

Epistemic Dilemmas Defended*

Nick Hughes
University of Oxford
nickhowellhughes@gmail.com

§I. Introduction

Can there be epistemic dilemmas? I say yes (Hughes 2019, fc1, ms1, ms2). Daniel Greco (this volume) says no.¹ He suggests that the mark of a dilemma is *unavoidable residue*. If you face a dilemma, then no matter what you do you will rightly be the target of blame or resentment, owe an apology or reparations, feel guilt or remorse, or something along these lines. But whilst we can make sense of the idea of unavoidable moral residue – and hence moral dilemmas – there are, Greco argues, principled reasons to think there can never be unavoidable epistemic residue. If so, there cannot be epistemic dilemmas.

* Forthcoming in the Oxford University Press volume *Essays on Epistemic Dilemmas*. The final version might be a bit different. Nevertheless, feel free to cite this version. Got questions or comments? Email me!

¹ Actually, strictly speaking he's neutral-but-sceptical. I'll present him as a straight up 'no' for ease of exposition. Nothing important will turn on the difference.

Greco's argument is unsuccessful. In the first part of this chapter I'll explain why. Nevertheless, his essay poses an important challenge to those of us who maintain that there are epistemic dilemmas. Greco thinks that if we don't identify epistemic dilemmas with unavoidable epistemic residue, disputes between those who argue that there are epistemic dilemmas and those who argue that there aren't may turn out to be merely verbal. This is wrong; they don't. In the second part of the chapter I'll explain why.

I'll start, in §2, by saying a bit about what epistemic dilemmas are and why there can be such things. In §3 I'll present Greco's anti-dilemmas argument. In §4 I'll identify three problems with the argument, which should lead us to reject it. In §5 I'll argue that there could be epistemic dilemmas even if the mark of an epistemic dilemma was unavoidable epistemic residue. In §6 and §7 I'll discuss the matter of verbal disputes. I'll show how diagnostic tools other than residue can be used to adjudicate between dilemmic and non-dilemmic views.

§II. Dilemmism

I call the view that there are epistemic dilemmas 'dilemmism'. An epistemic dilemma is a conflict in epistemic 'oughts'. A paradigmatic example is a situation in which you ought (epistemically) to adopt doxastic attitude D, and at the same time ought (epistemically) not adopt doxastic attitude D. I call these 'conflict cases'. In conflict cases, whatever you do, you'll do something wrong from the epistemic point of view.

Here's an example, which will be my focus in this chapter.² Consider these two norms:

KNOWLEDGE: One ought, epistemically, to believe that p only if one knows that p

RATIONALITY: One ought, epistemically, to be epistemically rational

In some circumstances they issue conflicting instructions. To see this, consider the following case (c.f. Hughes fc1, ms1):

BRUEGHEL: *The Crucifixion*, a painting by Pieter Brueghel the Younger, hangs in a church in a small town in Northern Italy. A gang of thieves intends to steal it. After weeks of planning, late one March night they quietly disable the church alarm system, break in through the apse door, snatch the painting from its frame, and make their escape. Back at the boss's house, they celebrate; they expect to negotiate a large ransom from the government for its return. Meanwhile, the local police are also celebrating. After being tipped off about the thieves' plan, they set up a hidden camera in the church and replaced the painting with an identical-looking replica. Now they can use the camera footage to identify the thieves. The actual Brueghel is sitting in a vault in the basement of the Uffizi.³

² This isn't the only kind of epistemic dilemma. For discussion of others, see Hughes (2019, fc1, ms2).

³ This is based on a true story. Check it out: <https://www.bbc.co.uk/news/world-europe-47568323>

Let 'p' = 'the thieves have a painting by Brueghel'. According to KNOWLEDGE the thieves shouldn't believe that p, because it's false. But it would be irrational for them not to believe it. So according to RATIONALITY they should believe that p. Hence, if KNOWLEDGE and RATIONALITY are both genuine epistemic norms, the thieves should believe that p, and, at the same time, shouldn't believe that p. That is, they face an epistemic dilemma. According to dilemmism, that's just how it is; we shouldn't try to avoid the dilemmic conclusion by denying one of the norms, or by claiming that one is weightier than the other, or anything like that.

Why accept the dilemmic view? Mark Nelson (2020) suggests that dilemmists might be driven by a perverse delight in making philosophical mischief, or perhaps by a kind of Romantic misologism. Now, I enjoy mischief and misologism as much as the next person, but my reasons are altogether less exciting. There are three of them. First, dilemmism allows us to embrace simple and straightforward interpretations of KNOWLEDGE and RATIONALITY. *Prima facie*, conformity with them is non-optional from the point of view of epistemic normativity. Dilemmism accommodates this observation. That is a good reason to accept it. The price to be paid is accepting epistemic dilemmas. It is a price worth paying. Second, as I have argued elsewhere (Hughes 2019, fc1, ms1), objections to the dilemmic view aren't nearly as persuasive as they might first seem. Third, (as we'll see), there is no plausible non-dilemmic account of these norms. So we're stuck with dilemmism whether we like it or not.

§III. Greco's Anti-Dilemmas Argument

§3.1. *Epistemic Blame*

Greco thinks dilemmism is untenable. His anti-dilemmas argument starts with the idea that the mark of a moral dilemma is unavoidable moral residue. According to this view, a situation is a moral dilemma just in case, whatever you do in that situation, you will rightly be a target of blame or resentment, owe an apology or reparations, feel guilt or remorse, or something along these lines. Consider the famous case of Sartre's student. He was torn between staying home with his mother, for whom he was the only source of comfort and solace in life, and leaving home to fight fascism with the Free French. It's not implausible to suppose there would have been some kind of moral residue whatever he did. If he stayed home to look after his mother, he would have, not unreasonably, felt guilty for not joining the struggle. But if he left home to join the struggle, he would have felt guilty for abandoning his mother – once again, his guilt would not have been unreasonable.

If the mark of a moral dilemma is unavoidable moral residue, it's natural to think that the mark of an epistemic dilemma – if there can be such a thing – is unavoidable epistemic residue. But what could this residue be? One possibility is that it is just the same thing as moral residue. If so, an epistemic dilemma is a situation in which guilt, remorse, blame, resentment, apologies, reparation, or some such is unavoidable whatever you believe. As Greco points out, if that's right, then it is implausible to think there are epistemic dilemmas – at least of the kind I think there are. That's not to say that this kind of residue never attaches to beliefs. If I

believe that you committed some heinous crime, purely on the basis of a prejudice against you, and without any evidence, then it would be appropriate for you to feel resentful. I should also feel guilty, and I owe you an apology too. But it is surely a mistake to think that these sorts of reactions are *always* appropriate in response to people believing what they don't know or believing irrationally. They wouldn't be appropriate reactions to have towards the thieves in the BRUEGHEL case when they falsely believe that they have a Brueghel, for instance.

