curious implication of Smithies' view, namely that ideal rationality sometimes requires agents to exhibit a certain kind of *epistemic "akrasia."* Consider the following two combinations of attitudes that Dr. Smart might adopt towards her own epistemic situation after having received the higher-order evidence (where "PJ" stands for "Propositional Justification" and "DJ" stands for "Doxastic Justification"):

> **PJ-akrasia:** I'm confident that *L* is not supported by my total evidence, but *L* is no doubt true.
>
> **DJ-akrasia:** I'm confident that I didn't form my high credence in *L* in a rational manner, but *L* is no doubt true.

On Smithies' view, Dr. Smart should exhibit DJ-akrasia, but not PJ-akrasia. She should exhibit DJ-akrasia, because she possesses strong (albeit misleading) higher-order evidence indicating that she was cognitively impaired while carrying out the proof. But she shouldn't exhibit PJ-akrasia, because she is perfectly sensitive to the fact that *L* is a logical truth.

What to make of this curious implication? According to Smithies, we should embrace it. As he sees it, the fact that you formed your belief in an irrational manner does not, in and of itself, create a rational pressure to revise that belief. What ultimately matters is whether you are *propositionally* justified. As long as you don't have any reason to doubt that you are propositionally justified, no level of doubt about whether you are *doxastically* justified will create a rational pressure to revise your belief.

That concludes my summary of Smithies' proposal. I now want to raise two worries about it.

The first worry is that Smithies' view clashes with what I think is a plausible connection between propositional and doxastic justification:

> **PJ/DJ-principle:** An agent is propositionally justified in believing *p* only if it is possible, at least in principle, for the agent to believe *p* in a doxastically justified way.

Put differently, according to the PJ/DJ-principle, it must be in-principle possible for agents to convert their propositional justification into doxastic justification by responding properly to the propositional justification (say, by basing their beliefs properly on the evidence). The exact content of this principle obviously depends on how we interpret the relevant notion of "in-principle possibility," but I won't try to make this notion fully precise. For present purposes, we can get enough of a sense of what the PJ/DJ-principle is supposed to capture by considering a paradigmatic example of "improper basing:"

> **Wishful Thinking:** Fred, a dedicated baseball fan, believes that the Red Sox will take home the trophy this season. As it happens, he has strong evidence for this belief. But his belief isn't based on the evidence. Rather, it is based on wishful thinking.

This sort of case is often used to motivate the distinction between propositional and doxastic justification. The idea is that Fred is *propositionally justified* in believing that the Red Sox will win, because he possesses strong evidence for this proposition. But he is not *doxastically justified* in believing as he does, since his belief is not based on the evidence in the right kind of way. Of course, it is notoriously difficult to spell out this "basing requirement" in precise terms.[16] But let us just assume that there *is* a suitably clear and precise way to draw the distinction between propositional and doxastic justification.

We can then ask: Could Fred in principle have been doxastically justified in believing as he does, even if he is not? It certainly seems so. After all, Fred *would* have been doxastically justified if only he had based his belief properly on the evidence; and there seems to be no principled obstacle to his doing so. It might well be highly cognitively demanding for Fred to avoid basing his belief on wishful thinking. It might even be psychologically impossible for him. But nothing about his evidential situation as such makes it in-principle impossible for him to convert his propositional justification into doxastic justification. In other words, his evidence does not, in and of itself, prevent him from believing in a doxastically justified way. This is what the PJ/DJ-principle is supposed to capture.

Now, if Smithies' view is correct, not all cases follow this pattern. On his view, there can be situations in which an agent is propositionally justified in believing a proposition, although it is not even in-principle possible for the agent to convert this propositional justification into doxastic justification. This is exactly the kind of situation that non-ideal agents will find themselves in when they receive defeating higher-order evidence.

To illustrate, consider again the version of the logic case in which a non-ideal agent believes that *L* is a logical truth on the basis of an impeccable proof, but then learns that he was drugged while carrying out the proof. On Smithies' view, this agent will end up in a situation in which he has the strongest possible level of propositional justification to believe that *L* is a logical truth, but in which it is in-principle impossible for him to convert his propositional justification into doxastic justification.

[16] A recent attempt to do so can be found in Neta (2019).

The crucial point here is that it is the agent's evidential situation *itself* that prevents him from believing in a doxastically justified way, not a flaw in the reasoning process that led him to conclude that *L* is a logical truth. This is what sets the present case apart from standard cases of improper basing like Wishful Thinking; and it is why, to my mind, Smithies' view obscures the distinction between propositional and doxastic justification.

