Inquiry, 2013
Vol. 56, No. 4, 400–415, http://dx.doi.org/10.1080/0020174X.2013.816087

Discussion

Williamson on Gettier Cases in Epistemic Logic and the Knowledge Norm for Rational Belief: A Reply to a Reply to a Reply

STEWART COHEN AND JUAN COMESAÑA

University of Arizona, USA

(Received 30 May 2013)

I. Gettier Cases in Epistemic Logic

In our reply to Williamson, we offered models where traditional Gettier cases arise but which differ from Williamson's models in three related respects: first, non-traditional, fake-barn style Gettier cases do not arise in our models; second, knowledge iterates in our models; and, finally, our models do not have the Moorish consequences of Williamson's models.[1] Predictably, Williamson thinks that our models have no value whatsoever. Our reasons for preferring them 'evaporate', he says, and they have 'epistemologically absurd consequences'. We think that Williamson exaggerates.

We first address the allegedly absurd consequences our models have. These are based on the fact that they allow for what Williamson calls 'cliff-edge knowledge': 'Say that one has *cliff-edge knowledge* of the real value of the parameter when it is in fact *e* if and only if one either knows that it is at least *e*

Correspondence Address: Juan Comesaña, Department of Philosophy, University of Arizona, Tucson, AZ 85721, USA. Email: comesana@email.arizona.edu

[1]Cohen and Comesaña, 'Williamson on Gettier Cases'. That paper is a reply to Williamson, 'Gettier Cases in Epistemic Logic'.

or knows that it is at most *e*.'[2] Williamson says that '[g]iven normal limits on human powers of perception and estimation, in a vast range of cases none of us have such cliff-edge knowledge', and later he calls the existence of cliff-edge knowledge an 'epistemologically absurd' consequence.[3]

As is to be expected after that negative characterization, Williamson's models do not display cliff-edge knowledge. By contrast, they satisfy the following principle, which is incompatible with cliff-edge knowledge:

> MFE: If the real value of some quantity is v, then one doesn't know that it is not v-1 and one doesn't know that it is not v + 1.

That principle is related to the following one by Weatherson, which Williamson cites approvingly: 'You cannot come to know something about the real value of some quantity on the basis of a suboptimal measurement that you could not also know on the basis of an optimal measurement.'[4] It might seem that cliff-edge knowledge runs afoul of Weatherson's principle. And, indeed, in many cases it will be that way. Suppose that a certain thermometer has 2 degrees of margin of error—that is to say, if the thermometer measures m, then it is at least $m - 2$ and at most $m + 2$. The thermometer is of course not *logically* (or even physically) guaranteed to satisfy that constraint, but it is well manufactured and it does indeed satisfy it. Now suppose that on one occasion the thermometer measures 60 degrees, although it is in fact 58 degrees. On the basis of the thermometer's measurement, I believe that it is not 57 degrees. Do I thereby know that it is not 57 degrees°? The argument to the contrary based on Weatherson's principle goes as follows. Had the thermometer given a perfectly accurate reading (had it read 58), I would not have known that it is not 57 degrees. Therefore, by appeal to Weatherson's principle, I do not know in the actual case, where the thermometer gives an inaccurate reading, that it is not 57 degrees. However, Weatherson's principle is not applicable to all the cases of cliff-edge knowledge, and it is in any case open to doubt.[5]

Suppose that you buy a thermometer that, instead of displaying a precise temperature reading, displays an interval of, say, 5 degrees. The thermometer may work in a couple of different ways. First, it may be that it differs from regular thermometers only in how it displays the temperature: maybe it makes a point-value estimation of the temperature and then adds 2 degrees on either side of that estimation. But this is not how this thermometer works,

[2]Williamson, 'Response to Cohen, Comesaña', 81. The 'if' after the 'only' is missing in Williamson's text. Beware: Williamson's definition is ambiguous. The main connective is the 'if and only if', not the 'when'. Thanks to Timothy Williamson for discussion of this point.

[3]Williamson, 'Response to Cohen, Comesaña', 81, 83.

[4]Weatherson, 'Margins and Errors', 65.

