

The norm of assertion and blame. Draft of Nov 2017.

This article will appear in a revised form in the Oxford Handbook of Assertion (editor Goldberg), Oxford University Press. © Oxford University Press.

Jessica Brown. Arché, University of St Andrews.

1. Introduction.

One of the most popular suggestions for the norm of assertion is the knowledge norm. Some hold that knowledge is necessary for proper assertion – one ought: assert that *p* only if one knows that *p*. Others endorse a bi-conditional norm: one is in a good enough epistemic position to assert that *p* if and only if one knows that *p*. However, the knowledge norm faces intuitive counterexamples. In some of these alleged counterexamples, it seems fine for a subject to assert that *p* even though she doesn't know that *p*; in other alleged counterexamples, it seems problematic for a subject to assert that *p* even when she does know that *p*. Defenders of the knowledge norm have responded by appealing to the distinction between whether one conforms to a norm and whether one is blameworthy: one can violate a norm and yet be blameless or conform to it but nonetheless be blameworthy. Furthermore, some suggest that it's futile to search for a norm such that whether one conforms to the norm aligns with whether one is blameless. They appeal to the failure of luminosity to defend the pessimistic conclusion that, for any norm whatsoever, the notions of norm conformity and being blameworthy can come apart.¹

¹ Thanks to Dan Greco for suggesting I should connect my work on blame with the argument for the pessimistic conclusion. Thanks for useful comments on various versions of this paper from Sandy Goldberg, Mona Simeon, my colleagues at St Andrews and the audience at the Joint Session of the Mind and Aristotelian Society 2017.

In this paper, I assess this defence of the knowledge norm of assertion. I criticise the argument from the failure of luminosity to the pessimistic conclusion and argue that there is no easy way to reformulate it successfully. The discussion of the pessimistic conclusion distinguishes a variety of different explanations of how norm-conformity can come apart from whether one is blameless, and the limitations of these explanations. Thus, it also provides some guidance on when it is reasonable to defend the knowledge norm of assertion against counterexample by appeal to the distinction between norm-conformity and whether one is blameless.

2. Defending the knowledge norm.

It's become recently popular to defend the idea that knowledge is the norm of assertion. Some defend the idea that knowledge is necessary for appropriate assertion: one ought to assert that *p* only if one knows that *p*. Others defend a bi-conditional version of the knowledge norm: one is in a good enough epistemic position to assert that *p* if and only if one knows that *p*. The knowledge norm is motivated by a variety of data. This includes the infelicity of asserting “*p* but I don't know that *p*”; the propriety of challenging the assertion that *p* by the questions, “How do you know that *p*?” and “Do you know that *p*?”; and, the infelicity of asserting lottery propositions (Williamson 2000).

However, the idea that knowledge is the norm of assertion faces objections to both its necessity and sufficiency direction. In some cases, it seems fine for a subject to assert that *p* even though she doesn't know that *p*, such as Gettier cases and cases in which a subject justifiably falsely believes that *p*. In other cases, it seems problematic for a subject to assert that *p* even though she does know that *p*. For instance, in Lackey's second-hand knowledge case, it seems problematic for a surgeon to flat-out assert to her patient that she has cancer when she hasn't looked at the patient's notes but has merely learnt this information from the

testimony of a junior. Another class of potential counterexamples involve cases in which a subject plausibly knows but is nonetheless in such a high-stakes situation that it seems inappropriate for her to assert the relevant claim. For instance, even if a doctor knows on the basis of the morning clinic that it is the patient's left kidney which is diseased, it seems that she should check the patient's notes again rather than straight out assert this in the context of the afternoon operation to remove the diseased kidney (Brown 2010; for similar cases see Reed). A last set of cases are generated by the controversial view that one can maintain knowledge that *p* despite receiving what many regard as defeating evidence, say evidence that not-*p*, evidence that the process which produced the belief that *p* is unreliable, or evidence that one's evidence does not support that *p* (e.g. Lasonen-Aarnio 2010, forthcoming; Hawthorne and Srinivasan 2013). In these cases, too, it seems problematic for one to assert that *p* even though one (allegedly) knows that *p*.

In response to these putative counterexamples, defenders of the knowledge norm have appealed to the idea that whether one conforms to a norm can come apart from whether one is blameworthy (Williamson 2000, Hawthorne and Stanley 2008, Hawthorne and Srinivasan 2013). To the extent that one can violate the knowledge norm blamelessly, perhaps this explains our intuitions about cases in which a subject lacks knowledge that *p* but it nonetheless seems fine for her to assert that *p*. To the extent that one can conform to the knowledge norm in a blameworthy way, perhaps this could explain some of the cases in which a subject knows that *p* but it seems problematic for her to assert that *p*. Further, some argue for the pessimistic conclusion that, for any norm whatsoever, whether one conforms to that norm can come apart from whether one is blameworthy (Hawthorne and Srinivasan 2013).

For this defence of the knowledge norm to be legitimate we need an account of when one is blameless for violating a norm or blameworthy for conforming to it. For otherwise, it's

hard to assess whether the right response to the putative counterexamples is to retain the norm and chalk them up to the distinction between norm conformity and whether one is blameworthy or, instead, revise the formulation of the relevant norm. Indeed, without an account of when one is blameless for norm violations or blameworthy for norm conformity, the appeal to the distinction between norm conformity and blameworthiness could be used irresponsibly. For example, consider the surprising suggestion that in polite English society it's a norm that one ought to belch in appreciation of a good meal. Of course, this putative norm faces apparent counterexamples: cases in which people who belch after a good meal are criticised and cases in which people who don't are not criticised. Without any account of when one is blameless for a norm violation or blameworthy despite norm conformity, somebody could attempt to defend the existence of the putative belching norm by appeal to this distinction, chalking up the apparent counterexamples to the distinction between norm conformity and blameworthiness.

Similarly, we need an account of blameless norm violation and blameworthy norm conformity in order to assess the argument for the pessimistic conclusion that, for any norm whatsoever, norm conformity and being blameless can come apart. The classic argument for the pessimistic conclusion proceeds from an epistemic claim, namely Williamson's anti-luminosity result, to the conclusion that, for all norms whatsoever, norm conformity and being blameless can come apart. But, any such argument implicitly relies on a bridging claim which connects the relevant result, here anti-luminosity, to the conclusion that norm conformity and being blameless can come apart. In other words, the argument implicitly relies on some account of blameless norm violation and blameworthy norm conformity. Thus, to assess both how defenders of the knowledge norm reply to particular counterexamples, as well as their more general pessimistic conclusion, we need an account of blameless norm violation and blameworthy norm conformity. In the rest of the paper, I examine this issue,

starting with an assessment of the argument from anti-luminosity to the pessimistic conclusion.

