

NON-PICKWICKIAN BELIEF AND ‘THE GETTIER PROBLEM’

John BIRO

ABSTRACT: That in Gettier's alleged counterexamples to the traditional analysis of knowledge as justified true belief the belief condition is satisfied has rarely been questioned. Yet there is reason to doubt that a rational person would come to believe what Gettier's protagonists are said to believe in the way they are said to have come to believe it. If they would not, the examples are not counter-examples to the traditional analysis. I go on to discuss a number of examples inspired by Gettier's and argue that they, too, fail to be counter-examples either for reasons similar to those I have urged or because it is not clear that their subject does not know.

KEYWORDS: Edmund Gettier, knowledge, belief

I. Few things are more widely agreed upon by philosophers today than that the traditional analysis of knowledge as justified true belief was dealt a damaging, perhaps fatal, blow by Edmund Gettier in his famous paper published more than fifty years ago.¹ Since that time, most discussion of the topic has centered on how to repair the analysis either by beefing up the justification condition or by adding a fourth one. Few have questioned Gettier's claim that in his alleged counter-examples to the analysis all three conditions claimed to be necessary are satisfied. That the truth condition is, cannot be doubted: that is a matter of stipulation. There has been some debate, though not much, and usually only with respect to the first example, whether Gettier's protagonist really satisfies the justification

¹ Edmund Gettier, “Is Justified True Belief Knowledge?” *Analysis* 23 (1963): 121-123. Examples of the many confident pronouncements to this effect: “Gettier's counter-examples leave the justified-true-belief theory stone dead.” (David Owens, *Reason without Freedom: The Problem of Epistemic Normativity* (London: Routledge, 2000), 41); “Gettier described two cases that decisively refute the analyses of knowledge as justified true belief.” (Mathias Steup, “The Analysis of Knowledge,” in *The Stanford Encyclopedia of Philosophy* (Summer 2012), ed. Edward N. Zalta, URL=<http://plato.stanford.edu/archives/sum2012/entries/knowledge-analysis>). As Williamson, who regards Gettier's examples as paradigms of thought experiment, remarks, “...his refutation of the justified true belief analysis was accepted almost overnight by the community of analytic epistemologists.” (Timothy Williamson, *The Philosophy of Philosophy* (Malden: Blackwell, 2007), 180) I shall consider some of Williamson's own Gettier-style cases below.

condition.² Even fewer have questioned – and, again, only in connection with the first example – whether the belief condition is satisfied.³ Indeed, even the possibility of defending the traditional analysis in this way has been denied. According to Floridi, one can try to defend it *only*

by strengthening/modifying the only flexible feature of the account, namely the justification condition; or by adding at least one more condition that would prevent the Gettier-ization of the required justified true beliefs or, alternatively, allow their de-Gettierization; or by combining (a) and (b). *No other general strategies are available* (my emphasis).⁴

Here I aim to show that in neither of Gettier's cases are both the belief condition and the justification condition satisfied and thus that the cases do not constitute counter-examples to the traditional analysis. I shall also discuss a number of examples modelled on those of Gettier's to show that they fail for similar reasons.

I take it that the following characterization of a standard 'Gettier' case would be widely accepted: S believes that *p*, S is justified in believing that *p*, *p* is true, but S's justification for believing that *p* is rooted not in the fact that makes *p* true but in some false proposition S is justified in believing from which *p* follows. It is then claimed that it is intuitively clear that S does not know *p*.⁵ While Gettier's claim that in the cases he describes the belief condition is satisfied has gone virtually unquestioned, I think there is reason to question it. I aim to show

² Christopher New, "Some Implications of 'Someone,'" *Analysis* 26, 2 (1965): 62-64 and "Someone' Renewed," *Analysis* 28, 3 (1968): 109-112; Charles Pailthorp, "Knowledge as Justified, True Belief," *The Review of Metaphysics* 23, 1 (1969): 25-47; Irving Thalberg, "In Defense of Justified True Belief," *The Journal of Philosophy* 66 (1969): 794-803.

³ Christoph Schmidt-Petri, "Is Gettier's First Example Flawed?" in *Proceedings of the 26th International Wittgenstein Symposium. Austrian Ludwig Wittgenstein Society* (Kirchberg am Wechsel, 2003), 317-319; Benoit Gaultier, "An Argument Against the Possibility of Gettiered Beliefs," *Logos & Episteme* V, 3 (1914): 265-272.