[5]Thanks to Brian Weatherson for discussion about the following paragraph.

for it differs from regular thermometers more radically. The only estimations that it makes are intervals, and these intervals are not derived from more precise estimations. Moreover, there are no assurances that values in the center of the interval are any more likely than off-center values. The thermometer is guaranteed to work correctly, in the sense that the real temperature is guaranteed to be within the interval the thermometer displays (again, this guarantee is of course not *logical*). You take a look at this thermometer and it reads [58, 62]. The temperature is 58 degrees. Based on the thermometer's reading, you conclude that it is not 57 degrees. Do you know this? In this case, at least, a positive answer seems warranted. Moreover, notice that Weatherson's principle is not applicable to this thermometer. We cannot say that, had the thermometer displayed [56, 60], its measurement would have been more accurate. For this thermometer, off-center measurements are as accurate as center measurements. When it is 58 degrees, any of the five readings compatible with this temperature is as good as any other. Therefore, the fact that, had the thermometer read [56, 60], we would not have known that it is not 57 degrees does not impugn the fact that we do actually know this. If the deliverances of our perceptual organs work like this thermometer, then Weatherson's principle is not applicable to them, and the existence of cliff-edge knowledge is a welcome result of our models (and their incompatibility with Williamson's a defect of his). Now, of course, it is not likely that our organs work this way *over this particular interval*. When it feels like it is between 60 and 70 degrees, it is usually more likely to be 65 than 60 degrees. But this does not mean that values towards the middle of an interval will be more likely to be correct *for every interval*. And, of course, if there is *any* interval over which this indifference holds, our point stands: there will be cliff-edge knowledge and our models better capture it.

Moreover, it is not at all clear that Weatherson's principle, and Williamson's MFE, are true. Cliff-edge knowledge may be surprising, perhaps even counterintuitive, but some consequences of MFE are at least as counterintuitive as cliff-edge knowledge. Take an example from Williamson himself. Suppose that we are judging by naked eye the height of a tree which is in fact 13.793 m high. If cliff-edge knowledge is allowed, then there can be cases where we can know that it is not 13.792 m or that it is not 13.794 m. Of course, that would happen only at the edges of our discriminatory range—that is to say, that will happen only if we are dealing with quasi-skeptical scenarios. Williamson thinks that this fact—that cliff-edge knowledge arises only when we are dealing with quasi-skeptical scenarios—is a further indictment of our models. On the contrary, we think that this is the expected behavior. But what we want to remark on is the fact that Williamson's models show the flip-side of this behavior. According to Williamson, when the tree both is and appears to be 30 m high, we know, say, that it is not 13.793 m high (let this be the limit of our perceptual capacities). However, when the tree has the same appearance but is 29.999 m high, we no longer know that it is not 13.793 m high. Say

that one has 'extrinsically fragile' knowledge of the real value of a parameter *e* just in case one can know that *e* is not *n* when *e* is *m*, but not when *e* is $m - \varepsilon$, where the distance between *n* and *m* is dictated by your discriminatory capacities and ε is arbitrarily small. Of course, we all agree that knowledge is *intrinsically* fragile—our knowledge that the real value is not *n* disappears as soon as the real value *is n*. But Williamson's models allow for extrinsic fragility, whereas ours do not, and extrinsically fragile knowledge strikes us as at least as counterintuitive as cliff-edge knowledge. Moreover, a model allows for cliff-edge knowledge if and only if it does not allow for extrinsic knowledge. The interest of our models reside precisely in the fact that they allow us to weigh the relative (de)merits of cliff-edge vs. extrinsically fragile knowledge.

Williamson approvingly cites Goodman as saying that the point of the margin for error is to achieve safety, and thus that the margin for error in our models does not deserve the name. We ourselves doubt the feasibility of a safety constraint on knowledge, at least when the constraint is thought of as being understandable independently of knowledge.[6] Williamson will not want to understand safety in that way, though, resting content with understanding safety in terms of knowledge. If so, then our judgments about whether a belief is safe cannot be made independently of our judgments about whether it amounts to knowledge. In that case, to say that the point of margins for error is to achieve safety is just to say that the point of margins for error is to avoid generating knowledge where there is not. We of course have no objections to that constraint, nor do our models violate it.

Williamson has a second, related complaint about our models, namely that they display luminosity. But notice that our models do not really display luminosity about the real temperature. Rather, what is true in our models is that there are conditions and appearances such that, if those conditions obtain and the appearances remain the same, we know that those conditions obtain. This result, however, is due as much to the assumption (which we take from Williamson) that appearances are luminous as to the fact that in our models knowledge iterates. Moreover, notice that Williamson's models exhibit 'coziness':[7] whenever the real value is within a certain interval, we know that it is within a larger interval.

So much for the allegedly absurd consequences our models have. Let us now turn to why Williamson thinks that our reasons for preferring our models to his 'evaporate'.

Our first reason was that our models do not display the fake-barn style of Gettier cases that Williamson's do. Before turning to Williamson's observations regarding this, one minor point. Williamson claims in his reply that the generation of fake-barn cases 'requires the minor assumption #1 in GCEL,

[6]See Comesaña, 'Unsafe Knowledge'; Cohen, 'Luminosity, Reliability, and the Sorites'.