For this reason, Greco argues that dilemmists are better off looking for a distinctive kind of epistemic residue. He proposes that it consists in *epistemic blame*, which he understands as the withdrawal of trust in a person as a reliable source of information:

EPISTEMIC BLAME AS TRUST WITHDRAWAL: To epistemically blame someone is to withdraw trust in them as a reliable source of information.⁴

Greco argues that if epistemic residue is understood along the lines of EPISTEMIC BLAME AS TRUST WITHDRAWAL, it cannot be unavoidable. If the mark of an epistemic dilemma is unavoidable epistemic residue, it follows that there cannot be epistemic dilemmas.

§3.1. *The Reflection Argument*

I'll call Greco's argument the 'Reflection Argument', because it starts from the assumption that the following is a requirement of rationality:

⁴ As Greco emphasises, withdrawal may only be *partial* (you still give their word *some* weight) and *local* (it may only concern certain topics).

REFLECTION: Your present beliefs should be a weighted guess as to what your future beliefs will be.

REFLECTION is motivated by the idea that if you know it's likely that you'll get compelling evidence for p tomorrow, this should be reflected in your beliefs *today* – you should already be confident that p is true.⁵

The Reflection Argument is simple. However, a word of caution: in order to best understand the argument, it will be helpful to put the KNOWLEDGE vs. RATIONALITY conflicts described above out of mind for the time being (I'll have much more to say about them shortly) and just consider *in the abstract* the idea that there could be unavoidable epistemic residue in the form of EPISTEMIC BLAME AS TRUST WITHDRAWAL. That is, consider the general idea that there can be situations in which, whatever someone believes, it will be appropriate for you to withdraw trust in them as a reliable source of information. Done that? Good. Now for the argument.

It goes like this. Suppose that at time t^1 you treat Dylan as a reliable epistemic surrogate, meaning that you're prepared to defer to her judgement on a wide range of questions. At t^1 you believe that p on Dylan's say-so. Let's say that $p =$ 'vitamin C prevents the common cold'. However, you know that at t^2 she will face what you will regard as an epistemic dilemma. As a result, whatever attitude she forms at t^2 you will judge her to be epistemically blameworthy. Given EPISTEMIC BLAME AS TRUST WITHDRAWAL, this means that you know at t^1 that at t^2 you will no

⁵ As Greco notes, the principle may need to be refined to handle tricky cases. For instance, we don't want to say that if you know that you'll believe that p tomorrow because you will have been brainwashed, you should believe that p today. We can ignore such complications here.

longer be prepared to defer to Dylan's judgement about p . But of course you don't blame her at t^1 , as she hasn't encountered the dilemma yet, and so hasn't yet done anything to merit blame. So, you are still prepared to defer to her at t^1 by believing that p on her say-so. But this means that you violate REFLECTION – you believe that p (= vitamin C prevents the common cold) on Dylan's say-so at t^1 whilst at the same time knowing that at t^2 you'll regard yourself as having insufficient evidence to believe this, Dylan having revealed herself to be untrustworthy. So, the existence of epistemic dilemmas is incompatible with REFLECTION. Since REFLECTION is true, it follows that there cannot be epistemic dilemmas.

§IV. Why the Reflection Argument Doesn't Work

Should we conclude on the basis of the Reflection Argument that there cannot be epistemic dilemmas, then? No. There are three serious problems with the argument.

§4.1. Blame and Quasi-Blame

The first problem is this. The argument depends on the assumption that you cannot rationally, or perhaps even literally, "blame" Dylan at t^1 for what she will do at t^2 by withdrawing trust in her at t^1 .⁶ For suppose that you could. In that case, you would not be doomed to violate REFLECTION. Upon coming to know that she will face what you regard as a dilemma at t^2 , you could, rationally, revise your belief that p at t^1 . And if you did, you

⁶ The scare quotes around "blame" are deliberate, for reasons that will soon become clear.

wouldn't violate REFLECTION. That is, you wouldn't be in a state of believing that p even whilst knowing that soon you will regard yourself as having insufficient evidence to believe that p .

But is it true that you cannot, or at least should not, "blame" Dylan at t^1 for what she will do at t^2 by withdrawing trust in her at t^1 ? I don't think so. What might well be true is that you cannot blame someone for ϕ -ing before they have ϕ -ed in the ordinary sense of 'blame'. If I know that you'll betray me next year, then whilst I might feel angry and resentful about the fact that you *will* betray me, it is tempting to say that I cannot blame you *for* betraying me until you've actually done the deed. Or, even if I can, I shouldn't – preemptive punishment is wrong. But the kind of "blame" that Greco's argument makes use of quite clearly isn't blame in the ordinary sense. One can withdraw epistemic trust in a person without any of the characteristic properties of blame (anger, resentment, etc.) being present. Here's an example. Suppose your kindly grandfather develops dementia. You begin to notice a decline in his cognitive performance – he starts saying strange things for which he has no evidence. You downgrade epistemic trust in him as a result. But do you *blame* him? Certainly not, at least in the ordinary sense of blame. You harbour no anger or resentment towards him – only pity.

This suggests that epistemic blame, as Greco defines it in EPISTEMIC BLAME AS TRUST WITHDRAWAL, isn't really a kind of *blame* at all. Or, if it is, it is atypical. So, in order to avoid confusion, it will be useful to give it a different name. Let's call it 'quasi-blame'.

It might not be possible or rational to preemptively blame someone at t^1 for what they haven't yet done but will do at t^2 . But is it possible to preemptively quasi-blame someone at t^1 for what you know they will do

at t^2 by withdrawing trust from them at t^1 ? And if it is possible, can it also be rational? The answer to both of these questions is surely ‘yes’. To see why, let’s put dilemmas aside for a moment. Imagine that an episode of InfoWars has just been broadcast, propagating spurious evidence-free conspiratorial nonsense as usual.⁷ Dylan hasn’t seen it yet. But you know that she will, and that when she does, she’ll be taken in by it. Given this, it would be rational for you to withdraw epistemic trust in her now. There’s no need to wait, because you already know everything you need to know about her epistemic incompetence. This shows that preemptive quasi-blame is both possible and sometimes rational.

Now, back to dilemmas. You know at t^1 that Dylan will face what you regard as an epistemic dilemma at t^2 , and so no matter what she does you will withdraw epistemic trust in her. What is the rational thing for you to do now, at t^1 ? Well, given that preemptive quasi-blame can be rational, it might well be to withdraw trust in her *now*. Nothing Greco has argued shows us otherwise.⁸ But if you do withdraw trust in her now, you won’t

⁷ For readers fortunate enough not to be familiar with InfoWars, it is a far-right US website and “news” show which promotes New World Order, 9/11, chemtrail, QAnon, and false flag, conspiracy theories.