⁴ Luciano Floridi, "On the Logical Unsolvability of the Gettier Problem," *Synthese* 142 (2004): 62

⁵ In what I am calling a *standard* case, the (supposed) belief that turns out to be fortuitously true is a false but justified belief. Other examples that supposedly show the inadequacy of the traditional analysis, such as Lehrer's Grabit, Harman's assassination and Goldman's fake-barn cases, are often lumped together with Gettier's, even though they do not exhibit the standard pattern and – perhaps for that reason – are more controversial, with intuitions dividing on whether their subject knows or not. (Keith Lehrer and Thomas Paxson, Jr., "Knowledge: Undefeated Justified True Belief," *Journal of Philosophy* 66, 8 (1969): 225-237; Gilbert Harman, *Thought* (Princeton: Princeton University Press, 1973); Alvin Goldman, "Discrimination and Perceptual Knowledge," *Journal of Philosophy* 73, 20 (1976): 771-791. Goldman tells us that the example comes from Carl Ginet.) I shall say something about such cases later in the paper.

that in neither of the cases Gettier describes does Smith believe the proposition Gettier claims he does. If I am right, the cases Gettier describes make no trouble for the traditional analysis.^{6,7}

In the first challenge I know of to the claim that in the first case Smith does believe the proposition that turns out to be true, Schmidt-Petri argues that the definite description ('the man who will get the job') in the sentence that is supposed to express the proposition for which Smith has good evidence ("Jones is the man who will get the job, and Jones has ten coins in his pocket" – Gettier's (d)) is used referentially, whereas the same definite description in the sentence expressing the proposition Smith supposedly infers from it ("The man who will get the job has ten coins in his pocket" – Gettier's (e)) it is clearly used attributively. But surely, should Smith utter (e), he would be expressing a belief about Jones and Jones only.⁸ It is, in general, not difficult to recognize which use is in play. When

⁶ The claim that the belief condition is not satisfied must not, of course, be confused with the claim that it is not a necessary condition, as Radford suggested. (Colin Radford, "Knowledge by Examples," *Analysis* 27, 1 (1966): 1-11.) I agree that it is not, though not for the same reasons. But whether we are right in this has no bearing on whether in Gettier's examples the condition is satisfied. If it is not, the examples are not counter-examples to the sufficiency of the traditional analysis. Reasons for thinking that believing is not a necessary condition of knowing, either, may be found in J.L. Austin, "Other Minds," *Proceedings of the Aristotelian Society*, Supplementary Volume 20 (1946):171 and Zeno Vendler, *Res Cogitans* (Ithaca, NY: Cornell University Press, 1972), Ch.5; reasons for thinking that being justified is not a necessary condition of knowing, in Crispin Sartwell, "Knowledge is Merely True Belief," *American Philosophical Quarterly* 28 (1991): 157-65).

⁷ This is, of course, not enough to show that there are no genuine Gettier-style counterexamples to the traditional analysis. Stopped-clock and fake-barns cases may be thought plausible candidates. The difference is that in these the belief condition is – arguably – satisfied, whereas, as I shall argue, there is reason to think that in the former it is not. Some think that in one or the other or both of these cases the subject does have knowledge. (Stephen Hetherington, "Knowing Failably," *Journal of Philosophy* 96 (1999), 565-587; Igor Douven, "A Contextualist Solution to the Gettier Problem," *Grazer Philosophische Studien* 69, 1 (2005): 207-228; William G. Lycan, "On the Gettier Problem Problem," in *Epistemology Futures*, ed. Stephen Hetherington (Oxford: Oxford University Press, 2006), 148-168). Heathcote thinks this with respect to the fake-barn case but views the stopped-clock case as a "classical Gettier counter-example." (Adrian Heathcote, "Gettier and the Stopped Clock," *Analysis* 72, 2 (2012): 309). I can be argued, though, that in that case the belief and the justification conditions fail to be jointly satisfied. (John Biro, "Showing the Time," *Analysis* 73, 1 (2013): 57-62.) I discuss these cases further in section IV.

⁸ More recently, Gaultier has offered a different reason for denying that Smith believes the second proposition (Gaultier, "An Argument"). Because while Gaultier confines his discussion to Gettier's first case, his argument is general enough to be easily extended to the second, I reserve discussion of it until after I have considered the latter.

Dr. Johnson says “The man who is tired of London is tired of life” (as he is sometimes reported, inaccurately, to have done), we do not take him to have learned that a particular man was tired of London and also tired of life. If we did, the utterance would obviously lose the interest it now has for us. What Johnson actually said was “When a man is tired of London, he is tired of life,” which is clearly not about any particular man. (Johnson – of all people – would not have said what he is reputed to have said, knowing that it could be interpreted to mean what he did not.) He could have said “He who is tired of London is tired of life,” much as we say “He who hesitates is lost.” But we cannot substitute ‘he who’ for ‘the man who’ in “I am the man who broke the bank at Monte Carlo.” Yet this is, in effect, what Gettier is asking us to do in suggesting that we can infer the belief that is made true by Smith’s having ten coins in his pocket from our false belief that Jones will get the job and has ten coins in his pocket.