[7]The 'coziness' terminology is from Hawthorne, 'Knowledge and Evidence'.

but not the major assumption #2'. #1 is the assumption that knowledge is imperfect: not everything is known about the real value of the parameter even in the most epistemologically propitious world. #2 is the assumption that knowledge decreases as the real value and the apparent value diverge. Williamson's explanation of how a fake-barn case arises is as follows. First, Let <f, f> be a world where appearances and reality match. Second, let <f, g> be a world where appearances and belief are as in <f, f>, but where less is known. Crucially, <f, g> must also be a world that is doxastically possible at <f, f>—that is to say, if we are in <f, f>, then for all we believe we are in <f, g>. Then it can be shown that, at <f, g>, there is a true (and justified, as are all beliefs in these models) belief that is not known. Now, the crucial step is the second: the existence of a world that is doxastically but not epistemically possible at <f, f>. This goes clearly beyond #1, which, as Williamson knows, is satisfied in our models. What else is needed to guarantee this result? The major assumption #2 will do, and Williamson does not provide us with a weaker sufficient assumption. It is not true, therefore, that fake-barn cases can be generated by relying on #1 alone.

Williamson adduces three considerations against us: (1) that our observation that many epistemologists are not as convinced by fake-barn cases as they are by traditional Gettier cases is trumped by the observation that most of them are; (2) that such sociological observations do not in any case constitute an argument; and (3) that the kind of fake-barn cases that arise in his models are importantly different from the original fake-barn cases (that may be why Williamson coins the terminology of 'purely veridical' and 'impurely veridical' Gettier cases).

The first two points can be taken care of swiftly. Williamson is looking for an argument where there isn't any, and so he thinks he sees a particularly bad one. We of course do not mean to be saying that because some epistemologists (we disagree with Williamson about the proportions here, but that does not matter) do not find fake-barn cases as convincing as the original Gettier cases, then there are no fake-barn cases. We meant simply to point out that many epistemologists (ourselves included) will find our models interesting precisely because they do not find fake-barn cases as convincing as traditional Gettier cases. For a convincing case against fake-barn cases, we refer again to Gendler and Hawthorne.[8]

With respect to Williamson's claim that the kind of fake-barn cases that arise in his models are importantly different from the original fake-barn cases, we beg to differ. Certainly, there are many differences between them, but the question is whether those differences matter. As we see it, a crucial feature of both kinds of cases is that doxastic possibility outstrips epistemic possibility even in the most epistemically propitious worlds (and we must here remember

[8]Gendler and Hawthorne, 'The Real Guide to Fake Barns'.

that we are operating under the assumption that all beliefs are justified). The dialectical situation here gets complicated, because presumably Williamson does not himself think that this is so—rather, he thinks that *according to an internalist interpretation of doxastic possibility*, doxastic possibility outstrips epistemic possibility even in the best epistemic worlds. But another thing that our models show is that this need not be so.

Our second reason for preferring our models to Williamson's was that ours do not generate knowledge of Moorish propositions like '*p* and I do not believe I know *p*' and '*p* and I am almost certain that I do not know *p*'. In his reply, Williamson concedes that 'such conjunctions sound bad', but he goes on to argue that they can still be known. Assuming the knowledge account of the norm of assertion, those propositions

> are tantamount to '*p* and I do not believe that I have warrant to assert *p*' and '*p* and I am almost certain that I do not have warrant to assert *p*' respectively. They resemble the commands 'Stand to attention!—and I do not believe that I have authority to order you to stand to attention' and 'Stand to attention!—and I am almost certain that I have no authority to order you to stand to attention'.[9]

These assertions sound bad, Williamson says, even if the speaker is entitled to make them. So, the fact that they are generated by Williamson's models is no reason against them.

But we are not sure what Williamson means when he says that, for instance, '*p* and I do not believe I know *p*' is 'tantamount' to '*p* and I do not believe that I have warrant to assert *p*'. Those propositions are certainly not equivalent, not even if we assume that Williamson is right about the norm of assertion. Of course, if Williamson is right about what warrants assertions, then those propositions have the same truth-value—but Williamson cannot be assuming that simply by providing a paraphrase with the same truth-value as the offending proposition and then explaining why there is nothing wrong with the paraphrase he is thereby explaining away the offensiveness of the original proposition. Of particular notice here is the fact that we are assuming that Williamson is right about the norm of assertion, but not that the speaker is. Normally, speakers have no opinions about the norm of assertion, and even those who do are not all Williamsonian. It is still strange for those speakers to say '*p* and I do not believe I know *p*', and the explanation cannot be that what they mean is '*p* and I do not believe that I have warrant to assert *p*'. Maybe Williamson will reply that even those who do not believe in his knowledge norm have nevertheless somehow internalized it, and that this explains why assertions like that sound bad even for them. This still does not explain

[9]Williamson, 'Response to Cohen, Comesaña', 82.

why the assertions sound bad even to those with anti-Williamsonian convictions about assertion, unless they too have internalized the truth against their explicit views.

