

Doxastic dilemmas and the method of division

Clayton Littlejohn

Australian Catholic University

It seems that different epistemic norms can come into conflict. We might wonder what happens when they impose incompatible requirements upon us. According to the dilemmic view, they might sometimes generate sets of requirements that cannot be satisfied, ensuring that there is no rationally acceptable way for a thinker to deal with the predicament she's in. After reviewing the case for the dilemmic view, I introduce an alternative framework that accounts for the appearance of dilemma-like conflicts without committing us to the possibility of genuine dilemmas. I argue that this alternative view has some important explanatory virtues that the dilemmic view lacks and can account for the problematic intuitions that seem to support the dilemmic view.

§0. Introduction

Inge tells Agnes something in confidence. Agnes ought to believe her. Question. Could it also be that Agnes shouldn't believe her?¹

¹ Readers be warned. I shall use 'ought', 'should', and 'required' interchangeably because it feels quite unnatural to stick with just one of these terms. If readers worry that they are different, I'm primarily interested in 'ought' and this paper can be read as primarily a paper about what we ought to believe.

I don't think so. The defenders of the dilemmic view say that this is possible.² They think that Agnes might, through no fault of her own, find herself in a situation in which she's required to believe what she's required not to believe. There have been few defences of the dilemmic view in epistemology, but powerful arguments have been offered in support of the view. While I'm sceptical, I don't have an equally powerful argument for their non-existence. Instead of presenting the argument I don't have, I shall try to show that we can account for the data (or very similar data) that supports the dilemmic view in a dilemma-free framework. I shall argue that this framework is simpler than the dilemmic view and seems to explain the same (or very similar) data without positing dilemmas. While the defenders of the dilemmic view might take on board much of what I propose below, doing so might undercut their case for believing in dilemmas in the first place.

§1. Doxastic dilemmas

The situations that interest us are ones in which each of the following is true:

1. A epistemically ought to believe p .
2. A epistemically ought not believe p .
3. It's not A's fault that she's in a situation in which (1) and (2) would be true.

If (1) and (2) obtain because some epistemic norms required us to believe and not to believe, we'd have an epistemic dilemma (i.e., a doxastic dilemma that involves a conflict between epistemic requirements).³ I've included (3) because I think the

² For a defence of the possibility of epistemic dilemmas of this kind, see Hughes (2019, this volume) and Leonard (2020).

³ As Tessman (2017: 27) reminds us, in a genuine dilemma neither (1) nor (2) ceases to be true because the other is true. There might doxastic dilemmas that don't

interesting conflict cases are the ones that aren't a mess of a thinker's own making. Dilemmas are interesting, in part, because they give us reason to doubt that we can (in principle) always respond to the situations we face so as to avoid blame or fault. Suppose someone knows that she's in a conflict case (i.e., a case in which these three conditions hold). Given the *prima facie* plausible assumption that (barring duress) we can be blamed for responding to the situations we're in in ways that we know we shouldn't, it seems that there's no way for her to avoid being blameworthy. I find the idea that blame might be inescapable quite disturbing. If you've been blamed for something and you don't see why it's fair or fitting, you might reasonably ask, 'But what could I have done differently?' You might expect a better answer than, 'Nothing'. We often blame or tell people that they shouldn't have done something because we want them to do better next time, but this presupposes that there was a better option. The dilemmic view denies that this is always so.

Are there any genuine doxastic dilemmas? Hughes (2019) argues that there are. He draws our attention to cases where these (putative) norms come into conflict:

NO FALSEHOODS: One ought (epistemically) to
never believe falsehoods.

EPISTEMIC RATIONALITY: One ought
(epistemically) to be epistemically rational.

Suppose that all the evidence indicates that the mug in the dishwasher is your favourite mug (i.e., it's the one that says, 'You've been nominated for a 2017/18 teaching excellence award!' because your employer wanted to remind you that you didn't win.) Unbeknownst to you, a perfect replica of your mug was slipped into your dishwasher and your mug was taken away to parts unknown. After a careful search of your kitchen for this mug, you see your mug's double in the dishwasher

involve conflicting epistemic norms. For a discussion of non-epistemic factors that might help bring about an impure doxastic dilemma, see Basu and Schroeder (2019) and Rinard (2017).

(REPLICA). On the one hand, it seems irrational *not* to believe that this mug is your mug. (Add whatever details you need to make this plausible.) So, if you don't believe this mug is your mug, you violate EPISTEMIC RATIONALITY. On the other, if you believe this mug is your mug, you violate NO FALSEHOODS.

In situations like this, you either believe, disbelieve, or suspend. EPISTEMIC RATIONALITY rules out disbelief and suspension. NO FALSEHOODS rules out belief. As those options are exhaustive, each option is prohibited. You wouldn't necessarily be at fault for finding yourself in a situation like this, so it seems that *if* these are genuine epistemic norms, there are doxastic dilemmas.

§1.1. Responses

Let's consider some responses.

§1.1.1 *Logical considerations*

Some might argue (1) and (2) cannot both hold because we can derive a contradiction from the assumption that they do. If (1) is true, the argument goes, (2) must be false because A is permitted to believe what she ought to believe and A is permitted to believe iff it is not the case that A ought not believe.

The defenders of the dilemmic view won't be moved by this argument. While we can treat 'may' or 'permitted' as the dual of 'ought' (i.e., accept that PA iff $\sim O \sim A$), the defenders of the dilemmic view say we shouldn't assume (D), the assumption that OA implies PA.⁴ Asserting (D), they'll say, comes close to flatly denying that dilemmas are possible.⁵ What (D) tells us, in effect, is that there's always some

⁴ See Conee (1982) for a discussion of this assumption.

⁵ See Goble (2009), Hughes (2019), and van Frassen (1973).

permissible response to any situation we might find ourselves in. Dilemmas are supposed to be cases in which there's no way forward that isn't forbidden. If we're agnostic about dilemmas, we ought to be agnostic about (D).

Some might argue that we know that (1) and (2) cannot both be true because (1) and (2) lead to deontic explosion. We can derive the problematic consequence that everything is required from (1) and (2) and these two assumptions:

(DIST) If $\text{Nec}(A \rightarrow B)$, then if OA then OB

(AGG) If OA and OB then $O(A \& B)$.⁶

Goble (2009) shows us that deontic explosion can be avoided by restricting (DIST) as follows:

(PDIST) If PA and $\text{Nec}(A \rightarrow B)$, then if OA then OB .

(DIST) and (PDIST) tell us how deontic necessity transmits from the deontically necessary to the deontically necessitated (in the sense relevant to the interpretation of 'Nec'). The rationale for rejecting (DIST) and adopting (PDIST) is that distribution should only happen when a required response is unconflicted (i.e., it is not prohibited). This is in keeping with the spirit of the dilemmic view. Goble's framework blocks deontic explosion.⁷ I doubt the dilemmic view can be refuted on straightforward logical grounds.

⁶ Suppose (1) and (2). If A believes p , it follows that A believes p or $A\phi$ s (for any ' ϕ ' you like). If A doesn't believe p , it follows that A doesn't believe p or $A\phi$ s. Because of this, (DIST) tells us that $O(A \text{ believes } p \text{ or } A\phi)$ and $O(A \text{ doesn't believe } p \text{ or } A\phi)$ follows from (1) and (2) respectively. By (AGG), $O((A \text{ believes } p \text{ or } A\phi) \& (A \text{ doesn't believe } p \text{ or } A\phi))$. Since that's equivalent to $O(A\phi \vee (A \text{ believes } p \& \text{ doesn't believe } p))$, which is equivalent to $O(A\phi)$. Thus, given (AGG) and (DIST), we can derive $OA\phi$ (for any ' ϕ ') from (1) and (2). The argument is taken from Goble (2009).

⁷ Hughes (2019) denies (AGG) and in so doing gives us a way to avoid deontic explosion without denying (DIST). I don't discuss views that reject (AGG) and

§1.1.2 Normative transparency

Some philosophers believe REPLICA is not a genuine dilemma because Agnes couldn't know that these norms come into conflict in her present situation. This is because, the objector assumes, an agent should be aware of the conflict in a genuine dilemma.⁸ (Think, for example, about Sophie's choice in *Sophie's Choice*.)