⁸ Curiously, Greco seems to agree, but misses the significance of the point. He writes “...”punishment” for epistemic wrongdoing plausibly differs from punishment for moral wrongdoing. While it’s a matter of some controversy, it’s commonly held that it is immoral to blame or punish people for crimes that they haven’t in fact committed, even when they almost certainly will commit some crime, or we know that they would commit such a crime in the right circumstances. More generally, the idea seems to be that bad (moral) character isn’t enough—bad moral character must be manifested in some bad act for certain negative reactive attitudes to be warranted. But if epistemic “punishment” is just the withdrawal of deference, such a restriction isn’t plausible. Bad “epistemic character” is enough to warrant epistemic “punishment”, even in the absence of some manifestation of that bad character in a cognitive act of (irrational) belief. If I know you’re the sort of person who would believe in homeopathy, were you to consider the question, then I know you have poor judgment, and this provides me the same sort of reason to cease relying on your concerning medical matters

violate REFLECTION, as you won't be in a state believing that p even whilst knowing that soon you'll regard yourself as having insufficient evidence to believe that p . This shows that there is no tension between REFLECTION, on the one hand, and the idea that epistemic dilemmas should be thought of along the lines of unavoidable trust withdrawal, on the other. But if not, then the Reflection argument doesn't work.

§4.2. Overgeneration

The argument has another serious problem.

Epistemic dilemmas are situations in which one is subject to conflicting epistemic 'oughts', such that if one satisfies the first, one will violate the other, and vice versa, with the result that no matter what one does, one is guaranteed to do something wrong from the epistemic point of view. In order for the Reflection Argument to be effective against this view, we have to assume that a subject, S , has violated an epistemic ought in ϕ -ing only if it is appropriate to quasi-blame S for ϕ -ing. If we don't assume this, then the argument could not show that there cannot be epistemic dilemmas.

Now, consider the Dylan case again, but this time replace 'faces an epistemic dilemma' with 'violates epistemic norm N ', where N is some candidate epistemic norm purporting to give rise to epistemic requirements, like, say, 'believe what your evidence supports'. So, at t^1 you know that at t^2 Dylan will violate N , and you'll quasi-blame her as a result

as I'd have if I knew that you had already considered the question and adopted a belief in homeopathy".

by withdrawing trust. On the assumption that it's not rational (and perhaps not even possible) to quasi-blame someone for φ -ing before they have φ -ed – an assumption we have already seen is essential to the Reflection Argument – we can now derive a violation of REFLECTION: you believe that p on Dylan's say-so at t^1 even whilst knowing that you won't believe that p at t^2 because Dylan will have shown herself to be untrustworthy. Following Greco's reasoning, we can now conclude that N isn't a genuine epistemic norm. But N could have been pretty much any putative epistemic norm. It could have been 'believe what your evidence supports'. It could even have been 'don't believe in a way that makes you epistemically untrustworthy'. So, if we accept the Reflection Argument, we will have to conclude that there aren't any epistemic norms at all. Hence, the argument overgenerates disastrously.

§4.3. Inscrutable Cognizers

The third problem with the argument is what I'll call 'the problem of inscrutable cognizers'. Focus again on the assumption that S has violated an epistemic ought in φ -ing only if it is appropriate to epistemically quasi-blame S for φ -ing by withdrawing trust. Assuming that these oughts are generated by epistemic norms, it follows that only creatures we are capable of trusting as reliable sources of information can be subject to epistemic norms. For if a creature cannot be trusted as a reliable source of information in the first place, trust cannot be withdrawn – you can't take back what was never given. But we stand, and are only able to stand, in this kind of trusting relationship with a rather limited range of creatures. Our fellow humans, most obviously, and perhaps a few non-human animals such as guide dogs. But what about, say, octopuses and lobsters? Both are highly intelligent species. But it stretches credulity to suppose that we could rely on creatures so utterly alien to ourselves as reliable

sources of information. *Ipsa facto*, it is implausible to suppose that we could ever withdraw epistemic trust from them. It follows, given the assumption under consideration, that epistemic norms do not apply to them.

Should we accept this result? Many epistemologists will be happy to. It is often said that epistemic norms only apply to humans.⁹ That's wrong. There's no denying that some epistemic norms are distinctly human. 'Don't trust InfoWars' – there wouldn't be much point in trying to understand octopus cognition by apply *that* norm. But the most basic and fundamental epistemic norms apply to virtually *all* cognizers, human or otherwise. KNOWLEDGE is a case in point. All animals, including octopuses and lobsters, need to know what's happening in their environment in order to navigate it successfully.¹⁰ RATIONALITY is another example. Plausibly, epistemic rationality is a matter of exercising cognitive dispositions that are conducive to knowing.¹¹ The possession and exercise of such dispositions is just as important for an octopus or a lobster as it is for us.¹²

If epistemic norms apply to non-human animals, why do so many epistemologists think otherwise? Here's what I think is going on. Many of our everyday epistemic evaluations are infused with Strawsonian reactive

⁹ See, for instance, Alston (1988) and Brandom (1994). The idea that epistemic norms only apply to humans seems to be implicit in the work of many contemporary epistemologists insofar as they propose norms that could not possibly apply to most non-human animals.

¹⁰ C.f. Kornblith (2002).

¹¹ The pioneer of this way of thinking about epistemic rationality is Maria Lasonen-Aarnio (2010, 2020a, 2020b, fc1, fc2). See also Williamson (fc1, fc2) and Hughes (fc2, ms1).

¹² Of course, we may not feel comfortable using the words 'rational' or 'irrational' when talking about octopuses and lobsters. This is revealing, and I will have more to say about it momentarily. In the meantime, we can simply note that nothing forces us to use these words – they could be replaced by technical terms that don't have the same connotations. The substantive point would still stand.

attitudes – attitudes that it only makes sense to adopt towards people. The language of normative epistemology reflects this. One often hears talk (including in this chapter) of people being epistemically ‘responsible’ or ‘irresponsible’, ‘justified’ or ‘unjustified’, ‘reasonable’ or ‘unreasonable’, ‘rational’ or ‘irrational’. There is discussion of epistemic ‘accountability’ and epistemically ‘praiseworthy’ and ‘blameworthy’ behaviour. We talk of epistemic ‘virtues’, ‘duties’, ‘obligations’, ‘permissions’, ‘prohibitions’, and ‘requirements’, and of ‘excuses’ and ‘exemptions’. What is striking about these concepts is that very few of them, if any, can be comfortably applied to most non-human animals. It just sounds *weird* to say of an octopus or a lobster (or even a dog) that it is being unreasonable, or that it has a duty to be rational. To the extent that such language can be used, it has a metaphorical feel to it, or it is as if we are generously but temporarily (and not entirely seriously) extending the boundaries of personhood. Why does this way of talking sound strange? The obvious explanation is that we cannot enter into the kinds of sophisticated interpersonal relationships with these animals that the correct application of these concepts presupposes.

