

Uniqueness, Rationality, and the Norm of Belief

Nick Hughes, University of Oxford
Appears in *Erkenntnis*, 2019: Vol. 84 (1)

§I. Introduction

This article addresses the question: when are you epistemically permitted to believe that P? I'll argue that the answer is: when it is epistemically rational for you to do so. More broadly, I'll argue that when it comes to the 'coarse-grained' doxastic attitudes of belief, suspension, and disbelief, it is permissible for you to take one of these attitudes to P if it is rational for you to do so. Rationality is the norm of belief, and of coarse-grained doxastic attitude formation in general.

This claim may seem so obviously true as to be in no need of defence. Traditionally the question 'what is it rational for you to believe?' has been taken to be *synonymous* with the question 'what may you believe?'.¹ This is hardly surprising, as there is a strong intuitive pull to the claim that you may believe what it is rational for you to believe.² But the orthodoxy has recently come under fire. A growing number of epistemologists think that permissible belief requires truth³ or knowledge.⁴ Truth and knowledge normers (as we might call them) often disavow the rationality norm and employ an error theory to explain away its intuitive appeal. Roughly stated, the error theory claims that the appeal of the rationality norm is the result of a failure to distinguish between permissibility and excusability. Those who hold beliefs that,

¹ To give just one example, White (2005) moves freely between talk of what it is rational for one to believe and what it is epistemically permissible for one to believe, without further elaboration. This is typical. Even Goldman – arch reliabilist – was moved to revise his thinking in order to respect the apparent equivalence of rationality and permissibility (1986).

² Which Cohen (1984) relies on in the course of presenting the 'new evil demon' argument against externalism, and which many epistemologists have subsequently relied on. I'll discuss new evil demon style arguments in more detail later.

³ Whiting (2010, 2013), Littlejohn (2012).

⁴ Williamson (2000, 2013, forthcoming), Adler (2002), Sutton (2005, 2007), Littlejohn (2013, forthcoming).

whilst rational, are false, or not items of knowledge, should be excused for holding their beliefs, the theory argues. But excusability shouldn't be mistaken for permissibility; even though they are to be excused, it remains the case that the beliefs of non-knowers and false-belief-holders are impermissible. Once we recognise this distinction and its application to the epistemology of belief, the error theorists argue, the appeal of the rationality norm is undermined.

Unless the error theory can be refuted – and I think it is fair to say that it has not been so far – to continue to treat rationality as sufficient for permissibility on the grounds of intuitiveness alone would be an act of evasion. Rationality normers must either tackle the error theory head on by showing why it is mistaken, or, alternatively, put forward a *principled* argument for the view, one that isn't undermined by the error theory. Here I do the latter.⁵

One way to respond to truth and knowledge normers which I will not pursue here, but it is nevertheless worth mentioning before we go any further, is to challenge their assumption that there is a notion of epistemic permissibility available which is distinct from epistemic rationality. As Cohen and Comesaña (forthcoming) point out, 'epistemically permissible' (along with its cognate 'epistemically justified') is a technical term, so it must be given some specified content before we begin theorising about it. One might object – as Cohen and Comesaña do – that truth and knowledge normers have not adequately explained what they mean when they say that S's belief that P is 'epistemically permissible' (or 'epistemically justified'). Cohen and Comesaña argue that, just as acts can be morally permissible and prudentially permissible, they can also be rationally permissible. So talk of rationality should be thought of as picking out a *certain kind* of permissibility. But in that case it makes no sense to claim that it can be rational for someone to believe that P yet not permissible for them to do so, and truth and knowledge normers are simply confused.

Perhaps Cohen and Comesaña are right, perhaps not. I won't go into the matter here. Either way truth and knowledge normers evidently reject this way of thinking, and I want to address them on their own terms. I think that even if we go along with them in distinguishing between epistemic rationality and epistemic permissibility – that is, even if we treat the two as conceptually distinct – we should *still* cleave to the idea that it is epistemically permissible to believe that P if it is epistemically rational to do so. So I'll assume for the sake of argument that Cohen

⁵ In Hughes (forthcoming) I argued that we should reject the knowledge norm. This paper builds on that one by showing how similar considerations can be used to motivate a principled argument for the rationality norm.

and Comesaña are wrong. If they are right, so much the better for the rationality norm.⁶

Of course, this immediately raises the question of how we should think of epistemic permissibility, as distinct from epistemic rationality. Obviously I want to conform with truth and knowledge normers way of thinking about it, in order to avoid talking past them. The most I can glean from their work is that they identify it with what we pick out when we say, from the epistemic point of view, that *S may* believe that *P* (just as epistemic obligation is what we pick out when we say that *S must* believe that *P*). It would be nice to have something a bit more informative, but further details haven't been forthcoming so far. Nevertheless, I'll go along with this way of thinking, somewhat undercooked though it may be. Even so, I'll argue, truth and knowledge normers make a mistake in taking epistemic permissibility to require more than epistemic rationality.

It will be useful to begin with a rough, brief, overview of my overall argument before we go through it in detail. Here (roughly, briefly) is how it all goes. First I'll argue that there are cases – which I call 'clear-cut cases' – in which there is only one doxastic attitude, *D*, that it is rational for you to take towards a proposition *P*. Since a doxastic attitude is irrational if and only if it is not rational, it follows that in clear-cut cases it is irrational for you to take any attitude *other* than *D* towards *P*. I'll then argue that epistemic irrationality is epistemically impermissible. It follows that it is impermissible to take any attitude other than *D* towards *P* in clear-cut cases. I'll then argue that in every case there is at least one doxastic attitude that it is permissible for you to take towards *P*, and I'll show how it follows that rationality is sufficient for permissibility in clear-cut cases. I'll then argue that a theory that takes rationality to be sufficient for permissibility *in general* (i.e. not just in clear-cut cases, but in every case) is preferable on broad theoretical grounds to one that restricts the sufficiency of rationality only to clear-cut cases. A lot of what follows will involve spelling out the details of these claims more carefully, motivating them, and showing how the error theory gains no traction against them. On the matter of the error theory the crucial point is this: rather than relying on the intuitive appeal of the claim that rationality is sufficient for permissibility, the argument relies on the much less controversial claim that rationality is *necessary* for permissibility. As we will see, the error theory does nothing to undermine this claim.