I'll say two things in response. First, I doubt that it's a constraint on the possibility of the relevant kind of conflict that the subject can justifiably believe they're in a conflict case. I would deny that individual requirements only apply to agents who can know or justifiably believe such requirements apply to you in the case you're in:

LUSTROUS NORMATIVITY: If A ought to ϕ if C obtains (where C is the relevant norm's application condition), C obtains only if A can rationally believe C to obtain.⁹

If the individual requirements aren't lustrous, it's hard to see why the conflict of such requirements would have to be. Making conflicts salient makes for good drama, but it's bad theory to insist that all requirements must be accessible.

preserve (DIST) because, as Goble (2009) argues, such views appear to be too weak to account for some important data.

⁸ Nick Hughes told me that he had heard this point many times in conversation.

⁹ For defences of access requirements in the neighbourhood of LUSTROUS NORMATIVITY, see Kiesewetter (2017), Lord (2018), and Smithies (this volume). For criticism having to do with anti-luminosity, misleading evidence about norms, and normative uncertainty, see Dutant and Littlejohn (2018), Hughes (2018, 2021), and Littlejohn (2012, 2014).

LUSTROUS NORMATIVITY is false if NO FALSEHOODS is a genuine norm. It's less obvious but no less true that it's false if EPISTEMIC RATIONALITY is a genuine norm. Suppose that we cannot rationally believe p unless we're warranted in having a sufficiently high degree of confidence in p (i.e., one that is greater than threshold t where $1 > t > .5$). This threshold presents two problems for LUSTROUS NORMATIVITY. A thinker might be required by EPISTEMIC RATIONALITY to believe p even if she's not in a position to rationally believe that she's warranted in believing that she's warranted in believing that she has a sufficiently high degree of confidence. (This would be a case in which A ought to believe p but it's not the case that A ought to believe or may believe that she is required to believe p .) In another, a thinker might not be warranted in having a sufficiently high degree of confidence in believing p when she isn't warranted in having a sufficiently high degree of confidence in believing that. (This would be a case in which A ought not believe p where A is not permitted to believe that she's required not to believe p .) If rationality allows for any uncertainty about how well supported our beliefs are, LUSTROUS NORMATIVITY fails.¹⁰

The second problem with trying to dispense with our putative dilemma on the grounds that dilemmas have to be transparent is that there are transparent conflicts between EPISTEMIC RATIONALITY and NO FALSEHOODS. Consider an example.¹¹ An eye doctor has you look into a device that flashes slides and asks you to identify the number or letter you see. In each case, the letter or number appears quite clearly to you. The exercise is a bit tedious as it goes on for quite some time without any slide seeming to pose any difficulty. Your doctor says that you did very well and won't need glasses. Suppose you remember each answer and how things seemed when

¹⁰ For discussion of these levels connections, see Das (n.d.), Dorst (2020) and Williamson (2019). It is not clear whether the authors who endorse views in the neighbourhood of LUSTROUS NORMATIVITY defend the view that we can be rationally certain what support our beliefs have.

¹¹ The example is from Dutant and Littlejohn (2021).

you gave that answer. When you ask your doctor how well you did, he tells you that you only got one wrong (EYE EXAM).

It might be that no particular answer stands out as the one that you'd be most likely to get wrong, so I'm tempted to say this. In each subcase, you ought to believe that the letter you seem to remember seeing was the right letter (or believe the doctor's testimony if we want to think of this as another subcase). To do otherwise would be to violate EPISTEMIC RATIONALITY. If you believe in each subcase, you know that you violate NO FALSEHOODS because of the doctor's testimony. While there's no subcase such that you know that this case is a conflict between EPISTEMIC RATIONALITY and NO FALSEHOODS, you know that this conflict exists. Choices have to be made. If you suspend judgment in some cases but not others, we have another violation of EPISTEMIC RATIONALITY. If you believe in no case, we have a systematic violation of EPISTEMIC RATIONALITY. And if you meet the requirements imposed by EPISTEMIC RATIONALITY, you're guaranteed to violate NO FALSEHOODS.¹²

§1.1.3 Normative revision

One way to try to resist the argument for the dilemmic view is by denying that there are positive epistemic duties (i.e., requirements that some thinker believes something as opposed to requirements that they don't believe something). The argument for the dilemmic view assumes that there are falsehoods that we ought to believe where we ought to believe them because this is what's required to conform

¹² I assume that rationality can require you to believe in each case even though you know the collection of judgments contains a falsehood. For discussion of preface cases sympathetic to the line we take here, see Christensen (2004), Dutant and Littlejohn (2021), Easwaran and Fitelson (2014), Makinson (1965), and Worsnip (2016). For an opposing view, see Ryan (1991) and Roush (2010).

to EPISTEMIC RATIONALITY. If we interpret this norm as saying only that certain beliefs are forbidden, we don't have a situation in which (1) and (2) hold.

I see three problems with this strategy. First, suspending judgment is sometimes irrational. If we ought to conform to EPISTEMIC RATIONALITY, we sometimes ought to believe. Second, it seems independently plausible that there are some things that we should believe. It might be that there's no purely epistemic story to tell about why we should believe some propositions when it's fine not to believe others, but it might be the case that Agnes ought to believe what Inge told her. Third, this strategy is hopeless as a general strategy for avoiding dilemmas in the practical sphere. If we're forced to admit that there can be practical dilemmas, we shouldn't dismiss the possibility of doxastic dilemmas.

Instead of interpreting EPISTEMIC RATIONALITY as a purely negative requirement against some beliefs, many epistemologists will reject the idea that there are externalist norms like NO FALSEHOODS. This is not the place to offer a full defence of externalist norms, but I have defended such norms at length, trying to respond to every objection that I've found in the literature.¹³ We shouldn't get in the habit of denying that there are externalist norms.

§1.1.4 Deny nothing.

We could accept doxastic dilemmas.¹⁴ Here's a virtue of this response. There are good reasons to accept EPISTEMIC RATIONALITY and NO FALSEHOODS. We have some reason to think that these norms, in combination, generate doxastic dilemmas. So, we have some reason to believe the dilemmic view.

¹³ See Littlejohn (2012).

¹⁴ See Hughes (2019, n.d.).

This is an option to consider. This lands us with the dilemmic view. There is an alternative proposal to consider, however.

§2. The method of division

Let's consider a simple case (PROMISE). I don't think of this as a conflict case, but I hope it will be instructive. Agnes promised that she'd return Inge's book. This is her obligation. She has two ways of returning it. She might ask the charming Caradoc to drop it off the next time he heads past Inge's flat. She might ask the studious Sebastian to do this instead. This is the only difference between these two options: given her evidence, it's more probable conditional on asking Seb to return the book that the book will be returned than it is on asking Caradoc to carry the book for her. She ends up asking Caradoc to take it. Caradoc returns the book. If it matters, Seb would have returned the book, too.

Question. Did Agnes do what she should have done? Here's a case for thinking that she did. The book was returned. This means the obligation was discharged. Since the obligation was discharged, she didn't fail to do what she ought to have done. Here's a case for thinking that she did not do what she should have done. She had two options. The *only* difference between the options was that the probability of success was lower conditional on choosing Caradoc than it was on choosing Seb. The better-chances principle tells us that if we have two gambles with the same prizes, we should prefer the one that has the greater chance of giving us the better prize.¹⁵ It tells us that Agnes should have chosen Seb, not Caradoc.

¹⁵ See Resnik (1987: 90).

Here's another case (FACT CHECK). Agnes's massive collection of books contains two sets of encyclopaedias. The second set is known to be the same length as the first and to offer answers to the same questions as the first. The only difference is that the second set contains fewer errors than the first. It's compatible with this that Agnes can come to know what she reads in either set. She's asked whether the Spanish flu originated in Spain and, after consulting the first set of encyclopaedias, asserts that it wasn't.

Should she have asserted this having relied on the first set of encyclopaedias? Here's a case for thinking that it was okay for her to have done so. When she believes what she reads, she comes to know. This means that she believed what she ought to have believed and violated no norm in asserting what she did. Here's a case for thinking that she should not have. She believed on the basis of a source she knows to be less accurate. On the assumption that it's desirable to know and undesirable to believe without knowing, Agnes violates the better-chances principle.