Moreover, a crucial fact is that, in Williamson's models, the subject *knows* that the offending proposition is true—that is to say, the subject knows that: *p* and she does not believe she knows *p*. This consequence comes close to generating the so-called Fitch–Church knowability paradox ('I know that [*p* and I do not know *p*]'), whereas the corresponding paraphrase (that the subject knows that *p* and she does not believe she has the authority to assert *p*) comes nowhere near. This is a strong reason to think that Williamson's paraphrases do not get to the heart of the matter, and thus our models, which do not generate the offending propositions, remain interesting.

II. Rationality and the Knowledge Norm

In his original paper, Williamson claims that '*in a strongly externalist sense* a belief is *fully* justified only if it constitutes knowledge'.[10] In our response, we spent some time puzzling about what he was referring to by 'strongly externalist justification' and how to distinguish it from rational justification, i.e. rationality:

> it is worth asking what Williamson's notion of strong externalist justification comes to. Are there two different notions of genuine justification—rational justification, and strongly externalist justification? Perhaps, but once we are clear about rational justification, it is unclear what Williamson is referring to by 'strongly externalist justification'. This is technical vocabulary that he does not define.[11]

We argued that either strong externalist justification concerns what he calls 'a constitutive norm' or it concerns a norm of rationality. But if it concerns a constitutive norm, his view is consistent with the view he opposes. What he calls 'internalism' about justification is best construed as a view about rationality. In his response to us, Williamson says the following: 'For Cohen and Comesaña, justified belief in the relevant sense is rational belief. They claim my view is false of rational belief.'[12] While we do understand justified belief as rational belief, Williamson neglects to mention our argument that he himself has no interesting notion of justification distinct from it.[13] It is indeed true that we claim his view is false of rational belief simply because we claim

---

[10]Williamson, 'Gettier Cases in Epistemic Logic'. A similar view has been defended by Sutton, *Without Justification*.

[11]Cohen and Comesaña, 'Williamson on Gettier Cases', 21.

[12]Williamson, 'Response to Cohen, Comesaña'.

[13]For more on this point, see Cohen, 'Theorizing about the Epistemic'.

his view (if it is to be an alternative to internalism) is false. Williamson proceeds to defend his view as a view of rational belief, without responding to our argument that he must.

We argued that Williamson's view is false because rationality can require one to believe a falsehood. Williamson objects that we never give an argument for this claim. Suppose that evidence consists only of truths:

> since truths never entail a falsehood, they never *fully* support a falsehood (support it in the strongest way); they support it, if at all, only partially. Rationality sometimes requires one to adopt a false belief only if it sometimes requires one to adopt a belief not fully supported by one's evidence; Cohen and Comesaña give no argument that rationality ever requires that.[14]

Williamson is correct that we did not give an argument. We did not think it necessary. Suppose you notice what appears to be a red table staring you in the face, you have no evidence of deception, everyone else around you says they see a red table, yet you fail to believe there is a red table before you. In our view you are paradigmatically irrational, even if unbeknown to you, you do not see a red table.[15] Moreover, our dialectical context with Williamson did not seem to require an argument. In earlier work, Williamson himself presupposes that one can rationally believe a falsehood. Here are a few examples.

> Suppose I rationally believe myself to know there is snow outside; in fact, there is no snow outside.
>
> If I rationally believe today, that tomorrow, I will rationally believe *p*, I cannot deduce *p* today; for all I rationally believe today, tomorrow's rational belief will be based on misleading evidence.
>
> In more difficult contexts, believing truly becomes an aim and we fall back on the method of believing rationally. Rationality becomes a subgoal on the way to truth.[16]

So in *Knowledge and its Limits*, Williamson seems to hold, without argument, the very view he criticizes us for holding without argument.

If it is rational for one to believe *p*, one is at minimum rationally *permitted* to believe *p*. While we framed the issue in terms of rational requirement, our target was Williamson's view that a false belief cannot be fully rational. We could just as easily have made the weaker claim that when one's evidence

---

[14] Williamson, 'Response to Cohen, Comesaña', p. 92.

[15] One might hold that you are irrational only if you are taking some attitude toward whether there is a table.

[16] Williamson, *Knowledge and its Limits*, 262, 235, 179.

is misleading, rationality can permit one to hold a false belief. It is clear that Williamson would still have objected that we failed to give an argument for our view.[17]

More importantly, based on his reply to us, Williamson must deny his previous view that one can be rationally permitted to believe a false proposition. He tells us:

> In a sceptical scenario, beliefs only partially supported by one's evidence appear to be fully supported by it. One appears to be fully conforming one's beliefs to the evidence, but really is conforming only partially. Hence one is deviating from it, even if only slightly. Thus the subject in the bad case is being slightly irrational.[18]

Williamson's view can appear somewhat moderate. After all, he claims only that in the bad case, one is *slightly* irrational. On the contrary, his position entails that in the bad case, one is not rationally permitted to believe.

