

Higher-Order Evidence is the Wrong Kind of Reason

1. Introduction

Philosophers are narcissists. Or at least, we have tendencies towards narcissism in a special epistemic sense: We are preoccupied with the rationality our own beliefs. Sure enough, we try to discover the truth and to persuade each other of it where we can. But if truth is too much to ask for, we hope at least to have good reasons for our beliefs. And when persuasion fails, we fall back on at least exhibiting to each other the coherence and systematicity of the views we have stitched together.

A little narcissism is probably forgivable, given what we do for a living. And maybe there is even something admirable in valuing rationality for its own sake. But many philosophers take our idiosyncratic tendencies to reflect an ideal that all rational beings must aspire to. They build a preoccupation with one's rationality into what it is to be rational.

These defenders of epistemic narcissism of course do not characterize it that way. Instead, they defend what they call *reflection* (or critical reason), the process of reflectively appreciating what first-order rational norms require of us, and of adopting or revising attitudes to bring them into line with the requirements. In Tyler Burge's words:

Critical reason is a capacity to recognize and effectively employ reasonable criticism of or support for propositional attitudes and for propositional reasoning, guided by an appreciation, use, and assessment of reasons and reasoning as such. Critical reasoning is the sort of reasoning that evaluates, checks, refines ordinary first-level reasoning, using concepts and standards of reason.¹

Declan Smithies develops a similar view in a recent book, claiming:

[T]he point of reflection is to bring your beliefs into line with higher-order reflections about which beliefs you have justification to hold. The reflective process has two stages: first, reflecting on which beliefs you have justification to hold; and second, revising your beliefs in light of these reflections.²

Some authors take rationality to be impossible in the absence of reflection.³ But most accept that some degree of rationality is attainable even if your attention remains directed at the world around you. If you know that tigers are dangerous and learn that the creature in the bushes is a tiger, then you can, based solely on these worldly facts, adopt the belief that it is dangerous. But this kind of world-directed reasoning is still disparaged as exemplifying a kind of low-grade rationality. It is what we have in common with children and animals. It is “automatic” or “unconscious” or “offline.”

¹ Burge 2013, pg. 167.

² Smithies 2019.

³ E.g., Shoemaker 1996, pp. 32-33

In contrast, the reflective agent has some appreciation of what rationality's requirements are, and of how well her current beliefs live up to those requirements. And when she judges that her beliefs do not live up, she modifies them to bring them into line. This is what Smithies above deemed the second stage of the reflective process. Burge puts it this way:

[I]f, in critical reasoning, one correctly and with warrant judges that a lower-level state is (or is not) reasonable, then it rationally follows directly that one has reason to sustain (or change) the lower-level state.⁴

While most reflectivists do not think reflection is required for garden variety reasoning, they still consider it vitally important. It is what makes us the kinds of beings that we are, and separates us from the brutes. It is what makes us responsible for our attitudes or actions. It is the source of epistemic normativity.⁵

These are ambitious claims. While common, they are controversial. Critics have doubted that reflection has the significance these authors attach to it.⁶ But what is less controversial is that, if one does happen to consider whether one's belief is rational, and one judges that it is not, then one ought to revise it, or at least will be under some pressure to do so.⁷ This moderate claim often turns up in discussion with no broader reflectivist agenda, as a well-known truism cited in support of some other point. Here is Richard Moran discussing doxastic agency:

If a person is at all rational, his first-order beliefs will indeed be sensitive to his second-order beliefs about them, and will change accordingly. He may, for instance, discover some set of his beliefs is inconsistent, or suspect that a particular belief of his is the product of prejudice or carelessness, or, at the limit case, that it is just plain false. His first-order beliefs will then normally change in response to his interpretation of them.⁸

And here is Ralph Wedgwood, while discussing the aim of belief:

[T]he concept "rational belief" is normative for the practice of theoretical reasoning. It is a constitutive feature of this concept that if one judges that a certain belief would not be a "rational" belief for one to hold, this judgment commits one to not holding that belief.⁹

Discussing certain externalist responses to skepticism, Roger White takes a similar line about credences. White considers an externalist who says that we who can see our hands should be nearly 100% sure we have hands, while a brain in a vat, who merely hallucinates hands, should only be, say, 90% sure of it. White observes that if the brain in the vat is 90% sure

⁴ Burge 2013, pg 372. (??)

⁵ E.g., Alston 1989, pp. 221-236; Burge 2013; McGinn 1982, pg. 20; and Smithies 2019.

⁶ E.g., Arpaly 2002, Broome 2013, and Kornblith 2012.

⁷ E.g., Elga 2005, pg. 116; Scanlon 2007; and the authors quoted below.

⁸ Moran 2001, pg. 55.

⁹ Wedgwood 2002, pg. 271.

he sees hands, given externalism he also should be 90% sure that really he should be 100% sure he has hands. White objects:

Can this be a rational state to be in? If [a brain in a vat] has reason to think that most likely he is less confident than he should be in thinking that here's a hand, doesn't this give him a reason to increase his confidence? ... Shouldn't he think: 'Most likely I'm not as confident as I should be that here's a hand, so I'll boost my confidence'?¹⁰

But in recent discussions, moderate reflectivism has received the most sustained attention in connection with peer disagreement, epistemic akrasia, and higher-order evidence.¹¹ Here discussion centers on agents whose reflective beliefs are based not on introspection and *a priori* reasoning, as traditional reflectivists imagined, but instead on *a posteriori* evidence. We will consider below some examples of higher-order evidence like this. For now, just compare with traditional reflectivism what is sometimes said to be the upshot of these examples, such as by David Christensen in his early work on the topic:

Intuitively, there seems to be a connection between what one is rational to believe, and what one is rational to believe one is rational to believe. ... Putting the thought in terms of justification, the idea is that (justified) higher-level doubts about the justification of one's belief that p can defeat one's justification for believing p.¹²

While all these authors put the point in different ways, they share a common idea. It is that one's reflectively believing that a doxastic attitude is irrational will normally defeat the justification or rationality of the attitude, or commit one to revising it, or gives one a reason to do so. As I will put it, they think reflective beliefs ought to *guide* you.

This moderate reflectionist idea is not just limited to philosophers with some partial sympathies to traditional reflectivism. It is often taken for granted by critics, who instead just dispute particular claims about the strength or philosophical significance of reflective reasons.¹³ For instance Maria Lasonen-Aarnio, while opposing stronger views like Christensen's, still concedes:

Consider the reason *I am required, given my epistemic reasons, to believe p*, or simply *I am rationally required to believe p*. These are conclusive reasons to believe p. Similarly, *I am forbidden, given my epistemic reasons, to believe p*, or simply *it is irrational for me to believe p*. These are conclusive reasons not to believe p.¹⁴

¹⁰ White 2013.

¹¹ E.g., Christensen 2010a, 2010b, 2016; Elga 2005; Greco 2014; Horowitz 2014; Kelly 2010; Lasonen-Aarnio 2014 and 2020; Smithies 2019, Chs. 9-10; Weatherson 2019, Worsnip 2018

¹² Christensen 2010a. As I discuss below, Christensen has revised his view. But even the new view partly resembles reflectivism, and anyway the the reaction to higher-order defeat remains common (e.g., Elga 2005, Greco 2014, Horowitz 2014, Worsnip 2018).

¹³ E.g., Arpaly 2002, Bergman 2005, and Broome 2013.

¹⁴ Lasonen-Aarnio 2020.

In contrast, I think the *very idea* of reflective reasons and reasoning is misguided. I doubt that the reasons Lasonen-Aarnio cites are reasons at all, much less conclusive ones. But even if they are admitted as reasons of some kind, they are the *wrong kind of reasons* for and against belief. An agent who responded to them as the above authors propose would be reasoning irrationally.

A paradigmatic wrong kind of reason for belief is a moral reason. Suppose my friend is accused of a crime, and I think I am morally required to believe she is innocent. Does this give me a reason to believe she is innocent?¹⁵ Maybe so. If I can press a button that will cause me to believe she is innocent, it plausibly gives me a reason to press. Maybe that is enough for it to count as giving me a reason to believe. But even if so, this strikes us as in some sense the wrong kind of reason to believe. It would be irrational for me to believe for a reason like that.

My thesis here is that the same goes for reflective judgments about what it is rational for one to believe. I think it would be just as irrational to believe for the reason that it is rationally required as it is to believe for the reason that it is morally required. On my alternative view, which I call *deflationism*, a reflective belief that one is rationally required to believe that *p* gives one a reason to believe *p* only to whatever extent the apparent fact that one is rationally required to believe *p* amounts to ordinary evidence for *p*. The idea is that apparent facts about one's own rational requirements are no more germane to the question whether *p* than apparent facts about anyone else's rational requirements. If some arbitrary agent is apparently required to believe *p*, either morally or epistemically, that will be a reason for me to believe *p* only insofar as it is evidence supporting that *p*. And the same goes, under deflationism, if I am the one who is apparently required to believe *p*.

2. The Wrong Kind of Reason

The thesis that reflective beliefs give the wrong kind of reason invites a natural suspicion. You can be morally required to believe something that you have no evidence for. But typically you are not rationally required to believe things without sufficient evidence—i.e., without the right kind of reason to believe, an epistemic reason. If you seem to find yourself in the possession of sufficient evidence, the reflectivist says, that ought to motivate you to go ahead and believe. And how else is it supposed to motivate belief, if not by amounting to an epistemic reason? Is it supposed to amount to some other kind of reason, like a moral or prudential one?

But this talk of different kinds of reasons—epistemic and practical, or the right kind and the wrong kind—has the potential to mislead. The right and wrong kinds of reasons are not just two types of considerations bearing on the same deliberative question, like personal and professional reasons for moving to a new city. Instead, they are considerations relevant to distinct kinds of deliberation or reasoning. Let me explain, starting with an example:

Cash's Prize: Cash is offered a large cash prize conditional on his believing it will rain. Cash can press a button that will cause himself to believe it will rain.