In PROMISE and FACT CHECK we have a case for thinking that Agnes's responses were permitted and a case for thinking that she shouldn't have responded as she did. In making the case that her response was permitted, we seem to be relying on this idea that given how things turned out her salient obligation was not left undischarged. In making the case that an alternative was required, we seem to be relying on this observation that she took the riskier route to discharging her obligation without any reason to take the risk. I want a view that does justice to both sets of intuitions.

When I think about the ambivalence that I feel in trying to settle questions like, 'Did Agnes do what she should have done?', I can see the attraction of what Sepielli (2018) calls the 'divider's' view. A divider thinks that there are different readings or interpretations of normative vocabulary. In this case, we have a reading of 'ought' that seems to be sensitive to facts about how things turn out (i.e., an objective reading of 'ought') and one that's sensitive to facts about the agent's epistemic situation, such

as how probable she thinks it is that some means for discharging an obligation will be effective (i.e., a subjective reading of 'ought'). The view nicely accounts for intuitions about cases like PROMISE and FACT CHECK. We're ambivalent because we (rightly) register that there's a sense in which she responded as she should have and a sense in which she didn't.

The divider's view is opposed by the 'debater', someone who thinks that there's only one reading of 'ought' (e.g., an objective one, an information-relative one, a belief-relative one, etc.). I can't here provide much defence of the divider's view. My pitch is this. It does justice to conflicting intuitions and feelings of ambivalence when we consider cases like PROMISE and FACT CHECK without forcing us to say anything too terribly embarrassing. The debater's view, on the other hand, seems to be largely motivated by controversial assumptions about normativity and forces us to choose between intuitions. The divider's view seems preferable.

Let's note one small thing. I don't see any reason why a defender of the dilemmic view shouldn't be a divider and agree with my take on PROMISE and FACT CHECK. The dilemmic view seems to be silent on the question as to whether there can be different readings of 'ought'. For all that's been said, there might be cases where we subjectively should and should not return a book or objectively should and should not assert something. I'll assume that we want to be dividers and argue that the divider's view can account for the intuitions that support the dilemmic view without forcing us to embrace dilemmas.

In what follows, I'll develop the divider's view. In §2.1, I shall offer a general account of the subjective 'ought'. In §2.2, I shall briefly explain why we shouldn't dispense with the objective 'ought'. In §2.3, I offer an account of what we subjectively and objectively ought to believe.

§2.1 Explaining the subjective ought

The divider's view tells us that the arguments for thinking that Agnes did and did not do what she should have done might both be sound. In PROMISE and in FACT CHECK she did something that was objectively permitted but subjectively forbidden. Intuitively, it seems that these are not conflict cases. On its face, it seems that we could all, regardless of what we think about dilemmas, appeal to this view to explain why it seems there's a sense in which she discharged her obligations and a sense in which she failed to do what she ought to have done. In this section, I want to give an account of the subjective ought.

My proposal is that the objective ought should be taken as primary and that we should characterise the subjective ought in terms of some connection to it. I shall assume that normative reasons determine how we objectively ought to respond to the situations we find ourselves in. Think of reasons as facts about the situation and our relationship to the situation. At least some can do their ought-making work even if we're not fully aware of them. For each of our options, we can imagine giant lists of normative reasons that count in favour or count against these options. Insert your favourite complicated story about how some reasons render others redundant, how some can outweigh or neutralise others, and we can imagine the options can be ordered so that some option or options is the most strongly supported. To keep things simple, let's say that we objectively ought to choose amongst the options that are the most strongly supported and ought not choose an option when there's another that is more strongly supported.

Notice that we have two notions at play. The objective ought is an overall notion. The objective ought-makers are contributory. Dividers and their critics typically assume that dividers will characterise the subjective ought in terms of some connection to the overall notion (e.g., that we subjectively ought to ϕ if we believe that we objectively ought to ϕ , rationally believe that we objectively ought to ϕ , etc.), but I think that this is a mistake.

Consider:

K-MAY: If you know that you objectively may A,
you subjectively may A.

We should reject this. In PROMISE, we can fill in the details so that Agnes knows that Seb and Caradoc would succeed in returning the book if asked, but the point remains that Agnes subjectively shouldn't choose Caradoc because of the better-chances principle. If that's right, Agnes might know that she's objectively permitted to give the book to Caradoc but subjectively she shouldn't. The risk matters.¹⁶

This principle fails, too:

C-UGHT: If you are (rationally) certain that
you objectively ought to A, you subjectively
ought to A.

Its failure is instructive. Suppose 10 miners are trapped at the bottom of one of two shafts. You don't know which shaft it is. If you press a button guessing correctly their location, you can block their shaft off and redirect the water to the empty shaft saving all 10. If you press the wrong button, all the miners will die. If you do nothing, both shafts will only partially flood and only one will die (MINESHAFT). As Spencer and Muñoz (2021) observe, it's certain given the setup that you objectively ought to press the button that blocks the shaft that the miners are in.¹⁷ Because 'ought' is upward monotonic (i.e., if you ought to do X and doing X entails doing Y, you ought to do Y), it's also certain that you objectively ought to press a button. But subjectively, you shouldn't press a button.

We should reject principles like K-MAY and C-UGHT because they imply that we subjectively should act as if certain risks didn't matter. Because 'ought' is upward monotone, a complete description of any detailed plan that you objectively ought to implement will entail that you objectively ought to implement each fragmentary part of it (e.g., if you ought to add some molecule containing sodium

¹⁶ For arguments against K-MAY and against the idea that only known reasons can make a normative contribution, see Littlejohn (2019).

¹⁷ See also Fritz (2021).

and chlorine to the soup, you ought to add something containing chlorine to the soup). You might be certain that you objectively ought to implement some fragmentary part of a plan (e.g., that you ought to add something that contains chlorine to the soup) whilst being uncertain about which of the many detailed ways of carrying out the plan would be objectively permitted or required (e.g., you might be able to list any number of things that contain chlorine that you can find in your kitchen and not know which stuff should go in the soup). In such a state, insisting on only acting on plans that include this fragmentary part of the objectively right plan could be exceedingly dangerous.

The lesson to take from this is that information about what we objectively ought to do (which might be limited to fragmentary information about the plans that we objectively ought to implement) might give us little information about what we subjectively ought to do. Remember that if the theory of subjective ought is going to earn its keep, it should allow that the objective oughts and the subjective oughts can attach to different options. This makes it difficult to characterise the subjective ought in terms of an objective overall notion. We should characterise the subjective ought in terms of the objective factors that help determine the objective oughts.

We should characterise the subjective ought in terms of a combination of epistemic factors (e.g., the agent's rational degrees of belief) and the contributory factors that determine the objective oughts (e.g., the weights of the objective reasons that might obtain). This expectabilist approach to the subjective ought secures a role for objective factors whilst providing us with a framework for managing risk. Here is my proposal:

OBJECTIVE OUGHT: You objectively ought to ϕ iff
 ϕ -ing minimises objective wrongfulness.

SUBJECTIVE OUGHT: You subjectively ought to ϕ
iff ϕ -ing minimises expected objective
wrongfulness.¹⁸

The combined weights of the objective reasons determine an ordering of options. It's a fact-relative ordering, if you like, because it's done in light of all the facts. These weights should also tell us how much worse some options are than others. Once we give a cardinal representation of how much worse one option is than another, we can determine the expected objective wrongfulness of an option. The objective normative status of some option will be determined by the weights of the objective reasons that do obtain and the subjective normative status of an option will be determined by the weights of the objective reasons that might obtain (i.e., the obtaining of which isn't ruled out by the evidence) and the probability that they are actual. This is an expectabilist account of the subjective ought because we're using the expectation of a quantity of stuff to determine subjective normative status.

To see how this expectabilist account of the subjective ought might work in practice, let's see how it works in PROMISE. Agnes might conform to the norm or she might fail to. Assume that each failure would be equally wrongful (-3, let's say) and that each way of conforming to the norm wouldn't be wrongful at all (0, let's say). We'll assume that she's certain she can discharge her obligation to return the book, so we'll focus on three states that might obtain in which both couriers or only one courier will return the book:

	C&S	C&~S	S&~C
Ask C	0	0	-3

¹⁸ To simplify things, I'm bracketing complications that arise when we take account of prerogatives. This account is defended in Olson (2018), though there might be differences in the role that we assign to contributory reasons. A similar view is defended by Lazar (2017, 2020), though his full account is more sophisticated since it is formulated to take account of prerogatives.