According to Williamson

(1) In the bad case, believing p is partly irrational

But surely,

(2) One cannot be rationally permitted to believe partly irrationally.

Rationally does not permit partial irrationality. Certainly, some degree of irrationality can be excused, but it is not clear what it would mean for some degree of irrationality to be rational.[19]

From (1) and (2) it follows that

(3) In the bad case, one is not rationally permitted to believe p.

So W's view in *Knowledge and its Limits* cannot be squared with his new view that in a skeptical scenario believing p is to some degree irrational. If in the skeptical scenario one is only partially rational in believing p, then in the skeptical scenario one is not permitted to believe p. Equivalently, in the skeptical scenario, one is rationally required not to believe p. So on Williamson's view, if there are two people in a skeptical scenario where there appears to be a red table, the one who does not believe the table is there is conforming to the requirements of rationality, while the one who does believe the table is there, fails to conform. To put it mildly, this is a surprising result. It runs contrary to what virtually everyone else, including Williamson's former self,

[17]The broader context concerns Williamson's knowledge norm for belief, which clearly *requires* that one know in order to believe. It follows that one is not rationally *permitted* to believe unless one knows

[18]Williamson, 'Response to Cohen, Comesaña', p. 92.

[19]We do describe *people* as rational even if they are occasionally irrational. Our claim is that we do not know what it would mean for a partially irrational *belief* to be rationally permitted.

believes. So we think the pertinent question here is whether Williamson has an argument for *his* view.

At the end of the above quoted passage, Williamson says: 'Thus the subject in the bad case is being slightly irrational.' This certainly suggests that what precedes it is an argument for his view.

The argument appears to be this.

(4) In the bad case, the evidence does not fully support, i.e. entail, one's beliefs.

Thus,

(5) In the bad case, one does not fully conform one's beliefs to the evidence.

Thus,

(6) In the bad case one is (at least) irrational slightly irrational.

(4) is true on the assumption that one's evidence is factive.[20] (5) appears to be meant to follow from (4), though it is not explicitly marked as a conclusion.

Suppose it is a conclusion. What is it to conform fully one's belief to the evidence? In a graded belief framework, one could say that one fully conforms ones belief p to one's evidence e just in case one's degree of belief (credence) equals the degree to which e supports p. But here we are talking about outright (binary) belief.

So what is it for an outright belief to conform fully to the evidence? Conforming is here a normative notion. A belief fully conforms to the evidence just in case it is fully rational on the evidence. If (5) is supposed to follow from (4), then Williamson is presupposing that believing p fully conforms to one's evidence only if one's evidence fully supports, that is entails, p. Given that a belief fully conforms to the evidence just in case it is fully rational on the evidence, this assumption is equivalent to the claim that believing p is fully rational on the evidence only if the evidence fully supports, that is entails, p. But this is just the point at issue. On the view we defend, believing p can be fully rational, and so fully conform to the evidence, even if the evidence does not entail p. An inference from (4) to (5), rather than providing an argument for Williamson's view, simply assumes it.

As I noted, Williamson does not explicitly indicate that (5) is supposed to follow from (4). If, however, (5) is an independent premise, it is not clear what work (4) is doing in the argument. Moreover, (5) is surely too close to his conclusion (6) to serve as a premise in the dialectical context. That one does not fully conform one's belief to the evidence in the bad case is just what we deny.

As we noted earlier, if believing p is not fully rational, then believing p is not rationally permissible. Because of this, Williamson's talk of full and partial

[20] For nonfactive views of evidence, see Schroeder, 'Having Reasons'; Comesana and McGrath, 'Having False Reasons'.

rationality is somewhat misleading. Believing *p* is either rationally permitted or it is not. Here it is important to distinguish between two importantly different relations:

- *Rational permissibility*——a binary relation that holds between a body of evidence and a believing.
- *Evidential support*——a graded relation that holds between a body of evidence and the content of a believing, viz., a proposition.

One is rationally permitted to believe *p* just in case one's evidence supports *p* to the required degree. On Williamson's view, if the evidence does not fully support *p*, then believing *p* is slightly irrational. Given that rational permissibility is binary, this can be understood as meaning only that believing *p* falls just short of rational permissibility.[21] And if believing *p* falls just short of rational permissibility, then believing *p* is not rationally permissible. What we are still looking for is an argument that, if the evidence does not fully support *p*, that is support *p* to the greatest degree, then believing *p* is not rationally permissible.

