

Dilemmic Epistemology

Nick Hughes, UCD
Forthcoming in Synthese

'In the tragic perspective, acting, being an agent, has a double character. On the one side, it consists in taking council with oneself, weighing the for and against and doing the best one can to foresee the order of means and ends. On the other hand, it is to make a bet on the unknown and the incomprehensible and to take a risk on a terrain that remains impenetrable to you.'

– Jean-Pierre Vernant

Abstract: This article argues that there can be epistemic dilemmas: situations in which one faces conflicting epistemic requirements with the result that whatever one does, one is doomed to do wrong from the epistemic point of view. Accepting this view, I argue, may enable us to solve several epistemological puzzles.

§I. Ethics and Epistemology

Ethicists have long debated the question of whether there are genuine moral dilemmas: situations in which one faces conflicting moral requirements, with the result that whatever one does, one is doomed to do wrong. But the question of whether there are genuine epistemic dilemmas, in which one faces conflicting epistemic requirements, has received little attention from epistemologists.¹ I find this surprising. It is currently *en vogue* to tackle epistemological questions with the tools of deontology. According to this approach, epistemic normativity is to be understood in terms of sets of rules (aka norms). One is naturally led to wonder

¹ Little, but not none. Ross (2010) argues that there can be certain kinds of epistemic dilemmas. Srinivasan (2015), Christensen (2016), and (tentatively) Hawthorne and Srinivasan (2013), endorse views that come quite close to the idea. Moss (2013) argues that there can be epistemic dilemmas, but not of the sort that I am interested in here. Her view is that there are cases in which nothing speaks in favour of adopting one epistemic attitude over another incompatible one, but not that one is *required* to take each attitude. Moss's view can be thought of as a kind of Permissivism.

whether those rules sometimes conflict with one another and thereby give rise to epistemic dilemmas. I think they do, and that accepting this view may enable us to solve several epistemological puzzles. My goal here is to argue that a dilemmic approach to epistemology can solve one such puzzle, and to lay the grounds for potential dilemmic solutions to others.

In §2 I'll present the puzzle, describe the dilemmic approach to solving it, along with the virtues of taking this approach, and briefly catalogue some possible alternatives. In §3 I'll say something about my ambitions in defending the dilemmic approach in this article. In §4-6 I'll present and respond to what I take to be the three main objections to it. In §7 I'll briefly sketch some other possible applications of the general framework of dilemmic epistemology.

§II. Dilemmism

§2.1. The Puzzle

Here are two plausible looking claims:

T: One ought (epistemically) to only believe truths

R: One ought (epistemically) to be epistemically rational²

There is plenty of scope for debate about just how the 'oughts' in *T* and *R* should be interpreted, but it is hard to deny that there is at least some sense (or senses) in which both of them are true. Moreover, each looks to be in some sense *non-optional*.³ If they were merely optional then there would be nothing wrong with self-consciously disregarding them, and so nothing wrong with holding beliefs one knows to be false or irrational. But there is something very wrong, from the

² I have in mind here the kind of rationality associated with reasonableness, rather than means-end coherence. On a different note, some philosophers treat 'it is rational for you to ϕ ' as *synonymous* with 'you ought to ϕ '. I think this is a mistake. The question whether you ought to always be rational cannot be settled by definitional fiat.

³ Gibbons (2013). I also think that the norm *K*, according to which one ought (epistemically) to only believe that *P* if one would thereby come to know that *P*, is genuine and non-optional (Unger 1975 Williamson 2000, Adler 2002, Sutton 2005, 2007, Littlejohn 2013). I won't discuss it here because it is superfluous to our concerns. Given the factivity of knowledge, anyone who endorses *K* is committed to *T*, and *K* and *T* cannot conflict with one another in the way that I'll be arguing *T* and *R* can. The reader is welcome to extend everything I say about *T* to *K* if they are also inclined accept it.

epistemic point of view, with such beliefs.⁴

So far so good. We have an epistemology of doxastic attitude formation, cashed out in the coin of deontology. But there is a problem lurking. Consider this case:

MUG: You have just woken up and are in the kitchen making breakfast. You want to make a cup of tea, and you'd prefer to drink it from your favourite mug. Remembering that you left it in the dishwasher last night, you look in there and see what appears to be the mug sitting on the rack, just where you left it. On the basis of your visual experience you form the belief that your favourite mug is in the dishwasher. Alas, you are deceived. Last night while you were sleeping a thief broke in, stole your mug, and replaced it with an identical-looking replica.

Call the proposition that your favourite mug is in the dishwasher 'P'. According to *T* you ought *not* believe that P, since P is false. But doesn't *R* tell you that you ought *to* believe that P? Surely that's the rational thing for you to believe. After all, you have no reason to suspect any funny business. People don't normally go around stealing mugs in the dead of night and replacing them with replicas. In fact, wouldn't it be irrational for you *not* to believe that P? You're motivated to find out the truth about P - you want to use your favourite mug, rather than another one - so it would make no sense not to form any doxastic attitude about it. But to suspend judgement would be to practice an odd form of excessive epistemic cautiousness. We may suppose that your visual experience is as vivid as it could be, and that you have a clear head and excellent eyesight. As already stated, the thing looks absolutely identical to your favourite mug, and you have no reason to think that someone might have replaced it. We can also add, if we like, that the mug has an unusually distinctive pattern, so you couldn't have easily mistaken it for one of your others. If you still suspend judgement despite all this, you're not exactly a paragon of rationality. To disbelieve that P would be even worse, of course; disbelieving that P is equivalent to believing that not-P. In your circumstances that would be egregious.

Rationality appears to dictate that you ought to believe that P, then. And therein lies the problem. If that's right, then there are cases, like *MUG*, in which *T* and *R* issue conflicting instructions.⁵ *T* tells you not to believe that P, and *R* tells you to

⁴ You might think that since it is irrational to believe that which one knows to be false, only *R* is needed to explain what is wrong with each kind of belief, and so the wrongness of believing them does not motivate *T*. I disagree. If *T* were false, or merely optional, then it would be a mystery just *why* it is irrational to believe known falsehoods in the first place.

⁵ Gibbons (2013).

believe that P.⁶ But you can't do both of these things. It's logically impossible to both believe that P and not believe that P (at the same time, anyway). So surely something must give? Call any case in which *T* and *R* apparently conflict a 'conflict case'. The puzzle is: what should we say about the epistemology of conflict cases?

§2.2. The Dilemmic View

I say we should take them at face value. My view - dilemmism - is that *T* and *R* both express full-blown, bona fide, epistemic requirements. So you're required to believe that P in conflict cases, and at the same time required not to believe that P. And that's pretty much that. Neither of these requirements outweighs or takes priority over the other, and there is no hope for resolving the conflict by appealing to the idea that there are different senses of 'ought' at work or anything like that. All things considered you both ought to believe that P and ought not to believe that P.⁷ Sometimes, through no fault of your own, you can stumble into a situation in which you're damned if you do and damned if you don't. Just as there are moral dilemmas, in which you fall short of living up to the demands of morality whatever you do,⁸ so too are there epistemic dilemmas. Life is hard.

It is worth mentioning straight off the bat - and you've probably noticed this already - that whilst the analogy with moral dilemmas is useful, and for the most part accurate, it isn't *quite* perfect. Those who face moral dilemmas are typically in a position to know about it; they usually know that there are mutually

⁶ A referee has pointed out that in the *MUG* case the fact that P is false is not part of your evidence. They ask: in what *epistemic* sense ought you not believe it, then? Perhaps we could say that there are *alethic* grounds on which you ought not believe it, but *epistemic*? I suspect that the worry here is partly a terminological one. Some philosophers take phrases like 'epistemically, you ought to believe that...' to essentially make reference to what one ought to believe *on the assumption that one ought to believe what the evidence supports*. They take the phrase 'epistemically, you ought to believe that...' to be synonymous with 'in order to conform to your evidence, you ought to believe that...'. Other use the phrase more broadly - to pick out a particular domain of normativity; roughly, that domain of normativity which is neither moral nor practical, but rather (putting it roughly) involved in acquiring an accurate picture of the world. On this usage it is not assumed that only your evidence bears on what you ought to do, epistemically speaking (though that remains a possibility). Here I use the phrase 'epistemically, you ought to believe that...' in the latter sense. One may prefer to use it in the former sense, but any dispute about how it should be used would be merely verbal - it would not concern what one ought to believe, but rather what words we should use in articulating a theory of belief formation. Of course, it might be argued that we should reject *T* on evidentialist grounds, but to assume a form of evidentialism incompatible with *T* at this point would be to rule out the view that I will be developing by prior to giving it the chance at a hearing

⁷ For a discussion of all-things-considered requirements to believe, see Booth (2012) and, *contra* Booth, Stapleford (2015).

⁸ Or at least, *so say I*. I don't deny that this is controversial.

incompatible courses of action, each of which has a moral claim on them. But the kinds of epistemic dilemmas I claim exist aren't like that. Rationality doesn't require you to believe that P when you are in a position to know that P is false. So you will never be aware that you are facing an epistemic dilemma. From your point of view *inside* a conflict case, both *T* and *R* appear to say that it's fine for you to believe that P. I maintain that this makes no difference. You can be in a bind without knowing it. But for the same reason, I'm not claiming that life is quite as hard as you might think I am. Whilst I think that conflict cases give rise to dilemmas, I don't think that you are necessarily epistemically *blameworthy* whatever attitude you take towards P in them. The deontic and the hypological realms do not perfectly align: it's possible to fail to do what is required of you without being an appropriate target of blame, provided that you have a good excuse.⁹ And if you believe that P in a conflict case, you have an excellent excuse, namely that from your point of view it looked like you were in accordance with both of *T* and *R*.¹⁰ Still, that doesn't mean that you've done nothing wrong: only wrongdoing needs excusing.¹¹

In order to ward off the possibility of misunderstanding, let me mention another way in which I am not claiming that life is quite as hard as you might think I am. I'm obviously not suggesting that *every* case is a conflict case. And dilemmism isn't only a view about conflict cases: it also has things to say about non-conflict cases. Most of the time the evidence is not misleading, and so *T* and *R* *don't* conflict with one another. When it's rational for you to believe that P, and P is true, dilemmism tells you to believe that P.¹² When it's rational for you to disbelieve that P, and P is false, dilemmism tells you to disbelieve that P. And when it's rational for you to suspend judgement on P, dilemmism tells you to suspend. The dilemmic view is not that you are *always* in a dilemma; only that you can *sometimes* be. When *T* and *R* concord - that is, when you are *not* in a

⁹ 'Hypological' is a rarely used term. It comes from the Greek *hypologos*: 'to hold accountable or liable' (Srinivasan 2015). The deontic realm covers things like requirements, obligations, and permissions. The hypological realm covers things like blameworthiness, criticisability, and praiseworthiness.

¹⁰ If you suspend judgement on P, on the other hand, then you are, intuitively, epistemically *blameworthy* for doing so, even though you did the right thing by *T*. At this point you might think that I'm identifying rationality with blamelessness and irrationality with blameworthiness. I'm not. Whilst I think that you are rational and blameless if you believe that P in *MUG*, and irrational and - excuses notwithstanding - blameworthy if you suspend judgement on P or disbelieve it, I also think that there can be cases of blameless irrationality. My reasons will become apparent later.

¹¹ That there is a distinction to be drawn between permissible behaviour and merely excusable behaviour is by now widely recognised (see Austin 1957 and Gardner 1997, amongst others). In epistemology, it is often invoked by those who accept various knowledge norms (Williamson 2000, and Hawthorne and Stanley 2008, for instance). Feldman (2008) also discusses the distinction.