Ask S	0	-3	0
Keep the book	-3	-3	-3

If either courier would succeed if asked (C&S), we can see that there's no objective reason that distinguishes asking Caradoc or asking Seb, so, on my view, Agnes ought to ask one of them and ought not keep the book for herself. If only one would succeed, there might be objective reason to ask Caradoc and not Seb (if C&~S obtains) or vice-versa (if ~C&S obtains), but what Agnes subjectively ought to do regardless of which state obtains is ask Seb since the expected objective wrongfulness of asking Caradoc is greater than asking Seb whatever probabilities we attach to these states.

The expectabilist view helps us see why K-UGHT and C-UGHT will fail. We should expect K-UGHT to fail if we can be rationally uncertain of what we know. Spot us that assumption and it's not hard to cook up a case in which there's another option available that has less expected objective wrongfulness. For example, I think it's consistent with the original set up that Agnes knew that both Caradoc and Seb would return the book, but provided that it's consistent with this that she's rational in being more confident in Seb's success, she subjectively shouldn't ask Caradoc. In MINESHAFT, the certainty that the way to minimise objective wrongfulness is to press some button is compatible with the fact that pushing no button minimises expected objective wrongfulness.

Let's note an interesting feature of the divider's view. Recall REPLICA. REPLICA is supposed to be a case in which there is some ϕ such that Agnes ought to ϕ and Agnes ought not ϕ . PROMISE wasn't like this. PROMISE is a case in which there is some ϕ such that Agnes is objectively permitted to ϕ even though it is subjectively forbidden. In other words, in REPLICA something is forbidden and required whereas in PROMISE something is permitted and forbidden (in some sense). If we make one small change to PROMISE, it will better match REPLICA. Suppose we stipulate that Seb would have failed to deliver the book to Inge. On this variant, there would be decisive objective reason to give the book to Caradoc and yet giving it to Seb would

minimise expected objective wrongfulness so Agnes subjectively should give it to Seb. Now there's some ϕ that is (in some sense) required and forbidden. The divider's view predicts that there can be conflicts between the objective and subjective ought that look like the conflict in REPLICA. Still, the view doesn't seem to commit us to any genuine dilemmas. There's still one and only one option that you objectively ought to choose, the one that minimises objective wrongfulness. This one isn't objectively prohibited and doesn't conflict with another option that's objectively required. There's still one and only one option that minimises expected objective wrongfulness. This one isn't subjectively prohibited and doesn't conflict with another option that's subjectively required. Once we have something that can account for our intuitions about the fulfilment of objective obligation and the subjective requirement to avoid violating the better-chances principle, we have something that can explain the appearance of conflicts with the same apparent structure as the conflict in REPLICA without committing us to the dilemmic view.

§2.2. Retaining an objective ought

According to the expectabilist account of the subjective ought, possible objective states of the world and the probability of their obtaining determines what we subjectively ought to do. In some ways, this is a familiar idea. In your introduction to decision theory, you're told that possible objective states of the world and the probability that they're actual determine what would be rational to do. What's interesting about the expectabilist view on offer is that it tells us that some of the objective things that matter will be normative reasons, claims, rights, principles. We don't want to say that the objective stuff that matters is limited to evaluative stuff, stuff that figures in explanations as to why some outcome was bad, regrettable, suboptimal, unlucky, etc.

We should distinguish between two expectabilist views. The evaluative expectabilist thinks that the objective stuff that matters must be purely evaluative.¹⁹ The deontic expectabilist thinks that the objective stuff that matters might also be non-evaluative stuff (e.g., objective reasons and their weights, the existence of various rights and claims that don't supervene upon the relevant agent's evidence, the application conditions of principles that might ground duties even though these conditions are inaccessible to the agent at the time of action, and so on.)

One way to understand the difference between these views is to think about how they might rank an agent's options in light of all the facts. The evaluative expectabilist will say that this could be done in terms of evaluative terms only, so their conception of what would be optimal, ideal, or best would have to be understood in terms of some notion of goodness. The deontic expectabilist wouldn't have to explain their orderings in these terms. They might say, for example, that if certain rights would be violated if the agent chose certain options, options that brought about better outcomes might nevertheless not be optimal, ideal, or best. The fact that the facts that would reveal a rights violation were unknown to the agent at the time of action might not matter to the deontic expectabilist when making their ordering, but such rights couldn't exist if the evaluative expectabilists were right so they won't figure in the evaluative expectabilists' understanding of objective orderings.

There are three reasons why evaluative expectabilism seems to be a popular view. First, it might be unclear how we could rationally take account of uncertain non-evaluative stuff and still deliver plausible directives. We might worry, for example, that if some rights exist that don't supervene upon our evidence we'll have to make unreasonable recommendations for how to manage such risks (e.g., that nobody drives again because lives will be lost, that we must free all the prisoners because one of them must be innocent, etc.). Second, it might seem that the view

¹⁹ This seems to be the view of Gibbons (2013) and Zimmerman (2008). See also Lord and Sylvan (this volume) on the evaluative register.

might be attractive only to dividers. Scepticism about the divider's view might explain scepticism about deontic expectabilism. Third, we might not be thinking about the right kinds of cases. Cases like MINESHAFT cause trouble for people who think that the objective ought is the only ought. So, whatever else we need in our normative theory, it seems we need some ought that is information sensitive.

Let's consider two norms that are practical analogues of the norms we discussed in connection with REPLICA:

RESPECT RIGHTS: One ought (morally) to never violate rights.²⁰

MORAL CONSCIENTIOUSNESS: One ought (morally) to do what a rational and morally conscientious agent would do.

Let's consider some cases that are structurally similar to REPLICA. Let's suppose Agnes:

... fails to return a valuable book that belongs to Inge, fails to give a copy of this book to Jaja, but finds a third copy (PROMISE AND GIFT);
... votes to convict the defendant on the basis of solid but misleading evidence (WRONGFUL CONVICTION);

²⁰ Some authors (e.g., Bolinger (n.d.), Zimmerman (2008)) might try to reconcile conflicts between these two norms by adopting a view about our rights according to which we don't have the right not to be harmed or wrongfully convicted but only (e.g.,) a right that others do not impose certain risks upon us. My main complaint about this move is that it doesn't do justice to intuitions about reparative duties and intuitions about the legitimate use of force by bystanders.

... maces a jogger who she mistakes for a
mugger in an effort to defend Inge (IMPERFECT
DEFENCE).²¹

If we imagine that she's the same on the inside as someone who knew that they were doing the right thing, we can elicit the intuition that (in some sense) she's done what she ought to have done and conformed to MORAL CONSCIENTIOUSNESS but on at least a naïve conception of what kinds of rights we have (e.g., property rights, rights against wrongful conviction, rights against others that they don't assault us) she's violated RESPECT RIGHTS. Debaters who reject the divider's suggestion that we distinguish between different readings of 'ought' and want to avoid recognising dilemmas have to choose between such norms, but our deontic expectabilists want to resist that. They want to allow for the possibility that MORAL CONSCIENTIOUSNESS and RESPECT RIGHTS are both genuine norms.

If we focus on what's happening up until the time of action, we might not expect to see much difference between the two expectabilist views. The evaluative expectabilist can say that they account for the importance of rights, duties, normative reasons, etc. by connecting such things to the agent's evidence at the time of action. Deontic expectabilists don't have to say that such things supervene upon the agent's evidence at the time of action. We can see the advantage of the deontic expectabilist view by thinking about various kinds of informational asymmetries. These might be diachronic informational asymmetries (e.g., an agent acquires new evidence after she acts) or interpersonal informational asymmetries (e.g., another

²¹ For a discussion of these cases and their significance for debates between internalists and externalists about justification, permissibility, rightness, etc., see Littlejohn (2012, 2014). There I argued for externalism about justification using so-called 'morally loaded cases'. If we should be guided by our beliefs and we shouldn't violate the rights of others, if their rights don't supervene upon conditions internal to us at the time of action, neither does our right to believe (or the justification of our beliefs).

agent has information that the agent lacks at the time of action). The evaluative expectabilist will say that the information that goes beyond the agent's information at the time of action either has no normative significance or purely evaluative significance. For example, discovering that she's misplaced Inge's book is the discovery that something bad has happened and that Inge has suffered a loss. Discovering that the defendant was innocent, we learn that the defendant has suffered a harm and that we don't now have sufficient reason to continue punishing them. Seeing that Agnes is attacking a jogger but not a mugger we learn that an innocent person is danger.