Of course, one might doubt whether the PJ/DJ-principle holds in full generality, and although I take something like the PJ/DJ-principle to be implicit in many discussions of the basing relation, I don't want to claim that it admits of no exceptions.[17] In any case, the PJ/DJ-principle seems to capture a robust enough pattern to make it worrisome that Smithies' view clashes with it. At the very least, it would be good to have an explanation of why higher-order defeaters have the effect of disrupting an otherwise close connection between propositional and doxastic justification.

The second worry I want to raise about Smithies' view concerns the implication that ideal rationality sometimes requires DJ-akrasia.[18] To bring out the worry, suppose that Dr. Smart does indeed exhibit DJ-akrasia after having received the higher-order evidence. That is, suppose she thinks to herself: "I'm confident that my proof of *L* is incorrect, but *L* is no doubt true." Given this, she must consider herself to be in a somewhat unlikely situation, namely, the situation of having reached a true conclusion through an incorrect proof. After all, correct proofs guarantee their conclusions, whereas incorrect proofs do not (although incorrect proofs may, of course, result in true conclusions by accident). But if this is how Dr. Smart conceives of her own situation, she must consider herself to have been subject to a certain kind of "epistemic luck." She must think to herself: "I most likely made a mistake in my proof, but I

[17] For example, some philosophers have argued that it can sometimes be irrational to form a belief that *p*, even if one has strong evidence for *p*, because one knows that the very formation of the belief would falsify *p* (Conee 1987). Such cases of "self-undermining" beliefs are notoriously puzzling, but they arguably show that one can have propositional justification to believe *p* even though it is not even in principle possible to convert this propositional justification to doxastic justification. Thanks to Kevin McCain for drawing my attention to this potential class of counterexamples to the PJ/DJ-principle.

[18] Here I'm indebted to Horowitz (2014, §3) who raises a similar worry about views that permit epistemic akrasia, though she doesn't distinguish between PJ-akrasia and DJ-akrasia.

nevertheless ended up with a true conclusion. Lucky me!" Yet, such self-ascriptions of epistemic luck seem patently irrational.

We can make the irrationality even more apparent by imagining a long series of similar cases. In each case, Dr. Smart writes down a correct and competently produced proof of some logical fact, but then learns that she was drugged while carrying out the proof. Given standard ways of thinking about risk accumulation, Dr. Smart must then end up thinking to herself: "It is *exceedingly* unlikely for someone to reach this many true conclusions on the basis of as many flawed proofs. Yet, I find myself in one of these exceedingly rare situations. What an incredible stroke of luck!" Again, this strikes me as a patently irrational attitude to have towards one's own epistemic situation.

I don't expect everyone to be convinced by this anti-luck intuition. Some philosophers have recently urged that it can be rational to consider oneself "lucky to be rational" in cases where one learns that one's beliefs have been influenced, in part or whole, by epistemically irrelevant factors (say, factors about one's upbringing or cultural background).[19] Perhaps something similar can be said about the kinds of self-ascriptions of epistemic luck to which Smithies' view gives rise. But it seems to me that the worries raised in this section, taken together, should make us doubt that the view is on the right track.

3.2. Can rational ideals come into conflict?

The second response proceeds by conceding that the requirement of logical omniscience can come into conflict with the evidential requirement, but maintains that logical omniscience is nevertheless a requirement of ideal rationality. A version of this view has been put forth by David Christensen, who suggests that the requirements of ideal rationality may sometimes issue conflicting demands, thereby forcing agents to violate at least one rational ideal.[20] In the case of Dr. Smart, the idea is that she is doomed to violate one of the following three requirements (cf. Christensen 2007, §5):

Logical Omniscience: Your beliefs should obey the laws of logic.

---

[19] See, e.g., Elga (ms.) and Schoenfield (forthcoming). See also White (2010) for an illuminating discussion of this issue.

[20] Christensen (2007; 2010; 2016). See also Pryor (2018) for a sympathetic, though noncommittal, discussion this kind of response.

Evidence-responsiveness: Your beliefs should be proportioned to your evidence.

Inter-level Coherence: Your first-order beliefs should "line up" with your higher-order beliefs (i.e., your beliefs about what beliefs you *should* have).