3. The pessimistic conclusion.

There are a variety of ways in which one might violate a norm and yet be blameless. One might be exempt from a norm, or forced by another to violate it. Alternatively, one might be blameless for violating one norm if by violating it one conforms to a more important norm which overrides it. The argument from anti-luminosity to the pessimistic conclusion exploits a further way in which one might be blameless for violating a norm, namely due to the distinction between one's perspective on whether one is conforming to the relevant norm, and whether one is in fact conforming to it. I will focus on this last possibility in the rest of this paper, assuming implicitly that the subject is not exempt from the relevant norm, not forced, and that the norm is not overridden by another norm.

Hawthorne and Srinivasan (2013) defend the pessimistic conclusion by appeal to “non-transparency”, the claim that there is no condition such that we are always in a position to know whether it obtains (9). In more detail, non-transparency amounts to the view that, for any non-trivial condition, it's possible that it obtains but one is not in a position to know that it obtains; and, it is possible that it does not obtain but one is not in a position to know that it does not obtain. As a result, they claim that there is no possibility of finding epistemic norms which are normatively satisfying:

Intuitively, we expect epistemic norms to be *normatively* satisfying: that is, we expect them to track our intuitions about blameworthy and praiseworthy epistemic conduct.

An epistemic norm that ties what one ought to do to a non-transparent condition (e.g. knowledge) is an epistemic norm that will not satisfy this basic desideratum. To construct an epistemic norm that is normatively satisfying, then, we require an

epistemic “ought” that is tied to only transparent conditions; unfortunately no such conditions plausibly exist. (28)

Given that no non-trivial condition is transparent, the same conclusion would extend to norms in general. Indeed, Srinivasan (2015) argues that given that no non-trivial condition is transparent, there is no norm which does not allow for blameless norm violation².

It seems that the basic form of Hawthorne and Srinivasan’s argument is to combine a certain epistemic result, namely non-transparency, with a bridging principle which connects that epistemic result to divergence between norm-conformity and blameworthiness to arrive at the general conclusion that there is no norm which aligns norm-conformity and whether one is blameless. In more detail, we can formulate their argument as follows:

1. For any norm α iff β , it’s possible that β /not- β obtains but one is not in a position to know that. (Non-transparency.)
2. If β /not- β obtains but one is not in a position to know that and one violates the norm (α iff β), then one violates the norm blamelessly; if β /not- β obtains but one is not in a position to know that and one conforms to the norm, one does so in a blameworthy way. (Bridging claim)
3. So, for any norm, it’s possible that one violates the norm blamelessly or conforms to it in a blameworthy fashion.³

² Srinivasan (2015) focuses on showing that every norm has cases of blameless norm violation. She argues that "if Anti-Cartesianism is true, then there are no norms that are always action guiding or that are never blamelessly violable" (285), where Anti-Cartesianism is just non-transparency (274).

³ Srinivasan's argument is slightly more complex and goes through some sub-premises concerning the connection between control, luck and blamelessness. In more detail, she

In assessing this argument, it's important to bear in mind that, quite independently of this argument, we know that in many cases, norm conformity comes apart from whether one is blameworthy. So, the interest of the argument is in establishing a stronger conclusion than this, namely the universal claim that, for any norm whatsoever, norm conformity comes apart from whether one is blameworthy. However, I will argue that the original argument which appeals to anti-luminosity fails to establish this universal pessimistic conclusion and, further, that it is hard to reformulate the argument so as to establish the universal pessimistic conclusion.

In assessing the original anti-luminosity argument, I will take it that Williamson (2000) has established the non-transparency premise (1). However, I will argue that the bridge principle (2) fails: not being in a position to know is not sufficient to make a violation blameless or to make conformity blameworthy. We can illustrate this by appeal to Gettier cases.⁴ Consider the following norm: call the boss if and only if the appointee is Etonian.

argues that given non-transparency, there is no norm such that a competent agent who knows the norm is in a position to know of every basic action available to her whether it would be in conformity with the norm. As a result, she claims that, for any norm, it's not always under one's control whether one's basic actions constitute norm violations, but is instead sometimes a matter of performance luck. Further, she claims that if one violates a norm through performance luck, then one is blameless (2015, section 5). Just as with the simpler argument considered in the main text, Gettier cases cause trouble for this more complex argument. We may accept that, in a Gettier case, it is a matter of luck whether the agent's action is a norm violation. But, it does not follow that she is blameless if she violates it.

⁴ I use the standard description of a Gettier case as one involving a justified true belief which is not knowledge. Some (Sutton 2007, Littlejohn forthcoming, Williamson forthcoming)

Suppose that a subject is in a Gettier case with respect to the claim that the appointee is Etonian. She justifiably and truly believes that the appointee is Etonian on the basis of inference from a justified false belief that Jones will get the job and a justified true belief that Jones is Etonian. Thus, she is not in a position to know that the appointee is Etonian. Nonetheless, if she violates the norm by failing to call the boss she is not blameless. Similarly, if the subject violates the norm by calling the boss when it's not the case that the appointee is Etonian, then it's not sufficient to render her blameless that she is not in a position to know that the appointee is not Etonian. For, the subject could be in a Gettier case with respect to the proposition that the appointee is not Etonian: she could have a justified true belief that the appointee is not Etonian but not be in a position to know that. But if she calls the boss in such circumstances she is not blameless for violating the norm.⁵

In a similar way, appeal to Gettier cases can be used to show that not being in a position to know is not sufficient to make conformity to a norm blameworthy. Suppose that a subject conforms to the norm by calling the boss in the circumstance that the appointee is Etonian. Further, suppose that the subject is not in a position to know that the appointee is

dispute this standard description, arguing that a belief is justified only if knowledge.

However, nothing hangs on this description and Gettier cases equally undermine the anti-luminosity argument for the pessimistic conclusion even on the description on which they merely involve having evidence for a true claim which is not knowledge.

⁵ Other cases may also cast doubt on the bridging principle 2). Consider Sher's (2009) case of a forgetful dog owner who leaves her pet dog in a sealed car on a hot day. Surely, she has violated a norm for looking after her pet dog and is blameworthy. That remains so even if, having clean forgotten the dog, she no longer knows, and is no longer in a position to know that she's left the dog in the car.

Etonian since she is in a Gettier case. In particular, she has a justified true belief that the appointee is Etonian but doesn't know that. It wouldn't follow that she is blameworthy for conforming to the norm. A similar chain of reasoning shows that even if a subject is in a Gettier case with respect to the claim that the appointee is not Etonian, it doesn't follow that she is blameworthy for conforming to the norm by not calling the boss when the appointee is not Etonian. It seems that not being in a position to know is not sufficient to make norm violation blameless or norm conformity blameworthy.