Does Williamson present any other argument for his view? He does say the following:

> I defend the equation E = K, on which one's evidence is simply what one knows . . . . In their discussion, they do not mention the equation, let alone try to argue against it. Given E = K, the requirement to conform one's beliefs to one's evidence is the requirement to conform one's beliefs to one's knowledge. Believing only what one knows is then a natural understanding of that requirement.[22]

We did not mention E = K because we did not, and still do not, see how it is relevant (beyond the relevance of factivity). We are happy to grant E = K for the sake of argument. Williamson thinks he can parlay E = K into an argument that rationality never permits one to have a false belief. We are having trouble following the argument. Why is it that if one's evidence is knowledge, only knowledge conforms to that evidence? Williamson says that this is a natural understanding of what rationality requires given that one's evidence is knowledge. We confess that the naturalness eludes us. Perhaps this would be the natural understanding if one's believing conforms to the evidence only if the evidence entails one's belief. But, again, that's just the point for which Williamson has no argument.

---

[21] In everyday parlance, we do talk about one belief being more or less rational than another. Given that rational permissibility is binary, this talk can be interpreted as being about degrees of irrationality, i.e. degrees of falling short of rational permissibility, or degrees of evidential support.

[22] Williamson, 'Reply to Cohen, Comesaña', p. 92.

It is worth noting that our view that in the bad case believing is rationally permitted is consistent with Williamson's view that the evidence better supports one's belief in the good case. We need not hold the internalist view that rational permissibility supervenes on the appearances. One is rationally permitted to believe if one's belief is sufficiently well supported by the evidence. But sufficiently good evidential support need not be maximally good. Thus, one can be fully rational in believing p even though someone else's evidence better supports p. Compare rationality with morality. I am morally permitted (and so fully moral) to break a promise to drive you to the airport if I am sick enough to make it a hardship. This is consistent with my having a better reason to break my promise to pick you up, if my life depends on it. It does not follow that I am partially immoral, and so not morally permitted, in the former case. Similarly, my being rationally permitted to believe p in the bad case is consistent with my having better evidence for p in the good case.

Williamson has not provided any arguments for his view. Aside from its sheer implausibility, is there any reason to think it is false? The view that believing p on evidence e is rationally permitted only if e entails p is widely believed to have skeptical consequences. Williamson avoids this result in the case of perception. On his view, perceptual knowledge is not based on evidence provided by appearances. In the good case, one comes to know there is a table before one. Because one knows that that there is a table, it is part of one's evidence that there is a table. So the evidence entails what one believes.

Even if we grant this account of perceptual evidence, we do not see how it generalizes to rational beliefs acquired by ampliative inference. Inference is a diachronic psychological process. So when one comes to know by inference, one's knowledge of the premise is temporally prior to one's knowledge of the conclusion. Suppose it rains while I am asleep. I wake up and notice that the streets are wet. I come rationally to believe that it rained last night by inferring it from my prior knowledge that the streets are wet. My knowledge that it rained cannot be part of the evidential basis for the inference, since prior to the inference I did not know that it rained. Had I already known that it rained, there would have no need to make the inference. I come rationally to believe it rained by inferring it from my knowledge that the streets are wet. But that the streets are wet does not entail that it rained.[23]

Is there any way for Williamson to avoid this result? Williamson denies that rationality sometimes requires one to adopt a belief not fully supported by one's evidence. As we argued, Williamson is committed to denying the weaker thesis that rationality sometimes permits one to adopt a belief not fully supported by one's evidence. Suppose an inference occurs over an interval t_1-t_2. Let e_1 be one's evidence at t_1 and e_2 be one's evidence at t_2. We can now distinguish between two readings of the claim Williamson denies.

[23] Assaf Sharon and Levi Spectre make a similar point in a framed probabilistically. Sharon and Spectre, 'Epistemic Closure'.

(7) Rationality sometimes permits one to adopt a belief p on the basis of evidence e_1, where e_1 does not fully support p.
(8) Rationality sometimes permits one to adopt a belief p on the basis of evidence e_1, where e_2 does not fully support p.

Given E = K, it is possible to deny (8) without denying (7). Where one knows p by inferring it from e_1, ones evidence at e_2 will include p itself. In such a case, one's evidence at e_2 will fully support, that is entail, p. So Williamson can allow that cases of ampliative inference are rationally permitted just in case they yield knowledge. For example, I can rationally infer that it rained from my knowledge that the streets are wet, just in case I thereby come to know it rained.