To satisfy the first requirement, Dr. Smart must believe *L* (assuming, as we do, that *L* is a logical truth). To satisfy the second requirement, she must believe that she *shouldn't* believe *L* (assuming, as we do, that Dr. Smart's total evidence supports this higher-order belief). So, if she satisfies the first two requirements, she believes that "*L*, but I shouldn't believe *L*," which violates the third requirement. This means that it is impossible for Dr. Smart to satisfy all three requirements at once.

A natural reaction to this result would be to say that one of the requirements cannot be a genuine requirement after all. But Christensen hesitates to draw this conclusion. Instead, he entertains a different hypothesis: Perhaps the demands of ideal rationality do not fit together as nicely as one might have expected or hoped they would. Perhaps ideal rationality sometimes places demands on agents that cannot even in principle be met.

There is an immediate worry about this sort of dilemma view, which comes from the familiar principle that "ought" implies "can" (OIC). The driving intuition behind this principle is that there seems to be something misguided about subjecting agents to requirements that they cannot possibly meet. The exact content of this principle clearly depends on how we interpret the notions of "ought" and "can." But for present purposes, we can focus on the following two interpretations:

Psychological OIC: You are subject to a requirement of ideal rationality only if it is psychologically possible for you to meet that requirement.

Logical OIC: You are subject to a requirement of ideal rationality only if it is logically possible for you to meet that requirement.

The Psychological OIC can arguably be denied without flying in the face of normative common sense. For example, it might well be psychologically impossible for Fred to avoid basing his beliefs on wishful thinking. Yet, it doesn't seem obviously misguided to say that his beliefs are nevertheless, for that reason, irrational.

It is, however, the Logical OIC that Christensen's proposal calls into doubt; and this version of the principle is harder to deny. This is true in a trivial sense: The Psychological OIC entails the Logical OIC, but not *vice versa* (assuming, as is customary, that psychological possibility entails logical possibility, but not *vice versa*). But I think the difference runs deeper than that. Consider again Fred, who suffers from wishful thinking. The reason why he fails to meet the requirements of rationality is that he exemplifies a kind of cognitive malfunction. If only he had performed better cognitively, he would have met the requirements of rationality. By contrast, Dr. Smart does *not* fail to meet the requirements of rationality because of any cognitive malfunction on her part. She *wouldn't* have met the requirements of rationality, if only she had performed better cognitively. This strikes me as a substantial difference, one that makes it considerably harder to deny the Logical OIC than the Psychological OIC.

Now, Christensen's proposal is not straightforwardly inconsistent with the Logical OIC. After all, for each individual requirement issued by the aforementioned norms, it *is* logically possible for Dr. Smart to meet *that* requirement. What is *not* logically possible is for her to meet all three requirements *at once*. Thus, Christensen's proposal only forces us to reject the Logical OIC provided that we accept another familiar deontic principle, referred to by Bernhard Williams (1965) as the principle of *agglomeration*:[21]

> **Agglomeration**: If you are required to φ and you are required to ψ, then you are required to $(\varphi \,\&\, \psi)$.

Agglomeration captures the seemingly banal claim that, say, you ought to buy coffee and cake if you ought to buy coffee and you ought to buy cake. This principle has undeniable appeal and is in fact validated by most systems of deontic logic.[22] Nevertheless, some philosophers have called it into doubt precisely on the grounds

[21] Thanks to David Christensen for discussion here.

[22] The most familiar such system comes from Standard Deontic Logic (SDL) in the tradition of von Wright (1951), which treats notions like *ought* and *is required* as a necessity operator with all of the standard properties (and more) inherent in the class of "normal modal logics." In particular, SDL validates the material conditional "$O\varphi \,\&\, O\psi \rightarrow O(\varphi \,\&\, \psi)$," and so validates Agglomeration. For additional background on deontic logic, see McNamara (2019).

that it leads to paradox when combined with (1) the claim that normative dilemmas are possible, and (2) the claim that "ought" implies "can."[23]

This puts us in a somewhat difficult dialectical situation, where one philosopher's Modus Ponens is likely to be another philosopher's Modus Tollens. I won't here try to argue that the best way forward is to retain Agglomeration and the Logical OIC, and so reject Christensen's dilemma view. Properly doing so would take us much too far afield. But the fact that Christensen's proposal is inconsistent with the conjunction of Agglomeration and the Logical OIC is at least somewhat disconcerting.