It seems, then, that the original form of the argument from non-transparency to the pessimistic conclusion fails. To reinstate the argument, one needs to show that for any condition, β , which might figure in a norm α iff β , it's possible that β (or not- β) obtains in circumstances C , such that 1) if one violates the norm when β (not- β) and C , one is blameless; and, 2) if one conforms to the norm when β (not- β) and C , one's norm conformity is blameworthy. One might hope to reinstate the argument to the pessimistic conclusion using some epistemic condition other than the condition that one is not in a position to know. For example, we sometimes judge a subject blameless for violating a norm if she is ignorant that she is violating it. However, we will see that, for a range of epistemic conditions that one might use to restate the argument, the argument fails either at the first premise which records a certain kind of epistemic result or at the bridging premise which connects that result to divergence between norm-conformity and whether one is blameless. Let's start by examining problems for the epistemic result.

4. Problems for the epistemic result.

Suppose that one tries to reformulate the argument to the pessimistic conclusion by appeal to a variety of doxastic or epistemic properties such as lack of belief, justification/evidence to believe, or justified belief. For instance, one might notice that one sometimes judges a subject

blameless for a norm violation when she is ignorant that she is violating the norm. However, to use this kind of case to reinstate the universal pessimistic conclusion, it's not enough to show that, for many conditions which can figure in norms, they can obtain while one is ignorant that they obtain. Instead, one needs to defend the universal claim that, for any condition which can figure in a norm, it can obtain despite one's ignorance that it obtains. Alternatively, if the idea is that violating the norm (α iff β) by α ing when not- β is blameless if one has justification to believe that β , then one needs to defend the universal result that, for any condition, it's possible that the condition does not obtain even though one has justification to believe that it does obtain. However, it's not clear how one would defend the needed universal result. In particular, Williamson's anti-luminosity argument would be of no help. For, the anti-luminosity argument centrally turns on the safety condition for knowledge: that if one knows that p in a case α , then there is no nearby case, β , in which one fails to know.

To see this, recall the general structure of the anti-luminosity argument (Williamson 2000). We are to consider a subject who starts out feeling cold one morning and gradually becomes warmer. In particular, consider a case c in the middle of the spectrum between feeling cold and feeling warm which constitutes a borderline case of feeling cold. Even if, in c , the subject truly believes that she feels cold, she is not in a position to know this. For, there is a nearby case in which she would continue to believe that she feels cold even though she is not. Thus, even if, in c , she truly believes that she feels cold, her belief is not safe or knowledge. Given the way in which the anti-luminosity result depends on the safety condition, the result does not apply to any property which does not require a safety condition.

But, arguably⁶, none of the following epistemic conditions require safety: belief, justification to believe, justified belief etc. So, even if we could formulate an account of blameless norm violation using these epistemic notions, it's not clear how we would establish the epistemic result which is supposed to generate the conclusion.

Even though Williamson's anti-luminosity argument does not extend to belief or justification, one might suggest that it is anyway independently plausible that, for any condition, it could obtain yet one not believe it obtains or it could obtain yet one justifiably believes it does not obtain. Certainly, we should grant that, for many conditions, this is true. However, to defend the pessimistic conclusion in full generality, the epistemic result from which the argument starts needs to be universal. For example, if the bridging premise employs the idea that lack of belief is sufficient for excuse, the epistemic result needs to take the following form: for any condition, β , β could obtain yet one not believe it obtain. Establishing any such universal claim is likely to be highly controversial. After all, some suggest that certain conditions are self-intimating, where a condition is self-intimating if whenever it obtains, one believes that it obtains. For example, some suggest that pain is self-intimating. Similar difficulties affect the suggestion that for any condition β , β could obtain yet one believes falsely that it not obtain. For some suggest that certain beliefs are self-verifying: for some p , if one believes that p , then p is true. For example, consider the belief that I'm now thinking. Some argue that whenever I form this belief, I thereby make it true

⁶ On the controversial knowledge view of justification, a belief is justified if and only if known, and so justified belief requires safety (e.g. Littlejohn, Sutton, Williamson). I set this view aside here since it would not help the defender of the pessimistic conclusion find an epistemic condition distinct from lack of knowledge with which to reinstate her argument.

(e.g. Burge). If that is right, then there are some conditions such that one cannot believe falsely that they obtain since, by believing they obtain, one makes it true that they obtain.

Consider instead whether the argument could work by exploiting justification or evidence. On a traditional view, one can have justified false beliefs and misleading evidence about many things.⁷ However, it's controversial whether, for every condition, p , whatsoever, one can have evidence/justification not- p , yet p ; or evidence/justification that p , yet not- p . This issue has surfaced in an especially acute form in the debate over the possibility of epistemic akrasia. Titlebaum (2015) and Littlejohn (2015) have argued that one cannot have justified false beliefs about what rationality requires. Others have suggested that there are inter-level coherence requirements which would prevent certain combinations of beliefs being justified. In particular, some argue that even though one's evidence can support both p and that one's evidence does not support that p , one cannot be rational or justified in believing both that p and that one's evidence does not support that p since rationality/justification requires certain inter-level coherence requirements (e.g. Christensen 2007, Worsnip 2015). If that's right, then for the norm, α iff one justifiably believes that p , one cannot be in the situation in which one justifiably believes that p and justifiably believes that one's evidence does not support that p . While I don't want to take a stand on these issues here, it's worth recording that it is controversial whether, for every condition, p , it's possible that p and yet one has evidence/justification/justified belief that not- p . As a result, it's far from clear that one could easily defend the pessimistic conclusion by appeal to the suggestion that, for any condition whatsoever, it can obtain even though one has evidence supporting

⁷ Of course, those who hold that a belief is justified if and only if known reject the idea that one can have justified false beliefs. But, even they accept that one can have evidence that p when not- p .

that it does not obtain, or justification to believe that it does not obtain, or justified belief that it does not obtain.

We've seen, then, that there is no easy substitute property for non-transparency to employ in the argument to the pessimistic conclusion. For, Williamson's anti-luminosity argument centrally relies on the safety condition for knowledge. Thus, it does not extend to properties which do not require safety, such as belief or justification. Further, even if it is plausible that, for many conditions, β , β could obtain without one's believing it, having justification to believe it, or one's justifiably believing it, it's far from clear that this is so for all conditions whatsoever. But only a universal claim would support the universal nature of the pessimistic conclusion, namely that for any norm whatsoever, norm-conformity comes apart from whether one is blameworthy.

However, even setting aside these concerns about establishing the relevant epistemic result, concerns can be raised about whether the various candidate epistemic properties are sufficient for norm conformity to come apart from whether one is blameless. If they are not so sufficient, then even if the epistemic result could be established, it would fail to establish the pessimistic conclusion. But, as we will see in the next section, there are problems for each of the various candidate epistemic properties.

5. Problems for the bridging claim.

Earlier, we saw that not being in a position to know is not sufficient to make norm violation blameless nor norm conformity blameworthy. So, let's now consider whether there are alternative bridging claims which link some epistemic or doxastic standing to blameless norm violation or blameworthy norm conformity. A first thought might be to formulate a bridge principle by appeal to the widely held view that factual ignorance can excuse norm violation. For instance, if a doctor kills a patient by unwittingly giving her a blood product

contaminated by a malevolent nurse, then we might judge the doctor blameless. While it is widely agreed that factual ignorance can excuse, it is much more controversial whether normative ignorance excuses.⁸ In particular, some doubt that normative ignorance, that is ignorance of norms, excuses norm violation. For example, consider a dictator's son raised within the palace walls and never given a proper moral education. Even if he is ignorant of moral norms, it doesn't follow that he is blameless when he, for example, orders the children of dissidents to be tortured in front of their parents and watches the scene with evident enjoyment. So, for the purposes of discussion, I will set aside normative ignorance, and focus wholly on factual ignorance.