Now, it’s hardly surprising that many of our epistemic evaluations take this form. Given that we depend on each other as sources of knowledge, it is to be expected that a social practice of praise and criticism will have developed with which we can keep one another in line, epistemically speaking. And a practice that draws on the full range of reactive attitudes is much more likely to be effective than one that does not. You’ll do a better job of changing someone’s mind by treating them as a person rather than an object. Nor is the practice uninteresting for normative epistemology. However, it does bring with it a risk. We are liable to forget (or simply fail to notice in the first place) what the ultimate purpose of this practice is: to bring about conformity with basic norms of cognition, like KNOWLEDGE

and RATIONALITY. When this is overlooked, it is easy to confuse the social practice itself with the norms that it enforces, in effect treating the social practice as the bedrock of epistemic normativity. And since most non-human animals cannot enter into the social practice, it is natural to conclude that they cannot be subject to epistemic norms. But just because they cannot enter into the social practice, that doesn't mean that the basic norms the social practice emerged to enforce don't apply to them. Once we recognise this, it becomes clear that the assumption that S has violated an epistemic norm in ϕ -ing only if it is appropriate to epistemically quasi-blame S for ϕ -ing by withdrawing trust simply cannot be right. And if not, the Reflection Argument should be rejected, as it relies on an flawed diagnostic tool for identifying epistemic norms.

I want to be clear on how minimal a claim I need here. One might object that we don't know enough about octopuses and lobsters to be confident that KNOWLEDGE and RATIONALITY apply to them. Perhaps the fundamental epistemic norms that govern their cognition are different to those that govern ours? This is highly doubtful – we should not underestimate how fundamental and widespread the need for knowledge is.¹³ However, even if we were to accept the objection, it wouldn't affect my point. We only need *one* case of an inscrutable cognizer to whom KNOWLEDGE and RATIONALITY apply to make the point, and that case only needs to be *possible* – it need not be actual. So, if actual non-human animals don't persuade you, consider instead a fictional cognizer – I'll call it 'The Inscrutable Cognizer' – which, *by stipulation*, thinks in such a way that

¹³ As Kornblith (2002) emphasises, it is impossible to understand the behaviour of nonhuman animals without attributing knowledge to them. None of this is to deny that many of the epistemic norms that apply to octopus and lobster cognition will be radically different to norms that apply to human cognition. But KNOWLEDGE and RATIONALITY will still be there, grounding the rest.

KNOWLEDGE and RATIONALITY apply to it, but which is inscrutable in a way that makes it impossible for us to place trust in it as a source of information, and hence impossible for us to withdraw trust in it. The Inscrutable Cognizer is surely possible.¹⁴ This fact alone shows that we should not try to identify epistemic norms by the method Greco proposes.

The Inscrutable Cognizer is a powerful philosophical device. Another objection might appeal to the idea that even if we can't trust, and hence withdraw trust, in octopuses, lobsters, and the like, their conspecifics can. If so, then it might be argued that a thinker has violated epistemic norms only when it is appropriate for conspecifics to withdraw trust in them as reliable sources of information. However, we can imagine that The Inscrutable Cognizer is an extremely asocial creature which cannot stand in trust relations even with its own kind, nor with any other species. KNOWLEDGE and RATIONALITY apply to it all the same.

§V. Beyond the Reflection Argument

The Reflection Argument doesn't work, then. And the problem of inscrutable cognizers shows that we shouldn't try to identify epistemic

¹⁴ The 2016 film *Arrival* provides a vivid example. In the film mysterious extraterrestrial spacecraft arrive on Earth, containing Cephalopod-like creatures called 'heptapods'. Heptapods are obviously highly intelligent – they have mastered interstellar travel, after all – but virtually impossible to communicate with, owing to their morphology and radically alien form of language. Much of the plot of the film is taken up with scientist's struggles interpret them. A Hollywood ending requires, of course, that the communication barrier is eventually overcome, so the example isn't perfect, but it is easy to imagine a version of the film in which it is insurmountable. The film is based on the SF writer Ted Chiang's short story 'Story of Your Life'. The film is quite good. Chiang's story is sentimental rubbish.

norms by looking for residue, at least not if residue is understood in terms of EPISTEMIC BLAME AS TRUST WITHDRAWAL.¹⁵ Nevertheless, one might still be curious if there could be epistemic dilemmas *if* we take the mark of a dilemma to be unavoidable residue and residue is understood in terms of trust withdrawal. The answer is ‘yes’. Here’s an example.

§5.1. *Cognitive Penetration*

Jamie thinks he’s God’s gift to women. One of the many unfortunate consequences of this is that his perceptual experiences are cognitively penetrated by his inflated sense of his own attractiveness. In particular, when he sees women on the street, he often sees their facial expressions and gestures as flirtatious when in fact they are neutral. Once we learn this about Jamie, we will have reason to withdraw epistemic trust in him, at least when it comes to questions about who fancies him. But imagine that instead he were to simply ignore these perceptual experiences, even though he is completely unaware that his faculties have gone awry and that the experiences misrepresent reality. Jamie would be disregarding perceptual experiences without being able to cite any reason to do so. This would be baffling behaviour, behaviour that doesn’t speak well of his

¹⁵ It might be thought that we should instead try to understand epistemic norms and epistemic dilemmas in terms of some other kind of epistemic residue. But consider the kinds of residue that Greco mentions: blame, resentment, apologies, reparations, guilt, and remorse. All involve reactive attitudes that it is only appropriate to adopt towards persons. This suggests that residue, at least as Greco thinks of it, can only attach to the actions of persons. If so, trying to identify epistemic norms by looking for residue will be a futile endeavour, since, as we have seen, epistemic norms do not only apply to persons. The alternative is to shift to a more encompassing conception of residue. But if we do so, then for all we have seen unavoidable epistemic residue (and hence epistemic dilemmas) may well be possible. In any case, I will shortly argue that we have good reasons to posit epistemic norms even when violations of them leave no residue at all.

cognitive dispositions. For if he is disposed to ignore these perceptual experiences, then he will also be disposed to ignore perceptual experiences that are not the result of cognitive penetration. This means he will not be a reliable source of information about what's going on around him. As a result, we will have reason to withdraw trust in him once again. So, whatever Jamie does, we will have reason to withdraw epistemic trust from him. According to Greco's criteria, this makes the case an epistemic dilemma.¹⁶

§5.2. *KNOWLEDGE/RATIONALITY Conflicts*

What about conflicts between KNOWLEDGE and RATIONALITY? Is trust withdrawal appropriate whatever the agent does in these cases? Greco thinks not. I agree. Consider the BRUEGHEL case. If one of the thieves were not to believe that they have a painting by Brueghel, thereby thinking irrationally, we would have reason to withdraw epistemic trust from him, just as we would with Jamie were he not to believe his eyes. However, if he believes that they have a Brueghel, then intuitively this should not lead us to withdraw trust, even though his belief is false. Why is this? The most likely explanation, it seems to me, is that our willingness to continue trusting or to withdraw trust when a mistake has been made is sensitive to facts about what explains the mistake. The thief's mistake is best explained by the fact that he is in abnormal circumstances for belief formation. Had things been normal, he would have known that they have a Brueghel. The normality or abnormality of the circumstances in which a mistake was made is something we take into account when deciding whether to withdraw trust or not. And with good reason, it would be

¹⁶ For more discussion of these kinds of cases in relation to epistemic dilemmas, see Hughes (ms2).

unrealistic and impractical to demand infallibility from our sources; so long as they are reliable in normal circumstances, that's good enough.