Should Cash adopt the belief that it will rain? Will doing the rational thing mean believing this? Many of us are reluctant to give univocal answers to these questions. We want to respond with some sort of distinction, and that is where talk of different kinds of reasons

¹⁵ Cf. Stroud 2006 and Rioux MS.

comes in. On the one hand, we want to grant that the prize gives Cash a good reason to bring it about that he believes, by pressing the button. But on the other, we want to deny that this makes it a reason to believe, or at least a reason of the right kind. But why not? Cash prizes are perfectly good reasons for doing most anything else. Why not believing?

At least part of the answer appeals to a plausible connection between reasons and reasoning (or deliberation). Look at it this way. When Cash knows that there is a prize for belief, or even if he merely believes there is one, this might potentially lead him to believing it will rain. In principle, there are many ways this might happen. Cash might have a microchip in his brain that scans for beliefs to the effect that some attitude would be beneficial, and then produces the attitude. But if that is how Cash's knowledge of the prize causes the belief that it will rain, it will just be an incidental factor in the belief's etiology, irrelevant to its rationality. For his knowledge of the prize to make the belief rational, it needs to in some further way *guide* him. In what way? A natural answer is that it must be a *motivating reason*, a reason for which Cash believes it will rain.

One way for the prize to be Cash's motivating reason is the old-fashioned way. Rather than using the button, Cash might simply consider the question whether it will rain, and be persuaded that it will by the fact that he will get a prize for believing it. Put another way, Cash might *infer* that it will rain from the premise that he will get a prize for believing it. In doing so, he would reason in a way he might express by saying "I'll get a prize for believing it will rain, so (probably) it will rain."

But reasoning like that would be irrational. Why? Well, the fact that he will get the prize is not evidence that it will rain, at least by ordinary standards. It does not deductively entail that it will rain, or inductively or abductively support it, or anything like that. That will make at least this route to the belief irrational, if we think ordinary evidential standards like these govern inferential reasoning.

Even pragmatists, who think it is unambiguously rational for Cash to believe it will rain, do not deny this much. They are not charitably interpreted as saying Cash ought to infer it will rain from the premise that he will get the prize. For comparison, when Pascal advocated belief that God exists on the basis that it maximizes expected utility, he did not expect you to infer God's existence from this premise. Rather, he wanted to motivate you to act so as to produce the belief, for example by going to church. Likewise, modern day pragmatists like Susanna Rinard do not want Cash to achieve belief by considering whether it will rain, and inferring that it will from a premise about the prize. They want him to consider whether to believe it will rain, and to decide to do so based on that premise. In doing this, Cash would reason in a way he might express by saying "I'll get a prize for believing it will rain, so I'll go ahead and do it."

Surely there is nothing irrational in Cash's deliberating like this, or even in following through, and pressing the button. But that does not mean pragmatists are right that it is rational for him to believe. It might just mean it is rational for Cash to cause himself to believe irrationally, for no good reason. Sure enough, the prize will be Cash's reason for pressing, and his pressing will be a crucial part of the etiology of the belief. But that does not make it the reason for which Cash believes.¹⁶

¹⁶ Cf. Hieronymi 2006, Kelly 2002, Shah 2006, Rinard 2019, and Way 2016.

Let's sum this up as a valid argument that a cash prize is the wrong kind of reason for belief. First, there is the premise even pragmatists should accept, that it is irrational to infer it will rain from a premise about the prize, and more generally that:

(EVIDENTIALISM) The apparent fact that *r* is a premise relevant to *p* in inferential reasoning only insofar as it is ordinary evidence regarding *p*.

Second, there is the premise pragmatists like Rinard dispute, that motivating reasons for belief must be premises in inferential reasoning. More generally,

(INFERENCEALISM) The apparent fact that *r* can be a reason for which one rationally holds a belief about *p* only insofar as *r* is a premise relevant to *p* in inferential reasoning.¹⁷

Assuming the prize is not evidence of rain, it follows that it is not the right kind of reason to believe it will rain, the kind of reason that might rationally motivate Cash to believe it will rain. But these two premises alone do not mean that Cash cannot rationally believe that it will rain, just that he cannot do so for the reason that he will get the prize. A pragmatist might take a different tack from Rinard, and say that it does not matter whether the prize ends up qualifying as the motivating reason for his belief. Even if Cash ultimately believes for no reason at all, maybe there remains a sense in which his belief was *guided* by his knowledge of the prize. And that, this pragmatist says, is enough for the beliefs to be rational.

To rule out a view like that, we need one further premise. It says roughly that for one belief to guide a second belief in a way directly relevant to its rationality, it has to give one's motivating reason. More generally:

(MOTIVATIONALISM) One's belief that *r* can rationally guide one's belief about *p* only insofar as *r* can be a reason for which one rationally holds a belief about *p*.

Now a lot of these points are familiar. What I hope to do now is to use them to motivate an unfamiliar claim regarding reflective beliefs. Consider an agent who judges that one of her existing beliefs is not really supported by her evidence, or another who is persuaded that some doubts he harbors are irrational, and that really he ought to believe. How should these agents respond? The same three premises just considered imply the following thesis, which I will defend:

(DEFLATIONISM) One's reflective belief that one is required to (or prohibited from) believing *p* can rationally guide one's belief about *p* only insofar as the apparent fact that one is required to (or prohibited from) believing *p* is ordinary evidence regarding *p*.

Deflationism is so called because it deflates the epistemic significance of reflective beliefs. When I deliberate about whether *p*, and consider the facts as I see them, deflationism says the apparent facts about what this guy here ought to believe about *p* do not have any distinctive rational significance which is not shared by apparent facts about what other people ought to believe, or apparent facts about shoes and ships and sealing wax. They are

¹⁷ Cf. Pryor's (2005) discussion of the 'Premise Principle'.

all just apparent facts that might speak for or against the truth of p according to the same general evidential standards.

In contrast, I will claim reflectivists need reflective beliefs to provide *inflationary* reasons (or guidance), which goes beyond their ordinary evidential import. This is for two reasons. The first is that many reflectivists think the *normativity* of rationality is central to the epistemic significance of reflective beliefs. I am supposed to have a special reason or motive to believe p when I think I *rationaly ought* to do so. In contrast, ordinary evidential standards arguably do not assign a special role to normative premises or evidence. For example, the validity of a syllogism does not turn on whether the major term is a normative or non-normative predicate. Something similar seems plausible for non-deductive inference, though I will not argue for it here. If normative beliefs have special motivational force, it is not because normative premises amount to especially strong evidence.

The second reason, which will be my focus, is that reflectivists need *first-personal* reflective beliefs to have a significance that does not extend to one's beliefs about what other people are required to believe. In contrast, ordinary evidential standards, like those governing deductive, inductive, and abductive inferences, do not carve out a distinctive role for first-personal premises or evidence. As I will put it, these and other familiar candidate standards of evidential support are invariant under *third-personalization*.¹⁸

To third-personalize an agent's beliefs and evidence, we replace any allegedly *de se* beliefs or evidence with third-personal analogues. For example, if my evidence includes that I am uncertain whether it will rain, and my beliefs include one that I am a lousy forecaster, then a third-personalization of my beliefs and evidence includes the evidence that DB is uncertain whether it will rain and the belief that DB is a lousy forecaster. If it helps, you might imagine the evidence of a guardian angel who, with no regard to herself, has followed me throughout my life. She knows everything I know, including about my mind, and she believes everything I believe. But all of the knowledge and beliefs about me are cast in third-personal terms. Or instead, you might suppose there is no such thing as *de se* thought, or else that I am incapable of it. I might still know all the same facts about myself that I now know, but conceive of them just as facts about DB.

Here is the sense in which I claim ordinary evidential standards are invariant under third-personalization. Suppose that my background evidence is such that, when I learn that I allowed the murderer to return to the crime scene, I can deductively infer that the butler did it. If so, then if my background evidence were third-personalized, I still can deductively infer the butler did it if I learn that DB allowed the murderer to return to the crime scene. For my original evidence must entail that I allowed the butler to return to the crime scene, and so when third-personalized will entail that DB did so. Or suppose that when I learn the butler lied to me about his whereabouts, I can abductively infer that the butler did it. Then if my background evidence is third-personalized, I can abductively infer the same thing if I learn the butler lied to DB about his whereabouts. Evidential standards are invariant when they are like that.

¹⁸ Some might suggest Sleeping Beauty and related cases favor evidential standards that are not invariant under third-personalization (e.g., Elga 2000). While I am skeptical, I won't pick a fight. My argument in Section 5 will focus on the particular standards for inferential reasoning required to vindicate reflectivism, rather than this general claim about third-personalization.

I will argue that reflectivists need first-personal reflective beliefs to have a significance that does not survive third-personalization, and so is not afforded them by ordinary evidential standards. So reflectivists need to resist the argument for deflationism. But how? In Section 5, I discuss what I call *transparency reflectivism*, a version of reflectivism that denies EVIDENTIALISM. In Section 6, I turn to *management reflectivism*, which rejects INFERENTIALISM. And in Section 7, I consider *guidance reflectivism*, which instead denies MOTIVATIONALISM.

3. Motivations for Reflectivism

I have found that when I tell people I endorse deflationism, I encounter two immediate objections. The first is that deflationism is obviously false, and the second is that everyone already accepts it. Most of this paper responds to the first objection. But before getting back to it in Sections 5-7, let me address the second.

