

Williamson on Gettier Cases and Epistemic Logic

STEWART COHEN AND JUAN COMESAÑA

University of Arizona, USA

(Received 21 December 2012)

ABSTRACT *Timothy Williamson has fruitfully exploited formal resources to shed considerable light on the nature of knowledge. In the paper under examination, Williamson turns his attention to Gettier cases, showing how they can be motivated formally. At the same time, he disparages the kind of justification he thinks gives rise to these cases. He favors instead his own notion of justification for which Gettier cases cannot arise. We take issue both with his disparagement of the kind of justification that figures in Gettier cases and the specifics of the formal motivation.*

I. Gettier Cases, Excusability, and Justification

Why, according to Williamson, do Gettier cases matter? For most of us, they matter because they show that the gap between knowing p and being justified in believing p is greater than what one might otherwise have thought. What Gettier cases show is that one's justified belief that p can fail to be knowledge that p for reasons other than that p is false. This realization has led epistemologists to think of knowledge in new and sometimes surprising ways. Williamson takes a different view of the relation between knowledge and justification. For Williamson, a belief is genuinely justified only if it is knowledge. He notes that on such a view, Gettier cases cannot arise. But if this is true, then the answer to the question posed in the title of Gettier's 1963 paper, 'Is Justified True Belief Knowledge?' should have been a clear 'yes!'. Of course, Gettier posed the question rhetorically. His paper seems to show convincingly that knowledge is not justified true belief. And epistemology's subsequent preoccupation with the search for a fourth condition for

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

knowledge presupposed that Gettier had it right. But if Williamson is right about justification, Gettier did not show that those who thought knowledge is justified true belief were wrong about the nature of knowledge. Rather, if Williamson is right, those who think that he did show this are confused, and badly so, about the nature of justification. Many people now think the project of trying to solve the Gettier problem was misguided, but few think it was *that* misguided!

Some may think that we have overstated Williamson's position. What he says is that '*in a strongly externalist sense a belief is fully justified only if it constitutes knowledge*'.¹ While the Gettier problem cannot arise for strongly externalist justification, Williamson allows that it does arise for an 'internalist' sense of justification. But this does not entirely absolve epistemology of the aforementioned confusion. For what Williamson calls 'internalist justification' is a degenerate form of justification—'closer to excusability than genuine justification'. Excusability and justification are very different notions. One's Φ -ing is justified just in case one's Φ -ing meets the relevant normative standard. One's Φ -ing is excusable only if one's Φ -ing fails to meet the normative standard, but in a blameless manner. So if Williamson is right, Gettier cases are not genuine cases of justified true belief. Rather they are cases of excusable true belief.

Actually, if Williamson is right, the confusion runs deeper. This emerges when we consider what it takes, on Williamson's view, to be excused for having a belief that falls short of knowledge:

. . . the victim of a paradigmatic sceptical scenario is not to be blamed for forming false beliefs under the misapprehension that they constitute knowledge. The subject has a cast-iron excuse for having formed those beliefs.²

Presumably, a subject in a Gettier case has the same excuse for forming a belief that fails to constitute knowledge, namely he formed it under the misapprehension that it does constitute knowledge. So the subject in the Gettier case has an excusable true belief, which amounts to having a justified true belief in the internalist sense. Thus, if Gettier's question is interpreted as pertaining to internalist justification, it does not have a trivial affirmative answer. On this internalist interpretation, epistemologists were not confusedly denying the trivial truth that (strongly external) justified true belief is (sufficient for) knowledge. Unfortunately, on this internalist interpretation, Gettier's question has a trivial negative answer. That one can have a true belief

¹The first emphasis is ours, the second Williamson's. We come back to that second emphasis below.

²Williamson, 'On Being Justified'. Williamson cites this paper when discussing the kind of justification at issue in Gettier cases.

while misapprehending that one knows is hardly controversial. Most, if not all, cases of true belief are cases in which the subject believes he knows. We do not need fancy Gettier cases to demonstrate that, in some of these cases, the subject fails to know.