This suggests that epistemic residue, understood as trust withdrawal, attaches to violations of RATIONALITY but not to violations of KNOWLEDGE that aren't also violations of RATIONALITY. An obvious conclusion to draw from this – one that Greco seems to be keen to draw – is that KNOWLEDGE isn't a genuine epistemic norm. Only RATIONALITY is. If so, there are no KNOWLEDGE/RATIONALITY dilemmas.

Let's call this view RATIONALITY ONLY. Is it tenable? Greco suggests that dilemmists who reject it, arguing that KNOWLEDGE is a genuine requirement-generating norm too, may only have a verbal disagreement with those who hold this view. Both parties seem to agree on all *practical* matters, such as when epistemic blame in the form of trust withdrawal is appropriate. If so, what more is there to discuss?

Quite a lot, as it turns out. There are other ways of identifying epistemic norms than just by looking at the residue violations of them leave. In the next section I'll argue that they show the RATIONALITY ONLY view to be untenable. In §7 I'll argue that no other non-dilemmic take on conflict cases works either.

§VI. RATIONALITY Without KNOWLEDGE?

§6.1. The Explanatory Problem

As it stands, the RATIONALITY ONLY view faces a serious explanatory problem.

If epistemic rationality is a matter of exercising cognitive dispositions that are conducive to knowledge, then the thief's belief is rational because it is the manifestation of a knowledge-conducive disposition; in normal circumstances he would have known that they have a Breughel. Likewise, irrational beliefs are irrational because they are the product of cognitive dispositions that are not knowledge-conducive. The beliefs of incorrigible conspiracy theorists are irrational because they are the product of cognitive dispositions that fail to produce knowledge even in normal circumstances. Knowledge and rationality are inextricably linked.

Now, according to the RATIONALITY ONLY view it is simply not true that one ought to believe only what one knows. But if not – if there's nothing wrong with non-knowledgeable beliefs – then two questions immediately arise:

1. If there's nothing wrong with believing without knowing, then why is a belief rational only if it is the manifestation of a disposition conducive to knowing?

2. If a belief is rational only if it is the manifestation of a disposition conducive to knowing, but it's fine to believe without knowing, then why be rational?

Dilemmism can answer these questions. On (1): If KNOWLEDGE is a norm, then the existence of a derivative norm according to which one should manifest knowledge-conducive dispositions is exactly what we would expect, since a norm on action typically generates a norm to the effect that one should be the kind of person who complies with the first norm.¹⁷ On (2): If you should be the kind of person who complies with KNOWLEDGE, you should form beliefs in a way that it conducive to complying with it – that is, you should be rational.

As things stand, the RATIONALITY ONLY view cannot answer either of these questions. This is a serious explanatory shortcoming. If adequate answers cannot be provided, then we should reject RATIONALITY ONLY.

§6.2. A Solution to the Explanatory Problem?

Can proponents of RATIONALITY ONLY solve the explanatory problem? It's clear how they should approach it. The trick will be to explain the epistemological significance of knowledge in a way that answers our two questions but doesn't entail that one ought to believe only what one knows.

The most promising way to try to do this is to look for a positive normative or evaluative status that knowledge has which isn't itself a requirement-generating norm of belief. This is what Mona Simion, Christoph Kelp, and

¹⁷ See Williamson (fc1) for discussion.

Harmen Ghijsen (2016, hereafter 'SKG') do. Following Conor McHugh (2012), they draw a distinction between *prescriptive norms* and *evaluative norms*. Prescriptive norms are about what one ought to do. Examples include 'don't steal', 'don't lie', and 'drive under 70mph on the motorway'. Their primary social function, SKG argue, is to reinforce certain forms of conduct and to discourage others. The function of the prescriptive norm 'drive under 70mph on the motorway', for example, is to discourage dangerous driving. Evaluative norms are different. They don't tell us what to do. Rather, they concern what it takes for a token of a particular type to be good or bad with regards to its type. Examples SKG give include 'a good hospital is a clean hospital', 'a good knife is a sharp knife', and 'good driving is safe driving'.

Evaluative norms don't necessarily entail prescriptive norms. A good summer's day is a sunny summer's day. Thus, an evaluative norm applies to summer's days. But it would be a mistake to think this entails a prescriptive norm according to which someone ought to make a summer's day sunny. Nor do prescriptive norms necessarily entail evaluative norms. To take an example from SKG, suppose your boss has imposed a pointless rule for making coffee at the office, according to which you must scoop the coffee powder with the red teaspoon and the blue teaspoon alternately. If you only use the blue teaspoon, then you've violated a prescriptive norm. But that doesn't mean you've violated an evaluative norm too. You might well still be making good coffee and making it in a good way.

Nevertheless, SKG observe, there is often a relationship between evaluative and prescriptive norms. In particular, we can often *derive* prescriptive norms from evaluative norms. If a good hospital is a clean hospital, then it makes sense that there should be prescriptive norms in

place which encourage hospital cleanliness and discourage hospital dirtiness.

SKG's idea is that the evaluative/prescriptive distinction helps us to understand the relationship between knowledge and rationality. They maintain that there is an evaluative norm according to which a good belief is a knowledgeable belief, but no corresponding prescriptive norm according to which one ought to believe only what one knows. So they think we should replace KNOWLEDGE with KNOWLEDGE*:

KNOWLEDGE*: One's belief that p is a good belief only if one knows that p.

By contrast, they maintain that RATIONALITY is a prescriptive norm:

RATIONALITY: One ought, epistemically, to be epistemically rational

This is because RATIONALITY fulfils the social function of a prescriptive norm: it reinforces behaviour conducive to fulfilling the evaluative norm KNOWLEDGE*.

This view does not commit us to epistemic dilemmas arising from conflicts between KNOWLEDGE and RATIONALITY. It says there is only one thing the thief ought to do in the BRUEGHEL case: believe that p. Moreover, it appears to solve the explanatory problem. SKG can argue that the reason a belief is rational only if it is the manifestation of a disposition conducive to having knowledgeable beliefs is because the function of the RATIONALITY norm is to reinforce behaviour that is conducive to knowing. And the

reason one ought to be rational is because doing so is the best way of getting a good belief – that is, a belief that is knowledge.

§6.3. *Problems with this View*

There are several problems with this view.

First, whilst it's certainly true that a useful distinction can be drawn between evaluative and prescriptive norms, and that not all evaluative norms generate prescriptive norms, there is a large class of cases where they do and where the prescriptive norm has the *same content* as the evaluative norm. KNOWLEDGE is one of these cases. Let me explain.