§2 introduces three principles and shows how, if they are all true, it deductively follows that rationality is the norm of doxastic attitude formation in clear-cut cases. §3 argues that we should accept the three principles and show why appeals to the permissibility/excusability error theory gain no traction against the argument that relies on them. §4 argues that the best explanation for why rationality is the norm

⁶ The above notwithstanding, Cohen and Comesaña briefly sketch a line of argument similar to the one I pursue in this article.

of doxastic attitude formation in clear-cut cases is that it is the norm of doxastic attitude formation in general. §5 discusses some implications.

§II. Part One: A Deductive Argument

Now for the details. The view that I ultimately want to argue for is this:

Universalised Rationality Norm: Necessarily: it is epistemically permissible to take a doxastic attitude D to P if it is epistemically rational to take D to P

Call this 'URN'. I intend the claim to be restricted to the coarse-grained doxastic attitudes {belief, suspension of judgement, disbelief}. Before I argue for the universalised norm though, I'll first argue for a weaker claim:

Restricted Rationality Norm: There exists a set of cases – 'clear-cut cases' – in which, necessarily, it is epistemically permissible to take a doxastic attitude D to P if it is epistemically rational to take D to P

Call this 'RRN'. 'Clear-cut cases' are cases in which there is a unique rational doxastic attitude to take towards P of belief, suspension, and disbelief. That is to say, there is *at least and at most* one attitude that it is rational to take of these three.

My approach is as follows. First, I'll show how the restricted rationality norm deductively follows from the conjunction of three principles. Then I'll motivate the principles. Then I'll give an abductive argument for the conclusion that if the restricted norm holds, then so does the universalised norm.

The three principles are these:

Weak Uniqueness: There exist cases – 'clear-cut' cases – in which there is a unique rational doxastic attitude to take to P of belief, suspension, and disbelief

Impermissible Irrationality: It is epistemically impermissible to take epistemically irrational doxastic attitudes

Doxastic Permissibility: In every case there is at least one epistemically permissible doxastic attitude to take towards P of belief, suspension, and disbelief

Hereafter I'll refer to these as 'WU', 'II' and 'DP'. To see how they jointly entail RRN, first assume for reductio that RRN is false. In that case, (1) is true (where 'D' is a variable whose value can be either belief, suspension, or disbelief):

- (1) There is a possible subject S, who is in a clear-cut case, who rationally takes doxastic attitude D to P and impermissibly takes D to P

According to WU, in clear-cut cases there is a unique rational doxastic attitude to take. As a doxastic attitude is irrational iff it is not rational this gives us:

- (2) It is irrational for S to take any attitude other than D to P

Because II states that it is impermissible to take irrational attitudes, and by (2) it is irrational for S to take any attitude other than D to P, when we combine II with (2) we get the following:

- (3) It is impermissible for S to take any attitude other than D to P

And putting together (1) and (3) gets us:

- (4) It is impermissible for S to take D to P and impermissible for S to take any attitude other than D to P

But according to DP there is at least one doxastic attitude that it is permissible for S to take towards P. So we have:

- (5) It is either permissible for S to take D to P or permissible for S to take an attitude other than D to P

And (4) and (5) contradict one another. So if WU, II, and DP are all true then (1) is false. And if (1) is false then there is no possible subject in a clear-cut case who rationally takes D to P but impermissibly takes D to P. And that's just to say that RRN is true: rationality is sufficient for permissibility in clear-cut cases.

§III. WU, II, DP

We just saw that if WU, II, and DP are all true then there is a set of cases – 'clear cut cases' – in which rationality sufficient for permissibility when it comes to doxastic attitude formation (at least for belief, suspension, and disbelief). If that's right, then the restricted rationality norm holds. In this section I'll argue that each of WU, II,

and DP *is* true, and so that we should accept the restricted rationality norm. I'll also show how the error theory mentioned in §1 gains no traction against this argument.

§3.1. Weak Uniqueness

Let's start with WU. Recall, it states that:

Weak Uniqueness: There exist cases – 'clear-cut' cases – in which there is a unique rational doxastic attitude to take to P of belief, suspension, and disbelief

WU is similar to a principle that has received a lot of discussion in recent debates about the epistemology of peer disagreement:

Uniqueness: Given one's total evidence there is a unique rational doxastic attitude that one can take to any proposition⁷

Call this principle 'U'. WU is similar to U but differs from it in some respects. The most obvious of these is that the former makes an existential claim whereas the latter makes a universal claim. WU claims that *in some cases* there is a unique rational doxastic attitude. It doesn't claim, as U does, that in *all* cases there is a unique rational attitude. That means that WU is a weaker principle than U. We'll see shortly why this is important.

Why should we accept that there are clear-cut cases, and so accept WU? Because there obviously are such cases. Here are three:

MUG: S wakes up and goes downstairs to make breakfast. In normal lighting conditions and with a clear head she sees what appears to be her favourite mug sitting on the kitchen table in front of her, where she remembers leaving it last night.

COIN: S is about to flip a coin that she knows to be fair. She is thinking about whether or not it will land tails.

CAR: S parked her car outside her house last night. She is now thinking about whether or not it spontaneously turned into a giant lizard last night. She has no evidence to suggest that this is possible, let alone that it has happened.

⁷ Various wordings of this principle can be found in the literature. I take this one from White (2005).

In each of these cases it is, I suggest, perfectly clear that there is a unique rational doxastic attitude for S to take towards the relevant proposition. In MUG it would be rational for S to believe that P (= there is a mug on the table), and it would be *irrational* for her to suspend judgement on P or disbelieve that P. To suspend judgement on P would be overly cautious and to disbelieve that P would be an egregious instance of disregarding the evidence.⁸ So S is in a clear-cut case: there is a uniquely rational attitude for her to take towards P – belief. In COIN, prior to flipping the coin, it would be rational for S to suspend judgement on P (= the coin will land tails). It would be irrational for her to believe that P and irrational for her to disbelieve that P; her evidence doesn't support taking either of these attitudes. So S is in a clear-cut case: there is a uniquely rational attitude for her to take towards P – suspension. In CAR it would be rational for S to disbelieve that P (= my car has spontaneously transformed into a giant lizard overnight). To suspend judgement on P would be too cautious, and to believe that P would be an egregious instance of disregarding the evidence. Both would be irrational. So there is a uniquely rational attitude for her to take towards P – disbelief. So MUG, COIN, and CAR are all clear-cut cases. Many others could be given of course.