This view has the result that the rationality at t_1 of inferring p does not supervene on one's evidence at t_1. Consider an inductive good case where at t_1 one infers p from one's evidence and thereby comes rationally to believe p at t_2. In the bad case, one has the same evidence at t_1 but p is false. Thus, on the view we are considering, in the good case, but not the bad case, one can at t_1 rationally infer p, despite the fact that at t_1 one has the same evidence in both cases. In the good case, what makes it rational at t_1 to infer p is evidence one does not have and will not have until t_2, viz., p itself. In effect, the evidence that makes it rational at t_1 to infer the conclusion is the conclusion itself. But as we noted, it cannot plausibly be maintained that the conclusion is the evidence that makes it rational to perform the inference in the first place. The evidence of the conclusion at t_2 is not relevant to whether it is rational at t_1 to infer p. Only the evidence of the premises is relevant.

A related result concerns one's inability to know one is in the inductive bad case. What one is rationally permitted to believe depends on one's evidence.[24] In the perceptual bad case, one is not in a position to know what rationality permits because one is not in a position to know what one's evidence is. But there is nothing about the inductive bad case that prevents one from knowing what one's evidence is. In such a case, one will still not be in a position to know what rationality permits at t_1. One can know what the evidence is and still not know whether one will know p by inferring it from the evidence. Moreover, in the bad case, nothing prevents one from knowing the evidential standards for rationality. So on this interpretation of Williamson's view, one could know the evidential standards for rationally inferring p, know what one's evidence is, but still not be in a position to know whether inferring p is rationally permitted.

Could we distinguish between the rationality of the inferring and the rationality of the conclusion based on the inferring? Perhaps it is the evidence at t_1 that determines whether the inference is rational, while the evidence at t_2

[24]Here we bracket issues about pragmatic encroachment. If the phenomenon exists, it does not affect the point we are making.

determines whether the conclusion is rational. Such a view cannot make sense of the bad case. The evidence at t_1 is the same in the bad case as it is in the good case. So this view would entail that in the bad case, at t_1, one is rationally permitted to infer p, even though doing so would yield at t_2 an irrational belief. It is hard to make sense of such a view.

III. Williamson on Believing Irrationally but Excusably

Williamson's strategy is to argue that internalism about justification confuses conforming to the standard of rationality with excusably failing to conform. On Williamson's view, in a skeptical scenario, one is excusably misled into thinking one is conforming to the requirements of rationality, when in fact one is not. According to Williamson, '[i]nternalists wrongly identify justified belief with belief excusable in such a way'[25]

In our response, we objected to his account of why the subject in the bad case is excused. Williamson was ruffled by what we said:

> Unfortunately, in their response Cohen and Comesaña badly misunderstand what I take the excuse to be. They misread a passage from an earlier article of mine . . . as saying that in the bad case one's excuse for falsely believing p is that one believes that one knows p. The passage says no such thing. I wrote 'the victim of a paradigmatic sceptical scenario is not to be blamed for forming false beliefs under the misapprehension that they constitute knowledge', but what one is not to be blamed *for* is not generally the same as the excuse in virtue of which one is not to be blamed.[26]

After protesting that we have misread him, Williamson seems to misread us. He 'badly misunderstands' us if he thinks that our misreading is attributable to our holding the absurd view that identifies offenses with excuses. So let us take a closer look at the passage in question: 'the victim of a paradigmatic skeptical scenario is not to be blamed for forming false beliefs under the misapprehension that they constitute knowledge.' We assumed that the offense, what one is not to be blamed for, is *falsely believing p*, thereby violating the norm for belief. On Williamson's view, that norm prohibits one from forming false beliefs (or more generally, from believing without knowing). Presumably, one violates the norm whether or not one misapprehends that believing p constitutes knowledge. So what work is the phrase 'under the misapprehension that they constitute knowledge' doing? Notice that in the above passage Williamson himself refers to the excusable offense only as 'falsely

[25] Williamson, 22. Williamson, 'Response to Cohen, Comesaña', p. 91.

[26] Williamson, 'Response to Cohen, Comesaña', p. 91.

believing *p*'. Moreover, prior to his reply to us, Williamson never says exactly what the excuse is supposed to be. Given our assumption that Williamson would not have included the point about the subject's misapprehension that the beliefs constitutes knowledge, were it not doing any work, we thought it natural to interpret the passage as saying that the misapprehension excuses one's violating the norm. No doubt Williamson's testy response was elicited by the implausibility of the view we attribute to him. Admittedly, the view is implausible, though less implausible than the view he attributes to us.