Deontic expectabilists can say that the additional information that creates the asymmetry can include information about evaluative matters but also rights, claims, reasons, duties, applicable principles or side-constraints, and so on. This information might matter when it comes to questions about what we should do in light of what we've done and what should be done about us in light of what we might do. Think, for example, about PROMISE AND GIFT. Suppose Agnes finds a third copy of the book. She can either keep it for herself, give it to Inge, or give it to Jaja. The evaluative considerations might support each option equally, but it seems that Inge has a claim that breaks the tie. Intuitively, giving Inge the book wouldn't just be nice and wouldn't just be a matter of beneficence. It seems that she has stronger reason to give the book to Inge precisely because she has a duty to compensate for something that she failed to do previously, but that seems to presuppose that this new information about what she did previously revealed a failure to do what she had reason to do, a failure to discharge an obligation, and a failure to see to it that something that was rightfully Inge's was returned. Similarly in WRONGFUL CONVICTION, when some individual or group agent discovers that they've been responsible for a wrongful conviction and then have to decide how to dedicate resources to people who have been wrongly convicted, it seems that taking steps to improve the lives of those that the agent has harmed takes priority over a duty to take steps to improve the lives of those who have suffered similar harms at the hands of others and that, in turn, is because the information about the defendant's

innocence can be seen as information about the wrongs that Agnes helped perpetuate and not just information about the bad states of affairs that eventuated as a consequence of her decision. Finally, when we compare cases like IMPERFECT DEFENCE to cases of genuine defence of others it is striking that bystanders who know what's happening might have different duties about the ways that they might use force to interfere depending upon whether Agnes's mace is being used on an apparent mugger or a real one. In turn, this suggests that she might have a right to non-interference when faced with a genuine threat that she lacks when faced with a merely apparent one, but this suggests that the information outsiders have isn't just information about evaluative matters but also about the existence or non-existence of Agnes's rights that don't supervene upon her evidence at the time of action.

It's not just the practical cases where deontic expectabilism seems to do better accounting for our intuitions. Think about the ways that information asymmetries matter in cases involving testimonial dependence and belief regulation. Let's suppose Agnes...

... tells Inge that the mushrooms growing near her house are safe to eat but she's made a (non-culpable) mistake about the marks of the safe mushrooms and these are poisonous (MUSHROOMS);

... discovers that while nearly all the conclusions that she's drawn about the guilt of the students charged with plagiarism were correct she made a few mistakes (WRONG CONCLUSIONS).

The deontic expectabilist and evaluative expectabilist can agree that there's something bad or defective about the false beliefs or testimony. The world is filled with false beliefs, though, and people often tell other people things that aren't so. In discovering that these mistakes were made, it seems that there's reason for Agnes to give priority to rectifying the mistakes that she made as opposed to, say, rectifying

the mistakes of others or dedicating her time and attention to new questions so that she might form new beliefs that realise some additional goods. Discovering after the poisoning or after the mistake that the results of her efforts were suboptimal, she felt the need to make things right would not just be the felt need to make things better in terms of purely evaluative stuff. She could make things better in any number of ways that didn't involve her rectifying the things that she got wrong. The deontic expectabilist can explain the normative pressure to address these suboptimal outcomes in terms of the violation of a norm that explains why rectifying these errors should take priority over other ways of trying to bring about the good or eliminate the bad. This view fits better with intuitions about the responses that take priority in the wake of discoveries about the harms we've caused or the mistakes we've made.

§2.3. What we (subjectively/objectively) ought to believe

Deontic expectabilists tell us that we subjectively ought to believe in ways that will minimise expected objective wrongfulness. We derive our account of what we subjectively ought to believe from a theory of what we objectively ought to believe and a general proposal about how these oughts are connected. The most familiar account will be a truth-centred account, such as one that takes these to be the fundamental epistemic norms:

NO FALSEHOODS: One objectively ought
(epistemically) to not believe falsehoods.

TRUTH: One objectively ought (epistemically)
to believe truths.²²

²² For defences of a truth-centred approach, see Joyce (1998), Wedgwood (2002), and Whiting (2012). This truth-centred approach, in turn, is often taken to support a Lockean view of rational belief. See Dorst (2019), Easwaran (2016), and Sturgeon (2008) for discussion.

Assuming that every proposition is either true or false (but not both), Agnes is either in the situation in which there's objective reason to believe that the Spanish flu didn't originate in Spain (provided by TRUTH) or objective reason not to believe the Spanish flu didn't originate in Spain (provided by NO FALSEHOODS). If it were equally normatively bad to violate these norms, we would say that Agnes subjectively should believe this proposition iff she rationally assigns a higher degree of confidence to this proposition than its negation. For what it's worth, I think it's plausible that it's worse to violate NO FALSEHOODS than it is to fail to know what we can know, so I think that if you accept a truth-centred view, you should probably say that if Agnes subjectively ought to believe p , she's rational in being more confident in p than its negation.²³ Provided that you don't think it's the end of the world to believe falsehoods, she subjectively ought to believe some things that aren't certain. If the target proposition is true, we have a case in which she subjectively

²³ An anonymous referee expressed scepticism about this claim about one wrong being 'worse' than the other. If the scepticism is about the very idea of one wrong being worse than another, it seems unwarranted. If it's impossible to keep an appointment and save a child's life, it would be worse to prioritise the former over the latter and worse still to decide to do neither. If it's scepticism about the idea that the sins of epistemic commission and omission differ, that's a different matter. It might be that the idea of wrongs differing in gravity makes sense, but the scepticism is about the specific claim that it is worse to violate NO FALSEHOODS than TRUTH (or that it is worse to violate IGNORANCE than KNOWLEDGE). If we thought the weights of these reasons were comparable but neither reason was weightier, we'd have to say that the expected desirability of believing p matches or exceeds that of suspending if it's rational to have at least a .5 degree of confidence in p (or that p is knowable in the circumstance). I suppose the best evidence that we don't think these reasons have the same weight is that most people seem to think that it must be more likely than not that a belief is desirable if it's to be rational to form that belief. If I were to try to represent the believers that strike me as rational, I'd represent them as caring about knowledge and ignorance and displaying a stronger aversion to violating IGNORANCE than KNOWLEDGE. This is not an argument exactly that we should give these different weights, but it seems the same kinds of considerations that convince us to think it is worse to let someone die to keep an appointment than it is to break that appointment to save a life support the claim about the comparative weight of these epistemic norms.

ought to believe what she objectively ought to believe. If the target proposition is false, we get that she subjectively should believe things she objectively shouldn't.

I prefer a knowledge-centred view, so I'd take the objective reasons that determine what we objectively ought to believe to be reasons to conform to these norms:

KNOWLEDGE: One ought (epistemically) to believe p if this is a way of coming to know p .

IGNORANCE: One ought (epistemically) not believe p if one is not in a position to know that p .²⁴

It seems worse to violate IGNORANCE than KNOWLEDGE (e.g., it's worse to believe lottery propositions than it is to fail to believe some bit of trivia reported by a newspaper if you harbour some doubts but nevertheless could have come to know by trusting the source). If this is right, we minimise expected objective wrongfulness in the case of full belief only if we believe when it's more likely than not that we're in a position to know. Rational belief requires a high degree of confidence not just in the truth of the target proposition but in the proposition that we're in a position to know. Provided that we don't assign implausible weights to the reasons to conform to these two norms, it should be possible to rationally believe much of what we read, what we seem to see, what others tell us, what we seem to remember, and so on.

To get a feel for the knowledge-centred view on offer and how it differs from the truth-centred view, think about four cases:

- Lottery cases;
- Preface cases;
- Moorean absurdities;

²⁴ For defences of the knowledge-centred approach, see Williamson (2000). Dutant and Littlejohn (2020; 2024) and Dutant and Fitelson (n.d.) develop and defend a knowledge-centred account of rational belief that's analogous to the Lockean view.

- Williamson's unmarked clock.