So, for Williamson, one can be justified in either a strongly externalist sense or in an internalist sense. In either sense, the Gettier problem is based on a serious confusion. Williamson says that 'it would be reassuring to have some independent way of checking that in Gettier cases the subject does not know after all'. But if it is worth reassuring ourselves that Gettier cases can indeed arise, then there must be a notion of justification distinct from either Williamson's strong externalist sense (according to which, trivially, they cannot arise), or his internalist sense (according to which, uncontroversially, they can). For this reason, even Williamson should be dissatisfied with the options he has given us for understanding justification. We think there is a middle ground—a notion of genuine justification that can make sense of the Gettier problem.

Williamson illustrates his two notions of justification using an analogy with moral justification. He considers the case of

a competent and non-negligent surgeon who causes the death of a patient because a lab technician with a grievance switched the labels on two bottles: he has a cast-iron excuse for killing the patient, but it does not amount to a justification for killing the patient.

The surgeon was not morally justified in administering the lethal medication, but his action was excusable.

What was the excuse? Williamson says that in the skeptical scenario, the excuse is that the subject believed under the misapprehension that his belief was knowledge. Presumably, by analogy, the excuse for the surgeon was his misapprehension that, by administering the medicine, he was doing the right thing. So, for Williamson, when one meets the normative standard—knowing, in the epistemic case, and doing the right thing, in the moral case—one is genuinely justified. When one fails to meet the standard but does so under the misapprehension that one does meet the standard, one is to be excused. The former notion corresponds to Williamson's notion of strongly externalist justification and the latter to his notion of internalist justification.

But when one fails to meet a normative standard, is one's failure excusable merely because one misapprehends that one meets the standard? Suppose the surgeon is under the misapprehension that he is doing the right thing because he carelessly assumes that the bottle closest to him contains the right medication. He administers the lethal medication thereby failing to do the right thing. Surely he is not to be excused for this failure even though he did it under the misapprehension that he was doing the right thing.

Why is the action of the surgeon who administered the lethal medication because of the switched labels excusable, but the action of the surgeon who blithely assumed he was grabbing the correct bottle not excusable? The answer is straightforward. The first surgeon, unlike the second, is to be excused because his misapprehension that he was doing the right thing, unlike the misapprehension of the second surgeon, was *rational*. Although his action did not meet the normative standard, he rationally believed it did. Equivalently, we can say that his belief that he was doing the right thing met the relevant standards for rationality. And, just as when one's Φ -ing meets the relevant standards of morality, one's Φ -ing is morally justified, so too when one's Φ -ing meets the relevant standards of rationality, one's Φ -ing is rationally justified. So in the misleading evidence case, unlike in the case of careless assumption, the surgeon's excuse for violating the moral standard is that he was rationally justified in believing that he met the standard.

To complete the analogy, we can see that the epistemic case has the same structure as the practical case. Having an excusable belief that does not constitute knowledge does not amount to believing under the misapprehension that one knows. Just as one can inexcusably misapprehend that one is doing the right thing (as in the case of the careless surgeon), so one can inexcusably misapprehend that one knows. If the Ouija board tells me it will rain and I trust its deliverances, I will inexcusably believe it will rain. My misapprehension that I know it will rain does not absolve me.

Rational justification is not strong externalist justification. The surgeon's belief that he was doing the right thing did not constitute knowledge, nor did my belief that I know that it will rain.

Excusably Φ -ing, then, requires more than Φ -ing under the misapprehension that one's Φ -ing meets the relevant normative standard. Misapprehensions themselves are subject to excusability constraints. When one's misapprehension is rational, one's failure to meet the standard is excusable.

Rationally believing one meets a standard is sufficient to excuse one's failure to meet it, but it is not necessary. Had the second surgeon's careless behavior resulted from a minor stroke, his administering the lethal medication would have been excusable despite his lacking any rational justification for thinking he was doing the right thing. Similarly, in an epistemic case, one may be so emotionally distraught after hearing bad news that one fails to appreciate properly one's evidence. In such a case one might excusably believe that one knows without being rationally justified in believing one knows.

While rational justification is distinct from excusability, we have seen that the notion of rational justification is indispensable for explaining why some subjects can be excused for failing to meet a normative standard. The surgeon with the misleading evidence is to be excused for administering the lethal medication because he was rationally justified in believing he was doing

the right thing. The subject in a standard Gettier case is to be excused for having a belief that falls short of knowledge because he is rationally justified in believing he knows.