The reason the evaluative norm 'a good summer's day is a sunny summer's day' doesn't generate a prescriptive norm is because there is nothing any agent can do to prevent violations of the norm. What it evaluates isn't within the sphere of human action.¹⁸ By contrast, things are quite different in situations where an agent (or group of agents) has the power to ensure that an evaluative norm is complied with. Consider the evaluative norm 'a good knife is a sharp knife'. From this we can derive a prescriptive norm for knife-makers: they ought to make sharp knives. Consider the evaluative norm 'a good climate is a habitable climate'. Again, one can derive a prescriptive norm from it: we (i.e. humanity) ought to keep the climate habitable. Or consider 'a good hospital is a clean hospital'. From this we can infer that the hospital staff ought to make the

¹⁸ Although one might think that even this is up for grabs. Consider the conditional norm: 'You ought: if you can make a summer's day sunny, do so'. (Thanks to Daniel Kodsí for bringing this up).

hospital clean. Or consider ‘good driving is safe driving’. From this we can infer that one ought to drive safely.

In general, where N is an evaluative norm, and it is within one’s power to ensure that N is satisfied, then one ought to ensure that it *is* satisfied. Call this the ‘good-to-ought’ principle.¹⁹ Applying it to normative epistemology, it follows that if a belief that p is a good belief only if one knows that p, then one ought to believe that p only if one knows that p. In other words, KNOWLEDGE is a genuine prescriptive norm, just as dilemmism tells us.

Anti-dilemmists like SKG have to deny this, of course. But these inferences all *look* reasonable, so if we’re going to reject them, we will want good grounds for doing so. What could they be? The obvious answer is that these allegedly prescriptive norms don’t fulfil the social function of prescriptive norms, which is to encourage and reinforce certain forms of behaviour that are conducive to complying with evaluative norms. But, *prima facie*, this is an implausible claim. After all, what better way is there to get a sharp knife, a clean hospital, or a knowledgeable belief than to tell the relevant person that they should make the knife sharp, the hospital clean, or the belief knowledge? Moreover, actual practice, at least in my

¹⁹ One conclusion we shouldn’t draw from the good-to-ought principle is that, e.g., a knife-maker ought to make as many sharp knives as possible. Rather, the relevant prescriptive norm for the knife-maker is ‘you ought: if you make a knife, make a sharp knife’. I bring this up because McHugh (2012) appeals to this kind of dubious conclusion to reject the principle. The fact that the Louvre is a good museum, he points out, doesn’t entail that, before it existed, someone was required to create it – not even the person responsible for Parisian museums. This is true, but it’s irrelevant. The relevant prescriptive norm is not ‘One ought to create the Louvre’, but rather ‘one ought: if one is creating a museum, create a good museum’. (It is not ‘one ought: if one is creating a museum, create the Louvre’, because creating the Louvre is not the only way of creating a good museum).

community, suggests that KNOWLEDGE does in fact serve the function of encouraging and reinforcing behaviour. We frequently rebuke others by pointing out that they don't *know* that p with the hope of getting them to not say or believe what they do not know. Indeed, this kind of criticism seems to be just as common, and perhaps more common, than criticism that appeals to irrationality, unreasonableness, and the like.

Nevertheless, there may be something to SKG's idea. Consider our thief in the BRUEGHEL case. Were he to not believe that p, he would avoid having a non-knowledgeable belief on this occasion, but the disposition that he manifests would be a bad one, since it will usually result in him not acquiring knowledge when it is available. Hence, his behaviour is the sort of thing that should be discouraged, given KNOWLEDGE*. By contrast, when he believes that p, he believes without knowing, and so has a bad belief on this occasion, but he manifests a disposition that usually results in knowledge, and so his behaviour should be encouraged. For this reason, a norm like RATIONALITY, which encourages and reinforces knowledge-conducive *dispositions* makes more sense given the functional role of prescriptive norms than does a norm like KNOWLEDGE.

However – and this brings us to the second problem – whilst prescriptive norms *can* have social functions, it is not a necessary condition on the *existence* of such norms that they have social functions. Like Greco's proposal, this idea faces the problem of inscrutable cognizers. Let me explain.

We are thinking of rationality as a matter of exercising cognitive dispositions that are conducive to knowing. As I emphasized earlier, the possession and exercise of such dispositions is just as important for non-human animals like octopuses, lobsters, and other inscrutable cognizers,

as it is for us. But obviously there is no chance whatsoever of encouraging or reinforcing knowledge-conducive behaviour in an inscrutable cognizer like an octopus or a lobster by issuing prescriptions to them. The very idea is comical. Try telling a lobster it's being irrational and see how far you get. There are two conclusions we might draw from this. One is that it isn't true that inscrutable cognizers ought to have and exercise knowledge-conducive dispositions.²⁰ That is implausible. The other is that the existence of a prescriptive norm does not depend on the norm's capacity to encourage and reinforce behaviour. This is the right conclusion to draw. But if so, then the fact that KNOWLEDGE does not encourage and reinforce forms of behaviour that are conducive to knowing as well as RATIONALITY does is no reason to think it isn't a prescriptive norm. This is not to say that we cannot make *use* of prescriptive norms to encourage and reinforce certain kinds of behaviour. We can. But it is not a precondition on something *being* a prescriptive norm that we can make use of it in this way.

Thinking about inscrutable cognizers makes this point vivid. But strictly speaking, it isn't necessary. We can stay closer to home. Imagine a community of people are exceptionally epistemically stubborn – they refuse to (or perhaps simply can't) pay heed to one another's appraisals of their epistemic behaviour. In this community – let's call them 'stubborners' – the practice of judging beliefs to be rational or irrational makes no difference whatsoever to their behaviour (and presumably will soon die out). According to SKG's view, RATIONALITY doesn't apply to stubborners as a prescriptive norm, because it fails to fulfil the function of such norms: it doesn't encourage and reinforce certain kinds of

²⁰ The thought might be that only an evaluative version of RATIONALITY applies to them. But notice that, if so, epistemic dilemmas reemerge, this time as conflicts between evaluative norms.

behaviours. Yet it surely remains the case that RATIONALITY is an epistemic norm that applies to stubborners.

The third problem with SKG's proposal is that it is quite unclear why the reasoning they employ should lead us to accept their package, according to which KNOWLEDGE* is an evaluative norm and RATIONALITY is a prescriptive norm, rather than some other package. Good action is successful action – action that achieves its purpose. We can think of this as an evaluative norm. Acting successfully usually requires acting on knowledge.²¹ So we should expect a norm like KNOWLEDGE to exist, the function of which is to reinforce behaviour conducive to satisfying the evaluative norm of successful action. By SKG's reasoning, this makes KNOWLEDGE a *prescriptive* norm rather than an evaluative norm. Why prefer SKG's view to this one? I can't see any reason to.