Since WU only makes a weak claim: that there are *some* cases in which there is a uniquely rational doxastic attitude to take, MUG, COIN, and CAR are enough to vindicate it. And indeed, even those who reject the stronger claim made by U typically accept WU. As Kelly says in the course of arguing against U: "...A Permissivist [i.e. someone who denies U] might very well think that there are many cases in which...there is one and one response to the evidence that's the fully rational response" (2014: 298). Schoenfield agrees: "In a nutshell, my view is this: there are non-permissive cases (cases in which there is only one rational way to respond to a body of evidence) and permissive cases (cases in which there is more than one rational way to respond to a body of evidence)." (2014: 194). In summary, WU ought to be uncontroversial.

§3.2. Impermissible Irrationality

We just saw that there are good reasons to accept WU – the claim that there exist clear-cut cases in which there is a unique rational doxastic attitude to take. What about II? Recall, it states that:

Impermissible Irrationality: It is epistemically impermissible to take epistemically irrational doxastic attitudes

⁸ One can suspend judgement on P whilst having a high credence in it. To be clear, my claim is that do to that in this case would be too cautious.

As we saw in §1, some epistemologists explicitly disavow the claim that rationality is sufficient for permissibility and employ an error theory to explain away its appeal. But II only states that rationality is *necessary* for permissibility.⁹ This is far less controversial. Even when it is not explicitly endorsed it is implicitly assumed in a great deal of epistemology. Indeed, it is a near truism. I think this is unsurprising; without it we cannot adequately capture the kind of epistemic wrongdoing that comes with being epistemically irrational. Let's take a concrete case. Suppose in CAR that despite overwhelmingly strong evidence that it has not, S suspends judgement on the question of whether her car has spontaneously transformed into a giant lizard overnight. Even if, by a freak quantum–mechanical miracle, it turns out that such a transformation has in fact taken place, it is patently irrational for S to suspend judgement on the matter. But if II is false, then the possibility remains that suspending is an acceptable thing to do by the norms of doxastic attitude formation. Could this possibly be right? Clearly there is *something* bad about S's suspending on the matter. Yet it doesn't seem enough to say that it is merely suberogatory, inadvisable, or non–ideal. Such descriptions fail to capture the kind and weight of epistemic wrongdoing that comes with failing to live up to the standards of rationality. It's not as though when one is epistemically irrational one is in a situation analogous to that of a person who gets drunk before a job interview: it's not a good idea, but you're within your rights to do it. The natural judgement is rather that it is epistemically *forbidden*; rationality isn't optional, it's a requirement.

Given the intuitive support it receives, then, I think our default position should embrace II in the absence of good reasons to give it up. But perhaps there are such reasons? Some things that Williamson says in a recent paper (forthcoming) suggest that he thinks so. Since Williamson is one of the most prominent advocates of rejecting rationality–centric approaches to epistemic normativity, I will address what I think his concerns are likely to be.¹⁰

Taking inspiration from Lasonen–Aarnio (2010), Williamson argues for a norm 'ODK' which seems to be similar to II. According to ODK you are criticisable if

⁹ The claim that rationality is necessary for permissibility is the contrapositive of II.

¹⁰ Given the standard assumption that knowledge entails rationality it is natural to think that knowledge normers are themselves committed to the truth of II. After all, if knowledge entails rationality, and you are obligated to only believe what you know, doesn't it follow that you are obligated to only believe what it is rational for you to believe? And if so, isn't II true? It turns out that the matter isn't so straightforward. For one thing, the idea that knowledge entails rationality has recently been called into question (Lasonen–Aarnio 2010). But more importantly, knowledge normers are committed to saying that there are cases (like CAR) in which you are obligated to suspend judgement on P, even though it is irrational to do so. The upshot is that whilst they might want to agree that it is impermissible for you to believe that P when it is irrational for you to believe that P, they are not in a position to endorse II in full generality (i.e. as saying that for *any* doxastic attitude – suspension included – it is permissible to take that attitude just in case it is rational to take that attitude). Thanks to a referee for this journal for pointing out the need for clarification on this matter.

you fail to take the doxastic attitude towards P that an agent who has knowledge–conducive doxastic dispositions would take towards P in your circumstances. But Williamson suggests that this norm carries less weight than what he takes to be the 'primary' norm of belief: knowledge. One way of interpreting what he is getting at here is that although it is somehow bad to take irrational attitudes, there are nevertheless cases in which it is permissible (and indeed, obligatory). Applied to the CAR case the thought would be that if P (= my car has spontaneously transformed into a giant lizard overnight) is true, then S should suspend judgement, rather than disbelieving it, even though it would be irrational for her to do so. We can, the argument goes, explain why she is nevertheless criticisable for suspending by appealing to the observation that in doing so she violates ODK. *Normally* her belief that P is false would be an item of knowledge, so although she is obligated to suspend she is criticisable for doing so, because by suspending she manifests a non–knowledge–conducive disposition. If that's right, then it may be argued that in endorsing II and saying that epistemic irrationality is epistemically impermissible we have mistaken impermissibility for mere criticisability (in much the same way as it is alleged that those who think rationality is sufficient for permissibility have mistaken excusability for permissibility).¹¹

Is this a good reason to reject II? I don't think so. Williamson is keen to point out – correctly, in my view – that it follows from some general facts about the structure of normativity that if the knowledge norm view is correct, we will be inclined to criticise those who manifest a disposition to violate it, as S would were she to suspend judgement in CAR. But it is one thing to judge that someone is criticisable for doing something but not forbidden to do it (think again of the person who gets drunk before an interview), and it is quite another to judge that they were forbidden to do it. And the natural judgement remains that S does something epistemically forbidden if she suspends in CAR even when the miraculous transformation has taken place.

Are we prone to mistake criticisable–but–nevertheless–permissible behaviour for impermissible behaviour? It matters, because if we're not, then the fact that the Williamsonian view delivers the result that S is criticisable for suspending, but still permitted to do so, does nothing whatsoever to undermine the standing of the original judgement that she is forbidden to suspend in virtue of the fact that it would be irrational for her to do so. So it is incumbent on anyone taking the Williamsonian line to show that we are prone to make this kind of mistake. If that can't be done, then the thing to conclude here is not that we have a reason to reject II, but rather (insofar as it delivers the result that S is *merely* criticisable for suspending in CAR) that we have a reason to reject the knowledge norm.

¹¹ Williamson's argument here is a development of the idea that we must draw a distinction between 'primary' and 'secondary' norms. That distinction was first drawn by DeRose (2002).