In the lottery cases, it's nearly certain that you'll violate TRUTH if you don't believe (and that you won't violate NO FALSEHOODS if you do believe) but certain that you'll violate IGNORANCE if you believe. My own intuitions, for what it's worth, is that there's no sense in which we should believe the truths about the outcomes of fair lotteries if we have only statistical evidence to rely on.²⁵ With preface cases, I'm moved by the point that Foley (1992) makes, which is that it's nearly certain that the (true) beliefs you form in that case will constitute knowledge. That means that on the knowledge-centred account, we objectively and subjectively speaking ought to believe much of what a carefully researched work of non-fiction says. We objectively speaking shouldn't believe the falsehoods contained in the book, but that seems to be in keeping with our intuitions about objective suitability. The truth-centred view agrees on these points, but it fails to capture the intuitive difference in terms of what we subjectively ought to believe in lottery and preface cases where there aren't interesting differences in the ratio of true to false beliefs. Moorean absurdities (e.g., dogs bark but we don't know if they do, this ticket lost but we don't know that yet) are easy to handle in the knowledge-centred framework. Williamson's unmarked clock is an interesting case. If it's a case of improbable knowledge, we have a case in which there are some propositions that we are objectively permitted to believe that we subjectively shouldn't believe. This might explain some of the intuitive resistance to what Williamson (2014) says about such cases. People might rightly register that there's a sense in which we shouldn't believe things we're nearly certain not to know even if such beliefs are objectively suitable or fitting in virtue of the fact that they constitute knowledge.²⁶

§3. Divisions, not dilemmas

²⁵ See Buchak (2013), Nelkin (2000), and Smith (2016) for discussion.

²⁶ See Lasonen-Aarnio (2010) for discussion of unreasonable knowledge.

Let's return to the argument for the dilemmic view:

P1. In REPLICA, Agnes ought to believe that this is her mug (because of EPISTEMIC RATIONALITY).

P2. In REPLICA, Agnes ought not believe that this is her mug (because of NO FALSEHOODS).

C. Thus, there are situations in which, through no fault of her own, Agnes ought to believe and ought not believe the same proposition.

Both premises strike me as plausible, but I don't think they'll commit us to the dilemmic view. Even if we agree that there are readings of (P1) and (P2) are correct, we only have an argument for the dilemmic view if they're both correct on the same reading. Whilst EPISTEMIC RATIONALITY might be true on its subjective reading, NO FALSEHOODS is false on a subjective reading. Whilst NO FALSEHOODS is true on its objective reading, it's false on its subjective reading. The divider's view predicts this and thus explains the appearance of a conflict without committing us to the dilemmic view.

Those who defend the dilemmic view might have no problem with the divider's suggestion that there are different readings of 'ought' and so they might agree that there are different ways to read these norms. Still, they might insist that EPISTEMIC RATIONALITY concerns the objective normative status of our attitudes. In what follows, I'll explain why I think a dilemma-free divider's view is preferable to the dilemmic view. In the course of doing so, I'll explain why dividers should accept EPISTEMIC RATIONALITY only on its subjective reading.

§3.1. On Suspension

Suppose a fair coin has been flipped, it's landed heads, but it's landed somewhere where none of us can see it (COIN). Intuitively, it seems that we ought to suspend judgment. Question. Which view provides the best explanation of this fact?

On a truth-centred divider's view, NO FALSEHOODS and TRUTH determine objective normative status and subjective normative status might be determined by EPISTEMIC RATIONALITY. The subjective standard's content might be extracted from the theory of objective normative status by means of the theory of subjective normative status offered above. If we attach greater weight to the reasons to conform to NO FALSEHOODS than to TRUTH, the view predicts that we subjectively ought to believe (i.e., it's rationally required to believe) when believing minimises expected objective wrongdoing, something it cannot do if the probability of the target proposition isn't greater than .5. That's great. The view also implies that it's never the case that we objectively ought to suspend. That's not great. It also implies that for some chance devices (e.g., the roll of an n-sided die, a lottery) we subjectively ought to believe when the probability of the target proposition is sufficiently high even if we're certain that we couldn't know the outcome. That's not great.

The knowledge-centred divider's view delivers the pleasing verdict that we objectively ought to suspend judgment on the grounds that neither belief nor disbelief would constitute knowledge and that the remaining option must be required. Since it's certain that we couldn't know either way, the view also predicts that suspension minimises expected objective wrongfulness. We get the norms that determine the objective and subjective suitability of suspension for free. This speaks to the explanatory power of the view.

This explanation assumes the validity of a pattern of inference that isn't valid on the dilemmic view.²⁷ It assumes that if we have three options and two are

²⁷ More formally, the problematic pattern is this: (1) Necessarily, you either believe, disbelieve, or suspend. (2) You are required to neither believe nor disbelieve. (3) Thus, you are required to suspend. (4) Thus, you are permitted to suspend. The move from (3) to (4) assumes (D). The move from (1) and (2) to (3) assumes DIST.

forbidden, the remaining option is permitted and required. In our example, we have three options: believing that the coin landed heads (and not believing that it landed tails), believing that it landed tails (and not believing that it landed heads), and believing neither. On the view proposed here, we're can infer that we're permitted and required (subjectively and objectively) to suspend because we're required (subjectively and objectively) not to believe that it landed heads and not to believe that it landed tails.

How do the defenders of the dilemmic view explain our intuitions about suspension (e.g., that it's permitted, that it's required) if this kind of normative argument by elimination is invalidated? We saw above that the defenders of the dilemmic view reject (D). As such, they deny that it follows from the fact that we ought to suspend that we may suspend. Thus, an explanation as to why we *ought* to suspend wouldn't explain why we're *permitted* to suspend. This adds to the explanatory burdens the view must shoulder. The dilemmic view has to find additional explanatory resources to explain why suspension in this case is permitted. Moreover, we know from Goble (2009) that principles like (DIST) have to be restricted so that we cannot say that suspension is required when the alternatives are forbidden without also establishing that suspension is not also forbidden. To explain why we ought to suspend, we cannot simply explain that we shouldn't believe or disbelieve. The defenders of the dilemmic view have to add further norms that tell us when suspension is permitted, something that we wouldn't be able to derive from norms like KNOWLEDGE and IGNORANCE or TRUTH and NO FALSEHOODS. They can do this, of course, but it means that the set of norms needed to explain the intuitive data will be larger than the set of norms that we need on the knowledge-centred divider's view. This suggests, in turn, that the divider's view has greater explanatory power since it explains the data with fewer norms.

It makes sense to drop (D) or revise (DIST) to allow for dilemmas if the dilemmic view provides the best explanation of the apparent conflict in cases like replica, but if, as I think, the divider's view provides an equally good explanation of that apparent conflict by distinguishing between objective and subjective normative

status and offering norms that slot into these two parts of our theory, we might think that the divider's view has explanatory virtues that the dilemmic view lacks.

§3.2. Comparing explanatory virtues

We want our theory of what we ought to believe to tell us the truth. It should be simple but explanatorily powerful. I haven't tried to show that the divider's view delivers different verdicts than the dilemmic view, but I have argued that it's simpler and thus has an attractive explanatory virtue. I shall continue on this theme.

There's a further reason that the divider's view has greater explanatory virtue than the dilemmic view. EPISTEMIC RATIONALITY seems to be a derivative norm. The divider's view tells us that this norm belongs to the theory of subjective normativity and tells us how to derive it by means of the theory of subjective ought. To sustain the argument that there's a genuine dilemma in cases like replica, the dilemmic view has to treat EPISTEMIC RATIONALITY as part of theory of objective normativity.²⁸ But then it cannot just borrow the story that I've given for how to derive EPISTEMIC RATIONALITY because, on that view, the derived norm doesn't

²⁸ Perhaps, more cautiously, I should say that it concerns objective normativity or concerns the same ought that the objective norms are concerned with. This might be some sort of all things considered ought that is distinct from the object and subjective ought but might possibly take account of both of them. My own view is that dividers should probably deny that there is an all things considered ought that takes account of the objective and subjective ought. I don't deny that there are oughts that take account of, say, the demands of morality and of law and tell us what we ought to do. This ought, in my view, might be understood as an objective ought that takes account of many sources but depends upon matters that don't supervene upon the agent's epistemic state and there might be a further subjective all things considered ought that takes account only of that which supervenes upon the agent's epistemic state.

conflict with the objective norm. Without showing us how to derive EPISTEMIC RATIONALITY from more basic norms, the dilemmic view seems forced to try to account for the same intuitions as the divider's view I've sketched with a larger set of non-derivative norms.²⁹ Thus, just as the divider offers a simpler theory of when we ought to suspend, it offers a simpler theory of when we ought to believe.³⁰

²⁹ One possibility might be to try to derive epistemic rationality from more basic norms (e.g., IGNORANCE and KNOWLEDGE) by means of the dispositionalist story found in Lasonen-Aarnio (2010, n.d.). A proper evaluation of this proposal would take up considerable space, but I have two big picture questions about this view. The first is whether this view's verdicts differ from the view sketched here. The second is whether the best interpretation of that view is one on which the derivative oughts (e.g., those connected to good dispositions) and the non-derivative oughts are the same ought. Her view might be developed as a knowledge-centred divider's view.