Another test of whether a definite description is being used referentially or not is adding ‘namely’ followed by a name to see if that yields something the speaker or thinker can be plausibly supposed to believe.⁹ Adding ‘namely, Jones’ in (e) does. But if that is what Smith believes, *his* getting the job does not make his

⁹ The ‘namely’ test was suggested by New (“Some Implications”) in connection with Lehrer’s Nogot/Havit example, as a way of distinguishing between the valid inference from “A is F, therefore someone, namely, A is F” and the invalid one “A is F, therefore, someone who may not be A is F.” (Keith Lehrer, “The Gettier Problem and the Analysis of Knowledge,” in *Justification and Knowledge*, ed. George S. Pappas (Dordrecht: Kluwer/Reidel, 1979), 65–78.) New argues that the sentence “Someone is F” is ambiguous in that it can be used to “express either a statement to the effect that an *identifiable*” someone is F or “a statement to the effect that an *unidentifiable*” someone is F. This is clearly close to – and anticipates – Donnellan’s famous distinction (with the difference that on the latter the difference is not one of meaning but of use). (Keith Donnellan, “Reference and Definite Descriptions,” *The Philosophical Review* 75 (1966): 281–304). However, New’s interest is, as usual, in how overlooking the difference affects the claim that the justification condition is satisfied. So is Pailthorp’s and Thalberg’s in making similar complaints. Bernecker also thinks that the first example turns on ambiguity: “In Gettier’s first example Smith is said to have a true and justified belief to the effect that the man who will get the job has ten coins in his pocket. The proposition that Smith believes – the man who will get the job has ten coins in his pocket – is ambiguous. Smith takes the definite description to refer to Jones but it in fact picks out Smith. If the definite description refers to Jones, Smith’s belief turns out to be justified but false. If the definite description refers to Smith, the belief is true but unjustified. The example therefore fails to show that justified true belief is insufficient for knowledge.” (Sven Bernecker, “Keeping Track of the Gettier Problem,” *Pacific Philosophical Quarterly* 92, 2 (2011): 127–52). Note, again, the focus on justification and the assumption that the belief condition is satisfied. However, Bernecker and I (and Schmidt-Petri) do agree that there is no proposition such that all three conditions are satisfied with respect to it.

belief true. Adding 'namely, Smith' also yields a belief, but it is not one he can be plausibly supposed to have. The only remaining reading of the definite description in (e) is as attributive. If so, this blocks the supposed inference from (d) to (e). Not only would such an inference fail to transmit justification, but Smith, being rational (as we must obviously assume him to be) would not in fact make it. Being rational, he must believe that even though he has good evidence that Jones will get the job, that evidence does not warrant believing that whoever gets the job has ten coins in his pocket. But if Smith does not believe (e) understood attributively, the example shows nothing about the adequacy of the traditional analysis.

However, this objection to Gettier's first example focuses on what is in fact an accidental feature of it. Other so-called Gettier examples, including his own second, do not involve an ambiguity of the sort that infects the first example. And it is easy to amend even that one by re-phrasing it without the ambiguous definite description. "Someone will get the job and that person has ten coins in his pocket" follows from "Jones will get the job and he has ten coins in his pocket," and if Smith recognizes this, he will, it may be thought, surely infer the former from the latter. (But see below.) His justification for believing the second proposition then carries over to the first, and if that happens to be true in the way Gettier asks us to imagine, he has a justified true belief but, it seems, no knowledge.

Another version of the attempt to undermine the first example by appealing to an ambiguity is to invoke a distinction between what Heathcote calls the speaker meaning interpretation and the objective referent interpretations, respectively. He argues that "...when we disambiguate (e), we either get a justified false belief or we get a unjustified true belief – but in neither case do we get a justified true belief."¹⁰ But this objection, too, appears to be blunted by eschewing the definite description.

¹⁰ Adrian Heathcote, "Truthmaking and the Gettier Problem," in *Aspects of Knowing: Epistemological Essays*, ed. Stephen Hetherington (Oxford: Elsevier, 2006), 153. Heathcote concedes that Gettier's second example is immune to this objection. (See footnote 22 below.) A variant of the same strategy is deployed by Mizrahi, who, appealing to Kripke's notion of ambiguous designators argues that Gettier cases – Gettier's own and the other discussed here – are misleading in that, contrary to the usual understanding, they reveal a semantic mistake, rather than an epistemic failure. (Moti Mizrahi, "Why Gettier Cases Are Misleading," *Logos & Episteme* VII, 1 (2016): 31-44). What is not clear in Mizrahi's discussion is why he thinks that the semantic mistake does not carry in its train the epistemic shortcoming alleged by those who take Gettier cases to be counterexamples to the traditional analysis. The account offered here can give a natural explanation of this: if in saying or thinking that something has a certain property one has in mind something other than the thing of that actually has that property, one does not believe that the latter has the property. If so, the belief condition of the traditional analysis is not satisfied with respect to the proposition that turns out to be true. This account

In any case, neither Gettier's second example nor the many other Gettier-style examples inspired by him involve such ambiguities. Take Chisholm's well-known one, in which one mistakes a bush for a sheep, but one's belief that there is a sheep in the field still turns out to be true, since there is an unseen sheep behind a tree.¹¹ "There is a sheep in the field" contains no definite description. The same goes for Lehrer's Nogot/Havit case and for Turri's more recent Lamborghini case (both to be discussed below).¹²

There is, though, another way to object: a way that brings out that the right question to ask is not whether Smith would be justified in believing (e) but whether he would believe it in the first place. Grant that "X is F" entails that "Something is F" and that I, who believe that X is F, know this. Is this enough to lead me to form the belief that something is F, a belief that would be made true by Y's being F?