¹² If we endorse *K* along with *T* and *R*, this will have to be revised slightly, since the combination of *R* and *T* permits belief in Gettier cases, but *K* doesn't.

conflict case - dilemmism says that you are required to conform with each, and that doing so presents no special difficulties.

§2.3. The Initial Reaction

That said, I don't expect many epistemologists to find dilemmism intuitively appealing. Quite the opposite, actually. Whilst the view has not received much attention in the literature, many of those with whom I have discussed it have responded with utter incredulity. Surely, the response goes, whatever the solution to the puzzle is, it can't be *that*. To this they sometimes add that it doesn't even have the form of a solution in the first place. Gibbons's (2013) remarks are representative. Though he mentions the view almost only glancingly, in order to put it aside, he nonetheless singles it out for special opprobrium. Describing the view as 'nihilistic', 'absurd', a 'disaster' and 'an abyss', he says, amongst other things, that he "cannot take it seriously as an option", and that "...it is far more likely that we've made a mistake in the argument for [this] absurd view than it is that the absurd view is correct" (23). Not a very favourable assessment, then. But one that will, I think, be widely shared. Still, it is a striking fact that there is no actual *argument* against dilemmism to be found amongst Gibbons' invective. He simply dismisses it out of hand. This is a risky strategy; there are plenty of philosophical views now taken seriously that were once unreflectively rejected. If we're going to dismiss the view, we should do so with reasons firmly in hand. And if those reasons can't be found, then we shouldn't be so quick to dismiss it. My contention is that reasons to reject the view are surprisingly thin on the ground, counterintuitive though it may be.¹³

§2.4. Virtues of Dilemmism

That doesn't mean that there are no *apparent* reasons to reject it. A large portion of this article will be dedicated to presenting and discussing what seem to me to be the three main objections to dilemmism. But before we get to all that, I want to discuss its virtues. Potential problems aside, what can be said in its favour? The answer, I think, is not that it has lots of nice little benefits, but rather that it has a few big ones.

¹³ Some of those with whom I have discussed the view have suggested that, whilst it shouldn't be discarded out of hand, it should nevertheless be thought of as a last resort - a view to be adopted only if all else fails - on the grounds that it is somehow *pessimistic*. I find this mentality puzzling. I see no reason to rank potential solutions to the puzzle prior to investigating their pros and cons, and given what I judge to be the pros and cons of dilemmism, I think it should be very far from the last resort.

Dilemmism vindicates the idea that *T* and *R* are both genuine epistemic norms. That's a big plus for the view. Both are absolutely central to our understanding of how one ought to go about forming one's beliefs. We may find ourselves forced to give up one or the other if things get desperate, but we should not do so lightly. There ought to be a strong presumption in favour of each. Since dilemmism allows us to 'save the appearances' here, that makes it an attractive view.

Moreover, because it interprets them as issuing *requirements*, dilemmism also vindicates the idea that *T* and *R* are both non-optional. If you're required to ϕ , then ϕ -ing isn't optional in the way that it would be if it were supererogatory or merely recommended; you *must* do it. The fact that it vindicates the non-optionalness of *T* and *R* is another highly attractive feature of the view. There is something deeply problematic, from the epistemic point of view, with holding on to beliefs that one knows to be false or irrational, and again, whilst we may eventually be forced to conclude that it is sometimes acceptable to do so, this should be avoided if at all possible. Once again, dilemmism saves the appearances, and that is to its credit.¹⁴

Finally, dilemmism does all this in a very straightforward and simple way. It provides an epistemology of doxastic attitude formation developed with a well recognised and well understood normative property - that of requirement - whilst at the same time both avoiding the need to posit any kind of error theory for either of *T* or *R*, and avoiding the need to appeal to a multitude of normative properties or statuses.

Each of these points - the fact that it vindicates *T* and *R*, the fact that it vindicates the idea that both are non-optional, and the fact that it does these things in a straightforward and simple way - should, I suggest, carry substantial weight when we come to considering the pros and cons of dilemmism versus the alternatives. My inclination is to think that they should make it a presumptive front runner.

¹⁴ A referee has asked: do I take the dilemmic view to provide an argument for *T* and *R*? I do not. Rather, I take them to be data that must be accommodated by an epistemology of belief. The dilemmic view accommodates them, and that is something that speaks in its favour. The same goes for the fact that *T* and *R* are both non-optional. This is data that likewise must be accommodated, and in virtue of interpreting them as both issuing requirements the dilemmic view accommodates it, since whenever one is required to ϕ , ϕ -ing is not optional in the way that it would be if it were supererogatory or merely permissible. This is what I mean when I say that dilemmism 'vindicates' the non-optionalness of *T* and *R*. Arguments for *T* and *R* abound in the literature. See, for instance, Unger (1975), Williamson (2000), Wedgwood (2002), Shah (2003), Gibbons (2013), and Whiting (2013), amongst others, for arguments for *T*. See Wedgwood (2002), Gibbons (2013), Cohen & Comesana (forthcoming) and Hughes (forthcoming), amongst others, for arguments for *R*. The observation that conformity with each of *T* and *R* is non-optional is due to Gibbons (2013).

§2.5. Alternatives

I'll have more to say about the details of the view in due course. But before we get to those details, and to the inevitable objections, I want to catalogue some of the possible alternatives to it. Dilemmism isn't the only game in town when it comes to potential solutions to the puzzle. A range of other strategies may look promising.

Firstly, and most obviously, we might reject one of *T* and *R*. Some epistemologists think that the 'ought' in 'you ought only believe truths' does not express a requirement, but rather articulates an *aim* or *goal*.¹⁵ If we can make sense of the idea that *T* expresses an aim or goal without also issuing a requirement, then perhaps it is possible to solve the puzzle by arguing that, whilst one is required to be rational, one is not required to believe only truths. If so, then there will be no need to posit dilemmas. All things considered you ought to believe that *P* in conflict cases.

A different approach, which Lasonen-Aarnio (2010, forthcoming) and Williamson (forthcoming) advocate (or come close to advocating, at least), is to argue that *T* expresses a genuine requirement, and that the 'ought' expressed by *R* articulates something less than a requirement.¹⁶ They argue that an agent who fails to believe that *P* in conflict cases is criticisable in virtue of the fact that they thereby manifest a non-knowledge-conducive disposition. Evidence isn't normally misleading. So even if you 'get it right' in conflict cases from the point of view of *T* by suspending on *P*, in doing so you exhibit a tendency to miss out on an epistemic good: knowledge, since you will, presumably, also suspend on *P* in cases in which the evidence is *not* misleading. But, Lasonen-Aarnio and Williamson may maintain, it is one thing to say that someone is criticisable for manifesting a non-knowledge-conducive disposition, and quite another to say that they are *required* not to do so. If this line of argument can be made to work, then it may be possible to avoid dilemmas by accepting *T* but rejecting (or rather, reinterpreting) *R*. All things considered you ought not believe that *P* in conflict cases.

Taking a different tack, it might be argued that *T* and *R* don't really conflict in the first place. Clayton Littlejohn has suggested to me that *R* only expresses a *negative*

¹⁵ For good discussions of the idea that belief aims at truth see Williams (1973) Velleman (2000), Wedgwood (2002), Boghossian (2003), Shah (2003), Shah and Velleman (2005), and the essays collected in Chan (2013). (Note: I don't mean to suggest that these authors all reject the idea that *T* expresses a requirement - one might think that belief aims at truth *and* that one is required to only believe truths).

¹⁶ The reason they think that *T* expresses a genuine requirement is because they think that *K* does, and *T* follows from *K*.

requirement - one that can be satisfied by taking no attitude whatsoever towards P (not even suspension). The idea would be that you are never required to form a doxastic attitude just because it would be rational for you to do so. Instead you're only ever required *not* to form a doxastic attitude if it would be *irrational* for you to do so.¹⁷ If so, then you can do everything that *T* and *R* ask of you in cases like *MUG* by simply refusing to take any doxastic attitude towards P in the first place. Again, if that's right, then there is no need to posit dilemmas. In conflict cases you ought, all-things-considered, take no attitude at all to P, and that's what you're required to do.

Another way to avoid positing conflict would be to adopt some kind of disambiguation solution to the puzzle. The thought, roughly, would be that *T* and *R* don't really conflict because each expresses a different sense of 'ought'. In conflict cases you ought^{truth} not believe that P, and ought^{rationality} believe that P, but, the argument would go, there is no more conflict here than there is there between the claims that a 6' 3" man is 'tall' by the standards of the average philosopher, but 'not tall' by the standards of the average NBA player. According to this view, the question of what you ought to believe all-things-considered is ill-formed and has no answer.¹⁸

Some philosophers think that one can be subject to two conflicting requirements, but that one of them can outweigh the other, giving rise to a univocal answer to the question of what one ought to do all-things-considered (Ross 1930, Nussbaum 1986, Booth 2012). Once again, this may hold the promise of avoiding the need to posit dilemmas. Finally, one might simply reject deontological approaches to epistemology in the first place (Alston 1988). If so, then there isn't even a real puzzle to be solved here in the first place; instead there's only a misguided way of doing epistemology.

§III. Ambitions

Many readers will have their favourites from this (no doubt incomplete) list. I don't think any of them can be made to work, which is one of the reasons I'm drawn to dilemmism. But my intention is not to argue against them here. Properly doing so would take us *far* beyond the scope of a single article. I could, I suppose, take a few potshots at them, leave the whole matter inconclusive, and then move

¹⁷ Nelson (2010), Littlejohn (2013). For a reply to Nelson, see Stapleford (2013).

¹⁸ Although he doesn't propose it as a solution to the puzzle I am interested in here, Feldman (2000) is an example of someone who thinks that the question of what you ought to believe all-things-considered is ill-formed and has no answer.

swiftly on to the main agenda: developing and defending dilemmism. But it would be a facile and perfunctory enterprise, convincing no-one of anything. Better to concentrate on the crux of the matter: to see what can be said in favour of the dilemmic approach, and to respond to the main objections to it. Dilemmism has not been taken seriously as an option so far (just look at Gibbons's reaction to it). I think that it should be. My ambition here is modest: to get it on the table for serious discussion. A full defense of the view will have to wait.

What I will do, however, is look closely at three objections to the view: that it must be rejected because it leads to contradictions and explosions in deontic logic; that it must be rejected because it fails to give agents useful guidance; and that it must be rejected because it requires one to do the impossible, and thereby violates the principle 'ought-implies-can'. As we will see, each of these objections is multifoliate, and each is intertwined with the others to some extent. I've chosen to focus on these objections, rather than others, for two reasons. The first reason is because I think that they provide the most serious challenges to the view. Between them they cover the reasons why nearly all of those with whom I have discussed the view have found it so hard to countenance. So if they can be shown not to add up to much, then that should be enough to at least get the view taken seriously. The second reason is because I have an eye on the programmatic dimension of the view. As I mentioned earlier, I think there may be other epistemological puzzles that can be solved by positing dilemmas. But as we will see, *any* application of the general approach to these puzzles will face the problem of having to overcome these objections. So if they can be overcome, some ground will have been cleared for extensions of the general framework.

§IV. Logic

§4.1. Contradictions?

Let's start with the logic. The dilemmic view says that in conflict cases one ought to believe that P and that one ought not to believe that P. These claims may appear to be straightforwardly contradictory. Since there are no true contradictions (or so I'll assume; I'll work with classical logic throughout), the dilemmic view, one might think, cannot possibly be correct.