The same problem arises in the opposite direction: we can use SKG's reasoning to conclude that RATIONALITY should be recast as an evaluative norm rather than a prescriptive norm. Being rational isn't easy. Empirical work in cognitive and social psychology suggests that we are prone to a range of cognitive biases. Beliefs formed as a result of these biases are often

²¹ Of course, it's possible to act successfully whilst acting on non-knowledgeable beliefs, but in such cases, luck typically provides a helping hand. To see this, consider a Gettier case (based on Chisholm 1966). Robin is a farmer. She needs to get the sheep back in their pen before nightfall, so the wolves don't get them. She sees what looks exactly like a sheep in a field over the way, but which is in fact a dog which has been dressed up to look like a sheep by a local prankster. As a result of her perceptual experience, Robin forms the belief that there is a sheep in the field, and starts to walk over with the intention of bringing it in. Luckily for her, there is a sheep in the field, hidden behind a tree, so her belief is true but not knowledge. Her journey will not be wasted, but this is only because of the lucky coincidence. Had there not been a hidden sheep, she would have traipsed over there for nothing. By contrast, had she known that there was a sheep in the field, her success would not have been dependent on good luck.

the product of cognitive dispositions that are not conducive to knowledge.²² According to the theory of rationality we are working with, this means they are not rational. In situations where being rational is hard, prescriptive norms can emerge which serve to encourage compliance with RATIONALITY. An example is anti-bias training in the workplace, which attempts to ameliorate implicit bias.²³ Should we conclude that RATIONALITY is best recast as an evaluative norm, and that the prescriptive norms are things like ‘undergo anti-bias training’? Once again, it is hard to see what rules out this way of approaching the issue.

SKG do not provide us with a satisfactory theory of the relationship between knowledge and rationality. *Ipsa facto* they do not provide us with a satisfactory solution to the explanatory problem. So we should prefer dilemmism, which does not face the problem.

§VII. Other Approaches

Things aren’t looking good for non-dilemmic approaches to conflicts between KNOWLEDGE and RATIONALITY, nor for the claim that disputes between dilemmists and others are merely verbal. But we haven’t exhausted the options yet. In the rest of the chapter, I’ll look at some of them. I’ll argue that none of the possible alternatives to dilemmism is plausible.

²² See Hughes (fc2)

²³ Although only with limited success, according to the Equality and Human Rights Commission (Report 113).

§7.1. *Permissivism*

One common reaction to the claim that cases like that of Sartre's student are moral dilemmas is to argue that the agent doesn't face two conflicting requirements – to stay home and to go and fight, for example – but rather a single, non-conflicting, disjunctive requirement: to *either* stay home *or* go and fight.²⁴ If epistemic dilemmas are analogous to moral dilemmas, then this response might also look attractive when it comes to conflicts between KNOWLEDGE and RATIONALITY. The resulting view would be a kind of radical permissivism. It would say that our thief ought to either believe that they have a painting by Brueghel *or* not believe it. So long as he does one or the other, he's doing fine.

Whatever its merits when it comes to moral dilemmas, radical permissivism has little to recommend it as an approach to conflicts between KNOWLEDGE and RATIONALITY. Because it says there is nothing wrong with believing without knowing, it too faces the explanatory problem. As far as I can tell, it provides no solution. It also fails to accommodate the datum that each of KNOWLEDGE and RATIONALITY seems to be distinctly *non-optional*. For these reasons, it should be rejected.

§7.2. *The Weighing Approach*

A different approach holds that each of KNOWLEDGE and RATIONALITY generates genuine epistemic requirements, but one of these requirements is weightier than the other, and so takes precedence when the two conflict. This approach holds the promise of offering us a univocal, non-dilemmic, answer to what our thief should believe: either he should believe that p,

²⁴ Donagan (1984), Zimmerman (1996).

because RATIONALITY outweighs KNOWLEDGE, or he should not believe that p, because KNOWLEDGE outweighs RATIONALITY. And if he should believe that p, it isn't also the case that he shouldn't believe that p, and vice versa.

The idea that one can face two conflicting requirements, one of which outweighs the other, has some appeal in certain cases of moral conflict. If I've promised to meet a friend for lunch, but on my way over I come across someone having a severe asthma attack and I'm the only one around to help, then whilst I'll violate an obligation whatever I do, it's clear what I should do – break my promise and help the person.

But how should we apply the idea to conflicts between KNOWLEDGE and RATIONALITY? Difficulties arise immediately. On the one hand, it seems plausible that RATIONALITY cannot outweigh KNOWLEDGE, since it *depends* on KNOWLEDGE for its normative force in the first place. On the other hand, we tend to judge violations of RATIONALITY more harshly than we do violations of KNOWLEDGE that aren't also violations of RATIONALITY. If our thief mistakenly believes that they have a Brueghel, it's understandable – anyone would have done the same in his situation. By contrast, if he doesn't believe that they have a Brueghel, thereby complying with KNOWLEDGE but violating RATIONALITY, we'll usually be much less sympathetic – the man's an idiot. And this seems to suggest that RATIONALITY should take precedence over KNOWLEDGE.

The problem is that the badness of violating RATIONALITY is altogether *different* to the badness of violating KNOWLEDGE. Violations of the latter don't necessarily reflect poorly on the cognitive dispositions of the agent. Rather, they focus on the bad results their cognition produced on this occasion. By contrast, violations of the former don't necessarily tell us anything about how the agent's cognition fared on this occasion, but they

do reflect poorly on the agent's cognitive dispositions. There are two kinds of rights and wrongs here, not one, and they cannot be judged along the same, or at least sufficiently similar, dimensions of evaluation in a way that makes meaningful weighing possible. This incommensurability makes it impossible to apply the weighing approach.

§7.3. *The Excuses Approach*

Another approach is to accept KNOWLEDGE but reject RATIONALITY, arguing that those who accept RATIONALITY have confused excusable violations of one norm with the existence of another norm.²⁵ Since an agent who complies with RATIONALITY does what any reasonable person trying to comply with KNOWLEDGE would do in their situation, they should not be blamed or criticised for believing that p. But, the argument goes, this only shows that they have a good excuse for believing that p, not that they ought to have believed that p. Excuses mitigate wrongdoing. They don't make wrongdoing rightdoing.

If the claim here was simply that the fact that a belief is rational makes it an excusable violation of KNOWLEDGE, then I would agree. This is not in tension with the idea that RATIONALITY is a requirement-generating norm. Hence, it is compatible with the dilemmic view. But a stronger claim is being made: that rational beliefs are not merely excusable violations of KNOWLEDGE, but that there is no sense in which one *ought to have* rational beliefs – that is, RATIONALITY isn't a genuine epistemic norm. And this, it seems to me, is seriously misguided. If it were correct, we should not think anything is going wrong, epistemically speaking, when one believes irrationally, provided that one complies with KNOWLEDGE. This entails,

²⁵ Littlejohn (2013, fc1)

inter alia, that there would be nothing wrong with our thief suspending judgement on whether they have a Brueghel, despite the fact that he hasn't the slightest reason to think it's a fake, since in suspending judgement, he complies with KNOWLEDGE. This is extremely counterintuitive, to say the least. *Clearly* something has gone wrong. Dilemmism accommodates that observation. The excuses view does not. So, we should prefer dilemmism.