It sure does not sound to me like everyone accepts deflationism, judging by what they say. For example, the authors I quote in Section 1 do not seem merely to be observing that premises about one's own rational requirements can amount to evidence. And while I won't bombard you with more quotes, reflectivist thinking seems to me pervasive in many debates in epistemology, for example between internalists and externalists about justification. For what it is worth, as a deflationist I routinely feel at odds with most discussions I read about 'reflective beliefs', 'higher-order evidence', 'metajustification', and the like. If I manage to convince you of deflationism, then perhaps you will, too. In fact, an openness to inflationism seems to me implicit in the common practice of using special terms like these when discussing beliefs or evidence about one's own epistemic situation. No one bothers to name the phenomenon of having meteorological beliefs that match one's evidence about what another person's meteorological beliefs should be, because most everyone is a deflationist about evidence like that.

Not to be too dramatic about it, but reflectivism seems to me presupposed by a natural conception of the enterprise of epistemology, which is often pitched (at least to students) as the branch of philosophy trying to help us figure out what to believe. Such a 'meliorative project' was certainly what many historical figures in epistemology thought they were up to. Even today, it is natural to take, say, many conciliationists about peer disagreement as concerned with persuading us actually to be more conciliatory. Yet under deflationism, there is no apparent room for theoretical knowledge about epistemology to guide us in our other beliefs. Epistemological inquiry might help us figure out what to believe *about epistemology*, but not much else.

All that said, a contrary impulse towards deflationism also seems to me common, including among some of the authors I have cited as reflectivists.¹⁹ It may be that on further thought, many of them would prefer deflationism after all. But I do think the inflationist tendencies go beyond some stray remarks. Rejecting deflationism seems to me baked into the central motivations commonly cited for reflectivism.

This is perhaps clearest for ambitious reflectivists like Burge and Smithies. Take for instance their claim that reflection is part of what makes us persons, or what makes us responsible

¹⁹ E.g., Christensen 2016; Lasonen-Aarnio 2014; pp. 317-318; and Smithies in correspondence

for our beliefs.²⁰ It is hard to see how that could be true if reflective beliefs merely give us deflationary reasons. Why single out reflective beliefs, as opposed to beliefs about what other people ought to believe, or beliefs about cabbages and kings, which also can give reasons insofar as they give ordinary evidence? For that matter, if unreflective brutes are able to make ordinary deductive and inductive inferences from premises about other topics, why not from premises about their own rational requirements? In any case, under deflationism a capacity to do so seems unlikely to mark a deep difference between kinds of rational beings.

Hostility to deflationism is not limited to these ambitious claims. Setting aside the importance of reflective reasoning, many moderate reflectivists just find it intuitively incoherent to combine a belief that *p* with a belief that it is irrational for one to believe that *p*.²¹ They accept something like:

(REFLECTION REQUIREMENT) Rationality requires that if you believe that it is irrational for you to hold a given doxastic attitude, then you do not hold that attitude.

This intuition is sometimes reinforced by appeal to Moore's paradox, with the suggestion that it is irrational to believe Moorean conjunctions like *It will rain, but I am rationally required not to believe it will rain*.²²

Reflectivism offers a natural explanation of these intuitions. Suppose I believed the second Moorean conjunct, which says I am required not to believe it will rain. Under reflectivism, that gives me a reason (or otherwise ought to guide me) not to believe the first conjunct. So I cannot believe the Moorean conjunction without having a reason not to. More generally, I cannot combine a doxastic attitude with the belief that it is irrational for me to hold that attitude, without having a reason to stop.

But this motivation for reflectivism seems incompatible with deflationism. Just as I can rationally believe third-personal Moorean conjunctions like *It will rain, but it is irrational for DB to believe it will rain*, I might believe it is irrational for some character DB to believe it will rain having a reason not to believe it myself. And ordinary evidential standards are neutral under third-personalization. Facts about what I am required to believe are not *stronger evidence* about what is true than facts about what others are required to believe, at least under ordinary standards. So it seems the reflectivist here must be proposing that first-personal reflective beliefs have a significance that is not afforded them under these standards.

While I think we should reject REFLECTION REQUIREMENT along with reflectivism, I will not argue for that here. Doing so would require rejecting not just reflectivism, but also:

(LUMINOSITY) You are always in a position to know whether you are rationally permitted to hold a given doxastic attitude.

If you are in a position to know that *p*, then arguably it is irrational for you to believe that not-*p*. If so, then LUMINOSITY will guarantee the truth of REFLECTION REQUIREMENT in

²⁰ E.g., Alsoton 1989, Burge 2013, and Smithies 2019.

²¹ E.g., Smithies 2019, pg. 314 and Worsnip 2018.

²² E.g., Greco 2014; Horowitz 2014, pg. 725; and Smithies 2019.

an entirely “bottom up” fashion. Unless your first-order attitude is antecedently irrational, you will be guilty of irrationality by believing it to be irrational. So rationality will never permit you both to hold a first-order attitude and the belief that the attitude is irrational, just because your higher-order beliefs should track their first-order subject matter.

Reflectivism is consistent with LUMINOSITY, and perhaps even a natural compliment to it. But it’s a distinct thesis. It alleges a “downward” pressure that higher-order judgments or beliefs are supposed to exert on your first-order attitudes. If you believe an attitude to be irrational, the reflectivist says that ought to motivate or otherwise guide you to giving it up. It is that claim, not REFLECTION REQUIREMENT alone, that I say is incompatible with deflationism.²³

While I will not directly oppose LUMINOSITY, I hope to undermine one motivation for it. Often when LUMINOSITY is upheld, the thinking is that you had better be able to know whether your beliefs are rational, or else you will be in a bad situation.²⁴ If you rationally believe p but cannot tell whether you do, you might have suspicions or even the belief that your belief that p is irrational. But what is so bad about that situation? The obvious answer is a reflectivist one: You will have a reason to give up a belief that otherwise would be rational. You also might be thought to be in a bad situation if your belief is irrational but you cannot tell. Why? Well, you won’t know any better than to keep on believing. Here an even stronger reflectivist claim seems assumed, that reflective beliefs are necessary to guide you to revise beliefs. Reject reflectivism, and these motivations for LUMINOSITY go with it. It might still turn out to be true, but we should not say in advance that it had better be.

So much for REFLECTION REQUIREMENT and Moore’s paradox. Let’s turn to another motivation for reflectivism, which appeals to the normativity of rationality. It says that if I believe it is irrational for me to hold a belief, then I think I *ought not* to hold it. And if I am rational, thinking that I ought not to do something will motivate me not to do it. Sometimes this idea is reinforced by an analogy with practical akrasia, as exemplified by agents like the nicotine addict who continues smoking despite believing she should quit.²⁵

This motivation also seems incompatible with deflationism. If the smoker’s practical akrasia is irrational, it is not just because her being obligated to quit amounts to evidence for something she did not already know. Instead, the idea is that first-personal normative judgments have a special motivational significance. By believing you ought not to ϕ , you thereby commit to not ϕ -ing, in some sense requiring a corresponding motivation.²⁶

If that is the idea the reflectivist wants to invoke, she cannot also accept deflationism, for two reasons. First, ordinary evidential standards plausibly assign no distinctive evidential significance to normative premises. The premise that I *ought* to believe p need not be stronger evidence for p than non-normative premises about my believing p that might not rationally motivate the belief. Second, and more importantly here, ordinary evidential

²³ Note that some reflectivism-adjacent views really just combine luminosity assumptions with further claims the deflationist can accept (e.g., Briggs 2009, Egan and Elga 2005, Titelbaum 2015).

²⁴ E.g., Burge 2013.

²⁵ E.g., Greco 2014; Horowitz 2014; Smithies 2019, Ch. 9; Titelbaum 2015; and Wedgwood 2002.

²⁶ For review, see Rosati 2016, Sec. 3.

standards are invariant under third-personalization. The premise that *I* ought to believe *p* need not be stronger evidence for *p* than a third-personal premise that *DB* ought to believe *p*, which might not rationally motivate me to believe. While this gets less attention than the normativity of the smoker's judgment, I suspect most who think practical akrasia is irrational also think it is important that her judgment be first-personal. If I judge that *I* ought to see to it that you win our tennis match, that is supposed to motivate me to throw the match. Not so if I judge that *you* ought to see to it that you win our tennis match.

A final motivation for reflectivism comes from intuitions about particular examples, especially one's involving higher-order evidence. Here I think things are a mixed bag. Some common intuitions commonly taken to favor reflectivism are probably consistent with deflationism, while others might not be. I turn to them next.

4. Higher-Order Evidence and Deflationism

Along with cash prizes, I have been citing moral reasons as paradigmatic reasons of the wrong kind for belief. Maybe that was too quick. Suppose I think one can never be morally required to believe a falsehood, and an oracle tells me I am morally required to believe my friend is innocent. I might rationally infer she is innocent. Does this show that moral reasons can be epistemic reasons, the right kind of reasons for belief?

Not really. Even in this case, the moral reason is not the right kind of reason in virtue of being a moral reason. Rather, the apparent fact that I am morally required just amounts to ordinary evidence for my friend's innocence. The same would go if I were told that some other person was morally required to believe my friend is innocent. For that matter, if I instead thought that one is never morally required to believe a truth, my apparent moral requirement would be an epistemic reason against believing.

I also have been saying that reflective beliefs give you the wrong kind of reasons for belief. Maybe that was too quick, too. Apparent facts about what you are rationally required to believe can amount to ordinary evidence. And to whatever extent they do, everyone should agree that they can be good reasons for belief. I do not mean to deny that when I say that higher-order evidence is the wrong kind of reason.

This is important because rational requirements to believe routinely have significant evidential weight. In a familiar slogan, evidence of evidence is evidence. I think we must be careful to distinguish this boring truth from the interesting falsehoods reflectivists endorse. For the boring truth, consider:

Sherlock's Secret Evidence: While investigating a murder, Sherlock tells Watson that he has gathered such strong evidence that he, Sherlock, is rationally required to believe the butler did it. Watson accepts Sherlock's claim.