Another way to bring out the point may be to ask what proposition is expressed by the sentence "The man who will get the job has ten coins in his pocket." It is supposed to be one made true by Smith's getting the job and having ten coins in his pocket. If so, it must be the proposition – (e) – that whoever gets the job has ten coins in his pocket. That proposition is one about a relation between two properties. By contrast, (d) says only that they are both instantiated in one, possibly idiosyncratic, case.¹³ What reason is there to think that the possibly accidental co-instantiation of two properties reveals some interesting relation between them?

also has the virtue of not requiring reliance on controversial views about meaning and reference, such as Grice's and Kripke's.

¹¹ Roderick Chisholm, *The Theory of Knowledge* (Englewood Cliffs: Prentice-Hall, 1977), 23

¹² Lehrer, "The Gettier Problem," 23; John Turri, "In Gettier's Wake," In *Epistemology: The Key Thinkers*, ed. Stephen Hetherington (London: Continuum, 2012), 214-229.

¹³ It is, no doubt, possible for someone to infer (e) from (d), thinking that the latter follows from the former. But if (e) is taken to express a belief that is supposed to be made true by Smith's getting the job, that belief cannot be justified on the basis of (d). Someone who recognized this would not make the inference and would therefore fail the belief condition. Someone who failed to do so and made the inference would fail the justification condition. Either way, one of the conditions deemed necessary by the traditional analysis would be left unsatisfied. Recall that in Donnellan's example about the murderer of Smith, in the scenario illustrating the attributive use it is the nature of the crime that grounds the belief, whereas in the one illustrating the referential one it is something entirely independent of it. Nothing in the second suggests a relation between the relevant properties. (Suppose the crime not particularly grizzly, but the man in the dock behaving as Donnellan describes, and suppose him innocent and the real culprit insane.)

II. Of course, even if Gettier's first example fails to be a genuine Gettier case and thus a counter-example to the traditional analysis, his second may succeed.¹⁴ Here we are asked to grant that S believes that either Jones owns a Ford or Brown is in Barcelona. (Gettier's (h)). He is said by Gettier to have 'constructed' (h), along with (g) that either Jones owns a Ford or Brown is in Boston and (i) that either Jones owns a Ford or Brown is in Brest-Litovsk. According to Gettier, Smith believes all these disjunctions, since he believes the first disjunct. No-one, to my knowledge, has suggested that we should not grant this. But I think there is reason not to. The assumption on which the claim that Smith believes these disjunctions rests is that what we may call belief tables mirror truth tables. This, I suggest, is an assumption we should not grant.

Is it really the case that anyone could believe (h) in a non-pickwickian sense? I take it that believing something in a non-pickwickian sense means being prepared to assert it (seriously, not pretending to, as one may in logic class). Being prepared to assert a disjunction implies that one is prepared to assert *either* one of the disjuncts on finding the other to be false. This may seem a surprising and unreasonable requirement. After all, is not the truth of one of the disjuncts sufficient for the truth of the disjunction? Smith is prepared to assert one of the disjuncts; why would his not being prepared to assert the other disjunct be a barrier to being prepared to assert the disjunction? The reason is not unlike the one we saw blocking a rational inference from (d) to (e) in Gettier's first example. Unless I *know* one of the disjuncts, I cannot, if I am rational, rule out the possibility that the one I believe (and think myself, perhaps rightly, justified in believing) is false. Believing that this is possible, I will be prepared to assert the disjunction only if I am prepared, on learning that one of the disjuncts, namely, the one I believed to be true, is false, to assert the other. But this is not so with Gettier's 'constructed' disjunctions.

The suggestion is that being prepared to assert seriously that $p \vee q$ requires not only being prepared to assert either that p or that q but also being prepared to assert that $(\neg p \rightarrow q) \wedge (\neg q \rightarrow p)$. This is not quite to require that the disjunction be understood as exclusive, since the requirement is compatible with being also prepared to assert that $p \wedge q$. (One sometimes says: " $p \wedge q$, but at least $p \vee q$.")

¹⁴ There is considerable variation in what is counted as a Gettier case. In characterizing 'the Gettier problem,' Steup does not even discuss the first example. Schmidt-Petri, by contrast, claims that it is 'the Gettier example.' Pritchard describes the stopped-clock example as 'the paradigm Gettier-style counter-example' (Duncan Pritchard, *Epistemic Luck* (Oxford: Oxford University Press, 2014), 156) and uses it as his stalking horse, without even mentioning Gettier's own cases. Heathcote thinks that it, but not the fake-barn case, is 'in the classic Gettier mold.' (Heathcote, "Gettier and the Stopped Clock," 309)

What a serious assertion of " $p \vee q$ " is not compatible with is being prepared to assert *only* that $p \wedge p \vee q$, without being prepared to assert that if $\neg p$, then q . But, surely, that is Smith's situation: he is not prepared to assert for any of the disjunctions he is said to have constructed that if the first disjunct is false, the second is true. Were he prepared to assert this with respect to one of them, he would have to be prepared to do so with respect to all three (not to mention the indefinitely large number of others he could construct in the same way). Clearly, someone rational would not be prepared to do this.