Rational justification is weaker than strong externalist justification but stronger than Williamson's internalist justification. Subjects in a Gettier case are rationally justified in believing they know p , and so trivially, rationally justified in believing p . Viewing the subjects in Gettier cases as rationally justified allows Gettier's question to have a non-trivial answer. The subject in a Gettier case is not justified in a sense that guarantees he knows. Nor does a Gettier subject merely truly believe under a misapprehension that he knows. Nor does a Gettier subject merely truly believe under the excusable misapprehension that his belief constitutes knowledge. A subject in a Gettier case has a rationally justified true belief. His failure to know is excusable because he is rationally justified in believing he knows.³

What is the relation between rationally justified belief and strongly externalist justified belief? Are they distinct kinds of justification? Williamson says that 'a belief is *fully* justified only if it constitutes knowledge'. As justification is a scalar notion, this way of talking suggests that to be justified in a strong externalist sense is to be maximally justified. It would follow that to be rationally justified is to be situated at some less than maximal point on the scale. On this picture, the distinction between rational and strong external justification is one of degree rather than kind. Were this true, we could not say that in a Gettier case the subject has a fully justified true belief. We could say at most that the subject has a partially justified true belief.

This is an implausible consequence. It entails that one cannot have a fully rational false belief. But falsity is not a failure of rationality. Rationality requires one to conform one's beliefs to one's evidence. When one's evidence is misleading, rationality will require one to adopt a false belief. This is precisely the situation of the subject in the bad case. If he is not fully rational, then he is, to a certain degree, irrational. But it is hard to see how the subject in the bad case could be convicted of irrationality.

Could we say that the subject in the bad case is irrational—but excusably so? As we have seen, the explanation for why the subject in the bad case is excusable presupposes the rationality of his belief (that he knows). Perhaps Williamson could say that unless the subject's belief is fully rational he is not fully excusable. But since the view we are considering requires knowledge for full rationality, it would follow that the subject in the bad case is not fully

³Perhaps one can be rationally justified in believing p without being rationally justified in believing one knows p . For example, perhaps one can rationally believe that one will lose the lottery (on the basis of the odds alone) even though one cannot rationally believe that one knows one will lose the lottery. Whether such a case is possible depends on subtle issues about the nature of outright belief that go beyond the scope of this paper. If such cases are possible, then it's not clear there is a knowledge norm for belief.

rational and so partly to blame for believing falsely. On what basis could such a subject be blamed?

Is there another way to understand Williamson's claim that strong externalist justification is full justification? Perhaps he intends strong externalist justification to be a distinct kind of justification, normatively superior to other kinds. Consider his description of the good and bad cases:

the victim of a paradigmatic sceptical scenario is not to be blamed for forming false beliefs under the misapprehension that they constitute knowledge. The subject has a cast-iron excuse for having formed those beliefs. But the subject in the corresponding non-sceptical scenario whose beliefs constitute knowledge needs no excuse for having formed those beliefs; there is nothing to excuse. That the two subjects are equally blameless does not imply that there is no normative difference between them. The subject who blamelessly forms false beliefs has not done as well as the subject who gains knowledge.⁴

Williamson's point is that being (genuinely) justified amounts to more than mere blamelessness. Each subject is blameless. But the subject in the good case does not need to be excused, presumably because his beliefs meet the relevant standard for strong externalist justification, namely knowledge. We have seen that this description of the cases misses an important normative element. In addition to both subjects blamelessly believing, each subject has a rational belief that, in the bad case, explains the blamelessness of the subject. But according to Williamson, the subject in the bad case, because he fails to know, has not done as well as the subject in the good case. As Williamson puts it, 'excusable failure is not normatively equivalent to success'.⁵

There is much dispute about what counts as normative, and so what counts as a normative difference. Insofar as what Williamson says is correct, there is no reason to suppose that all normative differences are differences in justification. Suppose you shoot and kill a terrorist about to blow up a school. I shoot at a terrorist about to blow up a school but miss. I did my best, so we are equally blameless. Is there a normative difference between us? Although I blamelessly missed the terrorist, I did not do as well as you. You succeeded while I failed—blamelessly. Presumably, we were each justified in trying to kill the terrorist. Were you better justified because you did better? That seems wrong. Whatever normative differences there are between us because you did better than I, you were no better justified in shooting at the terrorist than I was. If you were fully justified, so was I.