In fact though, the appearance of contradiction is illusory.¹⁹ Dilemmism asserts (1) and (2) (where ‘O’ stands for ‘it ought to be / is required that’ and ‘B’ stands for ‘one believes that’):

1. OBP
2. O¬BP

But (1) and (2) do not contradict one another. The negation of (1) is not (2), but (3):

3. ¬OBP

Dilemmism does not assert (3). So no immediate contradiction arises from the view. Nevertheless, as we’ll see, given some principles of Standard Deontic Logic (‘SDL’) one can derive both contradictions and deontic explosions from the conjunction of (1) and (2).

Some philosophers may want to take this fact to be a decisive reason to reject dilemmism all by itself. That would be a mistake. SDL is rife with problems and paradoxes, and does a notoriously bad job of capturing the range of deontic inferences that we find natural (Sayre-McCord 1986, Hansen *et al.* 2007). Indeed, the epithet ‘standard’ arguably gives it a more elevated status than it really deserves. Given its fraught and tangled status, we should not be shy about proposing revisions. In saying this I don’t mean to suggest that no weight whatsoever should be assigned to it - it captures at least some of the deontic inferences we find natural in a simple and straightforward way - but rather that the bar to accepting revisions should not be set as high as it should be set for proposed revisions to, for example, classical propositional logic.²⁰ When substantive normative theory conflicts with SDL, the former should not always and automatically be subordinated to the latter.²¹

With that in mind, let’s look at the conflicts that arise between SDL and dilemmism. Most of what I’ll say in this section is already familiar from the literature on moral dilemmas, so I’ll try keep things brief.

¹⁹ McConnell (1978), Marcus (1980) make this point about moral dilemmas.

²⁰ The reader is free to substitute in here whatever logic they think is most robust.

²¹ My impression is that this is the general consensus in ethical theory.

§4.2. Deontic Explosion

It's logically impossible to both believe that P and not believe that P. As Horty (2007) (amongst others) points out, given three principles accepted in SDL, the hypothesis that there are logically incompatible requirements leads to deontic explosion, according to which *everything* is required. The three principles are:

Ought-Entailment (OE): $\Box (\varphi \rightarrow \psi) \rightarrow (O\varphi \rightarrow O\psi)$

Agglomeration (AGG): $O\varphi \ \& \ O\psi \rightarrow O(\varphi \ \& \ \psi)$

Ex Falso Quodlibet (EFQ): $(\varphi \ \& \ \neg\varphi) \rightarrow \psi$

OE says that if φ entails ψ , then if you ought to φ , you ought to ψ . *AGG* says that if you ought to φ and you ought to ψ , then you ought to perform the complex action φ -and- ψ . *EFQ* says that everything follows from a contradiction.²²

So to avoid deontic explosion the dilemmist must reject either *OE*, *AGG*, or *EFQ*. I reject *AGG*. I'll explain my reasons for doing so shortly.

§4.3. Consistency, Ought-Implies-May, Duality

Lemmon (1965) (amongst others) points out that the claim that there are normative dilemmas leads to a contradiction given the following principle of SDL:

Principle of Consistency (PC): $O\varphi \rightarrow \neg O\neg\varphi$

²² Here's how the explosion occurs:

- | | |
|---|--------------------------------|
| 1. $O\text{BP}$ | (assump.) |
| 2. $O\neg\text{BP}$ | (assump.) |
| 3. $O(\text{BP} \ \& \ \neg\text{BP})$ | (by <i>AGG</i>) |
| 4. $O(\text{BP} \ \& \ \neg\text{BP}) \rightarrow O\varphi$ | (by <i>OE</i> and <i>EFQ</i>) |
| 5. $O\varphi$ | (from 3 and 4) |

(1) and (2) describe the dilemmic view. (3) follows from (1) and (2), given *AND*. (4) follows from (3), given *OE* and *EFQ*, and (5) follows from (3) and (4)

PC says that if you ought to ϕ , then it is not the case that you ought not ϕ .²³ *PC* follows from two further principles (where ‘P’ stands for ‘one is permitted to’):²⁴

$$\begin{aligned} \text{Ought-Implies-May (OIM): } & O\phi \rightarrow P\phi \\ \text{Duality (D): } & P\phi \leftrightarrow \neg O\neg\phi \end{aligned}$$

OIM says that if you ought / are required to ϕ , then you are permitted to ϕ . *D* says that you are permitted to ϕ iff it is not the case that you ought / are required to not- ϕ . So the dilemmist must reject *PC*, which in turn requires rejecting one of *OIM* or *D*. I reject *OIM*. *D* strikes me as by far the most plausible principle of SDL, but only someone already committed to the impossibility of dilemmas would accept *OIM* (and so *PC*) along with it, or so I’ll argue.

§4.4. Ought-Implies-Can

Normative dilemmas require you to do the impossible. That’s what makes them *dilemmas*. So they are obviously incompatible with some ought-implies-can principles. Since the dilemmic view rejects *AGG*, however, it is compatible with ought-implies-can as it is usually understood. That is, as having the logical form:

$$\text{Ought-Implies-Can (OIC): } O\phi \rightarrow C\phi$$

(Where ‘C’ stands for ‘one can’). *OIC* is an axiom of SDL. According to dilemmism, in conflict cases you ought to believe that P and ought not believe that P. Now, *if* we kept hold of *AGG*, it would follow that the dilemmic view is incompatible with *OIC*.²⁵ But once we reject *AGG* there is no action which you are required to perform but cannot perform. After all, the following are both true:

- CBP
- C \neg BP

What is *false* is:

²³ Here's how the contradiction arises:

- | | |
|---|-----------------|
| 1. OBP | (assump.) |
| 2. O \neg BP | (assump.) |
| 3. OBP \rightarrow \neg O \neg BP | (by <i>PC</i>) |
| 4. \neg O \neg BP | (from 1 and 3) |

(1) and (2) describe the dilemmic view. (3) applies *PC*. (4) follows from (1) and (3), and (2) and (4) contradict one another.

²⁴ McConnell (2014).

²⁵ Williams (1965).

- $C(BP \ \& \ \neg BP)$

But dilemmism doesn't claim that in conflict cases you are required to *both* believe that P and not believe that P (i.e. $O(BP \ \& \ \neg BP)$). And without this claim there is no conflict between dilemmism and *OIC*.

The dilemmic view is not compatible with the following version of ought-implies-can, however:

Aggregating-Ought-Implies-Can (A-OIC): $O\phi \ \& \ O\psi \ \rightarrow \ C(\phi \ \& \ \psi)$

A-OIC doesn't state that if you are required to perform some action, then you can perform it. Rather, it states that in a world w in which you have a set of requirements they are jointly satisfiable (i.e. compossible) at w .²⁶

OIC entails *A-OIC* given *AGG*. But the entailment does not hold if *AGG* is rejected. The upshot is that *OIC* might be true and *A-OIC* false. So the dilemmist can keep hold of *OIC*, but must reject *A-OIC* when it comes to the 'ought' of requirement and the 'can' of logical possibility. Since the 'can' of logical possibility is the weakest sense of can (i.e. every true statement of the form 'S can ϕ ' entails that it is logically possible for S to ϕ) that means the dilemmist must reject *A-OIC* for all senses of 'can'. I'll discuss my reasons for rejecting *A-OIC* in §6.

§4.5. On *PC*, *OIM*, and *AGG*

Unlike ought-implies-can principles, not very much has been offered by way of argument for *PC*, *OIM*, and *AGG*. Those who endorse them seem by-and-large to think that they are just obviously true and so in need of no defence.

That's not right. Given *D*, *OIM* merely represents a rejection of the possibility of dilemmas, disguised as a logical law (van Frassen 1973). As van Frassen points

²⁶ To avoid confusion I should be clear that *A-OIC* should not be read as only claiming that if you have *two* conflicting requirements, they are jointly satisfiable. The number of requirements may be arbitrarily large. So *A-OIC* says that when you have 10, or 100, or 1000...and so on, requirements, they are jointly satisfiable. Here's why dilemmism is incompatible with *A-OIC*:

- | | |
|--|-------------------|
| 1. OBP | (assump.) |
| 2. $O \neg BP$ | (assump.) |
| 3. $\neg C(BP \ \& \ \neg BP)$ | (assump.) |
| 4. $(OBP \ \& \ O \neg BP) \rightarrow C(BP \ \& \ \neg BP)$ | (by <i>OIC*</i>) |
| 5. $C(BP \ \& \ \neg BP)$ | (by 1, 2, and 4) |

(1) and (2) describe the dilemmic view. (3) follows from the fact that BP and $\neg BP$ are logically impossible. (4) follows from (1) and (2), given *A-OIC*, and (5) contradicts (3).

out, the conjunction of *D* and *OIM* (i.e. *PC*) is straightforwardly *equivalent* to the claim that there are no dilemmas. He conjectures that 'the original devisors of deontic axioms had a certain ethical bias...[against dilemmas]' (12). If *OIM* seems to be a natural and intuitive rule of inference to those who accept *D*, this is only because most of our deontic reasoning does not concern apparently conflicting requirements. But whilst *OIM* undoubtedly operates smoothly for most of our deontic reasoning, it breaks down in cases involving apparent dilemmas (assuming *D*, at least), and given the controversial status of *SDL* we cannot straightforwardly appeal to it in order to reject the possibility of dilemmas. To do so would be to give *SDL* unwarranted priority over substantive normative theorising (Sayre-McCord 1986). The matter of which direction to go - whether to keep hold of *OIM* along with *D* and thereby reject the possibility of dilemmas, or to reject *OIM* in order to make space for the possibility of dilemmas - isn't so easily decided. What we need to do is to come to a decision based on an evaluation of the costs and benefits of each course.

I'm also rejecting *AGG*. In doing so I follow Williams (1965), van Frassen (1973), and Marcus (1980), all of whom reject it in order to accommodate moral dilemmas. I likewise don't deny that *AGG* operates smoothly for most of our deontic theorising. But again, it doesn't always. Suppose you promise Jack that you will meet him but not Jill tomorrow, and also promise Jill that you will meet her but not Jack tomorrow. It is hardly obvious that you are not subject to two conflicting requirements in this situation: a requirement to keep your promise to Jack and a requirement to keep your promise to Jill. But it is difficult to even make sense of the idea that your requirements agglomerate to give rise to a single complex requirement: to meet Jack and not Jill and Jill and not Jack. There is no logically possible world in which you satisfy *that* would-be requirement. Such a state-of-affairs isn't even *conceivable*, so it is hard to see on what grounds we could base a judgement that you are required to bring it about. Moreover, when we bring in *OE* and *EFQ*, a consequence of the hypothesis that you are subject to the complex requirement is that *everything* is required, which is absurd.

Can we simply and unhesitatingly appeal to *AGG* to conclude that, contrary to appearances, you must not have the two conflicting individual requirements in this situation, then? No. As before, that would be to give *SDL* unwarranted priority over substantive normative theorising. Again, the matter of which direction to go isn't so easily decided. We need to come to a decision based on an evaluation of the costs and benefits of each course.

§4.6. Rejecting *PC*, *OIM*, and *AGG*

I'll discuss *A-OIC* later. But what about *PC*, *OIM*, and *AGG*? On what grounds do I reject them?