There is, however, a second worry, which makes me doubt that Christensen's proposal can ultimately be right. The source of the worry is that if we maintain that Dr. Smart is forced to violate at least one rational requirement, we haven't thereby said anything about *which* requirement she should violate. In particular, we haven't said that she should violate Logical Omniscience rather than Evidence-responsiveness or Inter-level Coherence. This might not seem like much of a problem. After all, why should we think that there is a fact of the matter as to which ideal Dr. Smart should violate? The reason, as I see it, is that there is a robust intuition that Dr. Smart should indeed violate Logical Omniscience, not Evidence-responsiveness or Inter-level Coherence. This is what I take the anti-dogmatism argument outlined in §2 to suggest, which is what motivated Christensen's proposal in the first place. Thus, it seems that Christensen's proposal is extensionally incomplete: it fails to give the right normative verdicts in cases like that of Dr. Smart.

Of course, an extensionally incomplete theory might still be correct as far as it goes. But I do think that the extensional incompleteness in the present case creates a real problem. To see why, let us suppose that we manage to augment the dilemma view in a way that lets us maintain that Dr. Smart should violate Logical Omniscience rather than Evidence-responsiveness or Inter-level Coherence.[24] We can then ask: How does

[23] See, e.g., Williams (1965), Marcus (1980), Brink (1994), Horty (2003), and Hughes (2019).

[24] What might such an augmentation look like? One possible option is to say that not all rational requirements are on a par, but that some requirements take priority over others, thereby creating a kind of "hierarchy" of requirements. In particular, the idea would be that Evidence-responsiveness and Inter-level Coherence somehow "overrule" Logical Omniscience, at least in cases like that of Dr. Smart. Of course, such a proposal runs an immediate risk of arbitrariness. After all, on what principled grounds

the resulting view differ from the view that ideal rationality does not require logical omniscience? It's not so clear. The idea is supposed to be that Dr. Smart still falls short of ideal rationality, if she violates Logical Omniscience. But it is not so clear how to make sense of this shortcoming in light of the fact that, by hypothesis, the augmented dilemma view says that Dr. Smart *should* violate Logical Omniscience. Perhaps there is a way to make sense of this by reconsidering the role of rational ideals in our theory of epistemic rationality. But as far as I can see, the strategy of augmenting the dilemma view in a way that lets us maintain that Dr. Smart should violate Logical Omniscience rather than Evidence-responsiveness or Inter-level Coherence—even if successful—will result in little more than a notational variant of the view that ideal rationality does not require logical omniscience.

4. Conclusion

Needless to say, I have considered by no means all imaginable responses to the conflict brought out by the anti-dogmatism argument. For example, I haven't considered the possibility of accepting a kind of "pluralism" about epistemic normativity, according to which the requirement of logical omniscience and the requirement to proportion one's beliefs to the evidence pertain to fundamentally different normative notions, which cannot be unified in a single "all-things-considered" notion of normativity, and which should not be expected to issue compatible demands.[25] Nevertheless, if the criticisms put forth are correct, it suggests, at least in a tentative way, that ideal rationality does not require logical omniscience.

None of this means that logic has *no* role to play in a theory of epistemic rationality. But it means that we need to reconsider what this role might be. I shall not

might we claim that one rational requirement takes priority over another? But for the sake of argument, we may simply set this kind of worry aside.

[25] A view along these lines has been suggested by Worsnip (2018) in response to a similar, but distinct, alleged conflict between the evidential requirement and the *enkratic requirement*, which says (roughly) that one shouldn't believe propositions of the form "*p*, but I shouldn't believe *p*." I have some misgivings about Worsnip's particular way of arguing for this alleged conflict, which I've articulated elsewhere (Skipper 2019). But if we grant that the enkratic requirement can indeed come into conflict with the evidential requirement, and if we grant that this conflict should be resolved by accepting a kind of pluralism about epistemic normativity, it's natural to think that a similar pluralist view might help resolve the conflict between the evidential requirement and the requirement of logical omniscience.

try to do so in any detail here; an undertaking of this magnitude is best left for another occasion. But let me close with a few preliminary remarks in this direction.

In light of what has been said, it might be tempting to think that logical omniscience can still serve as a rational ideal for what David Chalmers has called "insulated" agents: that is, agents who possess no evidence about their own cognitive reliability (Chalmers 2012, pp. 101-7). Such agents, by stipulation, never find themselves in situations like that of Dr. Smart in which their logical beliefs are defeated by misleading higher-order evidence. As such, it might be natural to think that nothing could prevent *them* from being rationally certain of all logical facts.