Sherlock's being required to believe the butler did it is evidence that the butler did it, just like a fingerprint match might be. People can be required to believe falsehoods when their evidence is misleading, but we think this is uncommon. So when Watson accepts that Sherlock is required to believe the butler did it, Watson gets a reason to believe it himself. But the explanation of this is arguably a boring, deflationary one. Watson gets a reason only insofar as he gets evidence that supports the butler did it in the usual way, by providing deductive, inductive, or abductive support for it, or something else of that kind.

More generally, when will some other agent's being apparently required to believe that p be a reason for one to believe p ? Here I will mostly assume that we are all deflationists when it comes to other people's rational requirements. If you think someone else ought to believe p , that will rationally motivate you to believe p only insofar as the premise that the person ought to believe p is ordinary evidence for p . But this assumption is not totally uncontroversial. For example, some fans of epistemic entitlement might think one is entitled to believe that rational beliefs are usually true, even without evidence or argument. For what it is worth, I suspect that if reflectivism is rejected, this view will seem unmotivated. Why would we need entitlement to accept *that*, unless one is supposed to be guided by beliefs about what is true via beliefs about what is rational? But in any case, my target is reflectivism, not the view that you can accept without evidence that rational beliefs are usually true.

Assuming we all are deflationists about Sherlock's Secret Evidence, compare it to:

Sherlock's Admonition: While investigating a murder, Watson is confronted with inconclusive evidence that his good friend, the butler, did it. Watson rationally withholds belief, proposing to himself various alternative explanations of the evidence. But Sherlock admonishes Watson to stop rationalizing, telling him that rationality requires belief given Watson's own evidence. Watson accepts Sherlock's claim.

Here it is plausible that Watson should be motivated to believe the butler did it. But that is compatible with deflationism. The deflationist can just say Watson gets the boring, deflationary kind of reason to believe that he got in Sherlock's Secret Evidence.

But some reflectivists might not be satisfied that this gives a complete accounting of the case. They might think that when Watson thinks *he* is required to believe, this has a special relevance to his deliberations. If he treats his apparent requirement as just one more bit of evidence, as he would the apparent requirements of other agents, then Watson is missing something. For here the person doing the deliberating is the very same person who is apparently required to resolve those deliberations by believing the butler did it. If Watson is rational, this will lead him to resolve those deliberations as he apparently is required to, by adopting the belief.

So while reflectivists and deflationists might disagree about why Watson should believe that the butler did it, they might agree that he should. There are other cases, however, where they might not even agree on that much, such as:

Watson's Apparent Mistake: Sherlock and Watson gather a large body of clues regarding a murder. Watson concludes that the butler did it via a long chain of deductive reasoning, and reports his conclusion to Sherlock without recounting the reasoning. Sherlock reminds Watson of his track record of committing fallacies in long chains of reasoning. Though Watson's reasoning was in fact flawless, he accepts that the objective chance of him reasoning flawlessly is low. So does a bystander, who was eavesdropping on their conversation.

It is intuitive that Watson should seriously reduce his confidence that the butler did it, once he speaks to Sherlock. Can the deflationist accommodate this intuition, or must she bite the

bullet and resist it? I will not take a stand on the question here, but will instead just recount why bullet biting might seem necessary.

The deflationist's best hope to accommodate the intuition is to compare Watson's situation to the eavesdropping bystander's. To explain why this bystander should reduce her confidence, we do not need to assign inflationary significance to anything the bystander knows or believes about Watson. Without taking a stand on the details, the explanation plausibly can be given in terms of general evidential standards. For example, it might be that the best explanation of why Watson believes the butler did it is no longer that it is true, once she learns of his track record. Or it might be that her credence that Watson's belief is true should follow her beliefs about the objective chance it is.

So if Watson's situation is broadly comparable to the bystander's, the deflationist can be accommodate our intuitions. The problem for the deflationist, or perhaps instead for our intuitions, is that Watson knows the clues. And thus unlike the bystander, he has total evidence that arguably still supports that the butler did it, at least by ordinary standards. Why? One answer is that Watson's evidence still deductively entails that the butler did it. But this assumes that deductive entailment is sufficient for support, which is plausible but not universally accepted. A better answer appeals to the invariance of support under third-personalization. For Watson's third-personal evidence is the same as Sherlock's, whose evidence surely supports that the butler did it. The clues by themselves support that the butler did it. And beyond this, both characters know that Watson inferred from these clues that the butler did it, and that Watson's track record is poor. Of course, Watson also knows that *he* is Watson, and thus that *his* track record is poor. But this will not make room for his total evidence to fail to support that the butler did it by ordinary standards, if these standards are invariant under third-personalization. Suppose for instance that the simplest explanation of Sherlock's third-personal evidence is that Watson hit on the right answer by dumb luck. That also will be the simplest explanation of Watson's evidence, if abductive standards are invariant.

So deflationism at least has trouble accommodating common intuitions about higher-order defeat, as in Watson's Apparent Mistake. This is roughly what led Christensen to suggest, as quoted in Section 1, that there is a more immediate connection between our reflective beliefs and our first-order rationality.²⁷ The idea is to distinguish Watson's situation from Sherlock's, by granting special significance to Watson's first-personal evidence that *he* has a poor track record. And this idea naturally lends itself to reflectivism, if we think Watson should conclude that *his* belief is probably irrational.

Now some recent discussions including Christensen's say it is Watson's first-personal beliefs about his *reliability*, rather than his rationality, that matter.²⁸ Perhaps this is a step away from conventional reflectivism, which stresses the normative content of reflective beliefs about rationality. But still, I think it is not on its own sufficient to uphold deflationism. So long as it is supposed to be first-personal beliefs about his reliability that motivate Watson to revise, this view still has Watson motivated by first-personal evidence in a way that is not licensed by ordinary evidential standards.

²⁷ Cf. Dutant and Littlejohn (MS), who think the modal condition of safety is normatively relevant, because necessary for knowledge.

²⁸ Christensen 2016; Weatherson 2019, Ch. 8; and perhaps Horowitz 2014.

Where does this leave us? Deflationists might follow common intuitions, and say that Watson's situation is like the bystander's, and that he should suspend just as she should. Or they might reject these intuitions, and say Watson's epistemic situation is instead like Sherlock's. As I see it, reflectivists do neither of these things. They agree with the intuition of defeat, but at the same time they grant that Watson, unlike the bystander, has total evidence supporting the butler's guilt. So they say that his first-personal track record evidence, in some way or other, motivates revision in a way it could not under ordinary third-personal evidential standards.

In these familiar cases it is at least debatable whether common intuitions can be explained by the ordinary evidential significance of reflective beliefs. Are there cases where it is clearer that reflective beliefs fail to provide the agent with enough evidence to license belief revision? I think they are easy to come by, if we allow for dramatic failures of self-knowledge. Just compare:

Guy's False Lemma: Gal knows that a weather forecast calling for rain is based on faulty radar data. Unbeknownst to Gal, Guy knows this, too. But Gal takes Guy to believe the false lemma that the forecast is reliable. She knows Guy has read the forecast, and so believes him to be rationally required to believe it will rain.

Gal's False Lemma: Gal knows that a weather forecast calling for rain is based on faulty radar data. But Gal cannot tell that she knows this, and takes herself to believe the false lemma that the forecast is reliable. She knows she has read the forecast, and so believes herself to be rationally required to believe it will rain.

When Gal thinks Guy is required to believe it will rain, this gives her no reason to believe it will rain. She thinks this because she takes him to believe a false lemma, and so his apparently being required to believe it will rain is not for her evidence that it will. Likewise, I say that when Gal takes herself to believe the false lemma, that means her apparently being required to believe it will rain also gives her no reason to believe. If she apparently is required just because of believing a false lemma, this is not evidence it will rain, and it would be irrational for it to motivate her to believe. The deflationist says, I think correctly, that Gal should not be motivated to believe.

This example turns on a dramatic failure of self-knowledge. Real people are not that bad at knowing about their own beliefs and evidence. But I am not so sure that prevents us from rendering judgment about it. For my part, it seems intuitive that if someone was misinformed the way Gal is, she should be unmoved by apparently finding herself required to believe it will rain. Indeed, even ordinary cases of higher-order evidence, like Sherlock's Admonition and Watson's Apparent Mistake, require some failure of self-knowledge. Even supporters of LUMINOSITY, who regard those failures as rational failures, have to admit it is possible for an agent to engage in them. And if so, it seems intelligible to ask how they should proceed once they do. Given that Watson in fact accepts Sherlock's testimony, we want to ask, what should he think about whether the butler did it? The same question seems intelligible to me in Gal's case, and the deflationist's answer seems intuitively correct.

Still, I realize Gal's failure of self-knowledge might be a bridge too far. More dialectically effective examples must limit themselves to realistic failures of self-knowledge, even if that

means adding further epicycles. I think there already are several examples in the literature that do just this.²⁹ Here is one more:

Sherlock's Preface: Watson rationally suspends judgment on who committed a murder, based on clues that leave it open. Watson writes a book detailing all the conclusions he has managed to draw, noting his failure ultimately to determine who did it. Sherlock writes a preface for the book, which asserts: "Watson's evidence indeed supports each claim made in this book. But I know his evidence is misleading in one respect. One of these claims, while supported, is false. And Watson has failed to notice that from this false lemma it follows that the butler did it." Watson accepts what Sherlock says, and concludes that he must even now be rationally required to believe that the butler did it.

As with Gal, Watson takes himself to be required to believe, but only because he thinks his evidence supports a false lemma. An agent's being required to hold the belief in this way is not evidence that the butler did it, at least not necessarily. So deflationism says, I think correctly, that Watson does not get a reason to believe the butler did it.