Herein lies the problem for the Williamsonian line: there is, as far as I can tell, no reason to think that we are prone to make the alleged error. Presumably the thought must be that even when a subject does not have an obligation not to manifest a disposition to violate another obligation, O, (rather, they only have an obligation not to *actually* violate O), we have a tendency to mistakenly believe that they do have such an obligation. But it is hard to see what support there is for that claim. For one thing, when it comes to a great many of our obligations it is rather plausible that insofar as one has an obligation O one *also* has an obligation not to manifest a disposition to violate O. Insofar as one is morally obligated not to punch people in the face, for example, it is plausible that one is also morally obligated not to *try* to punch people in the face. Insofar as one is morally obligated not to steal from others, it is plausible that one is morally obligated not to *try* to steal from others...and so on.

Clearly no support for the error-theoretic explanation can be found from *these* kinds of cases. Rather, if any support is going to be found for it, then it must come from cases in which it can be shown that one has a particular obligation O, but no obligation not to manifest a disposition to violate O. Since their rules are a matter of stipulation, games provide good examples of this phenomenon.¹² Within the rules of rugby it is forbidden to tackle a player around the neck (and so obligatory not to tackle a player around the neck). It is not, however, forbidden to manifest a disposition to violate this obligation by *trying* to tackle a player around the neck, provided that the attempt is, as a matter of fact, unsuccessful (it's unsportsmanlike, of course, but that's a different matter – the rules of rugby don't demand sportsmanlike behaviour). So if Billy tries to tackle George around the neck, but fails because George ducks at the last second and Billy goes sailing over the top of her, then whilst Billy is undoubtedly criticisable in virtue of her unsportsmanlike attempt at an illegal tackle, she nevertheless has not done anything that is forbidden by the rules of the game. But – and here is the crucial point – would anyone familiar with the game be prone to mistakenly think otherwise? I very much doubt it. But if not, then the error theory receives no support from this case either. And the same is true, it seems to me, of *all* other such cases in which there is an obligation O but no obligation not to manifest a disposition to violate O. What the error theory requires is cases in which people are prone to mistakenly believe that there is an obligation not to manifest a disposition to violate O when in fact there isn't. But, I suggest, there are no such cases. We just don't seem to be prone to making this kind of mistake. In that case the Williamsonian line is unmotivated: it relies on a hypothesis – that we have a tendency to mistakenly think that there is an

¹² Lackey (2007) appeals to game cases when discussing applications of the primary/secondary norm distinction to theorising about the epistemic norms of assertion. She does so for a different purpose to mine, though. Lackey wants to argue that the primary/secondary norm distinction is a spurious one. I only want to argue that we are not prone to mistakenly conflate the manifestation of disposition to violate an obligation, O, with an actual violation of O.

obligation not to manifest a disposition not to violate a different obligation when in fact there is no such thing – that is not supported by the evidence. In light of this, I do not think that we should be troubled by the challenge it appears to present for II. Given that, as I pointed out before, II has considerable intuitive appeal, we should hold on to it.

§3.3. Doxastic Permissibility

We just saw that there are compelling reasons to accept II. But we are not yet in a position to conclude that rationality is sufficient for permissibility in clear-cut cases. For it may be that, in addition to it being impermissible to take any attitude *other* than the rational one in clear-cut cases, it is *also* impermissible to take the rational attitude (if, for example, the rational attitude is belief, the relevant belief is false, and it is impermissible to have false beliefs). This is where the last principle of the deductive argument, DP, comes in. It is incompatible with the idea that there can be dilemmas of this kind. Recall, DP states that:

Doxastic Permissibility: In every case there is at least one permissible doxastic attitude to take towards P of belief, suspension, and disbelief

There are two good reasons to accept this principle. Firstly, suspension of judgement is normally understood as being the state an agent is in just in case she has considered the question of whether P and neither believes that P nor disbelieves that P. Since one either believes that P or does not believe that P and either disbelieves that P or does not disbelieve that P, on this view it is *impossible*, upon having considered whether P, to fail to take one of the attitudes belief, suspension, or disbelief, towards it. If that's right, then rejecting DP means claiming that there are cases in which there is no epistemically permissible course of doxastic action for an agent to take *whatsoever*. Thus, if DP is false there can be tragic epistemic dilemmas in which there is nothing one can do to avoid falling into epistemic disrepute. That seems very odd indeed.¹³

¹³ As Cohen and Comesaña (forthcoming) point out, given the pervasiveness of misleading evidence, a theory that posits dilemmas in order to rescue the truth or knowledge norms will be committed to saying that each of us faces dilemmas of this kind on a regular basis, which seems even more odd a position to take. On a different note, Friedman (2013) argues against the traditional view that one suspends on P just in case one has considered the question of whether P and neither believes nor disbelieves that P. She argues that it is possible to have *no attitude at all* towards a proposition even after having considered the question of whether or not it is true. Even if Friedman right, however, it's hard to see how that might help those who would reject DP. If each of the attitudes belief, suspension, and disbelief, is impermissible, then taking no attitude towards P would be the only permissible course of action, and would thereby be obligatory. Knowledge normers who reject DP wouldn't be committed to the possibility of epistemic dilemmas in the conventional sense, then, but the view that there can be cases in which one is obligated to take no attitude towards a proposition is, it seems to me, just as odd as a view on which there are epistemic dilemmas. Moreover, there is a fairly strong sense in which normal

The second reason to accept DP comes from reflection on the *raison d'être* of epistemic norms.¹⁴ One of the main reasons epistemic norms are useful is because they enable one to form the appropriate doxastic attitudes with the result one is then able to go on to use these doxastic attitudes in planning, decision-making, and action. In order to form and execute a plan of action one needs both desires *and* doxastic attitudes about how those desired can be fulfilled. To give a simple example, if I desire a glass of water then I have to make use of my doxastic attitudes about the location of water in my vicinity in order to plan and execute a course of action that fulfills my desire. If I have no doxastic attitudes about the location of water in my vicinity, then I simply *cannot* form and execute a plan of action. According to a theory that rejects DP there are circumstances in which I am not permitted to take any doxastic attitude towards P. So I am forbidden in these circumstances to be in a state of mind which it is necessary for me to be in to plan and execute a course of action to fulfill my desires (at least, when the fulfillment of those desires depends on the question of whether or not P is true). The result is that epistemic norms are of no practical use in these circumstances.¹⁵ So an epistemic theory that rejects DP commits itself to the view that epistemic norms do a bad job of fulfilling one of their main purposes – to assist subjects in navigating their way through the world. That's a good reason to accept DP.

In light of these considerations we should accept DP. As I argued in §3.1 and §3.2, we should also accept WU and II. And as we saw in §2 the conjunction of these principles entails RRN. So we should accept RRN. Rationality is sufficient for permissibility in clear-cut cases.