With respect to belief, unlike with truth and even, perhaps, assertion, ordinary 'or' *is* to be interpreted as always standing for exclusive disjunction. While $p \vee q$ is true even if both p and q are true (difficult enough to make a non-philosopher see at first), to say that someone believes that p or q (in a non-pickwickian sense), is not only not to say that he believes that p and also believes that q ; it is at least to imply that he believes only one of them. The correct way to report that someone believes both is to say that he believes that p and q . There is no analogue in ordinary belief talk of the inclusive 'v' of the truth table. Believing that p or q amounts to believing that one or the other is true but being undecided as to which that is. That is obviously different both from believing that p , as well as from believing that p and q .¹⁵

Even if this were not so generally, there is another reason why at least in Gettier's second example 'or' must be interpreted as exclusive. Presumably Smith would sign on to the conjunction of the three propositions he is said to have constructed only if he so took it, knowing that they could not all be true by way of both their respective disjuncts' being true. So, we need to take him as believing that either Jones owns a Ford or *else* he is in Barcelona. But then we must do the same with his other 'constructs' and take him to believe that either Jones owns a Ford or *else* he is in Boston, etc. Doing so would require that on learning that Jones does not own a Ford he believe – be prepared to assert – the conjunction of the second disjuncts of the constructed propositions.

Thus for Smith to believe (h), (g) and (i), he would have to believe that if Jones does not own a Ford, Brown is in Barcelona *and* if Jones does not own a Ford Brown is in Boston *and* if Jones does not own a Ford Brown is in Brest-Litovsk. He obviously does not believe all this. But if he is to satisfy the conditions of the traditional analysis with respect to (h), we must think of him as believing it. We

¹⁵ It may be asked, can one not believe that p and possibly also q ? Yes, but so understood, all the 'constructed' propositions (and many more) are true, and Brown's actual whereabouts are irrelevant to their truth. In any case, the propositions Jones 'constructs' and supposedly believes are categorical, not modal.

must thus think that there is something about (h) to make him believe *it* and be justified in believing it, rather than any of the others he has constructed (or could construct). What could that be?

Suppose, further, that believing that Jones owns a Ford, Smith constructs (j) "Jones owns a Ford or the moon is made of green cheese" or (k) "Jones owns a Ford or 4 is prime." Should we think of him as seriously believing these? Even if, as we are supposing, he has good evidence for the first disjunct, he presumably realizes that it could be false. And he knows that $p \vee q$ entails $\neg p \rightarrow q$. Can we think of him as seriously believing that if the first disjunct of (j) or (k) is false, their respective second disjuncts are true? (j) is unlikely to be made true by its second disjunct and (k) cannot be. Thus neither disjunction is something of which it can be said that Smith seriously believes it, hence neither satisfies the traditional analysis. It is no different with (h), which is 'constructed' in the same way.

In fact, Smith could 'construct' by addition any disjunctive proposition whatever and any of these could be made true by the truth of the added disjunct. Smith knows this. What is the difference between his believing the three he is said to believe and his not believing all the others? Just that those three happened to occur to him? Or should we say that he believes all of them? Neither choice strikes me as attractive.

III. Showing that Gettier's own examples are not genuine 'Gettier' cases and thus do not refute the traditional analysis is, of course, not enough to show that none of the many similar examples in the literature succeed where his fail. Obviously, I cannot review them all here. But a look at a few well-known ones may help bring out a general reason for thinking that they are all likely to be flawed in the same way. In Lehrer's Nogot/Havit case, the inference is by way of existential generalization. On the face of it, such an inference is unexceptionable. If Nogot, who is in the room, owns a Ford, it does follow that someone in the room owns a Ford. So, why balk at saying that someone who believes the first and infers the second from it believes the latter? Indeed, Lehrer himself does not, seeing the problem as one having to do with justification. He proposes a fourth condition, failure to satisfy which would "...have the effect of blocking... the transmission of justification" and thus underwriting the intuitively correct verdict that one would not know that someone in the room owned a Ford.¹⁶ Thus, again, it is the

¹⁶ Lehrer, "The Gettier Problem," 25. The details of how the fourth condition is supposed to do this are not important for present purposes. (As mentioned earlier, sometimes adding a fourth

justification condition that is seen as not satisfied, with the satisfaction of the belief condition assumed.

However, adding a fourth condition (or strengthening the justification condition) is conceding that, unimproved, the traditional analysis *is* inadequate. We are not forced into doing this if it can be shown that in this case and in similar ones the belief condition is not satisfied. And we can do this if we look more closely at whether the inference as described yields a seriously held, non-pickwickian, belief.

Here the inference has the form $Fa \wedge Ga$, therefore $Ex (Fx \wedge Gx)$, which is obviously valid, unlike the one from (d) to (e) in Gettier's first case. Yet even here it does not follow from one's believing that Nogot owns a Ford that one believes that someone or other in the room (that is, if not Nogot, then someone else) owns a Ford, which is what one would have to believe, if one were to believe something that Havit's owning a Ford could make it true. For one to have such a belief, one would have to believe that if Nogot does not own a Ford, someone else in the room does. In Lehrer's story, one has no reason to believe that. Thus, again, either one does not make the inference or, if one does, one is not justified in doing so.