The kind of ignorance which most plausibly excuses norm violation is a justified false belief or a false belief supported by the evidence. For merely failing to believe some truth is not plausibly sufficient to excuse a norm violation. For instance, suppose that a gardener empties a bottle of herbicide into a river, thereby violating the norm: dispose of liquids in the river only if they are harmless.⁹ It would be no excuse if she does so while failing to truly

⁸ Some argue that blameless moral ignorance excuses just as blameless factual ignorance does, including Calhoun 1989, Rosen 2003 and 2004, Fitzpatrick 2008, Zimmerman 2008, Levy 2009. Others deny this including Harman 2011, Mason 2016 and Weatherson forthcoming.

⁹ Here, I consider a norm which only specifies a necessary condition for appropriate action rather than a necessary and sufficient condition. Although I stated the original argument for the pessimistic conclusion with a bi-conditional norm, I take it that Hawthorne and Srinivasan want their pessimistic conclusion to also apply to one-way norms. This is particularly important in the context of the debate about the knowledge norm of assertion since some only

believe that the liquid is poisonous because she dogmatically persists in believing that the liquid is harmless despite reading the explicit warnings on the label. Further, it would be no excuse if she fails to believe truly that the liquid is poisonous because, as she realises, the evidence is inconclusive one way or the other so she suspends on the issue.

While lack of belief does not provide a sufficient condition for blameless norm violation, it might seem more plausible that we can find such a condition by appeal to a justified false belief or a false belief supported by the evidence (see Peels 2010, Harman 2011, Weatherston forthcoming).¹⁰ For instance, the gardener might be excused if she justifiably believes that the liquid is harmless on the basis of the information provided by the manufacturer. So, consider the idea that one is blameless for violating the norm (α only if β) by α -ing when not- β if one justifiably falsely believes that β ; and blameworthy for conforming to the norm (α only if β) by α -ing when β if one justifiably falsely believes that not- β .

Appeal to justified false belief looks initially promising for blameworthy norm-conformity. For example, if the gardener has a justified false belief that some liquid is poisonous, then she would plausibly be blameworthy for pouring that liquid into the river,

defend the claim that knowledge is necessary for proper assertion rather than a bi-conditional version of the knowledge norm.

¹⁰ Indeed, a great many of the examples used to suggest that not being in a position to know generates cases in which norm conformity comes apart from whether one is blameworthy involve having misleading evidence (see eg. Hawthorne and Srinivasan and 2013). Those who hold the knowledge view of justification on which a belief is justified only if known (e.g. Sutton, Williamson), would prefer the evidence formulation since they deny that there can be justified false beliefs.

even though this action conforms to the relevant norm. However, the suggested condition for blameless norm violation is less plausible. To see this, consider how much justification one needs in order to count as being justified tout court. On one possible view, one needs a maximal justification in order to count as being justified tout court. But, we will see that this makes the conditions for blameless norm violation implausibly tough. On an alternative view, in order to be justified, one's degree of justification needs to exceed a threshold which is less than maximal. But, we will see that, on this threshold view, having a justified false belief is not sufficient for blameless norm violation. Thus, we will see that it is difficult to formulate a condition plausibly sufficient for a norm violation to be blameless by appeal to the notion of a justified false belief. This is problematic for the purpose of establishing the overall pessimistic conclusion. After all, the original pessimistic conclusion stated that, for any norm, whether one conforms to that norm comes apart from whether one is blameworthy in both directions: one can conform to it in a blameworthy way or violate it blamelessly.¹¹ So, to establish the pessimistic conclusion we need not only a condition plausibly sufficient for blameworthy conformity but also a condition plausibly sufficient for blameless violation. So, even if appeal to justified false beliefs does provide a plausible account of blameworthy norm-conformity, if it doesn't provide an account of blameless norm-violation, then it doesn't establish the pessimistic conclusion.

¹¹ Recall that Hawthorne and Srinivasan (2013) argue that no norms are normatively satisfying i.e. "track our intuitions about blameworthy and praiseworthy epistemic conduct" (28). They stress the problem of normative divergence facing norms: "they generate instances of both blameworthy right-doing and, conversely, virtuous wrongdoing"(20). Srinivasan (2015) argues that any norm gives rise to cases of blameless norm violation.

Let's start by considering what degree of justification one needs in order for one's justified false belief to excuse a norm violation. One possible answer is that one needs to be have the maximal level of justification for a claim in order for it to render a norm violation blameless. But, this suggestion has the implausible result that we are hardly ever excused for norm violations since, plausibly, there are very few claims with respect to which we have the maximal level of justification. Indeed, on certain assumptions, we can never have the strongest possible evidence for falsehoods. For example, some defenders of the knowledge norm of assertion hold that evidence is factive (e.g. Williamson). But, falsehoods can never entail truths. So, falsehoods can never have probability 1 on one's evidence. Thus, on a factive view of evidence, one can never have the strongest possible amount of evidence for a falsehood.

Given that requiring maximal justification is too demanding, it might be suggested that a norm violation is blameless if one's degree of justification exceeds a high threshold of justification, where the threshold is less than the maximal amount. But, as soon as it is non-maximal, then we seem to face a threshold problem: what level of justification is sufficient, and what makes that level non-arbitrary? To illustrate, suppose that we understand justification in terms of probability on the evidence. Then someone might suggest that the requisite level of justification is a probability on the evidence greater than .9. What fixes this as the required level rather than .8999, or .911? If the level required is arbitrary, then it makes whether one is blameless for a norm violation arbitrary.

Further, wherever this level is fixed, then at least on a classic invariantist account of justification, it seems that justified false belief is not always sufficient for excuse. For example, suppose again that the required level of justification is probability on the evidence which exceeds .9. Even so, we can generate cases in which justification is insufficient for excuse. It would not excuse a teacher's feeding a known to be peanut-allergic child peanut-

containing curry that it was .9 likely that it didn't contain peanuts. The problem generalises to any plausible high threshold less than 1. For any invariant threshold of probability on the evidence less than 1, we can imagine a situation in which the stakes are so high that meeting that threshold is not sufficient for excuse. (Of course, analogous reasoning would apply to the alternative idea that justification is not a matter of probability on the evidence but instead the reliability of the relevant belief-forming process. First, we don't want to require maximal reliability for blameless norm violation, but only a sufficient threshold of reliability. But, second, on any invariant non-maximal threshold, there will be cases in which meeting that threshold is not sufficient for blameless norm violation.)