⁴Williamson, 'On Being Justified', 116.

⁵Ibid, 117.

Just as we each aimed at the terrorist, according to a familiar metaphor belief aims at truth. According to Williamson, a better metaphor is that belief aims at knowledge. As he says, ‘Mere believing is a case of botched knowing.’⁶ In effect, when one believes p , one tries to know p . When one fails, one merely believes p . Suppose both you and I believe p and you succeed in knowing p , but I blamelessly fail because I was rationally justified in believing (that I know) p . Williamson would say that, while I may be in some sense justified in believing p , only you are fully justified. You have done better than I.

But here the analogy with the case of shooting at the terrorist is telling. We noted that you were no better justified in trying to kill the terrorist than I was, even though you succeeded but I did not. By analogy, you are no better justified than I in trying to know p , i.e., believing p , simply because you succeeded and I did not. Whatever normative differences there are between success and failure, it is not in the justification for trying. We see no reason to think that Williamson’s strongly externalist justification is full justification in a way that rational justification is not. So we see no reason not to say that when one is rationally justified, as in a Gettier case, one can be fully justified.

At this point, 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. If all we know is that a belief is justified in a strongly externalist sense if and only if it is knowledge, then for all we know, to say a belief is justified in a strongly externalist sense is entirely redundant, just another way of saying that the belief constitutes knowledge. In our epistemological theorizing, we include both rational justification and knowledge. Have we left something out if we do not include strongly externalist justification? This is hard to say until we know what it is. Justification is relative to a domain of normativity, e.g., rational justification, moral justification, prudential justification. Strong externalist justification is distinct from each of these. So what is the relevant normative domain for strong externalist justification? Perhaps by ‘strong externalist justification’ Williamson means ‘epistemic justification’. But ‘epistemic’ is another technical expression whose meaning in this context would be quite unclear. Distinguishing strong externalist justification from rational justification by classifying the former as epistemic would yield the surprising result that rationally justified belief is not an epistemic notion.⁷ So Williamson would have to be using ‘epistemic’ in a non-standard way.⁸ What it would mean is hard to know. Allowing that

⁶Ibid, 47.

⁷Of course rational belief must be distinguished from rational action which is not an epistemic notion.

⁸For more on the problems associated with the use of ‘epistemic’, see Cohen (forthcoming).

both rational and strong externalist justification are epistemic would not be progress. It would leave us unable to answer our original query—how to distinguish strong externalist justification from rational justification.

A final option would be for Williamson to hold that ‘strongly externalist justification’ refers to a constitutive norm for belief, on the model of his knowledge norm for assertion. This would mean that it is in the nature of belief that the norm applies. As he notes in his discussion of the constitutive norm of assertion ‘believing p stands to asserting p as the inner stands to the outer’.⁹ So perhaps just as the norm of assertion is constitutive, so is the norm of belief.¹⁰ Nonetheless, it remains true that the subjects in a Gettier case are rationally justified. The fact that subjects in these cases fail to be constitutively justified does nothing to undermine our characterization of them as genuinely justified true beliefs.

We see no reason to modify our view that Gettier cases illustrate how knowledge and justification differ in substantial, and often surprising ways. We thus think that it is indeed interesting that the existence of Gettier cases can be motivated formally—although we remain puzzled as to why Williamson finds it interesting. We disagree, however, about the choice of logic on Williamson’s part.

II. Margins for Error and the Iterativity of Knowledge

In his paper, Williamson shows that the existence of Gettier cases can be motivated formally, for cases of justified true belief that do not amount to knowledge arise in models for epistemic logic where both knowledge and justified belief are represented. Williamson’s models, however, have some uncomfortable consequences. We show that stronger logics than those adopted by Williamson also exhibit the existence of Gettier cases without having those consequences.