However, even such a weakened proposal is too strong. To see why, we can ask what the relevant kind of insulated agent—that is, someone who is cognitively infallible, and who possesses no evidence about her own cognitive reliability—should think of her own cognitive reliability. Presumably, she shouldn't be *absolutely certain* of her own cognitive infallibility. In particular, she shouldn't be certain of her own infallibility when it comes to logical reasoning. After all, contingent claims about the reliability of one's own cognitive faculties are presumably not the kinds of things one can be rationally certain about on purely *a priori* grounds (as also pointed out towards the end of §2). But if so, this means that even insulated agents should exhibit the kind of uncertainty about their own cognitive reliability, which, on pain of dogmatism, should lead them to have uncertainty about logical matters.

A more promising hypothesis, to my mind, is that logic does not place any non-trivial constraints on what *belief states* it can be rational to be in, but only places constraints on what *reasoning processes* it can be rational to perform. Such a view might still allow us to say that an ideally rational agent would never base her beliefs on fallacious logical reasoning. But it wouldn't say of any particular belief state that *it*, necessarily, falls short of ideal rationality in virtue of not conforming to the laws of logic. Obviously, much more would need to be said to properly formulate and evaluate this hypothesis. But whatever we end up saying about the connection between logic and rationality, it will have to depart from the standard view that ideal rationality requires logical omniscience.

Acknowledgments: For helpful comments and questions, I'd like to thank David Christensen, Maria Lasonen-Aarnio, Declan Smithies, Eyal Tal, and the editors of this volume.

References

Alchourrón, C., P. Gärdenfors, and D. Makinson (1985): "On the Logic of Theory Change: Partial Meet Contraction and Revision Functions." *Journal of Symbol Logic* 50, pp. 510–30.

Brink, D. (1994): "Moral Conflict and its Structure." *The Philosophical Review* 103, pp. 215–47.

Chalmers, D. (2012): *Constructing the World*. Oxford University Press.

Christensen, D. (2004): *Putting Logic in Its Place*. Oxford University Press.

——— (2007): "Does Murphy's Law Apply in Epistemology?" *Oxford Studies in Epistemology* 2, pp. 3–31.

——— (2010): "Higher-Order Evidence." *Philosophy and Phenomenological Research* 81, pp. 185–215.

——— (2016): "Conciliation, Uniqueness, and Rational Toxicity." *Noûs* 50(3), pp. 584–603.

Coates, A. (2012): "Rational Epistemic Akrasia." *American Philosophical Quarterly* 49, pp. 113–24.

Conee, E. (1987): "Evident, But Rationally Unacceptable." *Australasian Journal of Philosophy* 65, pp. 316–26.

Conee, E. and R. Feldman (2004): *Evidentialism: Essays in Epistemology*. Oxford University Press.

——— (1985): "Evidentialism." *Philosophical Studies* 48, pp. 15–34.

Dorst, K. (2020): "Evidence: A Guide for the Uncertain." *Philosophy and Phenomenological Research* 100, pp. 586-632.

Elga, A. (ms): "Lucky to be Rational." Unpublished manuscript. URL = http://www.princeton.edu/ ~adame/papers/bellingham-lucky.pdf

Field, H. (2009): "What is the Normative Role of Logic?" *Proceedings of the Aristotelian Society* 83, pp. 251–68.

Foley, R. (1992): "The Epistemology of Belief and the Epistemology of Degrees of Belief." *American Philosophical Quarterly* 29, pp. 11–24.

Frege, G. (1893/1903): *Grundgesetze der Arithmetik*, Jena: Verlag Hermann Pohle. Band I-II.

Friedman, J. (2018): "Junk Beliefs and Interest-Driven Epistemology." *Philosophy and Phenomenological Research* 97, pp. 568–83.

Greaves, H. (2013): "Epistemic Decision Theory." *Mind* 122, pp. 915–52.

Harman, G. (1986): *Change in View*. MIT Press.