On the grounds that the price to be paid in rejecting them is one worth paying. The cost will be a revision to, and probably a complication of, our deontic logic. But, to my knowledge at least, no-one has shown that the consequences of rejecting them will be especially bad, and in any case deontic logic already needs to be overhauled anyway.²⁷ In return we get a simple and powerful solution to the puzzle at hand. One that vindicates both the idea that *T* and *R* are both genuine epistemic norms, and the idea that each is non-optional. The trade, I say, is worth making. Naturally I don't suppose everyone will see things this way. In order to convince those who don't it would have to be shown that dilemmism is sufficiently superior to the alternatives, and that the required revisions to deontic logic are sufficiently painless, that the price is right. Whilst I think that this case can be made, I can't make it here, so for now I'll flag it as one of the central tasks for the development of dilemmic epistemology down the line. But it is not, I think, necessary to get the view taken seriously, and that's my goal here. Some forbearance is needed.

§V. Guidance

Let's move on to the second big objection to dilemmism. It is widely thought that an adequate normative theory must be followable, useable, or - to adopt a popular piece of jargon - 'action guiding'. Though the exact details are often left unclear (almost always, actually), the basic idea seems to be that the role of normative theories, or a primary role at least, is to provide useful guidance to agents in their deliberation about what to do. What actions to perform, in the case of moral and prudential theories, and what beliefs to form in the case of epistemological theories of doxastic attitude formation. Given that the very purpose of a normative theory - its *raison d'être* - is to guide deliberation, a theory that is incapable of doing so is, the thought goes, thereby inadequate and must be rejected. It is for this reason that Hudson (1989) rejects objective utilitarianism. "For human agents" he says, "the theory is not really 'action guiding': it does say what one should do, but it gives this information in a practically unusable way" (221). Why?

²⁷ A number of logicians have already developed dilemma-permitting logics. Goble (2005) and Horty (2007) are two recent examples.

Because in an unpredictable world one is rarely (if ever) in a position to know which course of action will maximise utility in the long run.

Of all the objections to dilemmism that I have heard, one of the most frequently and vociferously voiced is that it fails to give adequate guidance (*useful* guidance, really).²⁸ It's easy to have sympathy with this complaint. The 'advice' given by dilemmism in conflict cases is that you ought to believe that P and at the same time ought not believe that P. It's natural to respond with exasperation: if *this* is the advice, critics are liable to think, it's no better than none at all. And since it is an essential feature of an acceptable normative theory that it gives agents adequate guidance, the dilemmic view is a non-starter. But is that right? What exactly *is* adequate guidance? And is it true that dilemmism fails to offer it? In order to answer these questions, we need to spell out the notion of 'adequate guidance' in more detail. Little work has been done on this task.²⁹ But as I see it, there are at least three interconnected ways in which a normative theory might be thought to give - or fall short of giving - adequate guidance.

§5.1. Varieties of Guidance

First we have what we can call the *transparency sense* of adequate guidance. This is what Hudson has in mind when he takes exception to objective utilitarianism.³⁰ When a normative theory allows for the possibility of circumstances in which you are required to ϕ , but not in a position to know about it, then it is, in some intuitive sense, not followable or guiding; you can't consult the theory to decide what to do, because its prescriptions are opaque to you. If, for example, someone teaches you the Highway Code, but neglects to tell you that you are required to stop at red lights, then you cannot use the fact of this requirement to guide your deliberation about what to do when you are confronted with one. This suggests a desideratum on adequate guidance:

The Transparency Desideratum: A normative theory is adequately guiding only if, whenever it requires you to ϕ , you are in a position to know that it requires you to ϕ

Secondly, we have what we can call the *possibility sense* of adequate guidance. When a normative theory allows for the possibility that you can be required to ϕ even though you are unable to ϕ , then it is again in some intuitive sense not followable or guiding; if the theory prescribes ϕ -ing, but you cannot ϕ , then you

²⁸ Rinard (forthcoming). The same objection is often made against the possibility of moral dilemmas.

²⁹ A notable exception is Smith (2012).

³⁰ See also Gibbard (1990), and Jackson (1991), amongst others.

cannot use this prescription to choose between your actual options.³¹ Suppose that you are a nurse dealing with a patient in severe pain. You ask the doctor what you ought to do. She tells you to give the patient an injection of morphine. If there is no morphine available, then the doctor's advice isn't much use. This gives us another desideratum:

The Possibility Desideratum: A normative theory is adequately guiding only if, whenever it requires you to ϕ , you are able to ϕ

Third is the *causal sense*. The thought here is that if you are to be guided in your behaviour by the fact that you are required to ϕ , then you must be able to ϕ *because* you are required to ϕ . If you know that you are required to ϕ , and you are able to ϕ , and you *do* ϕ , but you do so for reasons entirely unrelated to the fact that this is required of you, then you haven't really been *guided* in your behaviour by the fact of the requirement any more than someone who is entirely unaware that they are required to ϕ but happens to do so anyway. This gives us a third desideratum:

The Causal Connection Desideratum: A normative theory is adequately guiding only if, whenever it requires you to ϕ , you can ϕ because you are required to ϕ .

Clearly, if an adequate normative theory must satisfy each of these desiderata, then dilemmism should be consigned to the dustbin. One is not always in a position to know whether P is true or false or (I will argue shortly) whether or not it is rational for one to believe that P. So dilemmism violates the transparency desideratum. It also violates the causal connection desideratum. As several epistemologists have pointed out, one can ϕ for the reason that P - that is to say, *because* P - only if one knows that P (Unger 1975 Hyman 1999, Williamson 2000, Hornsby 2007). So if you're required by dilemmism not to believe some proposition because it is false, but you are in no position to know this, then you cannot refrain from believing it *because* you are subject to such a requirement.

Does it also violate the possibility desideratum? Not straightforwardly. Since I am rejecting *AGG*, the dilemmic view as I am developing it is compatible with ought-implies-can as it is usually understood. However, as we have already seen, it is not compatible with *A-OIC*:

$$A-OIC: O\phi \ \& \ O\psi \rightarrow C(\phi \ \& \ \psi)$$

³¹ Philosophers who have argued that an adequately guiding theory must not violate *OIC* include Hare (1963), Driver (1983), Griffin (1992), Copp (2003), and Andric (2015).

And there is a variation on the possibility desideratum that is just as intuitively attractive as the original:

The 2nd Possibility Desideratum: A normative theory is adequately guiding only if, whenever it requires you to do each of a set of actions ϕ^1, \dots, ϕ^n , you are able to do the conjunction of ϕ^1, \dots, ϕ^n

After all, insofar as it is unhelpful, guidance-wise, to be told that you are required to ϕ when you cannot, it is equally unhelpful to be told that you are both required to ϕ and at the same time required not to ϕ , for even if you are capable of ϕ -ing and capable of not ϕ -ing, if you are incapable of *jointly* ϕ -ing and not- ϕ -ing, you cannot use the set of requirements issued by dilemmic views to decide whether or not to ϕ .

So, guidance considerations look like they cause a major problem for dilemmism. There appear to be desiderata on an adequate normative theory that it falls woefully short of fulfilling. How serious a threat to the view does this pose? It depends on how plausible it is that a normative theory must meet these desiderata.

§5.2. Genuine Desiderata?

I don't think it's plausible. Starting with the transparency desideratum, it is worth noting that any view that takes T to express an epistemic requirement will face the objection arising from it with just as much force as dilemmism does. One is not always in a position to know whether a given proposition is true or false, after all. Moreover, unless one subscribes to an 'access internalist' conception of epistemic rationality, according to which we have privileged access to the facts about what it is rational for us to believe (Bonjour 1985, Chisholm 1988, Smithies 2011), the same will also go for any view that takes R to express an epistemic requirement. So it can hardly be said that there are reasons to reject dilemmism *in particular* and adopt a rival view here. If the transparency desideratum is genuine, then pretty much every view is in trouble, with the possible exception of access internalism.

But the transparency desideratum is deeply problematic in any case. Williamson (2000) argues that *no* non-trivial condition is *luminous* for creatures with epistemic limitations, such as ourselves.³² For every non-trivial condition, he argues, it can obtain even though one is not in a position to know that it obtains. Williamson's argument is well known, so I won't rehearse it here.³³ Although it is not universally accepted, I (along with many others) find it convincing, and I'll

³² A condition is luminous just in case whenever it obtains one is in a position to know that it obtains.

³³ See Williamson (2000), Chap. 4.

assume it here.³⁴ It follows from it that *no* possible normative theory satisfies the transparency desideratum. No matter how ‘internalist’ we go with the normative requirements that we posit, there will still be cases in which one is required to ϕ but not in a position to know about it.³⁵ The sensible conclusion to draw here isn’t that an adequate normative theory is impossible, but rather that transparency isn’t a genuine desideratum on an adequate normative theory in the first place.³⁶

The causal connection desideratum goes down with the transparency desideratum as collateral damage. If one can ϕ for the reason that one is required to ϕ only if one knows that one is required to ϕ , and no non-trivial condition is luminous, then it follows that there must be cases in which one is required to ϕ , but unable to ϕ *because* one is required to ϕ , because one does not know that one is required to ϕ . So the causal connection desideratum is likewise not a genuine desideratum.

Surprisingly, perhaps, anti-luminosity considerations also undermine the second possibility desideratum (and the first one, for that matter). The issue here isn’t that rival views face the same problems with it that dilemmism does. A view according to which *R* expresses a requirement but *T* doesn’t have the problem of issuing jointly unsatisfiable requirements in conflict cases. Likewise for a view that says that *T* expresses a requirement, but *R* doesn’t. Rather, the problem is that *even if* one is able to ϕ , there is no legitimate sense in which one can really be said to have been *guided* in deliberation by the fact of a requirement to ϕ if one is not in a position to know about it. To see why, suppose again that you are in the nursing scenario, and so you are subject to a rule stating that if one of your patients is in severe pain, they should be given morphine. This time, however, suppose that whilst there is morphine available on this occasion, you are entirely unaware that the patient is in pain (perhaps they are in such a bad way that there are no external signs of their internal torment). Whilst you *can* do what is required of you - i.e. give the patient an injection of morphine - given that you are not in a position to know that this is what is required of you, you can’t be *guided* in your deliberation about what to do with the patient by the requirement. If you happen to do what is required of you by giving the patient morphine anyway, that will be a *coincidence*: a case of conformity, but not compliance.³⁷ In order to be guided by a requirement, you have to *both* know what it instructs you to do, *and* be able to do it. If either

³⁴ Readers who are on the fence about the anti-luminosity argument are welcome to read much of the argument that follows as conditional: if no non-trivial condition is luminous, then (as we will see) many of the objections one might have to dilemmism lack bite.

³⁵ Thus, access internalism is an untenable position.

³⁶ Srinivasan (2015) also makes this point.

³⁷ Smith (2012) also makes this point

element is missing, then it's not as though you can be somehow partially guided by it - you get no guidance whatsoever.³⁸

The upshot of this is that the first and second possibility desiderata are left unmotivated. In cases in which you are unaware that you are required to ϕ - and anti-luminosity considerations dictate that there must be such cases - you cannot be guided in your deliberation by the fact of the requirement even if you are able to conform with it. But in that case, what could the motivation be for thinking that, in order to be adequately guiding, a normative theory must only issue requirements that you are able to fulfill? It can't be because the theory is thereby guaranteed not to give rise to cases in which you cannot use it for guidance. Such cases are inevitable if no non-trivial condition is luminous. But it is very hard to see what *else* the motivation could be. That's not to say that there are no good reasons to think that *OIC* and *A-OIC* are true. There are a number of reasons why one might want to embrace them, which I'll discuss in §6. The point is rather that once we reject the transparency desideratum, they aren't motivated by *guidance* considerations in particular.