In epistemic logics, that a proposition p is known at a world w is modeled by the fact that, in all worlds accessible from w , p is true. For Williamson, the logic of knowledge is T. The accessibility relation for T is reflexive, which results in the claim that knowledge is factive. Other candidates to be considered are S4 and S5. In S4, the accessibility relation is reflexive and transitive, which in addition to factivity yields also the iterativity of knowledge. In S5, the accessibility relation is symmetric in addition to reflexive and transitive, which yields the claim that, if in scenario A for all the agent knows she is in scenario B, then in scenario B for all the agent knows she is in A. We agree with Williamson that S5 is not the right logic for knowledge. For all they

⁹Williamson, *Knowledge and its Limits*.

¹⁰One problem with this way of viewing matters is that Williamson views constitutive norms as applying to game-like phenomena, where these are contrasted with natural phenomena. But while assertion may be game-like, surely belief is a natural phenomenon.

know, brains in vats are normal human beings; however, it is not true that for all we know we are brains in vats. The epistemic possibility relation is not symmetric. But what about S4? Might epistemic possibility be transitive?

Before trying to answer that question, let us take a closer look at the model that Williamson uses to illustrate the existence of Gettier cases. In this model, the worlds are ordered pairs of temperatures and appearances of temperatures, where both of these are real numbers. Thus, if in w it is 70 degrees and it appears to be 70 degrees, then $w = \langle 70, 70 \rangle$. As usual, a subject knows a proposition p in a world w if and only if p is true in all the worlds accessible from w . Williamson then characterizes the accessibility relation R for this model as follows:

$$\langle e, f \rangle R \langle e^*, f^* \rangle \text{ if and only if } f = f^* \text{ and } |e^* - f^*| \leq |e - f| + c$$

So, w^* is accessible from w if and only if w^* has the same appearance as w and the absolute value of the difference between appearance and reality in w^* does not exceed the absolute value of the difference between appearance and reality in w by more than the factor c . Let us suppose that the value of c is 5. In that case, if it is 70 degrees but it appears to be 71, then for all the agent knows it is as warm as 77 and as cold as 65 degrees; if it is 70 degrees and it also appears to be 70 degrees, then for all the agent knows it is as warm as 75 and as cold as 65; and if it is 71 and appears to be 70, then for all the agent knows it is as warm as 76 and as cold as 64. R so defined is reflexive, non-transitive and non-symmetric. To model beliefs, Williamson adds that, in any world $\langle e, f \rangle$, the subject believes all and only those propositions that she knows in a world where appearances and reality match, $\langle f, f \rangle$. Williamson claims that the subject thus has only justified beliefs (at least in what he calls the ‘internalist’ sense of justification, though not in any epistemologically interesting sense, a distinction which we have challenged).

Let us now return to our question: Could the epistemic accessibility relation be transitive? Isn’t the obvious answer ‘no’? Haven’t we learned that knowledge does not iterate? There are reasons for thinking that knowledge does not iterate that can be easily accounted for. We can assume, for instance, that the subjects that we are interested in possess the concept of knowledge. More serious reasons for doubting the iterativity of knowledge have been given by Williamson himself. He argues that the iterativity of knowledge conflicts with plausible margin for error principles.¹¹ Let us assume that we cannot discriminate differences in temperature smaller than 5 degrees. How best can we capture this discriminatory failure? Williamson would say that we should capture it as follows:

¹¹Williamson, *Knowledge and its Limits*, Ch. 5.

MFE(W): If the temperature is n , then the strongest proposition that we can know is that the temperature is between $n - c$ and $n + c$.

Notice that MFE(W) does not say that if the temperature is n then the strongest proposition that we *will* know is that the temperature is between $n - c$ and $n + c$. In a radical deception scenario where the temperature is 30 degrees but it appears to be 70 degrees, the strongest proposition that we know in Williamson's model is that the temperature is between $30 + c$ and $110 + c$. Rather, we know that the temperature is between $n - 5$ and $n + 5$ in the best case scenario, where appearances and reality match exactly.