One way to try to simultaneously solve both the threshold problem and the sufficiency problem would be to appeal to a shifty view about justification such as contextualism or pragmatic encroachment. According to such shifty views, whether a subject's belief counts as justified depends not just on the level of her evidence but also the stakes. As the stakes get higher, she needs a higher level of evidence in order to count as justified. Thus, it might be suggested that in the peanut-case, even having .9 probability on the evidence that the candy does not contain peanuts is not sufficient for the subject to count as justified.¹² While this

¹² Someone might object that, at least on pragmatic encroachment, it's not clear that the peanut case involves high stakes, as opposed to serious moral norm violation. According to pragmatic encroachment, the relevant stakes are the stakes for the subject of the epistemic attribution. But the fact that the situation is high stakes for the child does not entail that it is high stakes for the teacher. For instance, it's possible that a teacher has to decide what lunch to give to a known to be peanut-allergic child where the teacher doesn't care about the child, nor about whether she does a morally wrong action, and can expect not to face any serious consequences for feeding the child peanut-containing food. In such circumstances, the level

suggestion would be congenial to defenders of shifty views, it would be very unappealing to those who find shifty views problematic. Thus, although shifty views offer a way of defending a bridging principle linking justification to believe that one's not violating a norm with being blameless for violating it, it is not a solution which would be happily embraced by all. Thus, the argument to the pessimistic conclusion if filled out in this way would be very controversial.

A defender of the pessimistic conclusion might attempt to repair the argument in two different ways. First, she might employ the claim that, for any norm, there is some degree of justification which would excuse violations of it. She might attempt to defend the pessimistic conclusion by combining this first claim with the second claim that for any norm (α iff β), and for any degree of justification d , one can have a justified false belief that β of degree d . However this reformulated argument employs a much stronger epistemic claim than the original formulation (namely that for any norm (α iff β) one can have a justified false belief that β). It's not clear why we should endorse the stronger claim. The problems we encountered in section 4 for the original claim apply to the strengthened claim as well. Further, on certain views – such as Williamson's account of evidential support--one cannot have the maximal degree of evidential support for a falsehood. Given his factive account of evidence, and that truths don't entail falsehoods, a falsehood cannot have probability 1 on the evidence. Thus, the first strategy of repair looks unpromising.

of justification required for excusably feeding the child peanut-containing food would still intuitively be high, even if we cannot understand this in terms of the stakes for the teacher being high. To the extent that such cases are possible, we might conclude that pragmatic encroachment should be understood so that justification depends on the degree of wrongness of the relevant action as well as the stakes for the agent.

Second, the defender of the pessimistic conclusion might attempt to repair her argument by restricting her attention to low-stakes norm violations and arguing that justified false belief excuses low-stakes norm violations. For instance, she may suggest that: for any norm (α only if β), whenever it is not a serious matter whether that norm is violated, if the subject violates the norm by α ing when not- β and has justification to believe that β , she is blameless. However, even if this revised principle is plausible, it could be used to support the pessimistic conclusion only in conjunction with the following epistemic result: for any norm, there are cases in which it is not very serious whether that norm is violated. However, it's far from clear how we could show that the required epistemic result is true. In particular, certain norms seem in their nature to concern serious matters. This is the case with norms that prohibit torture or killing innocents except in certain special circumscribed situations. Given the significance of human life, it's simply not clear how there can be non-serious violations of these norms.

Let's summarise the discussion so far. We started by examining the original argument for the pessimistic conclusion which appealed to a certain epistemic result in conjunction with a bridge principle to conclude that, for any norm, whether one conforms to that norm comes apart from whether one is blameless in both directions: one can conform in a blameworthy way or violate it blamelessly. The original argument appealed to the notion of not being in a position to know. We saw that that argument failed since even granting the epistemic result, the bridge principle is implausible. Not being in a position to know is neither sufficient for a norm violation to be blameless nor for norm conformity to be blameworthy. We then examined whether we could reformulate the argument by appeal to some other doxastic or epistemic condition, such as belief, justification to believe or evidence. This reformulation faced two main problems. First, the relevant epistemic result is harder to establish. Second, we faced problems formulating plausible bridge conditions. True, appeal to

justified false belief provides a condition plausibly sufficient for blameworthy norm conformity: if one's conforms to the norm (α only if β) by α ing when β and one's evidence supports not- β , then one conforms in a blameworthy way. But, it's much harder to appeal to the notion of justification to provide a condition plausibly sufficient for blameless norm violation. Without appeal to controversial shifty views, such as contextualism or pragmatic encroachment, one can have justification to believe that β , violate the norm (α only if β) by α ing when not- β and still not be blameless. But unless we can formulate a plausible sufficient bridge principle for blameless norm violation as well as blameworthy norm conformity, then we cannot defend the original conclusion that, for any norm, whether one conforms to the norm comes apart from whether one is blameless in both directions. At best, we might be able to show that, for any norm, one can conform to it in a blameworthy way but not that, for any norm, one can violate it blamelessly. So, we cannot defend the original pessimistic conclusion in both directions.

6. Habit.

So far, the argument has focused on bridge principles using epistemic or doxastic conditions. But, it might be wondered if we could instead reformulate the bridge principles using other conditions. Indeed, some suggest that norm conformity can come apart from whether one is blameless because whether one conforms to a norm can come apart from whether one employs a good habit, that is a habit generally conducive to norm conformity, or to conforming to the relevant norm.

For instance, Lasonen-Aarnio (2010) considers a subject who knows that p despite having higher-order evidence which suggests that she doesn't know that p . If the subject asserts that p then she conforms to the knowledge norm of assertion but Lasonen-Aarnio suggests that she uses a bad habit and should thus be negatively evaluated. For even though

the evidence is misleading in this particular case, evidence is not typically misleading. Thus if the agent has a habit of ignoring evidence that she is not conforming to the knowledge norm of assertion, then this will generally tend to lead her to not conform to the knowledge norm. However, even if appeal to habits can plausibly provide an account of blameworthy norm-conformity, it has more difficulty providing a plausible account of blameless norm violation. In particular, that one employs a habit which is conducive to norm-conformity is not sufficient to make a norm-violation blameless. For, even if one knows that the habit is generally conducive to conformity with the norm, it might leave open too high a chance that one is failing to conform on this occasion. For instance consider the following norm: shoot if and only if it's an enemy plane. Even if in shooting a certain plane, one employs a habit which is generally conducive to shooting only enemy planes, this might leave open a significant chance that the relevant plane is in fact a passenger plane full of tourists going on holiday. That one knows that the habit is generally conducive to norm conformity would not necessarily render one blameless for shooting the jet out of the sky when one knew the risks. In general, any notion of a good habit as conducive to norm conformity but which is fallible is likely to be insufficient to render norm violations blameless when the stakes are high enough. For reasons similar to those discussed above, we can't hope to avoid this problem by appeal to the idea that when it is not a serious matter whether one conforms to the norm, a good habit is sufficient for a norm violation to be blameless. For such a bridging claim would yield the pessimistic conclusion only when combined with the claim that every norm has cases in which it's not a serious matter whether one conforms to it. But, we have already seen that that's implausible.