§7.4. *The Dividing Approach*

Like dilemmists, dividers (e.g. Littlejohn fc2) think KNOWLEDGE and RATIONALITY are both genuine epistemic norms, and that conformity with each is non-optional. However, they maintain that the two norms don't really conflict.²⁶ This, they claim, is because 'ought' means two different things in KNOWLEDGE and RATIONALITY. Our thief ought^K not believe that they have a Brueghel, and ought^R to believe that they have one. But according to the divider there is no more of a conflict here than there is in saying that I can^{LEGAL} buy a house in Hampstead (no law prohibits it), but also can't^{FINANCIAL} buy a house in Hampstead (I don't have that kind of money), or that a guy is tall^{ACADEMIC} for an academic, but not tall^{BASKETBALL} for a basketball player.

However, the claim that distinguishing between two meanings of 'ought' shows that there is no conflict looks wildly implausible when applied to analogous cases. Consider Sartre's student's dilemma again. In keeping with their approach to conflicts between KNOWLEDGE and RATIONALITY, a divider might argue that he ought^{FAMILY} to stay home, and ought^{PATRIOTISM} to

²⁶ At least, this is something one often hears. I'm not sure if it's Littlejohn's view. At one point he suggests that dilemmists can be dividers. So it may be that he sees the dividing approach as a way of semantically implementing dilemmism, rather than as a competing view.

go and fight. This may be true. But the idea that this shows there is no dilemma after all is absurd. Or suppose we invent a silly game with two rules. Rule 1: If it's a Monday, you must wear a hat. Rule 2: If it's an even-numbered calendar day, you must not wear a hat. On Monday the 2nd, these rules conflict, giving rise to a dilemma. It would be ridiculous to try to deny this by pointing out that you ought^{RULE-1} to wear a hat, but ought^{RULE-2} to not wear a hat. Again, this may be true, but it does absolutely nothing to show that there isn't a conflict.²⁷ If the divider's claim isn't plausible in these cases, there is no reason to think it is any more plausible when it comes to conflicts between KNOWLEDGE and RATIONALITY.

§VIII. Summing Up

Greco argues that there cannot be epistemic dilemmas; the mark of an epistemic dilemma is unavoidable epistemic residue, and there cannot be such a thing – the Reflection Argument shows as much. We have seen that the Reflection Argument doesn't work. It relies on the false assumption that we cannot rationally engage in preemptive epistemic quasi-blame, it overgenerates disastrously, and it faces the problem of inscrutable cognizers. The problem of inscrutable cognizers also shows that epistemic residue is not a good diagnostic by which to identify epistemic norms – and hence epistemic dilemmas. But even if it was, *contra* Greco, epistemic dilemmas would be possible, as the Jamie case shows. Greco is also concerned that if we don't identify epistemic dilemmas with unavoidable epistemic residue, we run the risk of merely verbal disputes. But as we have seen, this is incorrect – there are other diagnostic tools available with

²⁷ C.f. Gibbons (2013)

which to distinguish dilemmic from non-dilemmic views. Finally, we have seen that non-dilemmic approaches to KNOWLEDGE and RATIONALITY face serious problems, problems which dilemmism does not face. For this reason, amongst others, we should be dilemmists.²⁸

§IX. References

Alston, W. 1988, "The Deontological Conception of Epistemic Justification" *Philosophical Perspectives* 2: 257–299

Brandom, R. 1994. *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Harvard University Press.

Chisholm, R. 1966. *Theory of Knowledge*. Prentice Hall

Donagan, A. 1984. "Consistency in Rationalist Moral Systems" *The Journal of Philosophy* 81: 291-309

Equality and Human Rights Commission Research Report 113: *Unconscious bias training: An assessment of the evidence for effectiveness*

Gibbons, J. 2013. *The Norm of Belief*. Oxford University Press

Greco, D. forthcoming. "On the Very Idea of An Epistemic Dilemma"
Hughes, N. (ed.) *Epistemic Dilemmas*. Oxford University Press

²⁸ Thanks to Dan Greco, Mike Hannon, Daniel Kodsi, and Tim Williamson for comments and discussion.

Hughes, N. 2019. "Dilemmic Epistemology" *Synthese* 196: 4059-4090

Hughes, N. forthcoming-1. "Whose Afraid of Epistemic Dilemmas?"
McCain, K., Stapleford, S., & Steup, M. (eds) *Epistemic Dilemmas: New Arguments, New Angles*.

Hughes, N. forthcoming-2. "Evidence and Bias" Lasonen-Aarnio, M. & Littlejohn, C. (eds.) *The Routledge Handbook of the Philosophy of Evidence*

Hughes, N. manuscript-1. "Epistemology Without Guidance"

Hughes, N. manuscript-2. "Epistemic Feedback Loops (Or: How Not To Get Evidence)"

Kornblith, H. 2002. *Knowledge and its Place in Nature*. Oxford University Press.

Lasonen-Aarnio, M. 2010. "Unreasonable Knowledge" *Philosophical Perspectives*. 24: 1-21

Lasonen-Aarnio, M. 2020a. "Enkrasia or Evidentialism? Learning to Love Mismatch" *Philosophical Studies* 177: 597-632

Lasonen-Aarnio, M. 2020b. "Dispositional Evaluations and Defeat"
Brown, J. & Simion, M. (eds.) *Reasons, Justification, and Defeat*. Oxford University Press.

Lasonen-Aarnio, M. forthcoming-1. "Virtuous Failure and Victims of Deceit" Dutant, J. & Dorsch, F. (eds.) *The New Evil Demon Problem*. Oxford University Press

Lasonen-Aarnio, M. forthcoming-2. "Perspectives and Good Dispositions" *Philosophy and Phenomenological Research*.

Littlejohn, C. 2013, "The Russellian Retreat" *Proceedings of the Aristotelian Society* 113 (3): 293–320

Littlejohn, C. forthcoming-1. "A Plea for Epistemic Excuses" Dutant, J. & Dorsch, F. (eds.) *The New Evil Demon Problem*. Oxford University Press

Littlejohn, C. forthcoming-2. "Doxastic Dilemmas and the Method of Division" Hughes, N. (ed.) *Epistemic Dilemmas*. Oxford University Press

McHugh, C. 2012. "The Truth Norm of Belief" *Pacific Philosophical Quarterly* 93: 8-30

Nelson, M. 2020. "Epistemic Tension and Epistemic Duty: A Rossian Account" McCain, K. & Stapleford, S. (eds.) *Epistemic Duties: New Arguments, New Angles*.

Simion, M., Kelp, C., & Ghijsen, H., 2016. "Norm of Belief" *Philosophical Issues*

Williamson, T. forthcoming-1. "Justifications, Excuses, and Sceptical Scenarios" Dorsch, F. & Dutant, J. (eds.) *The New Evil Demon: New Essays on Knowledge, Rationality, and Justification*. Oxford University Press

Williamson, T. forthcoming-2. "Epistemological Ambivalence" Hughes, N. (ed.) *Epistemic Dilemmas*. Oxford University Press

Zimmerman, M. 1996. *The Concept of Moral Obligation*. Cambridge University Press.