Now certain features of the example are debatable. It might be argued that Watson really should not believe that he is required to believe the butler did it, now that he has read Sherlock's preface. (But is this reflective belief so irrational that I cannot even stipulate Watson holds it, and ask how he should proceed?) It might even be claimed that reflective beliefs can only give inflationary reasons for suspending judgment, never for believing. (But suppose Watson instead initially rationally believes the butler did it, and Sherlock's preface speaks not of a false lemma supporting this but of a false defeater for it. Doesn't that seem like no reason to suspend?)

In any case, I do not want too much to hang on this or any other particular example. For any given case, the reflectivist might claim that the reasons usually provided by reflective beliefs are defeated or outweighed. Now the reflectivist had better not say this about every case where apparent facts about one's rational requirements fail to be evidence. That risks collapsing the view into deflationism. And reflective reasons had better not be vanishingly rare, either. Reflective reasoning is supposed to be something we actually do, not just something an agent could do in some byzantine example. But reflectivists can still allow reflective reasons to be defeasible, and this gives them wiggle room.

So here I will not try to refute reflectivism by counterexample, or even take a definitive stand on which examples are supposed to be the ones where reflective beliefs offer guidance beyond their weight in evidence. Instead, I will focus on defending the general argument for deflationism sketched in Section 2. Now that we have seen why deflationists need to resist this argument's conclusion, let's consider how they might go about it.

5. Transparency Reflectivism

Start with *transparency reflectivism*, which denies EVIDENTIALISM. This was the premise that evidence governs inference, in the sense that a premise *r* is relevant to *p* in inferential reasoning only insofar as *r* is ordinary evidence regarding *p*. This premise is plausible, but

²⁹ See, e.g., Christensen 2016; Schoenfield 2015; Turri 2012; Worsnip 2018; and Horowitz's (2014) and Weatherson's (2019) discussions of Williamson's (2011) unmarked clock case.

not trivial. It relates a kind of reasoning, exemplified by inference, to a relationship between propositions, exemplified by deductive entailment and probabilistic support. Now it could be claimed that really our only grip on the notion of inferential reasoning is that it is the kind of reasoning governed by evidential support, or that our only grip on evidential support is that it is the relation governing inference. But I will assume for the sake of argument that we have enough of an independent grip on these two notions to at least make the denial of EVIDENTIALISM intelligible.

When I infer that it will rain from the premise that the weather forecast says so, I reason in a way I might express by saying “The weather forecast says it will rain, so (probably) it will rain.” In other words, I treat the forecast as *bearing on the question* whether it will rain. In contrast, if I decide to believe it will rain because I know I will get a cash prize, I reason in a way I might express by saying “I will get a prize for believing it will rain, so I’ll go ahead and believe it.” In doing that, I do not treat the prize as bearing on whether it will rain, but instead on the question whether to believe it will rain. Only in the former case do I arrive at the belief it will rain by inferential reasoning.

While I use natural language expressions of inferential reasoning to illustrate the kind of reasoning I have in mind, I do not suppose that it always is explicit or conscious. I mean for ‘inferential reasoning’ to be a very inclusive notion, covering most any way an agent might be persuaded that something is true by reasons or premises. Nor for all this talk of inference and persuasion is inferential reasoning always supposed to terminate in belief, as opposed to other doxastic attitudes. Suppose an agent already believes it will rain based on a weather forecast. If she then learns defeating evidence—that there are dark clouds outside, or that the forecast was based on faulty data—she might reason her way to suspended judgment in a broadly inferential manner. For example, she might express her reasoning by saying: “The forecast was based on faulty data, so maybe it won’t rain after all.” Contrast this with a case where she is offered a cash prize for suspending judgment on whether it will rain. If she deliberates about whether to suspend, for example by pressing a button, she might rationally decide to do so. But if so, it will be through a distinct kind of deliberation, which she might express by saying “I’ll get a prize for suspending, so I’ll go ahead and do it.”

That’s it for inferential reasoning. Next up is ordinary evidential standards, like those of deductive, inductive, and abductive support. There is of course room for disagreement about precisely which evidential standards or inference rules belong in this company. With some possible views, it might be unclear whether they reject EVIDENTIALISM, or merely propose unorthodox standards of evidential support. But I do not see the reflectivist as proposing a borderline evidentialist view like that. She needs rules or standards for inferential reasoning that are quite unlike these familiar ones. For as discussed in Section 2, these are invariant under a third-personalization of one’s evidence. If I can inductively infer that I will be bored at tomorrow’s faculty meeting from the premise that I have been bored at every previous faculty meeting, then I can inductively infer that DB will be bored at tomorrow’s faculty meeting from the premise that DB has been bored at every previous faculty meeting.

In contrast, reflectivists want reflective beliefs to have a rational significance that does not survive third-personalization. Reflectively believing that I ought to believe *p* is supposed to motivate me to believe *p* in a way that does not go for believing that DB ought to believe *p*. The transparency reflectivist in particular is one who thinks that reflective beliefs motivate by broadly inferential reasoning. She wants to let me infer that *p* from the premise that I am

required to believe p , even when my third-personalized evidence would not license the same inference from the premise that DB is required to believe p . And that is why she finds herself at odds with EVIDENTIALISM, which makes no such allowance. For the transparency reflectivist, the rationality of reasoning “I am required to believe p , so p ” is not supposed to come from antecedent evidence I have that I would not be required to believe p if it weren’t true.

Transparency reflectivism is so called because it might be likened to transparency accounts of how we know our beliefs.³⁰ These accounts say that the first-personal question whether I believe it will rain is not for me rationally independent of the question whether it will rain. Instead, I somehow can conclude that I believe it will rain by considering only the apparent facts about the weather. At least as developed by Alex Byrne, the idea is that I can infer the conclusion that I believe it will rain from the premise that it will rain.³¹

It is natural to motivate transparency reflectivism in an analogous way. Indeed, it often is suggested that the question whether it will rain is not from my perspective independent of the question whether it is rational for me to believe it will rain, the way it is independent of the question what it is rational for other people to believe.³² Maybe the idea is that I can answer the one question based on my answer to the other, as the transparency reflectivist supposes.

But transparency accounts of self-knowledge also face a natural objection. Apparent facts about the weather give me reason to ascribe beliefs to other people only insofar as they amount to evidence. If I know you are a weather expert, then maybe I can infer from the fact that it will rain that you believe it. But this will not be true as a general rule. People are routinely ignorant of matters of fact, and I know it. The objection for the transparency account is that if in my own case I do in general reason this way, then in some sense I seem to be reasoning as if the same did not go for me. That is, I will be reasoning as if I chauvinistically considered myself less susceptible to *ignorance* than an arbitrary person.³³

I will press a related objection to transparency reflectivism. If I settle whether it will rain just by considering the apparent fact that I am required to believe it, in a way I would not by considering that DB is required to, then I seem guilty of an analogous chauvinism. This time, I will be chauvinistically treating myself as less susceptible than an arbitrary person to being *misled*—i.e., to being rationally required to believe a falsehood.

Just as I know other people can be ignorant of the truth, I know they can be misled into believing falsehoods. Even when it will not rain, some unfortunate soul could have evidence supporting that it will rain, for example because it includes data from a malfunctioning radar system. It surely would be irrationally chauvinistic for me to believe it is impossible for me to be misled in the same way. I can rely on a malfunctioning radar system as well as anyone.

³⁰ E.g., Byrne 2018, Fernandez 2013, and Moran 2001.

³¹ Byrne 2018

³² E.g., Gibbons 2006 and 2013.

³³ See Barnett 2016, Boyle 2011, and Valaris 2011.

It would be just as chauvinistic to believe that the objective chance of possessing misleading evidence is lower for me than for other people, absent any evidence that this is so. Transparency reflectivism had better not force me to adopt unfounded beliefs about my powers to avoid misleading evidence. (Of course, I might with sufficient evidence think that my radar system is especially reliable, but then my third-personalized evidence would support the same of DB's radar.)

It furthermore would be chauvinistic if for any proposition p , I am especially confident that I will not be misled about p , even while acknowledging that the objective chance is the same for me as for others. Suppose I think our two radar systems are equally reliable, and that we have equal chances of being misled. I cannot then be especially confident that I won't be misled that it will rain, and that I won't be misled that it will be sunny, and that I won't be misled that it will snow, and so on.

But the transparency reflectivist still wants to allow cases where for some proposition p , I am prepared upon adopting the relevant reflective belief to reason "I am rationally required to believe p , so p "—and yet where I would not do so if my evidence and reflective belief were third-personalized. For example, suppose I already have learned the radar data for tomorrow but have not yet analyzed it. The transparency reflectivist might say that if an oracle tells me that the evidence in my possession requires me to believe it will rain, this should convince me it will rain, even when DB's being apparently required to would not be convincing. The problem is that it is hard to insulate a willingness to reason in this way from the kind of outright beliefs that I have already claimed are chauvinistic.

There are a number of ways to press this general problem. Here I will focus on two. The first appeals to a plausible connection between categorical and hypothetical reasoning:

(HYPOTHETICAL-CATEGORICAL COHERENCE) If one can infer p from r , then one can accept p under the hypothetical supposition that r .

While HYPOTHETICAL-CATEGORICAL COHERENCE is plausible, in conjunction with transparency reflectivism it leads to trouble. By transparency reflectivism, if I become convinced that I am rationally required to believe it will rain, then I should reason as follows: "I am rationally required to believe it will rain, so it will rain." But even if I am not sure whether I am rationally required to believe it, I still can hypothetically entertain the possibility that I am. HYPOTHETICAL-CATEGORICAL COHERENCE says that if so, I can accept that it will rain under this hypothetical supposition. But now it is a short trip to chauvinistic beliefs. Here is one way to get there:

(RAMSEY TEST) If one can accept p under the hypothetical supposition that r , then one can categorically believe that if r , then p .