³⁰ The reason that I think the dilemmic needs EPISTEMIC RATIONALITY to be part of objective normativity is that the dilemmic view needs a conflict between norms that concern the same sense of 'ought'. If it concerned only the subjective reading of 'ought', conflict is avoided. Hughes (personal communication) noted quite rightly that if you believe in some kind of all things considered 'ought', this might restore the conflict. On this way of thinking about things, you all things considered ought to conform to EPISTEMIC RATIONALITY and to some norm like NO FALSEHOODS. I am somewhat sceptical of the assumption that if there is a subjective requirement (e.g., EPISTEMIC RATIONALITY) and an objective requirement (e.g., NO FALSEHOODS), these both should generate some all things considered ought. When I think about a case like PROMISE, I don't think it's right to say that the subject all things considered ought to give the book to Sebastian and all things considered ought to give the book to Caradoc. Giving the book to one but not the other doesn't, in my view, amount to a failure to do something that they all things considered ought to do. Perhaps dividers should be wary of this idea that the all things considered ought 'takes account' of

The second point I want to make concerns the *content* of EPISTEMIC RATIONALITY. The requirement that we ought to have the attitudes that rationality requires is empty. It's like the requirement that we do the right thing. We can try to say what rationality consists in by deriving the requirements of rationality from something else in the way that, say, epistemic utility theorists do, the way that reasons-firsters have tried, or by my suggestion that it's about minimising expected objective wrongfulness.³¹ If the dilemmic view doesn't tell us how the specific requirements that determine the content of EPISTEMIC RATIONALITY is derived, we would seem to have to treat some possibly large set of rational requirements as primitive. If we want to avoid positing a heap of unconnected requirements, a unifying story should be told.

If the defenders of the dilemmic view were to try to derive the content of EPISTEMIC RATIONALITY from objective norms, it isn't clear how they'd go about this. In deriving EPISTEMIC RATIONALITY from KNOWLEDGE and IGNORANCE via my preferred approach to subjective normativity, I have to assume that the reasons to conform to my two objective norms were comparable so that we can say that it's worse to violate some norms and that it's worse by some amount.³² If they're not comparable (e.g., if we couldn't give a cardinal representation the weights of the relevant reasons), we couldn't say that we have some quantity of normative stuff or say that agents ought to maximise some expectation. Our ability to say that some

the objective and subjective oughts in the way that, say, it takes account of reasons to look after those near and dear and reasons to care for strangers in need. This is a large and difficult issue to explore further. For discussion of the all things considered ought, see Hughes (this volume).

³¹ See Dorst (2019), Easwaran and Fitelson (2014), and Pettigrew (2016) for a discussion of epistemic utility theory's attempts to vindicate rational requirements. See Kiesewetter (2016) and Lord (2017) for a reasons-centred approach.

³² See Hughes (2019) and Leonard (2020) for discussions of comparability in the epistemic domain.

cases are clear cases where the risk isn't worth taking and where the risk is worth taking (e.g., preface cases where we vary the size of the propositions at play) is, in my view, evidence that we can make the relevant comparisons. What line should the defenders of the dilemmic view take with regard to the comparability of the reasons associated with the different norms?

To explore this, let's suppose that we could compare the weights of the reasons we have to conform to NO FALSEHOODS and EPISTEMIC RATIONALITY. If they could be compared (e.g., we could assign weights to the reasons there are to conform to these norms), we could tell a systematic story about how we ought to respond to the risk of violating these norms, but it seems we might be able to avoid having to recognise any dilemmatic conflict by saying that we ought to have those attitudes that minimise wrongfulness. What if we instead were to say that the reasons to conform to NO FALSEHOODS and EPISTEMIC RATIONALITY were not comparable? This might help us see how there could be genuine dilemmas, but then we seem to lose the thing we need to use the sorts of tools that I've been hoping to use to decide how to manage the potential risks of violating NO FALSEHOODS or EPISTEMIC RATIONALITY. We couldn't use the weights of the relevant reasons and our probabilities to determine what it would take to minimise expected objective wrongfulness in the face of risks of violating NO FALSEHOODS or EPISTEMIC RATIONALITY.³³

³³ Dividers don't face this problem if they don't think that no falsehoods and epistemic rationality concern the same ought. The problem of comparing the weights of the reasons to conform to these norms arises if they're taken to concern the same ought (and they must concern the same ought if they generate a dilemmatic conflict). An anonymous referee asked, quite reasonably, whether the divider really eliminates dilemmas. Whilst it eliminates a kind of dilemma (i.e., situations in which there is one reading of 'ought' and 'permitted' such that no option is permitted), it leaves room for dilemmas in a looser sense (e.g., situations in which it is impossible to respond without doing something 'forbidden' in some sense or other). If dilemmas are not defined strictly (e.g., situations in which every response is bad or

Let's suppose that Agnes and Inge are at the pond. Inge wants to practice holding her breath and Agnes is supposed to stand by in case she needs any help (DIVING PRACTICE).³⁴ If she stands by watching and does nothing whilst Inge hasn't emerged after ten minutes, she's clearly failed in her responsibilities. If she jumps in to pull Inge from the water after one second, she's clearly failed in her responsibilities. We can ask similar questions about Agnes's beliefs. Should she believe after one second that she ought to jump in? Clearly not. Should she have come to believe that she should have jumped in if Inge hasn't emerged after ten minutes? Certainly. At ten minutes, it's nearly certain that Agnes can know that she ought to have jumped in by then. At the one second mark, it's nearly certain that Agnes can know that she ought not jump in. As time progresses, it becomes less

undesirable in some sense), then of course nobody could deny that dilemmas are possible, but there might be definitions of dilemmas looser than mine and stricter than this according to which the divider view allows for dilemmas in this sense. Here are two ways that the view on offer might differ from a dilemmic view. First, the objective norms are explanatorily more fundamental than the subjective ones. We should not say that we ought to conform to KNOWLEDGE or TRUTH because we ought to conform to epistemic rationality; rather, we should say that the demands of EPISTEMIC RATIONALITY are what they are because we ought to conform to KNOWLEDGE or TRUTH. Second, violations of the more fundamental and more derivative norms might call for different kinds of response. The violation of the fundamental norm might typically involve a change to the state of things that is the source of undesirability. The violation of the derivative norm might involve no such change to the state of things and call for a change in method or procedure instead of a change to the state of things. It seems clear enough in the practical domain that there is the possibility of discharging our obligations as a result of an irrational choice (e.g., using the less reliable courier but nevertheless bringing it about that a book is returned to its rightful owner).

³⁴ Example taken from Dutant and Littlejohn (2018).

certain that Agnes can know that she ought not jump in. We eventually hit a band where she couldn't know that she ought to jump in and couldn't know that she shouldn't. We then hit a band where it seems that she could know she ought to jump in even though it's not very likely that she could know. And it becomes increasingly certain over time that she's in a position to know that she ought to jump in.

As the probability that Agnes can know increases or decreases, the probability that she'd conform to EPISTEMIC RATIONALITY should change, too. As the probability that we're not in a position to know increases, the pressure to suspend increases. At some points, belief might still be required or permitted but some risks won't be rationally taken. If there were objective reasons to conform to EPISTEMIC RATIONALITY and, say, we couldn't compare their weights to the reasons to conform to KNOWLEDGE and IGNORANCE, we wouldn't have a framework that could tell us that we ought to take some small risks of violating some of these norms to serve others but shouldn't take some large risks. Thus, the framework might be friendly to the possibility of dilemmas, but it wouldn't vindicate certain intuitions we have about when it is and when it isn't rational to take risks of violating one of these norms to conform to the other.