§3.4. The Permissibility/Excusability Error Theory

I just argued that we should accept the principles on which my argument for the restricted rationality norm is based, and so that we should accept the norm. In §1 I touted this argument on its immunity to the permissibility/excusability error theory. It's time to make good on that advertisement.

agents will be unable to fulfil this would-be obligation. So this line of argument runs afoul of the 'ought-implies-can' principle.

¹⁴ See also Hughes (forthcoming)

¹⁵ Note that even if suspension is the appropriate attitude to take towards P, this is still action guiding. If I want water, and I suspend judgement on what locations it is available in my vicinity, then I can use this desire/doxastic attitude combination to motivate seeking out more evidence on the question of where there is water available. By contrast, if I am obligated, as per the rejection of DP, to have no attitude at all, then I will be obligated to be in a state of mind in which I cannot even do this. Relatedly, it should be clear that this argument will still apply even if we reject the tradition account of suspension of judgement and replace it with Friedman's account.

To begin with we need to get clear on exactly what the error theorists argue. Here's how it goes. First they put forward their arguments for the truth or knowledge norms. They then acknowledge that there appears to be a good objection to their views. If one is obligated to only believe truths or things one is in a position to know, then it follows that subjects who rationally believe that P when P is false or unknown hold impermissible belief. And this runs against the intuitive judgement that one is permitted to believe that P if it is rational for one to believe that P. By way of response to this they put forward the error theory, which, they argue, blocks the objection. First, they observe that there is a distinction to be made between permissibly φ -ing on the one hand and excusably φ -ing on the other. Crucially, excusability doesn't entail permissibility, so the fact that one is to be excused for φ -ing gives us no reason to think that one is permitted to φ . They then do two things. They argue that excusability is easily confused with permissibility and that this explains why it would appear intuitive to us that rationality suffices for permissibility even if it doesn't. And they argue that, given the considerations they have offered in favour of the truth/knowledge norms, we should think of subjects who rationally believe that P falsely or non-knowledgably as merely excused, in virtue of their rationality, rather than permitted. Their rationality makes them *blameless*, but that's only good enough for an excuse, not for permission.¹⁶

The purpose of the appeal to the permissibility/excusability distinction, then, is not to *rebut* the rationality norm, in the sense of providing reasons to think that it is false. That's what the positive arguments for the truth/knowledge norms are supposed to do. Rather, the appeal seeks to *undercut* the argument from intuitiveness by showing that the intuition doesn't support the conclusion (we would have the intuition even if the conclusion was false).¹⁷ In doing so, it aims to block the objection stemming from the intuitiveness of the rationality norm. Littlejohn is particularly clear on this point. In the context of employing the error theory against Cohen's reliance on the intuitiveness of the rationality norm in the New Evil Demon argument (1984), he writes: "Once we have a better understanding of the difference between justificatory and excusatory defenses, we shall see that externalists have nothing to fear from the demon and the intuitions that talk of demons can scare up. The intuitions that underwrite the new evil demon objection cannot be used to determine whether belief is governed by some particular norm (e.g., the truth norm, the knowledge norm, or some evidentialist norm)" (forthcoming: 1)

¹⁶ This error theory is put forward by Williamson (2000, 2013, forthcoming), Sutton (2005, 2007), and Littlejohn (2012, forthcoming)

¹⁷ The distinction between rebutting and undercutting arguments is this: rebutting arguments are supposed to show that P is false. Undercutting arguments aren't supposed that P is false. Rather, they aim to undermine the putative evidential connection between a set of facts and the conclusion that is being drawn from them. (See Pollock 1986). I discussed the error theory in Hughes (forthcoming). I don't think I properly characterised it there, however. What I have said here better captures it.

Whether the error theory is plausible or not is a matter of controversy. But we don't need to go into that debate here. The question that we are interested in is: might the error theory undermine the argument I have presented here? If it does, then it must be because it undermines one of WU, II, or DP, since these are the only claims that the argument relies on.

Let's start with WU and DP. It should be clear that the error theory provides no reason to doubt either of these principles. WU states that there are clear-cut cases, in which there is a unique rational doxastic attitude to take. The observation that there is a distinction between permissibility and excusability and that we are (allegedly) prone to conflate the two simply has no bearing on whether or not this is true, and neither point does anything to undermine the arguments that I have put forward for the claim. The two issues are simply unrelated. DP states that there is always at least one permissible attitude to take of belief, suspension, and disbelief. Again, the observation that there is a distinction between permissibility and excusability and the claim that we are prone to conflate the two are irrelevant here. The issues are, similarly, simply unrelated.

What about II? Here one might be tempted to think that there is a connection. II claims that it is impermissible to take irrational attitudes. Or, to put it contrapositively, that it is only permissible to take an attitude if it is rational to take it. If, as a result of failing to distinguish between permissibility and excusability, we would treat rational belief as permissible belief, even if it wasn't, doesn't this give us some reason to be suspicious of II?

It does not. I as pointed out in §3.2 II is the *converse* of the principle that the error theorists seek to undermine. In claiming that we conflate excusability and permissibility the error theorists want to undercut the evidential force of the intuition that rationality is *sufficient* for permissibility (i.e. $\text{Rat} \rightarrow \text{Per}$). But II says that rationality is *necessary* for permissibility (i.e. $\text{Per} \rightarrow \text{Rat}$), and nothing that the error theorists say provides us with any reason to doubt this. An argument against the claim that P entails Q gives one no reason to think that Q doesn't entail P. An argument to the effect that belief doesn't entail knowledge, for example, gives one no reason to think that knowledge doesn't entail belief. So even if the error theorists are right that people would tend to intuitively infer permissibility from rationality even if it was a mistake to do so, that would give us no reason to think that II is false. And indeed, even if we *agreed* with the error theorists that false/unknown-but-nevertheless-rational beliefs are merely excusable, that *still* wouldn't give us a reason to think that II is false.¹⁸ The error theory is simply inert when it comes to II.

¹⁸ Though of course I don't think we should agree with the error theorists about this. That's the point of my overall argument.

To sum up, the permissibility/excusability error theory has nothing to say about any of WU, II, and DP. I conclude, then, that appeals to the error theory gain traction against the deductive argument for the restricted rationality norm.