If we know MFE(W) (as we do in Williamson's model), then knowledge cannot iterate. For suppose that we do know MFE(W) and that knowledge iterates. Let us further suppose that the temperature both is and appears to be 70 degrees. In that case, by MFE(W) we know that the temperature is not 64 degrees. By the iterativity of knowledge, we know that we know that the temperature is not 64 degrees. Now, we also know (thanks to our knowledge of MFE(W)) the following: that if the temperature is 69, then the strongest proposition we may know is that it is between 64 and 74 degrees. Contraposing, we know that if we know that the temperature is not 64, then the temperature is not 69. We know a conditional and its antecedent. Without committing ourselves to any strong version of closure (although in the models that Williamson and we consider such a version of closure does hold), we can conclude that, in at least these kinds of cases, we know its consequent: we know that the temperature is not 69. But that is wrong according to MFE(W): when the temperature is 70 degrees (as we are assuming it is), for all we know it is as cold as 65.

That is an interesting argument for choosing T over S4 as the logic for knowledge. We will argue, however, that there are other ways of capturing the margins for error claim which are at least as compelling as MFE(W), and that they do not conflict with the iterativity of knowledge. Insofar as that conflict motivates choosing T over S4, we should reconsider that choice.

But before making that argument we want to point out that there are additional reasons to reconsider the choice of T over S4 as the logic for knowledge. First, Williamson is committed to our having knowledge of certain Moorean-looking propositions. The original Moorean propositions are of the form ' p and I do not believe p '.¹² Now, in Williamson's model, the strongest proposition that we know at a possible world is one that we do not know that we know. As we said above, moreover, Williamson's subjects believe all and only those propositions they know at a world where their actual appearances match reality. Thus, where ' p ' is the strongest proposition known at a world, the subject will know a proposition of the following form: ' p and I do not

¹²This is the so-called 'omissive' version of Moorean propositions—the 'commissive' Moorean propositions are of the form ' p and I believe not- p '.

believe I know p '. This is not quite a Moorean proposition, but it is close enough to qualify as Moorish, and may be uncomfortable for some philosophers. For instance, it is something that would be quite strange to assert, and yet it can be asserted if we follow Williamson in thinking that knowledge is the norm of assertion. In Williamson's models,¹³ that proposition can be upgraded to ' p and I'm almost certain that I do not know p ', which is even more uncomfortable.¹⁴

There is also a second noteworthy consequence of taking T to be the logic of knowledge. As Williamson himself notes, in his models there are two kinds of Gettier cases: those that closely resemble the cases originally presented by Gettier himself, and those that more closely resemble 'fake-barn' style of Gettier cases. Thus, for instance, when the temperature is 60 degrees but appears to be 70, the subject falsely (but justifiably) believes that the temperature is between 65 and 75. Because his beliefs are closed under entailment, she also believes that either it is between 65 and 75 or it is 60. This latter is a justified true belief which does not amount to knowledge (for all the agent knows, it is below 65 or above 75), and follows closely one of the original Gettier cases. Suppose now that the temperature is 65 degrees but appears to be 70. As before, the subject (justifiably) believes that it is between 65 and 75. However, this is not something the subject knows: given MFE(W), for all she knows it is 60. This is a case of justified true belief that does not amount to knowledge where the belief in question is not derived from any false belief (indeed, in this case the subject has no false beliefs at all). Many philosophers who accept the existence of Gettier cases because they are convinced by the original cases, however, are not so convinced by fake-barn style of cases.¹⁵ This is another reason to take a second look at the choice of T over S4 and MFE(W).

Let us take a closer look at Williamson's models. Notice, in particular, two features that they have. First, the margin for error always grows with the distance between appearance and reality (this is constraint [#2] in Williamson's paper). In other words, if the difference between appearance and reality is greater in w than in w^* , then more is known in w^* than in w . Second, this margin for error grows symmetrically: as the real temperature drops while the apparent temperature keeps the same, for instance, more and more warmer worlds become accessible (this follows from Williamson's characterization of the accessibility relation for his model). We disagree with both features. Take the second one first. Why would the margin for error grow symmetrically? Why would it be the case that the warmer it gets the colder I know it is? Fortunately, this feature of Williamson's models does not play any essential role in the results that Williamson proves, and so he could easily get rid of it.

¹³Williamson, 'Improbable Knowing', 'Very Improbable Knowing'.

¹⁴The 'Moorish' terminology comes from Williams and Hajek, "'p, and I'".

¹⁵See, for instance, Gendler and Hawthorne, 'Real Guide to Fake Barns'.