Someone might try and get around this first problem by suggesting that whether a habit counts as "good" or "bad" depends not merely on its reliability and so whether it is generally conducive to norm-conformity, but also how serious the relevant norm violation

would be. Thus, the idea might be that what counts as a "good habit" for conforming to the norm of shooting only enemy planes must be more reliable than what counts as a "good habit" for conforming to the norms of etiquette. Whatever the plausibility of this idea, it doesn't look as if it could be used in an argument for the pessimistic conclusion. It seems hard to ensure that the relevant bridging principle is true unless we understand a "good habit" as a habit whose employment ensures that one's action is blameless; and a "bad habit" as a habit whose employment ensures that one's action is blameworthy. But, if this is how the notion of a good habit is understood, then we don't so much have an argument for the pessimistic conclusion, but rather just an un-argued assertion of it. For, the supposed "argument" would start with the claim that, for any norm, whether one conforms to it can come apart from whether one employs a good habit or a bad habit, where a good habit is simply a habit employment of which renders one blameless; and a bad habit is simply a habit employment of which renders one blameworthy.

A different way to try and get around difficulties in defending the pessimistic conclusion by appeal to the notion of habit would be to shift to a non-probabilistic notion of a good/bad habit. Some defenders of the knowledge norm employ a non-probabilistic notion of a bad habit to deal with cases in which a subject knows that β and conforms to the norm (α if and only if β) by α ing but nonetheless seems blameworthy. If a subject knows that β , and knowledge is sufficient for evidence, then there is no epistemic chance that she is violating the norm (α if and only if β) by α ing. Nonetheless, the subject might use a habit which is poor in the sense that it easily leads to disastrous consequences in other similar cases (Hawthorne and Stanley 2008, Lasonen-Aarnio 2010, Hawthorne and Srinivasan 2013). Even if the subject knows that β , she cannot always distinguish when β holds and when it doesn't. So, by α ing when she believes that β , she is employing a habit which she may also employ when she believes that β but in fact not- β , resulting in her violating the relevant norm. For instance,

Hawthorne and Srinivasan (2013) appeal to this notion of a bad habit in explaining why an engineer who knows that she knows that the bridge only needs 10 struts is nonetheless blameworthy in ignoring the disagreement of a more junior but still competent engineer that more struts are required. Since the engineer cannot always distinguish when she knows and when she doesn't, by dismissing the disagreement she is employing a habit which could easily lead to disastrous consequences in nearby cases.

One might wonder if this non-probabilistic notion of a bad habit could be used to support the pessimistic conclusion. For instance, someone could appeal to the notion of a bad habit to try and argue that every norm allows for cases of blameworthy norm conformity as follows:

1. For any norm, one could conform to that norm while employing a habit which could easily lead to disaster in similar cases.
2. If one conforms to a norm while employing a habit which could easily lead to disaster in similar cases, then one is blameworthy.
3. Thus, for any norm, one could conform to that norm and nonetheless be blameworthy.

In considering the plausibility of this argument, one might first worry that 1) is not true for every norm since there are some norms which don't seem disastrous to break. However, even if violating a norm is not in itself important, in some circumstances the consequences of violating it may nonetheless be very important. For instance, even if violating a norm of punctuation is not in itself important, in certain circumstances it could lead to disastrous consequences (perhaps, you have been kidnapped by a pedantic psychopath who swears to kill you if you make the least punctuation mistake). So, perhaps 1) is defensible by considering the further consequences of norm violations. Nonetheless, we might question the bridge principle 2).

At the heart of the bridge principle is the idea that employing a good/bad habit goes with one's being blameless/blameworthy. However, the relevant notion of habit doesn't fit our judgements about whether one is blameworthy. To see the problem, notice that, intuitively, there are cases in which someone violates a norm leading to very bad consequences but is nonetheless blameless. For example, suppose that an engineer violates the knowledge norm of belief. She believes that 12 struts are enough for a bridge, even though this belief is not knowledge since false. As a result, the bridge is not sturdy enough and quickly collapses causing many deaths. Nonetheless, we can imagine circumstances in which the engineer is nonetheless blameless: perhaps, she did all the calculations and doublechecking we could expect but some unknown and unpredictable problem had the result that 12 struts were not enough. In believing that 12 struts are enough it seems that she used a habit of belief formation. Furthermore, given human fallibility, she cannot always discriminate cases in which 12 struts enough and those in which that's not the case. Indeed, it's her failure to be able to perfectly discriminate cases in which 12 is enough and those in which it's not which leads to disaster in this case. Thus, in the relevant sense, her norm violation manifests a bad habit: a habit that might lead to disaster. Furthermore, there are nearby cases in which it leads to disaster—indeed it leads to disaster in the actual case. Thus, appeal to the relevant notion of bad habit has the problematic consequence that she is blameworthy.

In conclusion, it seems that appeal to neither a probabilistic nor a non-probabilistic notion of a bad habit provides a good defence of the pessimistic conclusion.

7. An induction?

We have seen that the original argument from anti-luminosity to the pessimistic conclusion fails and that it cannot be successfully reformulated by appeal to other doxastic or epistemic

conditions even the notion of habit. Of course, this is compatible with there being many cases in which whether one conforms to a norm comes apart from whether one is blameless.

Indeed, a defender of the pessimistic conclusion might attempt to use this fact in a different argument than the deductive ones we've considered so far. In particular, she may suggest an inductive argument to the conclusion that, for any norm whatsoever, whether one conforms to it can come apart from whether one is blameless. Thus, she might argue as follows:

1. For all norms so far considered, whether one conforms to a norm can come apart from whether one is blameless.
2. Thus (by induction), for any norm whatsoever, whether one conforms to it can come apart from whether one is blameless.

This inductive argument faces problems. First, the first premise is not established merely by noting that, in many cases, whether one conforms to a norm can come apart from whether one is blameless. Rather, it would need to be demonstrated that for all norms considered so far, one can violate them blamelessly and conform to them in a blameworthy way. Second, and more importantly, the induction would be undermined if we have reason to doubt that we can generalise the kind of explanation of why, in cases so far considered, norm-conformity can come apart from whether one is blameless. However, when we start to consider the different reasons why norm-conformity can come apart from whether one is blameless, these reasons often fail to generalise to all norms whatsoever. In particular, we will see that many of the explanations of why the violation of a norm governing action may be blameless don't easily apply to norms of belief.