By RAMSEY TEST it will follow that I can categorically believe that if I am rationally required to believe it will rain, then it will rain. And in the same way, I can believe that if I am rationally required to believe it won't rain, then it won't. But if so, I can chauvinistically believe I am not misled about whether it will rain, even if I could not believe this of DB under a third-personalization of my evidence.³⁴

³⁴ Cf. Barnett 2016, Chalmers and Hájek 2007, and Valaris 2011.

This result does not depend on anything suspicious involving indicative conditionals. The usual shenanigans with indicative conditionals involve the converse of RAMSEY TEST, and anyway other principles involving hypothetical reasoning lead to the same results. For example, if I suppose for *reductio* that I am misled that it will rain, then it follows that it will not rain. But since being misled that it will rain also entails being rationally required to believe it will rain, by transparency reflectivism and HYPOTHETICAL-CATEGORICAL COHERENCE I can also accept that it will rain. Thus I can derive a contradiction under the supposition that I am misled.

The second way to press the chauvinism problem appeals to closure and transmission principles. The first is:

(CLOSURE) It is irrational for one to believe that p unless it is rational for one to disbelieve that not- p and r .

Now if I am misled that it will rain, that just means that it won't rain, but I am required to believe it will. So by CLOSURE, it is irrational for me to believe that it will rain unless it is rational to disbelieve that I am misled that it will rain. If I am going to take the view that it will rain, I cannot remain agnostic on the further question whether I am required to falsely believe it.

Where r *eliminates* a proposition iff one can infer its falsity from r , the second principle is:

(ELIMINATION) The premise that r does not eliminate the possibility that not- p and r .

The rationale for ELIMINATION is this. In order for a conjunction to be true, both its conjuncts must be. So I cannot very well conclude that a conjunction is false based on the premise that one of its conjuncts is true. In particular, since I cannot be misled that p without being rationally required to believe that p , my apparently being so required is not any reason to disbelieve I am misled, much less a sufficient one.³⁵

This gives us another route from transparency reflectivism to chauvinistic beliefs. By transparency reflectivism, my apparently being required to believe p can be a reason to believe p even if I lack sufficient third-personalizable evidence that I am not misled. But if I am to believe p , then by CLOSURE I also must believe I am not misled that p . But what reason will I have to believe that? I cannot believe it for the reason that I am required to believe p , by ELIMINATION. So it must be antecedently rational for me to believe I am not misled, despite lacking adequate third-personal evidence that DB is not misled.

While I have been pushing these problems for reflective beliefs that I ought to believe something, the same goes if I reflectively believe I ought *not* to believe that p . In that case, the transparency reflectivist might have me reason "I am required not to believe p , so maybe not- p ". But that also involves chauvinism, if my third-personalized evidence would not recommend reasoning "DB is required not to believe p , so maybe not- p ."

Suppose for instance I am in a desert where it has not rained in many years. I am confident that it will not rain tomorrow based on the base rate, but I also have in my possession some unanalyzed meteorological data. I know that there is a small chance this data will provide

³⁵ Cf. White 2006.

some evidence for rain, in which case I should on balance suspend judgment that it will rain. If I were told by an apparent oracle that I am required to suspend judgment given my total evidence, transparency reflectivists might consider this sufficient inferential reason to suspend, even if it would not be if third-personalized. But if so, then I already can accept that if I am required to suspend judgment, then it might rain. If under third-personalization I would not accept the same thing for DB, then it seems I am chauvinistically less open to the possibility that my meteorological data misleadingly supports rain.

Similar points go for CLOSURE and ELIMINATION. If I would remain confident of no rain even if convinced that DB ought to suspend, then I must already consider it more likely for DB to have misleading meteorological data supporting rain than for him to have accurate meteorological evidence of rain. For after I think DB ought to suspend, I would violate CLOSURE by not believing that DB's meteorological evidence misleadingly supports rain. And I cannot discount the possibility that DB's meteorological evidence accurately supports rain on the basis that it supports rain, or else I violate a cousin of ELIMINATION. Yet if despite all this I will suspend when convinced that I ought to suspend, then for similar reasons I cannot consider it too unlikely that I have meteorological evidence supporting rain accurately, compared to its doing so misleadingly. So again I end up chauvinistically considering myself less likely to have misleading meteorological evidence for rain (or else more likely to have evidence accurately supporting rain).

Can the reflectivist just reject all these principles about hypothetical reasoning and the like? That is not so easy. It is one thing just to say an apparent requirement to believe p bears on whether p without amounting to inductive or abductive evidence. Maybe these ordinary evidential standards are too limited, and only give us an incomplete picture of inferential reasoning. But it is another thing to say that the premise that one is required to believe p is enough for you to infer p when it is believed, but not when it is hypothetically supposed. Or that one can infer the falsity of a conjunction from the truth of one of its conjuncts. These claims seem harder to square with a more general notion of reasoning in which one considers premises or considerations that might bear on what is true.

6. Management Reflectivism

Next up is *management reflectivism*. It denies INFERENTIALISM, the premise that an apparent fact can rationally motivate you to believe only by serving as a premise in inferential reasoning. This was the same premise that was rejected by pragmatists like Susanna Rinard, who said that when Cash comes to believe it will rain by pressing a button, the prize he was offered for believing it will rain will be his motivating reason not just for pressing but also for believing. If so, Cash's reason for believing it will rain does not need to bear directly on the question whether it will rain, as inferentialists hold. It can instead bear in some other way on the desirability or appropriateness of believing it will rain. Likewise for the management reflectivist, the upshot of rejecting INFERENTIALISM is to allow reflective beliefs to motivate you to adopt beliefs you deem required, even without bearing directly on their truth.

By not requiring reflective reasons to serve as premises, management reflectivists are especially well positioned to accommodate the anti-akrasia motivation for reflectivism, which emphasizes the normativity of rationality. For comparison, if I think I am morally required to believe my friend is innocent, it will normally be irrational for me to treat this as a premise, and infer that he must really be innocent. But perhaps I should be motivated to produce this belief in myself somehow, if not by direct force of will, then by indirect means

like talking myself into it, or pressing a button. If I instead thought I ought to get you to believe he is innocent, maybe that should motivate me to produce the belief in your mind. Plausibly the same should go for producing the belief in my own mind, when I think I ought to.

Now management reflectivists presumably do not want reflective reasons to come out looking too much like moral reasons or cash prizes. But they can cite further differences, without sacrificing the core idea that reflective reasons bear on whether to hold a belief without necessarily bearing on its truth. After all, one's apparently being required to believe still concerns the epistemic merits of the belief, rather than its moral or financial merits. That might at least make it a better candidate to be an epistemic reason, the right kind of reason for belief.

There are different ways the core idea behind management reflectivism might be developed. It could be claimed that all reasons for belief ultimately work by bearing on the appropriateness of the belief. Knowledge-first epistemologists sometimes appear to take this line, though they differ from conventional reflectivists in considering knowledge rather than rationality to be fundamental norm for belief.³⁶ For example, on Dutant and Littlejohn's account of higher-order defeat, evidence that one does not know p defeats one's belief that p , even when it does not amount to evidence against p itself.³⁷ The idea is that in general, evidence justifies or defeats belief that p by bearing on the question whether p can be known, and thus whether it is appropriate to believe it.

Other management reflectivists might grant that evidence directly bearing on whether p can motivate belief in a distinctive way, via inferential reasoning. But they might add that in addition to the kind of low-grade rationality that goes with inferentially responding to evidence, there is a higher-grade rationality that involves properly managing one's epistemic affairs. For example, suppose it is important to know whether p , and you easily could find out, but you simply do not bother. Your uncertainty about p might fit your evidence, but still you seem to have failed to manage your epistemic affairs. Perhaps reflective reasons could be reasons like the ones you fail to respond to here.

But however the management reflectivist fills in the details of her view, I think she ends up sharing a big-picture problem with conventional pragmatism, and facing an additional problem of her own. The big-picture problem is that management reflectivism makes us out to be implausibly *alienated* from our beliefs. In principle, you might be able to manage someone else's beliefs, for example via belief-producing buttons. You might make them believe it will rain in order to secure a cash prize, or just because you think it would further their epistemic interests. But your relationship to your own beliefs is not like that. When you reason and deliberate, you are not in the position of trying to figure out which beliefs to plant in your own mind.³⁸ The worry about management reflectivism is that it makes things out to be that way, by letting reflective beliefs bear merely on the desirability or appropriateness of your having the relevant beliefs.

³⁶ E.g., Dutant and Littlejohn MS, and Gibbons 2013.

³⁷ Dutant and Littlejohn MS.

³⁸ Cf. Hieronymi 2005.

Now many reflectivists might agree with this complaint, and distance their own views from what I am calling management reflectivism. I take this to be Moran's intention when he denies that doxastic deliberation is like organizing one's closet, for example. And while I am less sure what he had in mind, Burge has emphasized that reflection is not like looking after the beliefs of a student or child. But others seem more favorably disposed to the view, at least on a first pass.³⁹ For example, I read the passage from Roger White quoted in Section 1 this way. And in other work, White is explicit that the reflective management of beliefs and actions involves something like what I am calling alienation.⁴⁰

Suppose we follow White in thinking reflective reasoning should involve managing one's beliefs, perhaps from an alienated perspective. There is a further problem, involving the aims or goals of the reflective belief manager. To some of us it already seem suspicious to assume that an agent, just by virtue of being fully epistemically rational and reflective, needs to have any particular set of aims regarding the merits of her beliefs. But management reflectivists have various replies to this worry.⁴¹ So I will grant for the sake of argument that some sort of concern for the the epistemic merits of one's beliefs is rationally mandatory. What I object to are the particular concerns the management reflectivist must demand, if she is to accommodate the usual motivations for reflectivism. In particular, she would have us manage our beliefs based on a *fetishistic* concern for rationality, rather than a concern for truth.