The first feature, however, is central to those results, for it entails the non-transitivity of the accessibility relation. But that feature is not as natural as Williamson would have us think. More natural is to start with the assumption that our evidence does not allow us to discriminate perfectly: when it appears to be 70 degrees, for all we know it could be between 65 and 75, say. Williamson will agree with us so far. But then he would add that when it appears to be 70 and it is, say, 69, then for all we know it is colder than 65. Why would this be the case? We already made allowances for the fact that when it appears to be 70 degrees, then for all we know it is 69. If it really is 69, then that is just one of the cases that we would expect to occur when it appears to be 70. Notice too that the case in which the temperature appears to be n and it is, say, $n + 5$, is a quasi-skeptical case: our evidence is almost misleading. Those cases are extremely improbable. But if they are indeed actual, then the consequence is as it should be: we do know that it is at most 75 degrees when it appears to be 70 and it is 75. That is just a consequence of the fact that our margin for error is 5 degrees either way.¹⁶

Suppose that I go outside and it feels like 70 degrees. You are inside getting dressed up, and you ask me how cold it is. I reply that it feels like 70 to me, and so it is at least 65. You then acquire a thermometer and come to know that it is actually 69 degrees. If Williamson is right, I did not really know that it was at least 65 degrees—and so there is something wrong with my telling you that it could not be as cold as 64. In our judgment, both of these are counterintuitive consequences of Williamson's way of thinking about margins for error. However, it is noteworthy that the same feature of Williamson's models responsible for this consequence (to wit, the intransitivity of R) plays a role in their yielding Gettier cases of the fake-barn variety. Those who think that fake barn cases are genuine Gettier cases, then, may have a reason to adopt Williamson's conception of margins for error. Having no such clear intuition ourselves, we think it worthwhile to consider the consequences of adopting an alternative conception of margins for error.

Our disagreement with Williamson comes down to this: if the distance between appearances and reality increases but the real temperature is still within our discriminatory range, then Williamson thinks we know less and we think we know the same. Now, if the temperature falls outside of the range, then we agree that we now know less than when it falls inside—although we think that the range of accessible worlds should grow asymmetrically in the direction of the real temperature. Those considerations motivate the following characterization of the accessibility relation:

$$\langle e, f \rangle R \langle e^*, f^* \rangle \text{ if and only if } f = f^* \text{ and either (i) } e < f - c \\ \text{and } e \leq e^* \leq f^* + c; \text{ or (ii) } e > f + c \text{ and } f^* - c \leq e^* \leq e; \text{ or} \\ \text{(iii) } f - c \leq e^* f + c$$

¹⁶For an argument along these lines, see Stalnaker, 'On Hawthorne and Magrider'.

That is to say, a world w^* is accessible from a world w if and only if appearances are the same in w and w^* and the temperature in w^* is anywhere between the temperature in w and the appearance, or anywhere at most c degrees away from the appearance. So, when e is within $f - c$ and $f + c$, the strongest proposition known is that the temperature is within that range; when e is below $f - c$, the strongest proposition known is that it is between e and $f + c$; and when e is above $f + c$, the strongest proposition known is that it is between $f - c$ and e .

So characterized, R is obviously reflexive. It is also non-symmetric: if the temperature appears to be 70 degrees but it is 30, then for all I know it is (say) between 65 and 75; but if the temperature both is and appears to be 70, then I know that it is not 30. But what makes the constraint different from Williamson's is that R is transitive. Consider the set $R(<e, f>)$, which contains all and only the worlds accessible from $<e, f>$. It is easy to verify that every world accessible from any member of $R(<e, f>)$ is a subset of $R(<e, f>)$ itself, which means that R is transitive.

Given that R so characterized is transitive, it follows that MFE(W) cannot be satisfied in our models: for instance, if the temperature appears to be 70 degrees and is 65, then I know that it is at least 65, violating MFE(W) (assuming that $c = 5$). But this does not mean that we cannot capture a margins for error principle within our models. Notice that MFE(W) takes our discriminatory capacities to be limited by the *real* temperature: when the temperature is n , we can at best know that it is between $n - c$ and $n + c$ (the best case being that where it also appears to be n). By contrast, we propose that our discriminatory capacities are determined instead by our appearances: when it appears to be n , we can at best know that it is between $n - c$ and $n + c$ (the best case being that where the temperature also is between $n - c$ and $n + c$). Thus, we propose the following formulation of a margins for error principle:

MFE: If the temperature appears to be n , then the strongest proposition that we may know is that the temperature is between $n - c$ and $n + c$.