§IV. Part Two: An Abductive Argument

Let's take stock. In §2 I showed that if WU, II, and DP are all true then it follows that there is a set of cases – 'clear-cut cases' – in which it is sufficient for one to permissibly take a doxastic attitude D to P (of belief, suspension, and disbelief) that it is rational for one to take D to P. If so, then RRN is true. I then argued that we should accept each of WU, II, and DP, and so that we should accept RRN. I also showed that the permissibility/excusability error theory is inert when it comes to this argument.

For all I have argued so far though it might be that we should *only* accept RRN. That is to say, it might be that we should only accept that rationality is sufficient for permissibility in clear-cut cases. So the question arises: should we go further and accept the universalised rationality norm URN, according to which *in every case* rationality is the norm of doxastic attitude formation? In this section I will argue that we should and explain why.

To start off we need to think about cases that are not clear-cut. Call a case that is not clear-cut a 'complex case'. A case is complex just in case there is more than one rational doxastic attitude to take towards P of belief, suspension, and disbelief. Are there any complex cases? Two options present themselves:

Strong Uniqueness: There are no complex cases. Every case is clear-cut.

Weak Permissivism: There are complex cases. Not every case is clear-cut.

Strong Uniqueness ('SU') is roughly equivalent to the principle U that we discussed in §3.1. It should be clear that if SU is true, then we can argue straightforwardly for the universalised rationality norm. All that is required is that we replace the conjunction WU, II, DP with the conjunction SU, II, DP.

But SU will be controversial for most of the same reasons that U is. As such, many will have no inclination to accept an argument for the universalised rationality norm that relies on it. I don't propose to argue for SU here. Rather, I'll argue that even if it is false, and so Weak Permissivism ('WP') is true, there are still good reasons to accept the universalised norm if we accept the restricted norm.

If WP is true then there are complex cases. Now, according to the theory that accepts the universalised rationality norm URN, the following is true:

URN-THEORY: In every case, rationally taking D to P is sufficient for permissibly taking D to P

And according to the theory that accepts the rationality norm as restricted to clear-cut cases, but rejects it the universalised version of it:

RRN-THEORY: In clear-cut cases rationally taking D to P is sufficient for permissibly taking D to P, but in complex cases rationally taking D to P is *not* sufficient for permissibly taking D to P

Given the argument of §2–3 our choice is between URN-THEORY and RRN-THEORY. There is no third option. Which one should we accept? I suggest that *all else being equal* we should prefer URN-THEORY. URN-THEORY is a simple and unified account of the relationship between rationality and permissibility. The link between them that it posits is a straightforward one. By contrast, RRN-THEORY is less simple and less unified. The link that it posits between rationality and permissibility is complicated and disjointed (or, at any rate, more complicated and disjointed than URN-THEORY). In general we should prefer theories that are more simple and unified to those that are less simple and less unified. So we should prefer URN-THEORY to RRN-THEORY unless there are additional considerations speaking in favour of the latter over the former.

The problem for RRN-THEORY is that there *doesn't* appear to be any reason to prefer it to URN-THEORY. Those, like knowledge normers and truth normers, who typically deny that rationality is sufficient for permissibility, will of course want to reject URN. But insofar as the kinds of arguments that they put forward are taken to be persuasive they motivate rejecting *both* URN-THEORY and RRN-THEORY (though of course I don't think they do). What they don't do is give us a reason to *accept* RRN-THEORY and *reject* URN-THEORY. Are there any other reasons to prefer RRN-THEORY to URN-THEORY? As I see it there is only one line of argument that might be thought to motivate the former over the latter. The argument fails but it may seem tempting, so it is worth presenting it and showing why it fails.

The argument goes by an analogy. A doctor has a patient with a serious skin complaint. There is only one drug, 'Fixan', that can cure her. Unfortunately Fixan is known to cause depression in around a quarter of patients. Nevertheless, in this case the benefits outweigh the risk and so it is permissible for the doctor to go ahead and prescribe the drug. Later a new drug comes on to the market: Fixanplus. Is it still permissible for the doctor to prescribe Fixan? That depends. If Fixanplus has everything going for it that Fixan does (it will cure the patient) *and more* (it does not cause depression), then the doctor ought to prescribe Fixanplus, and it is no

longer permissible for her to prescribe Fixan instead. What cases like this show is that the permissibility of a given option can depend on what other options are available. Sometimes when there are multiple available courses of action the standards for permissibility go up.

Something similar, the argument goes, applies in the epistemic case. Let's suppose first that it would only be rational for you to believe that P (it would be irrational for you to suspend or disbelieve). In that case, in order to avoid irrationality, it is permissible for you to believe that P even if P is false or unknown. But now suppose that the evidence changes. A new piece of evidence is added to it. The new evidence set makes it rationally acceptable to believe that P *and* rationally acceptable to suspend judgement on P. So now you don't have to believe that P simply to avoid irrationality – you have options. It might turn out that one option is better than the other for *non*-rationality-related reasons. For example, it might be that P is false, and so if you believe that P then, whilst you would be rational, you would incur the cost of having a false belief. By contrast, if you suspend on P you don't incur this cost. So in this case, whilst both belief and suspension have rationality going for them, suspension has something going for it that belief doesn't. If so, the argument goes, then just as it is no longer permissible for the doctor to prescribe Fixan when Fixanplus is available, which has everything going for it that Fixan does *and more*, so too is it impermissible for you to believe that P when suspension has everything going for it that belief does and more. What this shows is that we should expect the requirements of epistemic permissibility to be more stringent in complex case than they are in clear-cut cases. And for this reason we should RRN-THEORY to URN-THEORY.

Does this argument give us a good reason to prefer RRN-THEORY to URN-THEORY? It does not. Whilst it may have some *prima facie* appeal, that appeal rests on a wrongheaded view of the relative desirability of being in complex cases versus clear-cut cases.

The argument represents agents in clear-cut cases as being in a bind. Hemmed in on all sides by the potential for irrationality, they are forced to take the least bad option and adopt the rational attitude whether or not it has anything else (like truth or knowledge) to be said for it. Any port in a storm. By contrast, agents in complex cases are represented as enjoying comparative luxury. They are in the enviable position of having multiple rational options and it is for this reason that they must not rest content with mere rationality and should instead shoot for more – for truth or knowledge, for example. In this respect the argument takes the two positions to be analogous to the two positions that the doctor finds herself in. In the first instance the doctor is in an unfortunate situation. If she does nothing, then the patient will continue to suffer. But the only cure available – Fixan – comes with a downside: it runs the risk of causing depression. Later on though, the doctor is in a

more favourable situation. She has two cures available – Fixan and Fixanplus – and as a result she can afford to be picky and demand more from her cure of choice.