As before, the fact that the truth of the premise entails the truth of the conclusion should not be taken to mean that believing the premise (however strong the evidence for it) and recognizing that it entails the conclusion suffices for someone rational to believe the latter (in a non-pickwickian sense). Believing it in that sense requires, as we saw in connection with Gettier's second case, that one be prepared to assert that if, despite one's evidence, Havit does not own a Ford, someone else in the room does. This is something our reasoner is presumably not prepared to assert, having no evidence that suggests it. Hence we have no reason to say that in this case the belief condition is satisfied and thus no reason to think that we have a genuine Gettier case.

In Chisholm's example, the presence of an unseen sheep does not make true my mistaken belief that what I am looking at is a sheep.¹⁷ It makes the sentence

condition and beefing up the justification condition come to the same thing. This is the case here.)

¹⁷ The example fails the analogue of the 'namely' test. City Slicker: "There is a bull in the field!" Farmer: "Where?" "There, by the creek!" "Nah, that's a cow." An unseen bull in the field would not make what City Slicker believes true, even though the proposition it would make true follows from the one he does believe. (Contrast: Farmer: "There is a bull in the field." City Slicker: "Where?" Farmer: "I don't know, but those cows are sure acting nervous.") Imagine that we are on safari, hoping to catch a glimpse of the rare and elusive grumpus. I whisper, excitedly, "There is one!" You, my guide, deflate me by saying, "No, that just looks like one, it is a common pumpus. But I know there is a grumpus somewhere in this area – I have seen its tracks." I am

"There is a sheep in the field" true, but, as is well known, belief is finer-grained than sentence meaning. (Consider, most obviously, sentences containing indexical terms.)

Williamson claims that the following example exhibits the same pattern:

A clever bookseller fakes evidence which appears to show conclusively that a particular book once belonged to Virginia Wolf; convinced, Orlando pays a considerable sum for the book. He has a justified false belief that this book of his once belonged to Virginia Wolf. On that basis alone, he forms the existential belief that he owns a book which once belonged to Virginia Wolf. The latter belief is in fact true, because another of his books in fact once belonged to her, although he does not associate that one with her in any way. Thus Orlando has a justified true belief that he owns a book that once belonged to Virginia Wolf, but he does not know that he owns a book which once belonged to Virginia Wolf.¹⁸

Williamson is right that the example follows a familiar pattern. It should not come as a surprise that it fails for a by now familiar reason: it does not pass the 'namely' test, as is shown by the fact that the object of Orlando's justified false belief is described by Williamson himself as '*this* book of his' (my italics), namely, the ringer. That belief is not made true by his unwitting ownership of another book that did belong to Virginia Wolf.

Williamson also offers what he calls a real-life Gettier example. Here he describes himself as apologizing to an unsuspecting audience for not giving a power-point presentation, saying, falsely, that the only time he had given one it was a complete disaster. Believing him, the audience 'competently deduced' and thus came to acquire the justified belief that he had never given a successful power-point presentation. That belief, while true, was made true by the fact that he has never given a power-point presentation at all. According to Williamson, the listeners,

basing their justified true belief that I had never given a successful power-point presentation on their justified false belief that the only time I had given a power-

asserting that that pumpus is a grumpus, you are saying merely that there is some grumpus in the vicinity. Some languages mark the difference by different words or different word order: German "Da ist ein..." v. "Es gibt ein...;" Hungarian "Ott van egygy..." v. "Van ott egygy." It is not the first time that the fly was lured into the fly-bottle by taking a quirk of English at face value. (John Biro, "What Is 'That'?" *Analysis* 71, 4 (2011): 651-653.) I am not denying, of course, that the proposition expressed by a sentence containing the first locution entails the one expressed by a sentence containing the second. I am claiming only that believing the first proposition does not entail believing the second, so that the latter's turning out to be true is not enough for the conditions of the traditional analysis to be satisfied.

¹⁸ Williamson, *Philosophy*, 183.

point presentation it was a complete disaster... they did not know that I had never given a successful power-point presentation.¹⁹

I think that if we allow that the audience's false belief is justified, we have to say that it does know that Williamson had never given a *successful* power-point presentation, which is the only belief it acquired from his apology and which is both justified and true. (In VI. below I question whether we should allow this.) The fact that it is also true that he had never given one at all, something that entails what the audience believes but is not entailed by it and something that the audience does not believe, is neither here nor there. Of course they do not believe *that*, given Williamson's pretence.