§4. Conclusion

The divider's view helps make sense of our ambivalent attitudes towards cases like PROMISE and FACT CHECK. It shows us how we might make sense of our intuitions about what we ought to believe with intuitions about how rational agents will try to see to it that they believe what they ought to by meeting rational requirements like the better-chances principle. The account of objective and subjective normativity gives us a way to understand how apparent conflicts like those that arise in REPLICIA or EYE EXAM might arise even if there aren't any genuine dilemmas. Because the knowledge-centred divider's view handles the cases that might seem to involve a

dilemmatic conflict without positing the existence of dilemmas, I think it is an alternative to the dilemmic view worth considering seriously.³⁵

References

- Basu, R. & Schroeder, M. (2019) Doxastic Wronging. In B. Kim & M. McGrath (ed.), *Pragmatic Encroachment in Epistemology*. Routledge, pp. 181-205.
- Bolinger, R. (2021) The Moral Grounds of Reasonably Mistaken Self-Defense. *Philosophy and Phenomenological Research*. 103, 140-156.
- Buchak, L. (2013) Belief, Credence, and Norms. *Philosophical Studies* 169, 1–27.
- Christensen, D. (2004) *Putting Logic in its Place*. Oxford University Press.
- Conee, E. (1982) Against Moral Dilemmas. *Philosophical Review* 91, 87-97.
- Das, N. (n.d.) The Value of Biased Information. *The British Journal for the Philosophy of Science*.
- Dorst, K. (2019) Lockeans Maximize Expected Accuracy. *Mind* 128, 175–211.
- Dorst, K. (2020) Evidence: A Guide for the Uncertain. *Philosophy and Phenomenological Research* 100, 586-632.
- Dutant, J. & Fitelson, B. (n.d.) Knowledge-Centred Epistemic Utility Theory.
- Dutant, J. & Littlejohn, C. (2018) Just do it? When to Do What You Judge You Ought to Do. *Synthese* 195, 3755-3772.
- Dutant, J. & Littlejohn, C. (2020) Justification, Knowledge, and Normality. *Philosophical Studies* 177, 1593–1609.
- Dutant, J. and Littlejohn, C. (2021) Defeaters as Indicators of Ignorance. In J. Brown & M. Simion (ed.), *Reasons, Justification, and Defeat*. Oxford University Press.
- Dutant, J. and Littlejohn, C. (2024) What is Rational Belief? *Noûs* 58: 333-59.

³⁵ Thanks to Nick Hughes and two anonymous referees for extensive and helpful written feedback on earlier drafts of this paper.

- Easwaran, K. (2016) Dr. Truthlove or: How I Learned to Stop Worrying and Love Bayesian Probabilities. *Nous* 50, 816-853.
- Easwaran, K. & Fitelson, B. (2014) Accuracy, Coherence, and Evidence. In T. Gendler and J. Hawthorne (ed.), *Oxford Studies in Epistemology*. Vol. 5. Oxford University Press.
- Foley, R. (1992) The Epistemology of Belief and the Epistemology of Degrees of Belief. *American Philosophical Quarterly* 29, 111–121.
- Fritz, J. (2021) Akrasia and Epistemic Impurism. *Journal of the American Philosophical Association* 7, 98-116.
- Gibbons, J. (2013) *The Norm of Belief*. Oxford University Press.
- Goble, L. (2009) Normative Conflicts and the Logic of ‘Ought’. *Nous* 43, 450-89.
- Hughes, N. (2018) Luminosity Failure, Normative Guidance, and the Principle ‘Ought-Implies-Can’. *Utilitas* 30, 439-457.
- Hughes, N. (2019) Dilemmic Epistemology. *Synthese* 196, 4059-90.
- Hughes, N. (This Volume) Epistemic Dilemmas Defended. In N. Hughes (ed.), *Essays on Epistemic Dilemmas*. Oxford University Press.
- Hughes, Nick. (n.d.) Epistemology without Guidance. *Philosophical Studies*.
- Joyce, J. (1998) A Nonpragmatic Vindication of Probabilism. *Philosophy of Science* 65, 575-603.
- Kiesewetter, B. (2016) You Ought to ϕ Only if You May Believe that You Ought to ϕ . *Philosophical Quarterly* 66, 760-82.
- Kiesewetter, B. (2017) *The Normativity of Rationality*. Oxford University Press.
- Lasonen-Aarnio, M. (2010) Unreasonable Knowledge. *Philosophical Perspectives* 24, 1-21.
- Lasonen-Aarnio, M. (n.d.) Perspectives and Good Dispositions. *Philosophy and Phenomenological Research*.
- Lazar, S. (2017) Deontological Decision Theory and Agent-Centred Options. *Ethics* 127, 579-609.
- Lazar, S. (2020) Duty and Doubt. *Journal of Practical Ethics* 8, 28-55.

- Leonard, N. (2020) Epistemic Dilemmas and Rational Indeterminacy. *Philosophical Studies* 177: 573-96.
- Littlejohn, C. (2012) *Justification and the Truth-Connection*. Cambridge University Press.
- Littlejohn, C. (2014) The Unity of Reason. In C. Littlejohn and J. Turri (ed.), *Epistemic Norms*. Oxford University Press.
- Littlejohn, C. (2019) Being More Realistic about Reasons: On Rationality and Reasons Perspectivism. *Philosophy and Phenomenological Research* 99, 605-27.
- Lord, E. 2018. *The Importance of Being Rational*. Oxford University Press.
- Lord, E. & Sylvan, K. (This Volume) Beginning in Wonder: Suspensive Attitudes and Epistemic Dilemmas. In N. Hughes (ed.), *Epistemic Dilemmas*. Oxford University Press.
- Makinson, D. (1965) The Paradox of the Preface. *Analysis* 25, 205–207.
- Nelkin, D. (2000) The Lottery Paradox, Knowledge and Rationality. *Philosophical Review* 109, 373-409.
- Pettigrew, R. (2016) *Accuracy and the Laws of Credence*. Oxford University Press.
- Resnik, M. (1987) *Choices*. Minnesota University Press.
- Rinard, S. (2017) No Exception for Belief. *Philosophy and Phenomenological Research* 94, 121-143.
- Roush, S. (2010) Optimism about the Pessimistic Induction. In P.D. Magnus and J. Busch (ed.), *New Waves in Philosophy of Science*. Palgrave-MacMillan.
- Ryan, S. (1991) The Preface Paradox. *Philosophical Studies* 64, 293-307.
- Sepielli, A. (2018) Subjective and Objective Reasons. In D. Star (ed.), *The Oxford Handbook of Reasons and Normativity*. Oxford University Press.
- Simion, M. (2019) Knowledge-First Functionalism. *Philosophical Issues* 29, 254-267.
- Smith, M. (2016) *Between Probability and Certainty: What Justifies Belief*. Oxford University Press.
- Smithies, D. (This Volume) The Unity of Evidence and Coherence. In N. Hughes (ed.), *Epistemic Dilemmas*. Oxford University Press.

- Spencer, J. & Muñoz, D. (2021) Knowledge of Objective ‘Oughts’: Monotonicity and the New Miners Puzzle. *Philosophy and Phenomenological Research* 103, 77-91.
- Sturgeon, S. (2008) Reason and the Grain of Belief. *Nous* 42, 139-165.
- Tessman, L. (2017) When Doing the Right Thing is Impossible. *Oxford University Press*.
- van Frassen, B. (1973) Values and the Heart’s Command. *Journal of Philosophy* 70, 5-19.
- Wedgwood, R. (2002) The Aim of Belief. *Philosophical Perspectives* 16, 267-297.
- Whiting, D. (2012) Does Belief (only) Aim at the Truth? *Pacific Philosophical Quarterly* 93, 279-300.
- Williamson, T. (2000) *Knowledge and its Limits*. Oxford University Press.
- Williamson, T. (2014) Very Improbable Knowing. *Erkenntnis* 79, 971-999.
- Williamson, T. (2019) Evidence of Evidence in Epistemic Logic. In M. Rasmussen and A. Steglich-Petersen (ed.), *Higher-Order Evidence: New Essays*. Oxford University Press.
- Worsnip, A. (2016) Belief, Credence, and the Preface Paradox. *Australasian Journal of Philosophy* 94, 549-562.
- Zimmerman, M. (2008) *Living with Uncertainty*. Cambridge University Press.