§5.3. Degrees of Guidance

Things aren't looking great for the guidance objection so far. But before we proceed it will be useful to step back a bit and look at the idea of adequate guidance from a slightly different point of view. So far we've taken it to be a binary, all-or-nothing, matter. The assumption has been that if a normative theory throws up a single case in which you cannot use its prescriptions to guide your deliberation, then it must be cast to the flames. This is far too simplistic. The followability of a normative theory isn't an all-or-nothing matter. It comes in

³⁸ Might it be said that even if you don't know that you're required to ϕ , but you believe that you are, and so you ϕ anyway, then you have still been guided by the requirement? I don't think so. What you've been guided by is your *belief* that you are subject to the requirement. And this is still the case even if it turns out to be true that you are required to ϕ . Perhaps instead it could be argued that even if it isn't a desideratum on adequate guidance that whenever one is required to ϕ one is in a position to know about it, it is nevertheless a desideratum that one is in a position to *truly* believe it? I'm skeptical, but even if it were, dilemmism would satisfy this desideratum, for when you are in a conflict case you are in a position to truly believe that you both ought to believe that P and ought not believe that P. You have no reason to believe that, of course, but that's a different matter; you can still *do* it. Perhaps instead could it be argued that even if it is not a desideratum on adequate guidance that you know that you are required to ϕ , it is nevertheless a desideratum that it is *probable* on your evidence (to some degree n) that you are required to ϕ ? I'm still skeptical. Williamson (2014) argues - again, convincingly in my view - that one can know that P even though it is arbitrarily improbable short of 0 on one's evidence that one knows that P. And as Hawthorne and Srinivasan (2013) point out, this argument will extend to all non-trivial conditions. For any non-trivial condition C, it can obtain even though it is arbitrarily improbable short of 0 on one's evidence that it obtains.

degrees. This is because normative theories don't traffic only in one-off requirements. More often they take the form of general rules articulated by conditionals with antecedents describing particular circumstances, and consequents issuing instructions about what to do in these circumstances. Consider, for instance, the morphine norm:

Morphine Norm: When a patient is in severe pain, they should be given an injection of morphine.

The antecedent of *Morphine Norm* describes a circumstance and the consequent issues an instruction about what to do if that circumstance arises. When a norm takes this form, the degree of followability or guidance that it exhibits is a function of the *extent* to which it can be used for guidance in one's deliberation. That is to say (roughly) *how frequently* one can use the rule's instruction as a guide in deliberation when the circumstances described in the antecedent obtain. *Morphine Norm* is not perfectly luminous: it is possible for a patient to be in severe pain even though you're not in a position to know about it. So if the transparency desideratum were a genuine one, then *Morphine Norm* would be inadequate. But the norm might nevertheless display a high degree of guidance, even in the transparency sense, for one may *usually* be in a position to know when a patient is in severe pain, and thus *usually* in a position to use the norm to guide one's deliberation about what to do. *Morphine Norm* may likewise not be perfectly guiding in the possibility and causal connection senses. There may be cases in which there is no morphine available, in which case one cannot use the norm for guidance in the possibility sense of that notion, and, again, if it is possible for a patient to be in severe pain even though you are not in a position to know about it, then there are cases in which you cannot use the norm for guidance in the causal connection sense of the notion. For all that, it may nevertheless display a high degree of possibility-guidance and causal-connection-guidance. If morphine is available most of the time, and you are in a position to know when a patient is in severe pain most of the time, then you can use it to guide your deliberation most of the time.

Once we move away from the simplistic idea of adequate guidance being an all-or-nothing matter and towards a more sophisticated degree-based approach, there is both good news and bad news for dilemmism. The good news is that it displays at least *some* degree of followability and guidance (quite a high degree, actually). As I pointed out earlier, when things go well - and most of the time they do - you can straightforwardly and unproblematically use the prescriptions of dilemmism to guide your deliberation about what doxastic attitude to form towards P. Suppose you're not in a conflict case, and suppose that you know that P is true, for example. If so, then dilemmism tells you that it's fine to go ahead and believe that P (and doesn't tell you to not believe that P at the same time). Since you know

what it requires of you, you can use it, in the transparency sense, to guide your deliberation about what to do (i.e. believe that P). Since you are able to do what it requires of you, you can use it in both the first and second possibility senses to guide your deliberation about what to do. And since it satisfies both of these conditions, you can believe that P *because* it requires you to do so. Similarly, suppose - and this is another normal case - that your evidence doesn't favour P over not-P. Let's say you are about to flip a fair coin and you're thinking about whether it will land heads or tails. In that case, the dilemmic view tells you to suspend judgement on P. And again, this is something you are able to do, in a position to know that you are required to do, and able to do *because* you are required to do it. So again, you can straightforwardly and unproblematically look to the prescriptions of dilemmism for guidance here. The dilemmic view may not be perfectly guiding, but it is guiding most of the time.

It is also guiding in a broader sense. Misleading evidence can make it rational to believe a falsehood - as is the case in *MUG* - but one ought in general avoid it as far as possible. The requirements posited by dilemmism combine to instruct you to do so. Misleading evidence makes it more likely that you will find yourself in a conflict case in which you cannot fulfill your joint requirements to be rational whilst at the same time believing only truths. Since you ought to fulfill your epistemic requirements, you can use your knowledge of the requirements of dilemmism to guide your deliberation about whether or not to engage in practices - seeking out reliable sources, and avoiding unreliable sources, for instance - that will lead you to maximise your chances of steering clear of it.³⁹

It is also good news for dilemmism that, as Srinivasan (2015) has pointed out, our ordinary understanding of what it is for a rule or norm to be adequately followable or guiding does not require it to be perfectly so. Srinivasan asks us to consider a Seder norm: when setting the table for Passover, one ought to set as many places as there will be Seder guests, plus one. As she points out, by everyday standards this is a useful, informative, and guiding, norm. But it is clearly not perfectly guiding. You might not know how many people will be attending (sometimes people don't show up), or you might not have enough cutlery (the kids have hidden it for a laugh). All the same, it hardly follows that the norm is useless; most of the time you *can* use it as a guide. The same goes for the set of norms posited by dilemmism.

But there is also bad news. The demand for perfectly guiding norms may be quixotic, and our ordinary standards may not make that demand in the first place,

³⁹ I don't mean to suggest that only dilemmism instructs you to avoid misleading evidence. Any view on which *T* expresses a requirement will also do so. All the same, it is to the credit of dilemmism that it does too.

but surely an adequate normative theory must meet some reasonably high threshold on the spectrum of guidance? And whilst the dilemmic view isn't useless on this score - again, most of the time it works fine - doesn't it fail to exhibit a *sufficient* degree of guidance? Putting the point more plainly: misleading evidence is a pervasive feature of our epistemic lives, so if dilemmism is correct we *frequently* stumble, through sheer bad luck, into epistemic dilemmas. Given this, doesn't it fail to give useful guidance, *all too often*?⁴⁰

I don't think so. I want to say a few things in response to this suggestion. Firstly, insofar as there is a problem for dilemmism here, there is no *more* of a problem than there is for any view that takes *T* to express a requirement. Recall, in order to be guided by a norm on a particular occasion, one needs *both* to know what it requires of one, *and* to be able to do what it requires. If either of these conditions is not met, then one cannot be guided by it *at all* on that occasion. Again, one does not get *partial* guidance from a norm when one is able to do what it requires, but not in a position to know what it requires (and vice versa). Since in conflict cases one is not in a position to know what *T* requires of one, the consequence of this is that every conflict case is one in which, for *both* dilemmism *and* other theories that endorse *T*, one gets no guidance at all. The fact that theories that endorse *T* but not *R* don't have the additional property - which is a feature of dilemmism - of not being guiding in what I have called the 'second possibility' sense makes no difference to the overall degree of guidance they give. So theories that endorse *T* fail to be guiding every bit as often as dilemmism does. (This is because the set of cases in which *R* requires you to believe that *P* but *T* requires you not to believe that *P* is a *subset* of the cases in which *T* requires you not to believe that *P* while you are not in a position to know about it. So the number of cases in which dilemmism fails to give you guidance cannot be greater than the number of cases in which *T*-centric theories fail to give you guidance.)

One might want to take this to be a reason to reject *T*-centric theories of course, but many epistemologists have gone in for them, and those views are taken seriously. If we're going to take them seriously, then the fact that dilemmism falls short of giving useful guidance on a fairly regular basis should not prohibit it from serious consideration, as it does no worse on this score.

Secondly, it's far from obvious that even theories that reject *T* but accept *R* do a much better job than dilemmism here. We flawed human creatures are subject to a host of unconscious biases impeding our ability to rationally respond to the evidence. Most of us unwittingly hold many irrational beliefs. Perhaps not *wildly* irrational beliefs, like the belief that, say, your car has spontaneously transformed into a giant lizard (despite all evidence to the contrary), but beliefs that fail to live

⁴⁰ Cohen & Comesana (forthcoming).

up to *R* all the same. You might, for example, have misinterpreted someone's gestures as expressing irritation, when in fact they were being affectionate, because you unconsciously, and without good reason, think that they dislike you. Cases like this are hardly uncommon. But if so, then even *R*-centric theories fall short of offering useful guidance fairly frequently. Since your unsupported belief that the person doesn't like you is unconscious, the fact that it is irrational for you to interpret their gestures as expressing irritation is opaque to you, so you're not in a position to know that *R* requires you not to interpret them in this way.⁴¹ And again, as with *T*-centric theories, the fact that, unlike dilemmism, *R*-centric theories at least prescribe a course of doxastic action that it is *possible* to act upon in conflict cases won't make a difference here. To repeat: one gets guidance from a norm only if one is both able to do what it requires of one *and* in a position to know what it requires of one. Bearing in mind that each of us has an *unfathomably* large number of beliefs,⁴² can we really be sure that for the average person the proportion of these in which rationality requires them to believe a falsehood greatly outnumbers the proportion that are, unbeknownst to the individual, irrational? I very much doubt it.⁴³

But even if we could, it would be a grave mistake to suppose that one normative theory is superior to another *simply* in virtue of the fact that it offers a higher degree of guidance. Take an epistemology of doxastic attitude formation according to which you should believe whatever you feel like believing. This may exhibit a very high degree of guidance. Even if anti-luminosity considerations dictate that you are not *always* in a position to know what you feel like believing, you may *more often* be in a position to know what such a theory requires of you than you are in a position to know what dilemmism or *T*- or *R*-centric theories require of you (or let's suppose, anyway). But 'believe whatever you feel like believing' is a *terrible* epistemology. Yet if we take degrees of guidance to be the be-all-and-end-all, then we will be compelled to rank it as superior to *T*- and *R*-centric views.⁴⁴ Once we recognise this point it becomes apparent that other

⁴¹ For more discussion of these kinds of cases see Siegel (2013). Srinivasan (2015) makes a similar point to the one I'm making here.

⁴² I have in mind here dispositional rather than occurrent beliefs.

⁴³ That said, unlike with *T*-centric theories the number of cases in which dilemmism fails to give you guidance can be larger than the number of cases in which *R*-centric theories fail, because dilemmism fails to give you guidance in all cases in which you don't know what the rational attitude to take is *and* in cases in which you know that *R* requires you to believe that *P* but (unbeknownst to you) *T* requires you not to. However, the question is whether the set of cases in which dilemmism fails to give guidance will be so much greater than the set in which *R*-centric theories fail to give guidance that we are thereby availed of a reason to reject dilemmism and instead opt for *R*-centric theories. I find it hard to see how such a claim could be substantiated.