Hercule Poirot to the assembled company: "Someone in this room is the murderer!" Poirot may or may not have a particular person in mind. Suppose he does. Are we to think of him as believing that if that person turns out to have an iron-clad alibi, someone else in the room is the murderer? That would be to see him as having reason to believe that it could not have been someone other than those present. He may well believe this, but, surely, not on the grounds that it follows from his belief that his prime suspect did the deed. Even though "The nephew, standing by the fireplace, is the culprit" entails "Someone in this room is the culprit," believing the former does not entail believing the latter in a non-pickwickian way. To be non-pickwickian, the latter belief would have to survive the demise of the former; and it requires different evidence – at once, more and less.²⁰

To think that Poirot would make such an inference would be to commit both the mistake we found in Gettier's first example and that we found in his second. We would have to see him inferring a general proposition from one essentially involving an individual, going from $Fa \wedge Ga$ to one of the form $(Ex) Fx \rightarrow Gx$. We would also have to think that since he believes Fa , and since Fa entails $Fa \vee Fb$ and $Fa \vee Fc$ etc., he would believe the disjunctions (and therefore believe that the nephew or the wife did the deed and that the nephew or the secretary did and so on). However, though both inference forms are obviously valid, thinking that believing the premise suffices for believing the conclusion assumes that

¹⁹ Williamson, *Philosophy*, 192.

²⁰ Suppose, by contrast, that Poirot does not have someone in particular in mind but is confident that one of the company is guilty and hopes that he or she will give him/herself away. Now it does not matter who the culprit turns out to be. As long as someone in the room *is* the murderer, Poirot has a justified (we have assumed) and true belief, and there is no reason to deny him knowledge. But in such a case, he is not inferring his true belief from a false – albeit justified – one, as does Smith in Gettier's first case.

(recognized) validity is enough for basing one belief on another, that truth-preservation is sufficient for belief-preservation. That it is sufficient for the preservation of justification – the so-called Principle of Deducibility for Justification – has been denied (Thalberg), though the prevailing view seems still to be that it is. What I am calling into question here is an analogous principle for serious, non-pickwickian, belief. Such belief does not simply track logical relations.²¹ While it is certainly constrained by them (for rational believers), it is also subject to other conditions. Thus neither existential generalization nor addition, valid forms of inference though they be, is enough to generate such belief. But since the belief condition of the traditional analysis is satisfied only if such a belief is present, Gettier's examples and those similar to it do not pose a threat to it.²²

What about the beliefs involved in what Sorensen calls junk knowledge?²³ Can one not believe, and be justified in believing, that p or q only because one believes that p , no longer believing the disjunction once one ceases to believe the first disjunct? According to Sorensen, Smith knows the disjunction in Gettier's second example if only the first disjunct is true but does not know it if only the second is. What he has in the former case is “not a useful type of knowledge but... nevertheless knowledge.”²⁴ But it is not the usefulness of what Smith has that is in

²¹ Heathcote makes a parallel observation about evidence: “The idea that warrant transfers from one belief to another is to be in the grip of a false analogy: the analogy between evidence for a proposition and the truth of a proposition. Logical implication preserves truth but it does not preserve evidence for, and if one tries to force the analogy then evidence for, and hence the notion of justification, will end up being an epistemic concept that has truth-like properties... This has done untold damage in the history of philosophy, creating bad doctrine along with un-meetable demands.” (Adrian Heathcote, “Truthmaking, Evidence of, and Impossibility Proofs,” *Acta Analytica* (2014): 373) However, he does not recognize that the Gettier ‘problem’ is a case in point.

²² Heathcote also says (surprisingly, in view of the passage just quoted) that to deny that Smith is justified in believing (h) requires “deny[ing] some tried-and-tested rules of logical inference.” (“Truthmaking and the Gettier Problem,” 164) This is to commit the very mistake he laments. We need not challenge the validity of the inference from p to $p \vee q$ to deny that it is sufficient to yield serious belief. Distinguish between “ p or q or possibly both” and “ p or else q .” The former is true if either p is true or q is true or both p and q are true, whereas the latter is true only if just one of p and q is true and not true if both p and q are true. In ordinary discourse, “ p or q ” is always understood in the second way, unless the speaker adds the third disjunct. Doing so is always odd (“She spoke in German or she spoke in French (or she spoke in both)”) and sometimes not even possible (“She is in Paris or she is in Rome (or she is in both)”).

²³ Roy Sorensen, “Dogmatism, Junk Knowledge, and Conditionals,” *Philosophical Quarterly* 38 (1988): 433-54.

²⁴ Sorensen, “Dogmatism,” 446.

question, but whether it is genuine knowledge at all. What more is he supposed to know in knowing the disjunction other than what he knows in knowing the first disjunct when that disjunct is true? If the answer is, as I suggest, nothing, the traditional analysis suffices for telling what Smith does and does not know. It is no surprise that junk belief would be enough for junk knowledge but not for the real thing. Just as calling something junk food conveys that it is not the genuine article, so calling something junk knowledge tells us that it is not real knowledge. It is only the latter, however, that the traditional analysis claims to capture.