The persuasiveness of the argument depends on looking at the respective situations of agents in clear-cut cases and complex cases in this way. It is *because* the agent in the clear-cut case is in rather desperate circumstances that – as though pity must be taken on her – she is permitted to settle for mere rationality. And it is *because* the agent in the complex case is in a position of comparative luxury that it is incumbent on her not to settle for mere rationality and instead seek something more, like truth or knowledge. Were this not the appropriate way of looking at their respective positions the claim that agent in the clear-cut case is permitted to settle for less than the agent in the complex case would be unmotivated.

But this is the wrong way to view the respective situations of agents in clear-cut cases versus agents in complex cases. It is a typical characteristic of complex cases that there is more than one rational doxastic attitude to take in them because the evidence is complex, unwieldy, ambiguous, or marginal. One of the core motivations for Permissivism is the thought that in such cases the evidence is often hard to interpret, and it can be unclear which way it points and to what extent. The Permissivists argue that to insist that if two people take different doxastic attitudes to the evidence in these cases one of them must be being irrational, is to postulate unrealistically demanding standards for rationality (Rosen 2001, Schoenfield 2014, Kelly 2014). I don't mean to suggest that every complex case is like this. But it is the normal situation. By contrast, clear-cut cases are cases in which, typically, one's evidence straightforwardly supports taking one particular attitude. The MUG, COIN, and CAR cases testify to this. Insofar as the view that the argument rests upon is correct we should, presumably, prefer to be in complex cases and indeed actively seek them out. But this is obviously perverse. Often, from an epistemic point of view, it is more desirable to be in a clear-cut case than a complex case.¹⁹ So the argument for RRN-THEORY over URN-THEORY rests on a mistaken view of epistemically desirable states. For this reason, we should reject it.

Whilst the argument we just looked at has some *prima facie* appeal, then, it is ultimately flawed. Are there any other reasons one might prefer RRN-THEORY to URN-THEORY? There doesn't seem to be to me. But in that case we should go for URN-THEORY instead of RRN-THEORY. As we have seen, it is the simpler and more unified view. The additional complexity posited by RRN-THEORY is unmotivated.

¹⁹ I don't mean to suggest here that one should *always* prefer, and seek, to be in a clear-cut case over a complex case. But equally one should not always prefer and seek to be in a complex case rather than a clear-cut case. One should go wherever the evidence leads.

§V. Implications

I have argued that URN is true. When it comes to the attitudes belief, suspension, and disbelief, necessarily it is epistemically permissible to take attitude D to P if it is epistemically rational to take D to P. Rationality is the norm of belief, and of coarse-grained doxastic attitude formation in general. What are the implications of this?

The obvious big one is this: insofar as they require us to be irrational in certain cases (such as when your car has spontaneously transformed into a giant lizard, or someone has crept in and stolen your favourite mug overnight) we should reject the truth and knowledge norms. Crucially, this conclusion is arrived at without a need to confront the permissibility/excusability error theory. No appeal to the intuitiveness of the rationality norm is needed to get URN – that's one of the main selling points of the argument of this article – and the judgement that it can be *rational* for you to believe that P even when P is false or unknown is one that the error theorists typically accept. So the argument for URN puts us in a stronger position to reject truth and knowledge norms.²⁰

URN only applies to the coarse-grained doxastic attitudes belief, suspension, and disbelief. I have remained neutral on the question of whether WU applies to credences and it is quite unclear whether the motivations I offered for it would support such an application. But that's not to say that, once we have URN in hand, it has no implications for credal epistemology. A common line on credences is that the rationality norm governs them: it is permissible to have credence c in P just in case it is rational to have c in P. The rationality norm for credences has, for the most part, not been pitted against knowledge and truth norm alternatives. This is unsurprising; it is widely accepted that credences don't have truth-values, so knowledge and truth norms don't look to be appropriate for them. Recently, however, this orthodoxy has been challenged. Moss (2014) proposes that we extend the concept of knowledge in such a way that, despite not having truth-values, credences can properly be described as items of knowledge or non-knowledge. She then floats the idea that we should replace the rationality norm for credences with a knowledge norm. If URN is true we should, I suggest, reject this idea and hold on to the rationality norm. The view that rationality is the norm of coarse-grained doxastic attitude formation but that knowledge is the norm of fine-grained doxastic attitude formation is an odd one. What would explain the differing normative roles of rationality and knowledge between the two domains? It is reasonable to expect unity.

²⁰ See also Hughes (forthcoming)

I've argued that rationality is the norm of belief. But an obvious question still remains outstanding: what should we say about truth and knowledge? After all, there certainly seems to be *something* to be said for the idea that you ought only believe truths, and in particular, those truths that you are in a position to know to be true

I think the question of how rationality normers should seek to explain the normativity of truth and knowledge is an important one, and that it gives rise to a number of rich, subtle, and difficult issues. Since the purpose of this paper has only been to show how we do not need to tackle the permissibility/error theory in order to motivate the rationality norm, won't attempt to give anything like a comprehensive answer to the question here. I will, however, close with a very brief sketch of one way of thinking that I find attractive.

Briefly, then, my preferred way of thinking (and this is by no means a new idea) is this: truth and knowledge are *aims* or *goals* of belief, and the claims that 'you ought only believe truths' and 'you ought only believe those truths you are in a position to know to be true' are best read as articulating these goals.²¹ But just because they are goals, that doesn't mean they are obligations. Indeed, there are cases in which it is natural to think that even if your goal is to φ , your obligation is to *not*- φ . For example, consider the following norm:

BEST BY MOTHER: you ought to do best by your unwell mother.

And consider the following case borrowed from Lord (2015):

Jack's mother is in the hospital. She needs an operation in order to survive past this week. Her insurance won't pay. Jack, being a fledgling art dealer, doesn't have the money. It looks like his mother is going to die. She would be extremely comforted by Jack's presence in her final days. She lives in California; Jack lives in New York. Jack needs to decide whether to go see her. As it happens, a pawn shop owner in Queens has just unknowingly (and legitimately) bought a rare Picasso. He's selling it at a fraction of the price it's worth. If Jack were to buy it, he would be able to use it as collateral for a loan that would pay for his mother's surgery. Unfortunately, he has no idea that this pawn shop exists.