Knowledge of MFE doesn't conflict with the iterativity of knowledge. Suppose that we do know MFE and that knowledge iterates. Let us also suppose that the temperature both is and appears to be 70 degrees. In that case, by MFE we know that the temperature is at least 65 degrees (assuming again that $c = 5$). By the iterativity of knowledge, we know that we know that the temperature is at least 65 degrees. But this time the conditional that we know (thanks to our knowledge of MFE) is innocuous: if we know that we know that the temperature is at least 65, then the temperature does not appear to be 71. We know, therefore, that the temperature does not appear to be 71, but that does not conflict with anything else we know.

In our models, then, knowledge iterates without conflicting with a plausible rendition of a margins for error principle. Moreover, our models do not have either of the uncomfortable consequences that Williamson's does. Given the iterativity of knowledge, if I know that p I know that I know it, and so our model does not have the consequence that I can know p without believing that I know it. For the same reason, everything that is known in a world w is also known in any world accessible from w , and so our models do not exhibit the fake-barn style of Gettier cases.

Recall the traditional Gettier case that arises in Williamson's model: if the temperature is, say, 30 degrees but it appears to be 70, then the subject has a justified true belief that it is either between 65 and 75 or it is 30, which does not amount to knowledge, because for all the agent knows it is between 30 and 65 (and also above 75 or below 30). Therefore, the subject has a justified true belief that does not amount to knowledge. The same case arises in our models: when the temperature is 30 degrees but it appears to be 70, the subject has a justified true belief that it is either between 65 and 75 or it is 30, but, for all the subject knows, it is between 30 and 65 (though in our models it is not true that for all the subject knows it is above 75 or below 30).

In summary, the existence of Gettier cases can be motivated formally by assuming that the logic of knowledge is S4 as well as T. Moreover, if we assume that the logic of knowledge is S4 then we can avoid Moorish propositions and fake-barn style Gettier cases, while at the same time doing justice to the existence of margins for error for knowledge.

III. Conclusion

Against Williamson, we have argued that the kind of justification at issue in Gettier cases cannot be identified with excusable or blameless belief. We have also argued that the iterativity of knowledge is not in conflict with plausible renditions of a margins for error principle, and that taking S4 to be the logic of knowledge allows us to see how the existence of Gettier cases can be motivated without having to commit ourselves to some of the uncomfortable consequences of taking T to be that logic.

References

- Cohen, Stewart. 'Theorizing about the Epistemic', in *Oxford Studies in Epistemology*, vol. 5, ed. John Hawthorne and Tamar Gendler. Forthcoming.
- Gendler, Tamar, and John Hawthorne. 'The Real Guide to Fake Barns: A Catalogue of Gifts for Your Epistemic Enemies'. *Philosophical Studies* 124 (2005): 331–52.
- Stalnaker, Robert. 'On Hawthorne and Magridor on Assertion, Context, and Epistemic Accessibility'. *Mind* 118, no. 470 (2009): 399–409.
- Williams, John, and Alan Hajek "'p, and I Have Absolutely No Justification for Believing That p": The Necessary Falsehood of Orthodox Bayesianism'. *Research Collection School of*

- Social Sciences* Paper 175, 2006. http://ink.library.smu.edu.sg/sooss_research/175. (accessed March 1, 2013)
- Williamson, Timothy. 'Improbable Knowing'. In *Evidentialism and its Discontents*, ed. Trent Dougherty, 147–64. Oxford: Oxford University Press, 2011.
- Williamson, Timothy. *Knowledge and its Limits*. Oxford: Oxford University Press, 2000.
- Williamson Timothy. 'On Being Justified in One's Head'. In *Rationality and the Good: Critical Essays on the Ethics and Epistemology of Robert Audi*, ed. Mark Timmons, John Greco, and Alfred Mele, 106–22. Oxford: Oxford University Press, 2007.
- Williamson, Timothy. 'Very Improbable Knowing'. *Synthese* (forthcoming).