It is standard to suggest that the violation of a norm governing action may be blameless if that norm is overridden by another norm, or the violation involved maximising expected utility or if violating the norm does not exhibit bad will. For instance, a failure to meet a friend for lunch as promised may be blameless if overridden by a more important

norm (to save the life of a drowning child), or if done without bad will (e.g. my watch is running slow and I don't notice I'm running late). And it is sometimes suggested that a failure to maximise actual value can be excused if one maximises expected value. But these kinds of reasons why a norm violation may be blameless fail to extend easily to violations of the norm of belief. It is not clear how to apply the notion of a good or bad will to violations of the norm for belief. For instance, if I fall for the gambler's fallacy in concluding that the next coin will be heads, I might nonetheless be blameworthy even if in no sense was my belief the product of bad will. Further, appeal to overriding norms or maximising expected value doesn't obviously render violation of the norm for belief blameless. On a standard view, the norm of belief is not overridden either by norms of prudence or morality. Instead, the norm of belief is incommensurable with the norms of prudence and morality (e.g. Kelly 2013). Thus, the standard view about Pascal's Wager is that I have a prudential reason to believe in God against the evidence and an epistemic reason not to do so where these are incommensurable. Relatedly, that believing in God against the evidence would maximise expected value does not make the violation blameless. Indeed this conclusion remains true even if we restrict expected value to expected epistemic value. In a variant of Pascal's Wager, suppose that the Demon promises that if you believe that p he will reward you with outstanding epistemic abilities and conducive conditions for their use whereas if you don't, he will ensure that you have poor epistemic abilities and that conditions are not conducive. Even if believing that p would maximise expected epistemic value, if there is no evidence for p , one is nonetheless epistemically blameworthy in believing that p . Thus, maximising even expected epistemic value is not enough to make a violation of the norm for belief blameless.¹³

¹³ One might wonder whether Williamson's account of excusable norm violation can be used to defend the argument to the pessimistic conclusion. On his account, one is excusable for a

In conclusion, it's far from clear that one could mount an inductive argument to the pessimistic conclusion from everyday cases in which norm conformity comes apart from whether one is blameless. For when we consider the different kinds of explanations why norm conformity can come apart from whether one is blameless, many of these explanations do not generalise to all norms whatsoever. This undermines the idea that we can easily infer from the observation that norm conformity regularly comes apart from whether one is blameless to the conclusion that, for any norm whatsoever, norm conformity comes apart from whether one is blameless.

8. Blame and replying to counterexamples.

Despite the problems facing arguments to the pessimistic conclusion, norm-conformity often does not align with blamelessness. As a result, the failure of arguments to the pessimistic conclusion still leaves open the possibility of defending some suggested norm against counterexample by appealing to the idea that norm conformity can come apart from blamelessness. In the light of our discussion then, let us examine whether putative counterexamples to the knowledge norm of assertion can be deflected by appeal to the way in which norm conformity comes apart from being blameworthy.

norm violation if one does what someone who is disposed to conform to that norm would do in the circumstances. Employing this to defend the pessimistic conclusion would involve defending 1) for any norm, one can violate it yet do what someone who is disposed to conform to that norm would do in the circumstances; and 2) whenever that occurs, one blamelessly violates the norm. While 1) might seem plausible, it's not entirely clear how to prove it. Further, I raise problems for 2) in Brown forthcoming.

First, let's consider examples in which, intuitively, it's fine for the subject to assert that p even though she doesn't know that p , such as Gettier cases and cases in which the subject falsely but justifiably believes that p . The standard way for defenders of the knowledge norm to deal with such cases is to say that they involve a blameless violation of the knowledge norm. In order to show that this appeal to the distinction between being blameless and norm-conformity is not ad hoc, it is preferable if the defender of the knowledge norm can appeal to a general condition sufficient for blameless norm violation. Our earlier discussion cast doubt on several putative such conditions, including appeal to not being in a position to know. But, we did find one plausible such condition: if one violates the norm (α iff β) by α ing when not- β where it's not a very serious matter whether one violates the norm and one justifiably believes that β , then one's violation is blameless. Of course to apply this condition, the putative violations of the knowledge norm of assertion need to be not very serious. However, that seems plausible. Or, at least it seems that the best kinds of examples to cast doubt on the necessity of knowledge for assertion involve low-stakes situations.

To see this, compare two different putative counterexamples to the necessity of knowledge, one low stakes and one high stakes. In both cases, I visited the bank several weeks ago and saw that it was open on a Saturday. As a result, I believe that the bank is open on Saturday and take it that this belief is knowledge. In the low-stakes version of the case, a friend casually asks me about the bank's opening hours explaining that since her business is non-urgent, she would prefer to go on Saturday if the bank's open then rather than wait in the long queue on Friday. In this case, it seems fine for me to reply by saying "Yes the bank is open on a Saturday" even though, as it turns out, what I said is false since the bank has recently unexpectedly changed its hours. By contrast, in the high-stakes case, my friend makes it clear to me that it's a high-stakes matter for her whether the bank is open on

Saturday because of the urgency with which she needs to pay money into her account. In this high-stakes case, it would not seem fine for me to say “Yes the bank is open on a Saturday” without qualification. Instead, I should have said something like “Well, I was there a couple of weeks ago and it was open, but I guess you better check”. Similar points hold for a variant case in which I’m instead in a Gettier case in which I justifiably and truly believe that the bank is open on Saturday on the basis of the recent visit although I fail to know because many local businesses have recently changed their hours although the bank has not.

It seems, then, that a defender of the necessity of knowledge for assertion could use the earlier discussion to draw on a plausible sufficient condition for blameless norm violation in attempting to deal with putative counterexamples in which a subject doesn’t know yet seems fine to assert. The defender will say that, in these cases, the subject violates the knowledge norm but is blameless because the violation is not a very serious matter and she justifiably believes that she knows. Of course, to point out that the defender of the necessity of knowledge for assertion can make this move is not to say that it is fully satisfactory. For some have argued that in such cases it is problematic to say that the subject is merely excused for her assertion rather than fully justified. I do not attempt to adjudicate that issue here, but rather simply point out that the defender of the necessity of knowledge for assertion can attempt to deal with such cases by appeal to the notion of blameless norm violation.

Now consider putative counterexamples to the sufficiency of knowledge for being in a good enough epistemic position for assertion. In these cases, one does know the relevant proposition but it nonetheless seems problematic for one to assert it. Defenders of the knowledge norm would like to deal with such cases by appeal to the idea that one can conform to a norm and nonetheless be blameworthy. In some of these cases, they can appeal to the plausible principle concerning blameworthy norm conformity defended above (S5): if one conforms to the norm (A if and only if β) by doing A when β and one justifiably believes

that not- β , then one is blameworthy for conforming. In some putative counterexamples to the sufficiency direction of the knowledge norm for assertion, one knows that p although one's evidence suggests that one does not, say because one has misleading evidence that not p, or misleading evidence that the process which produced the belief that p is unreliable, or misleading evidence that one's evidence does not support that p etc. For example, some suggest that in a long deduction in which one comes to believe the conclusion by competent deduction from known premises, one's belief in the conclusion constitutes knowledge despite the fact that one knows one has likely made a mistake given the length of the inference and so likely doesn't know the conclusion. In asserting the conclusion one in fact conforms to the knowledge norm of assertion even though one has justification to believe that one is not conforming. The defender of the knowledge norm of assertion would claim that this is a case of blameworthy norm conformity.