To motivate the idea, start with cases where one agent manages another's beliefs:

Patient's Lucky Mistake: Agent believes it will rain, and knows that Patient also believes this. But Agent also has reason to believe that Patient's belief is based on fallacious reasoning from evidence that does not really support rain. Agent can press a button causing Patient to instead suspend judgment.

How should Agent manage Patient's belief, if she has the appropriate concerns for its epistemic merits? It is at least tempting to say she should not press, despite knowing this would replace Patient's irrational attitude with a rational one. In pressing, Agent plausibly would be fetishizing rationality in an objectionable way. She would put the rationality of Patient's belief ahead of its truth. Even if we had no positive objection to Agent's fetishizing rationality like this, it surely is not mandatory.

Fetishism seems equally objectionable (or at least non-mandatory) when managing one's own beliefs, as in:

Guy's Lucky Mistake: Guy is told by an oracle that he has some unspecified belief that, while coincidentally true, is based on fallacious reasoning from evidence that does not support it. He can press a button that will cause him to give up whichever belief the oracle is referring to.

³⁹ E.g., Dutant and Littlejohn MS, and Schoenfield 2014, Secs. 3.1-3.2.

⁴⁰ 2005, Sec. 6.

⁴¹ Dutant and Littlejohn MS; Gibbons 2013; and Shah and Velleman 2005.

It is tempting to say that Guy should not press in this case, either. To do so would be to fetishize rationality, by prioritizing it over truth. We might imagine Guy thinking “As much as I like having true beliefs, rationality requires that I give this one up, so I’ll go ahead and do it.” If not objectionably fetishistic, these epistemic priorities are hardly mandatory for all rational agents.

Management reflectivists might accept what I say about these cases, but deny that it carries over to typical doxastic management. Typically we are not fetishistically prioritizing rationality over truth. Instead, we try to adopt rational beliefs as a means to adopting true ones. If an oracle tells us which buttons produce true beliefs and which produce rational ones, then we would not need to worry about rationality. But as it stands, our best shot at having true beliefs is to have rational ones. In circumstances like that, there is nothing fetishistic about focusing on rationality, and hoping that truth comes with it.

But I doubt the management reflectivist can really avoid the problem of fetishism this way. Belief management avoids fetishism only at the cost of redundancy. Consider:

Mystery Detective’s Lucky Mistake: Watson rationally believes that the butler did it, despite having been told by Sherlock that a Mystery Detective believes the same thing based on fallacious reasoning. Watson can press a button that will cause the Mystery Detective to withhold belief, but decides against pressing. Sherlock then informs Watson that he, Watson, is the Mystery Detective.

Reflectivists are united that when Watson learns he is the Mystery Detective, he somehow should end up withholding belief that the butler did it. But how? The transparency reflectivist had an answer with some intuitive appeal. He said that Watson gets a reason bearing on whether the butler did it, which he might employ in inferential reasoning like: “I am the one who is required not to believe that the butler did it, so maybe he didn’t do it after all.” But if Watson reasons like that, then he also ought to accept that maybe the butler didn’t do it under the merely hypothetical supposition that he is the Mystery Detective. And that is what I claimed in Section 5 leads to chauvinism.

The management reflectivist has a different answer, with the potential to avoid chauvinism. She agrees that Watson should be motivated to suspend, including by pressing the button. But that is not because he is antecedently willing to accept that maybe the butler didn’t do it under the hypothetical supposition that he is required to. Just supposing he is required to might give him no motivation to press a button causing hypothetical suspension, for example. The supposedly required attitude is *categorical* suspension, and anyway he is merely supposing that he is required to adopt it.

This gets management reflectivists out of chauvinism, but at the same time sticks them with fetishism. When Watson speaks to Sherlock, they cannot have him non-fetishistically reason “Since I’m the one who is required not to believe that the butler did it, maybe he didn’t do it after all. So I’ll go ahead and suspend judgment.” If he did, pressing the button would be redundant, for he *already* would have suspended via a chauvinistic inference. To avoid redundancy, the management reflectivist needs Watson instead to think, “I’m the one who is required not to believe that the butler did it, so even though my belief is (probably) true, I’ll go ahead and suspend judgment.” But that is as fetishistic as pressing in Guy’s Lucky Mistake.

7. Guidance Reflectivism

So far we have focused on versions of reflectivism that take reflective beliefs to motivate our first-order beliefs by providing reasons. Not so *guidance reflectivism*. It denies MOTIVATIONALISM, the premise that providing motivating reasons is the only way for beliefs to guide us. In merely denying that this is how reflective beliefs guide us, I take it to be more open-ended than transparency and managerial reflectivism. The challenge for guidance reflectivists is to give a positive characterization of the sort of guidance in question, which avoids the fetishism and chauvinism problems, but still captures the initial motivations for reflectivism.

To preserve reflectivism's motivations, the guidance reflectivist must be careful not to make reflective beliefs out to be merely incidental parts of the etiology of the beliefs they are supposed to guide. If I believe that *p* just because a brain-scanning microchip detected that I reflectively believed it was required, that presumably is insufficient for my belief to be rational. But at the same time, the guidance reflectivist cannot say that the reflective belief ought to motivate me to believe. So what is its role, if not to motivate?

Perhaps the best option for guidance reflectivists is to take reflective beliefs to be involved somehow in basing and inference. When I believe that my evidence supports *p*, perhaps that should lead me to believe *p* based on that evidence, rather than based on the further premise that my evidence supports *p*. This suggestion fits naturally with common views about inference and the basing relation, on which reflective beliefs are crucially involved.⁴² In a slogan, it says reflective beliefs are reasoning-guiding, rather than reason-providing.

But I think it is not obvious there is a non-incidental reasoning-guiding role available for reflective beliefs. When I believe that Fermat's Last Theorem is supported by Peano's axioms, that does not just automatically guide me to me believing the theorem based on the axioms themselves. Instead, I believe it based in part on the further premise the axioms support the theorem. For me at least, my belief about support seems to be reason-providing. Now perhaps the belief could motivate someone indirectly to construct a proof, as Andrew Wiles did. When he hit on the proof, Wiles arguably did end up believing the theorem based on the axioms themselves. But in that case, his earlier belief about support seems like a merely incidental part of the etiology of his ultimate belief in the theorem. For comparison, his belief in the theorem might be just as rational if he hit on the proof by accident, or in an attempt to collect a cash prize for time spent thinking about the theorem.

Now maybe I am missing some more immediate way for reflective beliefs to guide basing. But there still are risks of falling into the fetishism problem facing management reflectivism. Whatever else the role of reflective beliefs in basing might be, it cannot be to supply motives for seeing that you end up in the state of believing based on your evidence, as you might by pressing a button. That would again cast reflective beliefs as motivating belief management, albeit with some extra bells and whistles. Instead of just seeing to it that you have the required belief, you would be motivated to see to it that it is held on some particular basis.⁴³ But it is unclear how this further constraint would make belief management any less fetishistic.

⁴² See, e.g., Neta 2019, and Boghossian's (2014) discussion of the 'intention view'.

⁴³ Cf. Whiting 2017.

Nor can the role of reflective beliefs in basing be to supply suppressed premises, or else we are back at chauvinism. If inferring p from e just requires you to believe some further *premise* connecting e to p , then merely hypothetically supposing that the premise is true should allow you to accept that p . Knowing e , you should be able to say “If e is connected to p in that way, then (probably) p .” But on pain of chauvinism, you cannot in the relevant cases be allowed to say “If my evidence supports that p , then (probably) p .” So whatever reflective beliefs contribute to basing, it cannot be suppressed premises.

But suppose the guidance reflectivist can navigate these problems, and find a role for reflective beliefs that is not too incidental, but that also does not involve supplying suppressed premises or motivating belief management. Even so, I think she has trouble accommodating the original motivations for reflectivism, especially ones involving higher-order evidence.

Consider for instance Sherlock’s Admonition, where Watson is persuaded by Sherlock’s testimony that a given set of clues support that the butler did it, when in fact they do not. Just believing his clues support that the butler did it is not sufficient for Watson to believe it based on his clues, any more than my belief about Fermat’s Last Theorem is sufficient for me to believe it based on the axioms. For basing, Watson needs something else, in addition to the reflective belief. The problem is that it seems that something else will have to include fallaciously inferring that the butler did it from clues that do not really support that. Even if we agree Watson should come to believe the butler did it, we do not want him to get there by making fallacious inferences.

A similar problem goes for agents who are misled about what evidence they possess. Suppose an agent knows both that total evidence e_1 supports that p , and that e_2 fails to support that p . While in fact her total evidence is e_2 , she thinks it is e_1 , and so reflectively believes she is required to believe p . Some reflectivists might say that an agent like that still should be led to believe that p , on pain of akrasia and Moore-paradoxicality. But her reflective belief cannot guide her to hold this belief based on e_1 , which she does not possess. And it is also hard to see how it could guide her to rationally believe it based on e_2 , which she knows does not support that p .

Guidance reflectivists might respond with a concession. Perhaps believing one’s evidence supports an attitude is not sufficient for rationally basing that attitude on the evidence. But it still might be necessary. This concession will prevent the guidance reflectivist from handling the cases above. But it might leave her enough for cases of higher-order defeat, like Watson’s Apparent Mistake. This was the case where Watson rationally believed based on his clues that the butler did it, but was misled by Sherlock into thinking his own reasoning was fallacious. Since Watson no longer should reflectively believe that his evidence supports that the butler did it, reflectivists can claim, he no longer can rationally believe it based on his clues.

But the view that reflective beliefs are merely necessary for rational basing still struggles with closely related cases, like:

Watson and Lestrade’s Disagreement: Watson is told by Sherlock that his clues either strongly support the butler’s guilt or strongly support his innocence. In fact it is guilt, and after thinking it over Watson rationally believes both that the butler is guilty and that his own clues support guilt. He then encounters Lestrade, who takes the same clues to support the

butler's innocence. Watson considers Lestrade as reliable as he is, and suspends judgment on what the clues support.