⁴⁴ I don't mean to suggest that anyone would think that 'believe whatever you feel like believing' is a good epistemology; I'm not trying to knock down a strawman here. Rather, I only want to use it to illustrate my point.

benefits and costs of respective views must also come into play. Displaying a high degree of guidance may be a virtue of a normative theory, but it is at most a *ceteris paribus* reason to opt for one over another. As I argued earlier, the dilemmic view does well in a number of important respects - it captures the core data, and it does so in a straightforward and simple way. My view is the same as it is when it comes to the conflict between dilemmism and SDL: that even if dilemmism does worse than some of the alternatives on the degrees-of-guidance score, the benefits make the trade off one worth making. But in any case, as we have just seen it is very hard to establish the claim that dilemmism *does* in fact score much worse than the alternatives when it comes to guidance. So I don't think there is a good reason to reject dilemmism here.

§VI. Ability

§6.1. *OIC* and *A-OIC*

Let's turn to the third objection to dilemmism: that in requiring you to believe that P and requiring you not to believe that P in conflict cases, it requires you to do the impossible, which contravenes various widely accepted principles that go under the umbrella name 'ought-implies-can'. You *cannot* both believe that P and not believe that P, so if ought implies can, and if you are required to ϕ then you ought to ϕ , then you cannot be both required to believe that P and required to not believe that P. So the dilemmic view must be false.

We've already looked at one way of trying to motivate this objection: by appealing to the notion of adequate guidance. That argument may not be persuasive, but there are others out there. Perhaps one of these torpedoes the dilemmic view?

I'll argue that none of them does. But before we get going, it will be useful to clarify a few points. Firstly, ought-implies-can has received a lot more discussion in ethics than in epistemology.⁴⁵ For this reason, I'll focus on the arguments that have been put forward for it in ethics, in order to see whether and how they carry over to an application of the principle to the epistemology of doxastic attitude formation. There is a good reason to take this approach. It is easy to imagine that ought-implies-can might be true in ethics (i.e. when it comes to moral requirements) but false in epistemology (i.e. when it comes to epistemic requirements). But I very much doubt that anyone would think the opposite: that

⁴⁵ Though see Ryan (2003, 2015) for critical discussion of various epistemic ought-implies-can principles.

it is true when it comes to epistemic requirements but not when it comes to moral requirements.⁴⁶ So if the arguments for a *moral* ought-implies-can fail to carry over to support an *epistemic* ought-implies-can, we can be fairly confident that there is little reason to think that the epistemic ought-implies-can is well motivated.

Secondly, it is important to remember what the logical form of the relevant ought-implies-can principle is. As I pointed out in §3, since the dilemmic view rejects *AGG*, it is compatible with ought-implies-can as it is usually understood. That is, as having the logical form of *OIC*. Nevertheless, one might think, that doesn't get us very far, for even if dilemmism is compatible with *OIC*, it is not compatible with *A-OIC*, and it is natural to think that whatever considerations speak in favour of *OIC* will also speak in favour of *A-OIC*. As we will see, however, this turns out not to be true. The most compelling arguments for *OIC* do not carry over to *A-OIC*, even if they motivate *OIC*.

Thirdly, one might think that although there is a sense in which you are able to suspend judgement in conflict cases even though all your evidence supports belief (it is logically, metaphysically, physically, and physiologically possible) there is nevertheless also a genuine sense in which you *cannot* suspend on P. With the situation as it is, you may well be unable to *voluntarily* suspend judgement. This line of thought find support from the oft-made observation that our doxastic attitudes are not under our direct control (Williams 1973). If someone offers me £100 to willingly believe that the world is flat, I will not be able to collect the prize. In that case you might think that C-BP is false in conflict cases, at least for some agents. These are muddy waters, and I don't want to wade into the debate over doxastic voluntarism. What I will say, however, is this. Anyone who thinks that one can only be required to take a doxastic attitude towards P if one has the relevant kind of strong voluntary ability to take that attitude to P, and also thinks that we lack that ability, will presumably want to have nothing to do with deontological approaches to epistemic normativity in the first place. I take myself to be engaged in dialogue with those of us who *are* interested in deontological approaches to epistemic normativity. If you've read this far, presumably you are among our number, and I will not discuss the matter any further, except to say that just because we reject *that* kind of ought-implies-can principle, that doesn't mean we won't accept others.

With all that out of the way, we can now turn our attention to the question of whether or not there are good reasons to think that a compelling objection to dilemmism can be found stemming from *A-OIC*. I am not aware of any discussion

⁴⁶ Why? Because the consequences of failing to meet one's moral requirements are usually more serious than the consequences of failing to meet one's epistemic requirements.

of *A-OIC* in the literature, so what I'll do is look at arguments for *OIC*, and ask a series of questions about them: 1. Are they convincing? 2. Do they carry over to motivate *A-OIC*? And 3. Does the argument for them in the moral domain carry over to the epistemic domain?⁴⁷ I'll discuss four such arguments. They aren't the only arguments for *OIC* out there, but they are the ones I think are most plausible.

§6.2. The Best Explanation Argument

What seems to me to be the most powerful argument for a moral ought-implies-can principle is the 'best explanation argument'. I suspect that it is *the* main reason the principle is so popular. The argument is most clearly expressed in Graham (2011).⁴⁸ It goes like this. Consider the following pair of facts:

1. I am not morally required to snap my fingers and thereby end all suffering.
2. If I were able to snap my fingers and thereby end all suffering, I would be morally required to do so.

What's the best explanation of the joint truth of (1) and (2) and pairs like them? The obvious answer is that one is morally required to ϕ only if one can ϕ . *If I could*, I would be morally required to snap my fingers and thereby end all suffering. But since I *can't*, I am not subject to a moral requirement to do so.

This is an interesting argument. But do analogous considerations motivate *A-OIC*? And do they motivate an *epistemic* version of *A-OIC*? That's what's required to extract an objection to dilemmism from them. Let's take these questions in turn.

The (apparent) fact that *OIC* is the best explanation of the joint truth (1) and (2) by itself gives us no reason to think that *A-OIC* is true. We may agree that *OIC* is the best explanation of the joint truth of (1) and (2), but since one can reject *A-OIC* whilst accepting *OIC*, an argument for *A-OIC* cannot appeal to cases in which it is impossible for you to perform a *single* action - like snapping your fingers and thereby ending all suffering - as this argument does. In order to motivate *A-OIC* along the lines of the best explanation argument, what we would need is cases in which each of (a)-(d) is satisfied:

⁴⁷*A-OIC* entails *OIC* (just replace ' ψ ' in *A-OIC* with ' ϕ '), so we need not worry about the possibility of arguments for *A-OIC* that are distinct from arguments for *OIC*. Thanks to Tim Williamson for drawing this to my attention.

⁴⁸ Though Graham doesn't endorse the argument.

- a. One can do each of a set of actions ϕ^1, \dots, ϕ^n
- b. One cannot jointly ϕ^1, \dots, ϕ^n
- c. One is not required to do each of ϕ^1, \dots, ϕ^n
- d. If one could jointly ϕ^1, \dots, ϕ^n , one would be required to do each

There *may* be cases like this in the moral domain. Suppose that a medic in a warzone has 100 patients, all of whom have been poisoned with chemical weapons, and all of whom are at death's door. There is a drug that cures the afflicted, but unfortunately the doctor only has one dose of it. For each patient, the doctor can cure that patient. But, since she only has one dose of the drug, she cannot cure every patient. It seems quite plausible to say that the doctor is morally required to cure at least *one* of the patients, but not morally required for each patient to cure that patient. But if she *could* cure all of them - if she had 100 doses of the drug - then she would be required cure each of them. Thus we have a case pair that fits the description above.

So there may be reason to think that the best explanation argument motivates a moral *A-OIC* principle (note: I'm not saying that it *does*, I think there is more to be said about the matter). But does it carry over to the epistemic *A-OIC*?

I don't think so. The argument gets its intuitive clout from the fact that there is a large number of patients involved. When there are 100 patients but only one dose of the cure the claim that the medic is morally required to give it to each of the patients seems, intuitively, to be ridiculously demanding. But suppose that instead of there being 100 patients, there are only two. Under that supposition there is much less *intuitive* pull to the claim that the medic doesn't have an obligation to each of them (to my mind at least). Indeed, the case then becomes exactly like those in which many have thought agents *do* face moral dilemmas. Now, one might argue that this reveals a problem with embracing the possibility of moral dilemmas in cases in which there are only two mutually incompatible actions: one is forced on a march to absurdity, for if we accept that there can be dilemmas when there are two mutually incompatible actions, then we will have no principled way of blocking the conclusion that there can be moral dilemmas in which there are 100, or 1000, or 1 million, mutually incompatible actions. Hence, the argument would go, the moral version of *A-OIC* must be accepted, and the possibility of moral dilemmas rejected. But crucially, even if one is convinced by this line of reasoning, the same argument cannot be pressed against the possibility of dilemmas when it comes to the epistemology of doxastic attitude formation (I don't think it's a convincing argument when it comes to moral dilemmas either, but that's by the by). It is very hard to see how anything in the epistemology of doxastic attitude formation - and in the dilemmic view in particular - could require one to do a vast number of mutually incompatible things. It's not as if dilemmism claims that one must take 100 (or 1000, or 1 million) different doxastic attitudes

to a single proposition. Nor do I see how there could be cases in which a state-of-affairs like this arises. Dilemmism says you should do two things in conflict cases: believe that P and not believe that P. But clearly it won't do to rely on *these* kinds of cases in the course of mounting the 'best explanation' argument. To do so would be to simply *presuppose* that dilemmism is false; the very first premise of the argument would be that there are no epistemic dilemmas. But that premise would beg the question against the view. So there isn't a good objection to be found here.

§6.3. The Reasons Argument

Struemer (2007) argues for *OIC* like this:

1. If one is required to ϕ , then one has a reason to ϕ
2. One has a reason to ϕ only if one can ϕ
3. Therefore: one is required to ϕ only if one can ϕ

If this argument works in the moral case (which is where Struemer applies it), it presumably also works in the epistemic case. We should expect a certain amount of unity in the structure of moral and epistemic reasons: it would be rather odd if moral reasons were ability-dependent but epistemic reasons weren't. So, do we have a motivation for *A-OIC*, and hence a reason to reject dilemmism, here? I doubt it.

Even if we accept Struemer's argument, it does not motivate *A-OIC*. You *can* believe that P and you *can* not believe that P. What you cannot do is *both* believe that P and not believe that P. But since you are not required to do so (*AGG* is being rejected), there is nothing you are required to do that you cannot do. If so, then even if (1) and (2) are true, it doesn't follow that *A-OIC* is true. So there is no problem for dilemmism here.