——— (1973): *Thought*. Princeton University Press.
Hintikka, J. (1962): *Knowledge and Belief: An Introduction to the Logic of the Two Notions*. Ithaca: Cornell University Press.
Horowitz, S. (2014): "Epistemic Akrasia." *Noûs* 48, pp. 718–744.
Horty, J. (2003): "Reasoning With Moral Conflicts." *Noûs* 37, pp. 557–605.
Hughes, N. (2019): "Dilemmic Epistemology." *Synthese* 196, pp. 4059–90.
Joyce, J. (1998): "A Non-Pragmatic Vindication of Probabilism." *Philosophy of Science* 65, p. 575–603.
Kiesewetter, B. (2017): *The Normativity of Rationality*. Oxford University Press.
Kolodny, N. (2005): "Why Be Rational?" *Mind* 114, pp. 509–63.
Kripke, S. (2011): "Two Paradoxes of Knowledge." *Philosophical Troubles: Collected Papers I*, New York: Oxford University Press, pp. 27–51.
Lasonen-Aarnio, M. (2014): "Higher-Order Evidence and the Limits of Defeat." *Philosophy and Phenomenological Research* 88, pp. 314–45.
Leitgeb, H. (2014): "The Stability Theory of Belief." *The Philosophical Review* 123, pp. 131–71.
Littlejohn, C. (2018): "Evidence and Its Limits." C. McHugh, J. Way, and D. Whiting (eds.), *Normativity: Epistemic and Practical*. Oxford University Press.
MacFarlane, J. (ms): "In What Sense (If Any) is Logic Normative for Thought?" Unpublished manuscript. URL = http://johnmacfarlane.net/normativity_of_logic.pdf.
Marcus, R. (1980): "Moral Dilemmas and Consistency." *Journal of Philosophy* 77, pp. 121–36.
McNamara, P. (2019): "Deontic Logic." E. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*.
Neta, R. (2019): "The Basing Relation." *The Philosophical Review* 128, pp. 179–217.
Pettigrew, R. (2020): *The Dutch Book Arguments*. Cambridge University Press.
——— (2016): *Accuracy and the Laws of Credence*. Oxford University Press.
Pryor, J. (2018): "The Merits of Incoherence." *Analytic Philosophy* 59, pp. 112–41.
Ramsey, P.F. (1926): "Truth and Probability." H. Kyburg and H. Smokler (eds.), *Studies in Subjective Probability*, Huntington, NY: Robert E. Kreiger Publishing Co.
Rinard, S. (2015): "Against the New Evidentialists." *Philosophical Issues* 25, pp. 208–23.
Schechter, J. (2013): "Rational Self-Doubt and the Failure of Closure." *Philosophical Studies* 163, pp. 428–52.
Schoenfield, M. (forthcoming): "Meditations on Beliefs Formed Arbitrarily." *Oxford Studies in Epistemology*.
——— (2018): "An Accuracy Based Approach to Higher Order Evidence." *Philosophy and Phenomenological Research* 96, pp. 690–715.
——— (2015): "A Dilemma for Calibrationism." *Philosophy and Phenomenological Research* 91(2), pp. 425–455.

Shah, N. (2006): "A New Argument for Evidentialism." *The Philosophical Quarterly* 56, pp. 481–98.

Skipper, M. (forthcoming): "Does Rationality Demand Higher-Order Certainty?" *Synthese*.

——— (2019): "Higher-Order Defeat and the Impossibility of Self-Misleading Evidence." M. Skipper and A. Steglich-Petersen (eds.), *Higher-Order Evidence: New Essays*. Oxford University Press.

Smithies, D. (2015): "Ideal Rationality and Logical Omniscience." *Synthese* 192, pp. 2769–93.

——— (2019): *The Epistemic Role of Consciousness*. Oxford University Press.

Steinberger, F. (2019a): "Logical Pluralism and Logical Normativity." *Philosophers' Imprint* 19, pp. 1–19.

——— (2019b): "Three Ways in Which Logic Might Be Normative." *Journal of Philosophy* 116, pp. 513.

——— (2019c): "Consequence and Normative Guidance." *Philosophy and Phenomenological Research* 98, pp. 306–28.

Tal, E. (forthcoming): "Is Higher-Order Evidence Evidence?" *Philosophical Studies*.

Titelbaum, M. (forthcoming): *Fundamentals of Bayesian Epistemology*. Oxford University Press.

——— (2015): "Rationality's Fixed Point (Or: In Defense of Right Reason)." *Oxford Studies in Epistemology 5*, T. Gendler and J. Hawthorne (eds.), pp. 253–94.

von Wright, G. (1951): "Deontic Logic." *Mind* 60, pp. 1–15.

Weatherson, B. (2019): *Normative Externalism*. Oxford: Oxford University Press.

White, R. (2010): "You Just Believe That Because…" *Philosophical Perspectives* 24, pp. 573–615.

Whiting, D. (2020): "Recent Work on Higher-Order Evidence." *Analysis*. Online first.

Williams, B. (1965): "Ethical Consistency." *Proceedings of the Aristotelian Society*, Supplementary volumes 39, pp. 103–38.

Worsnip, A. (2018): "The Conflict of Evidence and Coherence." *Philosophy and Phenomenological Research* 96, pp. 3–44.