In his reply, Williamson says more about what the excuse is supposed to be:

> Clearly, in the sceptical scenario one's excuse involves the nature of the scenario, in which appearances are the same as in the good case in which the same beliefs constitute knowledge; one is in no position to know that one is violating the norm for belief. As a result of their misreading, Cohen and Comesaña labour the obvious point, which I of course have always accepted, that believing that one knows p is not by itself a good excuse for falsely believing *p,* since one's belief that one knows may itself be utterly irrational.[27]

While Williamson characterizes our point as obvious, he seems not fully to grasp it. As Williamson puts it, our point is that believing one knows *p*, if itself utterly irrational, will not excuse falsely believing *p*. Our point is stronger, viz., it is unclear how believing one knows *p* can excuse falsely believing p if believing one knows *p* is itself irrational.[28] Even if believing is not utterly irrational, it may still fail to be rational. This point is central to our objection to Williamson's view of the bad case. According to Williamson, in the bad case, one believes *p* without knowing *p*, thereby violating the norm for belief. This violation is excusable because, in the bad case, one is in no position to know one violates the norm. And one is in no position to know one violates the norm because in the bad case the appearances are the same as in the good case. On Williamson's view, this makes the appearances in the bad case misleading. One appears to have more evidence than one does. One appears to have sufficient evidence to satisfy the norm, when in fact one does not.

But one can be misled by evidence rationally or irrationally. Suppose, as Williamson says, that the appearances in the bad case misleadingly suggest that one has sufficient evidence to satisfy the norm. Do these same appearances rationally permit believing one has sufficient evidence to satisfy the norm? Williamson must deny that they do. On his view, one satisfies the norm only if one knows. So if the appearances rationally permit believing one satisfies the norm by believing *p*, then they rationally permit believing one knows

[27]Williamson, 'Response to Cohen, Comesaña', p. 91.

[28]Of course one need not actually believe one knows. It is sufficient that it is rational for the subject to believe.

p. In that case, one is rationally permitted to believe *p* itself and one needs no excuse. Thus Williamson must hold that the appearances do not rationally permit believing one satisfies the norm. If one believes one satisfies the norms on the basis of the appearances, one does so irrationally. But if believing on the appearances that one satisfies the norm is irrational, how exactly do the appearances excuse one for irrationally believing *p*?

Perhaps Williamson will say that, although the evidence of the appearances is not strong enough rationally to permit believing one satisfies the norm, it is strong enough to *excuse* falsely believing one satisfies the norm. Again, on Williamson's view, this entails that the evidence of the appearances is strong enough to excuse falsely believing one knows *p*. But the appearances provide no better evidence for one's knowing *p* than they do for *p* itself. So, if the appearances provide enough evidence for one's knowing *p* to excuse falsely believing one knows *p*, they also provide enough evidence for *p* directly to excuse believing *p*. Williamson's appeal to one's not being in a position to know whether one violates the norm does no explanatory work. For all he has said, the explanation for why in the bad case one is excused for falsely believing *p* is that, although the appearances do not rationally permit believing *p*, they do excuse believing *p*. This is more stipulation than explanation.

References

Cohen, Stewart. 'Luminosity, Reliability, and the Sorites'. *Philosophy and Phenomenological Research* 81, no. 3 (2010): 718–730.

Cohen, Stewart, and Juan Comesaña. 'Williamson on Gettier Cases and Epistemic Logics', *Inquiry* 56, no. 1 (2013): 15–29.

Cohen, Stewart, 'Theorizing about the Epistemic'. *Oxford Studies in Epistemology*, forthcoming.

Comesaña, Juan. 'Unsafe Knowledge'. *Synthese* 146, no. 3 (2005): 395–404.

Comesaña, Juan, and Matthew McGrath. 'Having False Reasons'. In *Epistemic Norms*, ed. John Turri and Clayton Littlejohn. Oxford University Press, forthcoming.

Gendler and Hawthorne. 'The Real Guide to Fake Barns: A Catalogue of Gifts for Your Epistemic Enemies', Philosophical Studies 124, no. 3 (2005): 331–52.

Hawthorne, John. 'Knowledge and Evidence'. *Philosophy and Phenomenological Research* 70 (2005): 452–58.

Schroeder, Mark. 'Having Reasons'. *Philosophical Studies* 139, no. 1 (2008): 57–71.

Sharon, Assaf, and Levi Spectre. 'Epistemic Closure under Deductive Inference: What Is It and Can We Afford It'. *Synthese*, forthcoming.

Sutton, Jonathan. *Without Justification*. Cambridge, MA: MIT Press, 2007.

Williamson, Timothy. 'Gettier Cases in Epistemic Logic'. *Inquiry* 56, no. 1 (2013): 1–14.

Williamson, Timothy. *Knowledge and its Limits*. Oxford: Oxford University Press, 2000.

Williamson, Timothy. 'Response to Cohen, Comesaña, Goodman, Nagel and Weatherson'. *Inquiry* 56, no. 1 (2013): 77–96.

Weatherson, Brian. 'Margins and Errors'. *Inquiry* 56, no. 1 (2013): 63–76.