It may be tempting to respond to this by saying 'well, if you are required to believe that P and you are required to not believe that P, then you have a reason to believe that P and a reason to not believe that P, and in that case you have a reason to *both* believe that P and not believe that P.' But that would be to assume an agglomeration principle for reasons, and that principle is false. I might have a reason to believe that John stole the diamonds (his fingerprints were found at the crime scene) and a reason to think that Jack stole the diamonds (an eyewitness has picked him out of a line-up). It does not follow that I have a reason to believe that both John and Jack stole the diamonds, for I may know that, whoever the thief is,

he acted alone.⁴⁹

§6.4 The Fairness Argument

It has been suggested to me by some of those with whom I have discussed the dilemmic view that it is somehow *unfair* in virtue of the fact that it requires one to do two impossible things, and that this is a reason to reject it. A similar line of thought can be found in the literature on *OIC* in ethics. Copp (2003) argues that a moral theory that requires one to do the impossible is morally unfair, and that, as a matter of conceptual necessity, morality cannot be morally unfair. Thus, he concludes, *OIC* is true.

The negation of *A-OIC* does not say that can be an act ϕ that a person is required to do and cannot do. Rather, it says there can be sets of actions, each of which is required, which cannot be jointly performed. On the face of it, this seems to be equally unfair. So if this argument motivates *OIC*, then it seems to also motivate *A-OIC*, and there is a problem for dilemmism. But *does* it motivate *OIC*? I don't think so

So far as I have a grip on the relevant notion of fairness at work in the argument,⁵⁰ it is, it seems to me, just as unfair for someone to be required to do something *without their being in a position to know about* it as it is for someone to be required to do something that they are unable to do. And this is true in both the moral and epistemic cases. So if we accept the fairness argument when it comes to the epistemic *OIC* and *A-OIC*, then we must conclude that the epistemic norms of doxastic attitude formation are luminous. But if the Williamsonian anti-luminosity argument is sound (as I think), they are not. So I think we should reject the argument on the grounds that it overgenerates. (Graham 2011 makes a similar point against Copp's argument for a moral *OIC*).

What does seem more plausible is that it would be unfair to *blame or criticise* someone, *qua* epistemic agent, for failing to do what they are unable to do (Graham 2011).⁵¹ I do not think that this is a problem however, since dilemmism

⁴⁹ Robin McKenna has suggested to me that the argument could be revised as follows: 1. If one is required to ϕ one has a conclusive reason to ϕ . 2. One cannot have a conclusive reason to ϕ and a conclusive reason to not ϕ . 3. Therefore, one cannot be required to ϕ and at the same time required to not ϕ . My worry about this revised argument is whether we have a sufficiently good grip on how conclusive reasons function to think that (2) is any more plausible than (3).

⁵⁰ Whatever it is, it is perhaps worth noting that it is clearly not the notion at work in discussions of epistemic injustice such as Fricker (2007).

⁵¹ Even that's not *always* right. If someone who should know better reads Breitbart so often that they find themselves psychologically incapable of believing that Donald Trump is prone to lying, it is quite

does not claim that subjects in conflict cases are epistemically blameworthy come what may. As I said in §2.2, my view is that in conflict cases the subject who believes that P is to be excused for violating T and K . Hence they are blameless.

§6.5. The Blamelessness Argument

The blamelessness argument goes like this:⁵²

1. If one is required to ϕ one is blameworthy for not ϕ -ing
2. One is only blameworthy for not ϕ -ing if one can ϕ
3. Therefore, one is required to ϕ only if one can ϕ

Again, this argument, if it works, works equally as well for A - OIC as it does for OIC , and equally well in the epistemic case as it does in the moral case. So if it works, it's a problem for dilemmism.

My response won't come as a galloping shock. I accept a separation of the deontic and the hypological. One can be excused for failing to fulfil a requirement, and when the excuse is good enough - as it is when one violates T in conflict cases - one is blameless. So I reject premise (1).

§6.6. Stepping Back a Bit

We've looked at a number of arguments for the epistemic A - OIC . None of them is persuasive. Yet I suspect that many will still find dilemmism highly objectionable precisely because it conflicts with that principle. They might want to stick to their guns and insist that the sheer strength of the intuitive pull of A - OIC is a reason to accept the principle and thus reject dilemmism.

I don't put much stock in this line of argument. Whilst dilemmism may be counterintuitive, there are plenty of counterintuitive truths, and we've already seen that the dilemmic view has a number of attractive features. I think the benefits outweigh the cost of counterintuitiveness. However, I also think we can say a bit more here. My view - and this is tentative - is that we should not find it surprising that people will be inclined to find dilemmism counterintuitive, even if it is true.

natural to think that - excuses notwithstanding - they are criticisable and blameworthy, morally *and* epistemically for getting themselves into such a mess. Still, to say that someone is epistemically blameworthy in *that* case is not to say that they are epistemically blameworthy in conflict cases no matter what they do.

⁵² To my knowledge Hintikka (1969) was the first to present this argument.

A recent empirical study by Turri and Blouw (2015) indicates that when an agent blamelessly breaks a rule, we tend to judge, paradoxically, that there was no rule to be broken in the first place (Turri and Blouw call this phenomenon 'excuse validation'). Another study by Chituc *et al.* (2016) indicates that our judgements about whether or not a subject is required to ϕ when they are unable to ϕ are strongly influenced by our judgements about whether or not the subject is blameworthy for not being able to ϕ . When we judge that a subject is blameless for their inability to ϕ , we are inclined to think that they are not required to ϕ , but when they are blameworthy for the fact of their inability, we display no such inclination.

The picture emerging is that we find it psychologically difficult to accept a separation of the deontic and the hypological. If so, then an explanation naturally arises for why principles like *A-OIC* strike many people as intuitively plausible even if no good argument can be found for them, and why dilemmism may seem counterintuitive even if it is correct. *Normally* one is blameless for one's inability to satisfy conflicting requirements, and hence also blameless for not satisfying them. So if we find the possibility of blameless violations of requirements hard to countenance, it should be no surprise that we find principles like *A-OIC* intuitively plausible. But once we recognise that the deontic and hypological must be separated, the intuition should no longer be trusted.

We've looked at a number of different ways in which one might object to dilemmism on the grounds that it violates *A-OIC*. None of them is compelling. There is no good objection to the dilemmic view to be found here.

§VII. Ahead

So far we've looked at three of the main objections to dilemmism. None has much bite. Still, some outstanding issues remain. The two big one's are to develop a deontic logic that allows for the possibility of dilemmas,⁵³ and to show that the dilemmic view is superior to the alternatives catalogued in §2.5. These are both large-scale projects for the future.

But putting them aside for now, are there other areas of epistemology to which the general framework of dilemmism could be fruitfully applied? There may be. In particular, to puzzles concerning assertion, reasoning, disagreement, and higher-order evidence. My goal here isn't to offer anything like a defense or

⁵³ As I said earlier, some logicians have already made a start on this task

endorsement of these further applications - I don't propose that we go around posing epistemic dilemmas willy-nilly - rather, it is only to put them forward as interesting possibilities for further exploration, and to show how the work done so far clears some ground for developing them if we think that a worthwhile enterprise. Let's start with assertion and reasoning.

Norms similar to *T* and *R* feature prominently in debates about the epistemology of assertion. Each of the following has its proponents:

T-Assertion: One ought (epistemically) to only assert truths.⁵⁴

R-Assertion: One ought (epistemically) to be epistemically rational in one's practice of making assertions.⁵⁵

The *MUG* case serves to show how *R-Assertion* can conflict with *T-Assertion*. Suppose that someone walks into the kitchen and asks you where your favourite mug is. According to *T-Assertion* you ought *not* assert that it is in the dishwasher. But given the situation, doesn't rationality require you to assert precisely this? Wouldn't it be irrational *not* to? A dilemmist about the epistemology of assertion would argue that you both ought to assert it and at the same time ought not assert it. That view will face the same objections as those faced by the dilemmic approach to belief formation; that it leads to contradictions and explosions in *SDL*, that it doesn't give adequate guidance, and that it violates *A-OIC*. To the extent that I have answered those objections in the belief formation case, I believe that the same answers will apply in the assertion case. If so, then some progress has been made in clearing the grounds for a dilemmic epistemology of assertion.

Pretty much the same goes for the epistemology of reasoning. Each of the following has its proponents:

T-Reasoning: One ought (epistemically) to treat *P* as true in one's reasoning only if *P* is true⁵⁶

R-Reasoning: One ought (epistemically) to treat *P* as true in one's reasoning if it is epistemically rational to do so⁵⁷

⁵⁴ Unger (1975) and Williamson (2000), amongst many others. Most of those who endorse *T-Assertion* do so because they endorse *K-Assertion*, according to which one ought (epistemically) to assert only what one knows. That includes Unger and Williamson.

⁵⁵ Douven (2006), Lackey (2007), Kvanvig (2009), amongst others.

⁵⁶ Hawthorne (2004), Stanley (2005), Hawthorne and Stanley (2008), Littlejohn (2012). As before, we could add a *K-Reasoning* norm to this list.

⁵⁷ Brown (2008), Fantl and McGrath (2009), Neta (2009), Gerken (2011).

It should be obvious what I'll want to say here about a dilemmic epistemology of reasoning (the same as I do about a dilemmic epistemology of assertion), so I won't bore the reader by repeating myself.

The question of how one ought to revise one's doxastic attitudes in the face of disagreement has generated a huge amount of discussion from epistemologists over the last few years. Hawthorne and Srinivasan (2013) argue for a knowledge-centric approach to it. They favour what they call the 'knowledge disagreement norm' (*KDN*), according to which, in cases of disagreement about whether *P*, where *A* believes that *P* and *B* believes that not-*P*:

1. *A* ought to trust *B* and believe that not-*P* if and only if were *A* to trust *B*, this would result in *A* knowing not-*P*, and
2. *A* ought to dismiss *B* and continue to believe that *P* if and only if were *A* to stick to her guns, this would result in *A* knowing that *P*, and
3. In all other cases, *A* ought to suspend judgement about *P*

However, they also acknowledge that there will be cases in which the rational thing for *A* to do is to flout *KDN*, because her evidence strongly indicates, contrary to fact, that she will know that *P* if she sticks to her guns. They float the idea - albeit rather tentatively - that we should embrace what they call a 'two-state solution' to the puzzle thrown up from the conflicting demands of rationality and *KDN*. The two-state solution is broadly similar to the kind of view that a dilemmist about the epistemology of disagreement may wish to endorse. That is, a view on which one ought to remain steadfast in one's belief in the face of disagreement in cases in which one 'got it right' the first time around by satisfying norms like *KDN* and *T*, but at the same time one ought to be conciliatory and revise one's belief because - sometimes at least - that is what rationality requires in these cases.

Again, such a view will face the objections discussed earlier. As before, insofar I have answered those objections in the belief formation case, I believe that the same answers will apply, to a large extent, in the disagreement case.⁵⁸ If so, then some progress has been made in clearing the grounds for a dilemmic epistemology of disagreement.

⁵⁸ We need to be careful though. How often will the demands of steadfastness and conciliationism pull in opposite directions? The answer might be very often. If so, then a dilemmic epistemology of disagreement may be vulnerable to the objection that it fails to be guiding all too often.