Jack would do best by his mother if he went to Queens rather than California. So if BEST BY MOTHER expresses an obligation, Jack is obligated to go to Queens rather than California. Yet surely the opposite is true. Given what he knows, he is morally obligated to go to California rather than Queens. Why is that? One obvious answer

²¹ This view is taken by Williams (1973), Railton (1994), Velleman (2000), Boghossian, (2003), Shah (2003), and Shah and Velleman (2005), amongst others.

is because the 'ought' in BEST BY MOTHER doesn't express an obligation, but rather articulates a goal to be pursued. Oftentimes when one ought to pursue a particular goal, one does not have an obligation to meet it, but rather to do one's best to meet it in light of the available evidence. That's why Jack is obligated to go to California.

The thought is that the same goes in the epistemic case. Truth and knowledge are epistemic goals that one ought to pursue. But one's *obligation* is to do one's best to meet them in the light of the available evidence. And that means one's obligation is to be rational, even if that sometimes leads to failure.

All of this needs a lot more spelling out, of course, and I certainly don't want to suggest that there are no difficulties with the view (see Gibbons 2013, amongst others). But it strikes me as promising nevertheless. Making a comprehensive case for it is beyond the scope of this paper. For now, I am content to have shown that one does not need to tackle the permissibility/excusability error theory in order to motivate the rationality norm.²²

§VI. References

- Adler, J. 2002, *Belief's Own Ethics* (Cambridge: MIT Press)
- Boghossian, P. 2003, "The Normativity of Content" in *Philosophical Issues*, Vol. 13 (1). Pp. 31–45
- Cohen, S. 1984, "Justification and Truth" in *Philosophical Studies* Vol. 46 (3). pp. 279–295
- Cohen, S. & Comesaña, J. 2013. "Williamson on Gettier Cases in Epistemic Logic and the Knowledge Norm for Rational Belief: A Reply to a Reply to a Reply" in *Inquiry*, Vol. 56 (4). pp. 400–415
- Cohen, S. & Comesaña, J. forthcoming. "Being Rational and Being Right" in Dutant, J. and Dohrn, D. (eds.) *The New Evil Demon* (Oxford: Oxford University Press)
- DeRose, K, 2002, "Assertion, Knowledge, and Context" in *Philosophical Review*, Vol. 111 (2). Pp. 167–203
- Gibbons, J 2013, *The Norm of Belief* (Oxford: Oxford University Press)
- Goldman, A, 1986, *Epistemology and Cognition* (Cambridge, MA: Harvard University Press)

²² Thanks to Matthew McGrath, Herman Cappelen, Jessica Brown, Patrick Greenough, Torfinn Huvenes, Stewart Cohen, Clayton Littlejohn, and two anonymous referees for this journal for discussion and helpful feedback on earlier versions of this article.

- Friedman, J. 2013b, "Suspended Judgment" in *Philosophical Studies* Vol. 162 (2). pp. 165–181
- Hawthorne, J. 2004, *Knowledge and Lotteries* (Oxford: Oxford University Press)
- Hughes, N. forthcoming, "No Excuses: Against the Knowledge Norm of Belief" in *Thought: A Journal of Philosophy*
- Kelly, T. 2014, "Evidence Can Be Permissive" in Steup, M., Turri, J., and Sosa, E (eds.) *Contemporary Debates in Epistemology* 2nd edition, pp. 298–313 (Blackwell)
- Lackey, J. 2007, "Norms of Assertion" in *Nous*, Vol. 41 (4). Pp. 594–626
- Lasonen–Aarnio, M. 2010, "Unreasonable Knowledge" in *Philosophical Perspectives* Vol. 24. pp. 1–21
- Littlejohn, C. 2012, *Justification and the Truth–Connection* (Cambridge: Cambridge University Press)
- Littlejohn, C. 2013, "The Russellian Retreat" in *Proceedings of the Aristotelian Society* Vol. 113 (3). pp. 293–320
- Littlejohn, C. forthcoming, "A Plea for Epistemic Excuses" in Dutant, J & Dorsch, F (eds.) *The New Evil Demon* (Oxford University Press)
- Lord, E. 2015, "Acting for the Right Reasons, Abilities, and Obligations" in *Oxford Studies in Metaethics*, Vol. 10 (Oxford: Oxford University Press)
- Moss, S. 2014, "Epistemology Formalized" in *Philosophical Review* Vol. 122. pp. 1–43
- Pollock, J. 1986, *Contemporary Theories of Knowledge* (Savage MD: Rowman and Littlefield)
- Railton, P. 1994 "Truth, Reason, and the Regulation of Belief" in *Philosophical Issues*. Vol 5. Pp. 71–93
- Schoenfeld, M. 2014, "Permission to Believe: Why Permissivism is True and What It Tells Us About Irrelevant Influences on Belief" in *Nous* Vol. 48 (2). pp. 263–277
- Shah, N. 2003, "How Truth Governs Belief" in *Philosophical Review*. Vol 112 (4). Pp. 447–482
- Shah, N. & Velleman, D. 2005, "Doxastic Deliberation" in *Philosophical Review* Vol. 114 (4). Pp. 497–534
- Stanley, J. 2005, *Knowledge and Practical Interests* (Oxford University Press: Oxford)
- Sutton, J. 2005, "Stick To What You Know" in *Nous* Vol. 39 (3). pp. 359–396
- Sutton, J. 2007, *Without Justification* (Cambridge, MA: MIT Press)
- Velleman, D. 2000 "On the Aim of Belief" in *The Possibility of Practical Reason* (Oxford: Oxford University Press)
- White, R. 2005, "Epistemic Permissiveness" in *Philosophical Perspectives* Vol. 19 (3). pp. 445–459
- Whiting, D. 2010, "Does belief aim (only) at the truth?" in *Pacific Philosophical Quarterly*. Vol. 93. pp. 279–300
- Whiting, D. 2013, "Nothing but the truth: on the aims and norms of belief" in T. Chan (ed.) *The Aim of Belief* (Oxford: Oxford University Press)
- Williams, B. 1973, "Deciding to Believe" in *Problems of the Self* (Cambridge: Cambridge University Press)
- Williamson, T. 2000, *Knowledge and its Limits*, (New York: Oxford University Press)
- Williamson, T. 2013, "Response to Cohen, Comesaña, Goodman, Nagel, and

Weatherson on Gettier cases in Epistemic Logic" in *Inquiry: An Interdisciplinary Journal of Philosophy* Vol. 56 (1). pp. 77–96

Williamson, T. forthcoming. "Justifications, excuses, and sceptical scenarios" in Dutant, J. and Dorsch, F. (eds.) *The New Evil Demon* (Oxford: Oxford University Press)