There remain a range of putative counterexamples to the sufficiency of knowledge for assertion which don't have this structure including Lackey's second-hand knowledge case and the high-stakes cases. In Lackey's case, it seems problematic for the consultant to assert straight out that the patient has cancer even though she knows this, given that she learned it from a junior colleague and she hasn't even looked at her notes. In this case, it still seems problematic for the consultant to assert that the patient has cancer even if she justifiably believes that she does know that the patient has cancer. Similarly, in Brown's surgical case, it can seem problematic for a speaker to assert that p even if she knows that p and knows that she knows that p.

In dealing with putative counterexamples to a knowledge norm in which a speaker knows and knows that she knows, some attempt to appeal to habits. Earlier, I cast doubts on whether appeal to habit could provide a satisfactory account of blameless norm violation. So, the defender of the knowledge norm of assertion might instead try to deal with such cases by

appeal to the idea that some norm overrides the norm of assertion. For example, it may be suggested that, in Lackey's case, there is some kind of professional norm according to which the consultant shouldn't assert straight out that the patient has cancer if she hasn't even looked at her notes. It's far less clear that all the high-stakes cases can be dealt with in the same way. For, they need not all involve special professional norms. For instance, in the affair case, there is no professional relationship governing the assertion. While one might think that there are certain special norms governing the relations between friends, it looks as if we can set up a variant affair case in which the informant is not a friend and so has no special such relations.

In sum, then, the discussion of the pessimistic conclusion has helped clarify when someone can reasonably attempt to defend the knowledge norm of assertion by appeal to the idea that blamelessness can come apart from norm conformity. Some defenders of the knowledge norm of assertion have tried to deal with putative counterexamples by appeal either to lack of a position to know or appeal to habits. In fact, neither of these is satisfactory. Instead, it seems that they need to appeal either to evidence about whether the relevant condition which figures in the norm is satisfied, or the norm being overridden by rival norms.

8. Conclusion.

I have been discussing the attempt to defend the knowledge norm of assertion by appealing to the idea that whether one conforms to a norm can come apart from whether one is blameless. This idea has been appealed to in responding to putative counterexamples to the knowledge norm of assertion. In addition, defenders of the knowledge norm have endorsed the pessimistic conclusion that, for any norm whatsoever, whether one conforms to it can come apart from whether one is blameless. As a result, defenders of the knowledge norm claim that it is no objection to the knowledge norm that it allows norm conformity to come apart from

whether one is blameless since, they say, every norm does so. I've challenged the argument to the pessimistic conclusion. Rejecting the pessimistic conclusion is, of course, compatible with accepting that norm-conformity often comes apart from whether one is blameless. So the rejection of the pessimistic conclusion still allows the defender of the knowledge norm to deal with putative counterexamples by invoking the distinction between norm-conformity and being blameless. However, the discussion of the pessimistic conclusion illuminated the conditions under which it is reasonable to claim that norm conformity comes apart from blamelessness, and so when it is legitimate to appeal to the distinction between norm-conformity and blamelessness in dealing with counterexamples to one's favoured norm.

Brown. Forthcoming. *Fallibilism: evidence and knowledge*. Oxford: Oxford University Press.

Brown. 2017. "Blame and wrongdoing". *Episteme* 14, 3: 275-296.

Brown. 2010. "Knowledge and assertion. *Philosophy and Phenomenological Research* 81,3: 549-566.

Calhoun. 1989. "Responsibility and reproach". *Ethics* 99:389-406.

Christensen. 2007. "Does Murphy's Law apply in epistemology? Self-doubt and rational ideals" *Oxford Studies in Epistemology* 2:3-31.

Fitzpatrick. 2008. "Moral responsibility and normative ignorance: answering a new sceptical challenge". *Ethics* 618, 4:589-613.

Harman. 2011. "Does moral ignorance exculpate?" *Ratio* 24:443-468.

Hawthorne and Stanley. 2008. "Knowledge and action". *Journal of Philosophy* 105:571-90.

Hawthorne and Srinivasan. 2013. "Disagreement without transparency: some bleak thoughts." In Christensen and Lackey, *The Epistemology of Disagreement*, Oxford: Oxford University Press, 9-30.

- Kelly. 2003. "Epistemic rationality is instrumental rationality: a critique". *Philosophy and Phenomenological Research* LXVI, 3:612-640.
- Lackey. 2007. "Norms of assertion". *Nous* 41:595-626.
- Lasonen-Aarnio. 2010. "Unreasonable knowledge". *Philosophical Perspectives* 24.
- Lasonen-Aarnio. Forthcoming. "Enkrasia or evidentialism: learning to love mismatch".
- Levy. 2009. "Culpable ignorance and moral responsibility: a reply to Fitzpatrick". *Ethics* 119:729-741.
- Littlejohn. 2015. "Stop making sense? On a puzzle about rationality". *Philosophy and Phenomenological Research*. DOI: 10.1111/phpr.12271.
- Littlejohn. Forthcoming. "A plea for epistemic excuses". In Dorsch and Dutant forthcoming.
- Mason. 2015. "Moral ignorance and blameworthiness". *Philosophical Studies* 172:3037-3057.
- Peels. 2010. "What is ignorance?" *Philosophia* 38:57-67.
- Reed. 2010. "A defence of stable invariantism" *Nous* 44: 224-244.
- Rosen. 2003. "Culpability and ignorance". *Proceedings of the Aristotelian Society* 103, 1:61-84.
- Rosen. 2004. "Scepticism about moral responsibility" *Philosophical Perspectives* 18, 1:295-313.
- Sher. 2009. *Who knew? Responsibility without awareness*. OUP: Oxford.
- Srinivasan. 2015. Normativity without Cartesian privilege. *Philosophical Issues* 25:273-299.
- Sutton. 2007. *Without Justification*. Massachusetts: MIT Press.
- Titlebaum. 2015. "Rationality's fixed point (or: In defence of right reason)." *Oxford Studies in Epistemology*, 5.
- Weatherson. Forthcoming. *Normative Externalism*.
- Williamson. 2000. *Knowledge and Its Limits*. OUP: Oxford.

Williamson. 2013. "Gettier cases in epistemic logic". *Inquiry* 56:1-14.

Williamson. Forthcoming. "Justifications, excuses, and sceptical scenarios". In Dutant and Dorsch (eds.), *The New Evil Demon*, Oxford: Oxford University Press.

Worsnip. 2015. "The conflict of evidence and coherence". *Philosophy and Phenomenological Research* DOI: 10.1111/phpr.12246.

Zimmerman. 2008. *Living with uncertainty: the moral significance of ignorance*. Cambridge: Cambridge University Press.

Zimmerman. 2014. *Ignorance and Moral Obligation*. Oxford: Oxford University Press.