Christensen (2016) argues that rationality can be ‘toxic’ in the sense that there are several rational ideals that cannot always be jointly satisfied. Sometimes agents have misleading higher-order evidence that they have misevaluated the first-order evidence or are poor reasoners. In cases where they have engaged in accurate reasoning to the effect that P is a logical truth, for instance, satisfying an evidentialist ideal, which appears to demand that they take the misleading higher-order evidence into account and so downgrade their credence in P, will require them to be at odds with an ideal of logical reasoning according to which one should assign a credence 1 to all logical truths in order to maintain probabilistically coherent beliefs. Whilst Christensen’s view is that it is rational *ideals* that conflict in cases like this, one might want to take the line that living up to these ideals is an epistemic *requirement*. If so, then Christensen’s view of the ‘toxic’ effects of higher-order-evidence will become a form of dilemmism. Once again, such a view will face the problems discussed earlier. And again, to the extent that those problems have been addressed in the kind of cases I have been interested in, the work done here will carry over to clear at least some grounds for a dilemmist approach to rational toxicity, should we wish to adopt one.⁵⁹

Finally, in extreme cases luminosity failure is in tension with the idea that one ought not be epistemically akratic (Horowitz 2013). Pushed far enough, luminosity failure gives rise to cases in which rationality requires one to believe that P even though it is highly probable on one’s evidence that it is not rational for one to believe that P (see fn. 33). If so, then there will be cases in which rationality requires one to believe that P and at the same time requires one to believe that it is probably irrational for one to believe that P - an akratic combination of attitudes. Must we accept that it is sometime permissible to be epistemically akratic, then? Not necessarily. We might instead adopt a dilemmic view according to which epistemic akrasia is forbidden, and at the same time sometimes required. Again, the work done here clears at least some of the grounds for the development of such a view.

§VIII. Conclusion

My goal here has been relatively modest: to get dilemmism on the table as a serious position in the epistemology of doxastic attitude formation. As we have seen the view has a number of attractive features, the main objections to it aren’t nearly as compelling as one might have first thought, and there are popular rival views - especially *T*- and *K*-centric views that accept luminosity failure - that

⁵⁹ But again, we need to be careful. How often will this view fail to be guiding?

aren't in a position to press many of these objections in the first place. Moreover, dilemmic epistemology is potentially programmatic; there are several other areas in which the general framework may be fruitfully applied. There is much more work to be done, but dilemmic approaches to epistemic normativity should no longer be discarded out of hand.⁶⁰

§IX. References

- Adler, J. 2002, *Beliefs Own Ethics* (MIT Press)
- Alston, W. 1988, "The Deontological Conception of Epistemic Justification" *Philosophical Perspectives* 2: 257–299
- Andric, V. 2015, 'Objective Consequentialism and The Rationales Of 'Ought' Implies 'Can'' *Ratio* 29: 1–16
- Austin, J. L., 1957, "A Plea for Excuses" *Proceedings of the Aristotelian Society*: 1–30
- Boghossian, P. 2003. "The normativity of content" *Philosophical Issues* 13 (1): 31–45
- Bonjour, L. 1980, "Externalist Theories of Empirical Knowledge" *Midwest Studies in Philosophy* 5 (1): 53–74
- Booth, A. R. 2012, "All things considered duties to believe" *Synthese* 187: 509–517
- Brown, J. 2008, "Subject-Sensitive Invariantism and the Knowledge Norm of Practical Reasoning" *Nous* 42: 167–189
- Chan, T. 2013, *The Aim of Belief* (OUP)
- Chisholm, R. 1988, "The Indispensability of Internal Justification" *Synthese* 74 (3): 285–296
- Christensen, D. 2016, "Conciliationism, Uniqueness, and Rational Toxicity" *Nous* 50: 584–603
- Chituc, V. et al. 2016, "Blame, not ability, impacts moral "ought" judgements for impossible actions: Towards an empirical refutation of "ought" implies "can"" *Cognition* 150: 20–25
- Cohen, S. & Comesaña, J. forthcoming, "Being Rational and Being Right" Dutant, J and Dorsch, F (eds.) *The New Evil Demon* (OUP)
- Copp, D. 2003, 'Defending the Principle of Alternate Possibilities: Blameworthiness and Moral Responsibility' *Nous* 31: 441–456
- Douven, I. 2006. "Assertion, Knowledge, and Rational Credibility" *Philosophical*

⁶⁰ Thanks to Daniel Greco, Jessica Brown, Robin McKenna, Tim Williamson, Clayton Littlejohn, Torfinn Huvenes, Finnur Dellsen, Michael Hannon, Sebastian Watzl, Peter Fritz, Juhani Yli-Vakkuri, Maria Baghramian, Olav Gjelsvik, Rowland Stout, audiences at University College Dublin, the University of Oslo, and KU Leuven, and two anonymous referees.

- Review* 144: 361–375
- Driver, J. 1983, 'Promises, Obligations, and Abilities' *Philosophical Studies* 44: 221–223
- Fantl, J. & McGrath, M. 2009. *Knowledge in An Uncertain World* (OUP)
- Feldman, R. 2000, "The Ethics of Belief" *Philosophy and Phenomenological Research* 60 (3): 667–695
- Feldman, R. 2008, "Modest deontologism in epistemology" *Synthese* 161: 339–55
- Fricker, M. 2007, *Epistemic Injustice* (OUP)
- Gardner, J. 1997, "The Gist of Excuses" *Buffalo Criminal Law Review* 2 (1).
- Gerken, M. 2011, "Warrant and Action" *Synthese* 178 (3): 529–547
- Gibbard, A. 1990 *Wise Choices, Apt Feelings* (OUP)
- Gibbons, J. 2013, *The Norm of Belief* (OUP)
- Gowans, C. 1987, *Moral Dilemmas* (OUP)
- Graham, P. 2011, "'Ought' and Ability" *Philosophical Review* 120: 337–382
- Griffin, J. 1992. "The Human Good and the Ambitions of Consequentialism"
- Paul, E. F., Miller, D. & Paul, J.F (eds.) *The Good Life and the Human Good* (CUP)
- Hansen, J, *et al.* 2007, "Ten philosophical problems in deontic logic". G. Boella, L. van der Torre, & H. Verhagen (Eds.), *Normative multi-agent systems. Dagstuhl Seminar Proceedings*, Vol. 07122.
- Hare, R. M. 1963, *Reason and Freedom* (OUP)
- Hawthorne, J. 2004. *Knowledge and lotteries* (OUP)
- Hawthorne, J. & Stanley, J. 2008, "Knowledge and Action" *Journal of Philosophy* 105 (10): 571–590
- Hawthorne, J. & Srinivasan, A. 2013, "Disagreement without Transparency: Some Bleak Thoughts" Lackey, J. & Christensen, D. (eds.) *The Epistemology of Disagreement* (OUP)
- Hintikka, J. 1969. "Deontic Logic and its Philosophical Morals" *Models for Modalities* (D. Reidel)
- Hornsby, J. 2007. "Knowledge in action" Leist, A. (ed.) *Action in Context* (De Gruyter): 285–302
- Horowitz, S. 2013, "Epistemic Akrasia" *Nous* 28 (3): 718–744
- Horty, J. 2003. "Reasoning with Moral Conflicts" *Nous* 37 (4): 557–605
- Hudson, J. 1989, "Subjectivization in Ethics" *American Philosophical Quarterly* 26 (3): 221–229
- Hughes, N. forthcoming, "Uniqueness, Rationality, and the Norm of Belief" *Erkenntnis*
- Hyman, J. 1999. "How knowledge works" *Philosophical Quarterly* 49 (197): 433–451
- Jackson, F. 1991. "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection" *Ethics* 101: 461–482

- Kvanvig, J. 2011, "Norms of Assertion" Brown, J. & Cappelen, H. (eds.) *Assertion* (OUP): 233–250
- Lackey, J. 2007, "Norms of Assertion" *Nous* 41 (4): 594–626
- Lasonen-Aarnio, M. 2010, "Unreasonable Knowledge" *Philosophical Perspectives* 24: 1–21
- Lasonen-Aarnio, M. forthcoming, "Virtuous Failure and Victims of Deceit" Dutant, J and Dorsch, F (eds.) *The New Evil Demon* (OUP)
- Littlejohn, C. 2012, *Justification and the Truth-Connection* (CUP)
- Littlejohn, C. 2013, "The Russellian Retreat" *Proceedings of the Aristotelian Society* 113 (3): 293–320
- Nelson, M. 2010, "We Have No Positive Epistemic Duties" *Mind* 119 (473): 83–102
- Neta, R. 2009, "Treating something as a reason for action" *Nous* 43: 684–699
- Nussbaum, M. 1986 *The Fragility of Goodness* (CUP)
- Marcus, R. B. 1980, "Moral Dilemmas and Consistency" *Journal of Philosophy* 77 (3):121–136
- Moss, S. 2013, "Epistemology Formalized" *Philosophical Review* 122 (1): 1–43
- Rinard, S. forthcoming, "Reasoning One's Way Out Of Skepticism" *Brill Studies in Skepticism*
- Ross, W. D. 1930, *The Right and the Good* (OUP)
- Ross, J. 2010, "Sleeping Beauty, Countable Additivity, and Rational Dilemmas" *Philosophical Review* 119 (4): 411–447
- Ryan, S. 2003, "Doxastic Compatibilism and the Ethics of Belief" *Philosophical Studies* 114: 47-79
- Ryan, S. 2015, "In Defense of Moral Evidentialism" *Logos and Episteme* 6 (4): 405–427
- Sayre-McCord, G. 1986, "Deontic Logic and The Priority of Moral Theory" *Nous* 20: 179–197
- Shah, N. 2003, "How truth governs belief" *Philosophical Review* 112 (4): 447–482
- Shah, N & Velleman, D. 2005. "Doxastic Deliberation" *Philosophical Review* 114 (4): 497–534
- Siegel, S. 2013, "The Epistemic Impact of the Etiology of Experience" *Philosophical Studies* 162: 697–722
- Smith, H. 2012, "Using Moral Principles To Guide Decisions" *Philosophical Issues* 22: 369–386
- Smithies, D. 2011, "Moore's Paradox and the Accessibility of Justification" *Philosophy and Phenomenological Research* 85 (2): 273–300
- Srinivasan, A. 2015, "Normativity Without Cartesian Privilege" *Philosophical Issues* 25 (1): 273–299
- Stapleford, S. 2013, "Imperfect epistemic duties and the justificational fecundity of evidence" *Synthese* 190: 4065–4075

- Stapleford, S. 2015, "Epistemic versus all things considered requirements" *Synthese* 192: 1861-1881
- Struemer, B. 2007, "Reasons and Impossibility" *Philosophical Studies* 136: 351–384
- Sutton, J. 2005, "Stick To What You Know" *Nous* 39 (3): 359–396
- Sutton, J. 2007, *Without Justification* (MIT Press)
- Turri, J. & Blouw, P. 2015, "Excuse Validation: a study in rule-breaking" *Philosophical Studies* 172 (3): 615–634
- Unger, P. 1975. *Ignorance* (Clarendon Press)
- van Frassen, B, 1973, "Values and the Heart's Command" *Journal of Philosophy*, 70 (1): 5–19
- Velleman, D. 2000. "On the aim of belief" Velleman, D. (ed.) *The Possibility of Practical Reason* (OUP)
- Wedgwood, R. 2002. "The aim of belief" *Philosophical Perspectives* 16: 267–297
- Whiting, D. 2013, "Nothing but the truth: on the aims and norms of belief" T. Chan (ed.) *The Aim of Belief* (OUP) pp. 184–204
- Williams, B. 1965, "Ethical Consistency" *Proceedings of the Aristotelian Society, Supplementary Volumes* 39: 103–138
- Williams, B. 1973, "Deciding To Believe" *Problems of the Self* (CUP): 136–151
- Williams, B. 1989, "Internal Reasons and the Obscurity of Blame" *Making Sense of Humanity* (CUP): 35–46
- Williamson, T. 2000, *Knowledge and its Limits*, (OUP)
- Williamson, T. 2014, "Very improbable knowing" *Erkenntnis* 79 (5): 971–999