If Watson should suspend in Watson's Apparent Mistake, he plausibly should suspend here. But if so, guidance reflectivism is in trouble. Watson reflectively believes that his clues do not support suspending judgment on whether the butler did it. So it seems guidance reflectivists cannot allow his suspended judgment to be well-founded.

The guidance reflectivist might instead appeal to a distinction between belief and suspended judgment. Maybe a reflective belief that one's evidence supports a belief is necessary for a well-founded belief, even if a reflective belief that it supports suspended judgment is not necessary for well-founded suspended judgment.

But I think this distinction is unmotivated. Suspended judgment can be ill-founded just like belief, when it is insensitive to one's evidence. If reflective beliefs figure into the account of proper basing for beliefs, they should for suspended judgment, too.

But in any case, distinguishing belief from suspended judgment like this will not do the guidance reflectivist any good. Consider:

Watson's Conflicting Evidence: On Monday, Watson believes the butler did it based on clues that inductively support his guilt. On Tuesday, Watson learns new evidence that abductively supports that really the maid did it, and so suspends judgment on whether the butler did it. On Wednesday, Watson is reminded by Sherlock that he has an abysmal track record at abductive reasoning, though he has never had any trouble with inductive reasoning.

Anyone who says Watson ought to suspend in an ordinary higher-order defeat case like Watson's Apparent Mistake should say he ought to believe here. Yet if guidance reflectivists hold that the appropriate reflective belief is only necessary for well-founded belief, and not for well-founded suspended judgment, they cannot explain why. So not only does guidance reflectivism have trouble avoiding the disadvantages of transparency and management reflectivists, it cannot claim their advantages over deflationism in explaining higher-order defeat.

8. Conclusion

It seems that of all the people in the world whose rational requirements I might learn about, I cannot single out for special treatment the one of them who is me. Can this really be right? I might learn about the errors and biases afflicting the reasoning of many people, but only one of them is the same person whose reasoning led to this very belief which I have right now. Can there really be no way for evidence about this agent's reasoning in particular to play a special role in how I manage that reasoning, or whether I trust it?

If there is a way, it is not by following ordinary evidential standards, like those governing deductive, inductive, and abductive inferences. Since these standards are invariant under third-personalization, I can conform to them without making any special use of information about myself. I can be logically consistent and probabilistically coherent, infer the best explanations of my evidence, and so on—all without allowing information about the reasoning I am performing to inform that reasoning in any distinctive way.

Even if conforming to ordinary evidential standards is all it takes to be rational, it might at least seem suboptimal. Wouldn't an ideal agent have some further capacity that allows it to coordinate third-personal information about its reasoning with how it conducts that reasoning? Maybe this is the kernel of insight that motivated reflectivists to propose reflective reasoning as the additional capacity that distinguishes ideal agents from the brutes. Maybe an ideal agent would have such a capacity, and maybe we do, too. But even if all that is so, reflectivism offers a false theory of what that capacity is.

References

- Alston, William (1989) *Epistemic Justification*. Ithaca: Cornell University Press.
- Arpaly, Nomy (2002) *Unprincipled Virtue: An Inquiry Into Moral Agency*. OUP.
- Barnett, David James (2016) 'Inferential Justification and the Transparency of Belief' *Noûs* 50(1): 184-212.
- Bergmann, Michael (2005) 'Defeaters and Higher-Level Requirements' *Philosophical Quarterly* 55(220): 419-436.
- Boghossian, Paul (2014) 'What is Inference?' *Philosophical Studies* 169(1): 1-18.
- Boyle, Matthew (2011) 'Transparent Self-Knowledge' *Aristotelian Society Supplementary Volume* 111(2pt2): 223-241.
- Briggs, Ray (2009) 'Distorted Reflection' *Philosophical Review* 118(1): 58-85.
- Burge, Tyler (2013) *Cognition Through Understanding*. OUP.
- Byrne, Alex (2018) *Transparency and Self-Knowledge*. OUP.
- Chalmers, David and Alan Hájek (2007) 'Ramsey + Moore = God' *Analysis* 67(2): 170-172.
- Christensen, David (2010a) 'Rational Reflection' *Philosophical Perspectives* 24(1): 121-140.
- (2010b) 'Higher-Order Evidence' *Philosophy and Phenomenological Research* 81(1): 185-215.
- (2016) 'Disagreement, Drugs, etc.: from Accuracy to Akrasia' *Episteme* 13(4): 397-422.
- Dutant, Julien and Littlejohn, Clayton (MS) 'Defeaters as Indicators of Ignorance'
- Egan, Andy and Elga, Adam (2005) 'I Can't Believe I'm Stupid' *Philosophical Perspectives* 19(1): 77-93.
- Elga, Adam (2000) 'Self-Locating Belief and the Sleeping Beauty Problem' *Analysis* 60(2).
- (2005) 'On Overrating Oneself... And Knowing It' *Philosophical Studies* 123(1-2): 115-124.
- Fernandez, Jordi (2013) *Transparent Minds: A Study of Self-Knowledge*. OUP.
- Gibbons, John (2006) 'Access Externalism' *Mind* 115(457): 19-39.
- (2013) *The Norm of Belief*. OUP.
- Greco, Daniel (2014) 'A Puzzle About Epistemic Akrasia' *Philosophical Studies* 167(2): 201-219
- Hieronymi, Pamela (2005) 'The Wrong Kind of Reason' *Journal of Philosophy* 102: 437-457.

- (2006) 'Controlling Attitudes' *Pacific Philosophical Quarterly* 87(1): 45-74.
- Horowitz, Sophie (2014) 'Epistemic Akrasia' *Noûs* 48(4): 718-744.
- Kelly, Thomas (2002) 'The Rationality of Belief and Other Propositional Attitudes' *Philosophical Studies* 110(2): 163-196.
- (2010) 'Peer Disagreement and Higher-Order Evidence', in *Social Epistemology: Essential Readings*, ed. Alving Goldman and Denis Whitcomb, OUP.
- Kornblith, Hilary (2012) *On Reflection*. OUP.
- Lasonen-Aarnio, Maria (2014) 'Higher-Order Evidence and the Limits of Defeat' *Philosophy and Phenomenological Research* 88(2): 314-345.
- (2020) 'Enkrasia or Evidentialism? Learning to Love Mismatch' *Philosophical Studies* 177(3): 597-632.
- McGinn, Colin (1982) *The Character of Mind*. OUP.
- Moran, Richard (2001) *Authority and Estrangement: An Essay on Self-Knowledge*. Princeton University Press.
- Neta, Ram (2019) 'The Basing Relation' *Philosophical Review* 128(2): 179-217.
- Pryor, James (2005) 'There is Immediate Justification' in *Contemporary Debates in Epistemology*, ed. Matthias Steup and Ernest Sosa. Blackwell.
- Rinard, Susanna (2017) 'No Exception for Belief' *Philosophy and Phenomenological Research* 94(1): 121-143.
- (2019) 'Equal Treatment for Belief' *Philosophical Studies* 176(7): 1923-1950.
- Rioux, Catherine (MS) 'The Rationality of Practical Hope'
- Rosati, Connie (2016) 'Moral Motivation' *Stanford Encyclopedia of Philosophy*, ed. Edward Zalta. <<https://plato.stanford.edu/archives/win2016/entries/moral-motivation/>>
- Scanlon, Thomas (2007) 'Structural Irrationality' in *Common Minds: Themes from the Philosophy of Philip Pettit*, ed. Geoffrey Brennan, Robert Goodin, Frank Jackson, and Michael Smith. OUP.
- Schoenfield, Miriam (2014) 'Permission to Belief: Why Permissivism Is True and What It Tells Us About Irrelevant Influences on Belief' *Noûs* 48(2): 193-218.
- (2015) 'A Dilemma for Calibrationism' *Philosophy and Phenomenological Research* 91(2): 425-455.
- Shah, Nishi (2006) 'A New Argument for Evidentialism' *Philosophical Quarterly* 56(225): 481-498.

- Shah, Nishi and Velleman, J. David (2005) 'Doxastic Deliberation' *Philosophical Review* 114(4): 497-534.
- Shoemaker, Sydney (1996) *The First-Person Perspective and Other Essays*. Cambridge University Press.
- Smithies, Declan (2019) *The Epistemic Role of Consciousness*. OUP.
- Stroud, Sarah (2006) 'Epistemic Partiality in Friendship' *Ethics* 116(3): 498-524.
- Titelbaum, Michael (2015) 'Rationality's Fixed Point (Or: In Defense of Right Reason)' *Oxford Studies in Epistemology* 253-294.
- Turri, John (2012) 'A Puzzle about Withholding' *Philosophical Quarterly* 62(247): 355-364.
- Way, Jonathan (2016) 'Two Arguments for Evidentialism' *Philosophical Quarterly* 66(265): 805-818.
- Weatherson, Brian (2019) *Normative Externalism*. OUP.
- Wedgwood, Ralph (2002) 'The Aim of Belief' *Philosophical Perspectives* 36(s16): 267-297.
- White, Roger (2005) 'Epistemic Permissiveness' *Philosophical Perspectives* 19(1): 445-459.
- (2006) 'Problems for Dogmatism' *Philosophical Studies* 131(3): 525-557.
- (2013) 'What is My Evidence that Here is a Hand?' in *Scepticism and Perceptual Justification*, ed. Dylan Dodd and Elia Zardini, OUP.
- Whiting, Daniel (2017) 'Against Second-Order Reasons' *Noûs* 51(2): 398-420.
- Williamson, Timothy (2011) 'Improbable Knowing' in *Evidentialism and its Discontents*, ed. T. Dougherty. OUP.
- Worsnip, Alex (2018) 'The Conflict of Evidence and Coherence' *Philosophy and Phenomenological Research* 96(1): 3-44.