IV. What, though, to make of the fake-barn and the stopped-clock cases? I have already noted (fn.4) that there is disagreement about whether these fit the Gettier mold. Take fake barns first. While it has been claimed that “it is almost universally accepted that the agent in the barn façade case lacks knowledge,” I am not alone in thinking that it is far from obvious that Henry, looking at a real barn in fake-barn country, does not know that he is seeing a barn, a view shared by at least Lycan and Hetherington.²⁵ After all, he has formed his belief that he is doing so in a non-defective way and without relying on any false assumption, explicitly or tacitly (as Gettier's Smith or the believer in the Lehrer's Nogot/Havit case have been alleged to have done.)²⁶ But wait: is he not assuming that he is not in fake-barn country? Yes, but that is different from assuming that Nogot owns a Ford on the way to concluding that someone in the office does. The latter assumption is a positive contributor to a belief that would not be formed without it. The assumption that things are normal, in the absence of any reason to think that they are not, plays no such role in the fake-barn case. It is merely an instance of a general non-skeptical assumption, just like the assumption that there is no evil demon or that one is not a brain in a vat.

Pritchard says about the fake-barn case that “...there are... a great many nearby possible worlds where Henry forms the same belief on the same basis (by simply looking at the ‘barns’) and yet his belief is false.”²⁷ The worlds Pritchard has in mind are ones in which Henry is in fake-barn country, does not know it, and comes to believe that he is seeing a barn on the basis of a glance at a mere façade. But why think these worlds relevant? In them, one of the three conditions of the traditional analysis is not satisfied. Why should that be thought to impugn

²⁵ William G. Lycan, “Evidence One Does Not Possess,” *Australasian Journal of Philosophy* 55 (1977): 114-126; Hetherington, “Knowing Failably.”

²⁶ For a catalogue of the ways in which these cases may be modified in such a way that the no-false-assumption condition is satisfied, see Lycan, “On the Gettier.”

²⁷ Pritchard, *Epistemic Luck*, 162

the sufficiency of the analysis when the conditions *are* satisfied? Why should the fact that Henry might have been deceived deprive him of knowledge when he is not? We are allowed to assume that there is no evil demon. Still, there might have been. Does that force us to be skeptics?

The supposition that Henry is unaware that he is in fake-barn country is crucial here. In supposing Henry to be ignorant of being in a world as remote from the actual world as is a fake-barn world, we render his being in that world irrelevant to whether he has knowledge. Should Henry have some reason to suspect that he may be in fake-barn country, things would be different.²⁸ Then he would fail to satisfy the justification condition if he formed his belief merely by glancing at a facade. Nor would he, if sensible, believe his eyes. (Here we have the same trade-off we saw with Smith in Gettier's first case. See fns. 7 and 9 above.)

This is where the stopped-clock case is different from the fake-barn case in which Henry has no reason to be suspicious. In the stopped-clock case there is a specific tacit assumption in play, not just a general non-skeptical one. We cannot be supposed to be ignorant of the fact that in the actual world there are inaccurate or non-working clocks, hence anyone who forms a belief, even a true one, on the basis of a mere glance, fails the justification condition. That the clock I have just glanced at is a working and accurate one is a far more vulnerable assumption than a no-evil-demon one, and a world in which it is not is far closer to the actual world than is a fake-barn one. We all know that clocks are often inaccurate and sometimes stop, and we know how to go about finding out whether they are or have. (How does one go about finding out whether an evil demon is at work?)

Thus with clocks, whether the justification condition is satisfied depends on whether one has taken sufficient care to make sure that one is looking at a working and accurate clock.²⁹ I would not think I had done so just by glancing at a

²⁸ We are in the days of the land rush. Put something on a section and you own it.

²⁹ What counts as sufficient care is arguably context-dependent. See Igor Douven, "A Contextualist Solution to the Gettier Problem," *Grazer Philosophische Studien* 69, 1 (2005): 207-228 and Biro, "Showing." On the other hand, it may be held that there is a matter of fact about whether someone is sufficiently justified to know that is independent of his (or his attributor's) context, the latter affecting only how much he does or should care about whether he does satisfy it. Much of the time it does not matter whether one acts on knowledge or reasonable belief.

Harman's assassination case admits of being treated along similar lines. It is assumed that Jill, unaware of the false retractions of a true report of an assassination, satisfies the justification condition, though she would have believed the retractions, and thus ceased to believe the original reports, had she seen them, as did others who did. But that depends. (As usual, the case is severely under-described.) True, "Don't believe everything you read in the papers!" is not based on worries about evil demons – remember "Dewey defeats Truman"! But

clock, even one that has worked and been found accurate in the past. I may well form a belief about what the time is on the basis of a casual glance, but would I claim to *know* it, especially, did something important – say, your detonating the charge at just the right moment – hang on my doing so? Indeed, would I claim to *seriously* believe it?

It has not been noticed, to my knowledge, that in one way the fake-barn case and the stopped-clock case are inverses of one another. In the former, beliefs based on glances at all but one of the barn façades are false. In the latter, beliefs based on all but one glance at the clock – that is, all the past ones, at various times, at the (let us suppose) working and accurate clock – are, other things being equal, true. (I say ‘other things being equal,’ to allow for mis-readings and the like.) This is to be distinguished, of course, from the fact, often noted, that glances at the stopped clock at other times would have yielded false beliefs.³⁰ We can then imagine a clock case that mirrors the fake-barn case, one in which my every past glance yielded a false belief, since each was, as it happens, a glance at an inaccurate (or non-working) clock. But the clock has been repaired and when I glance at it again it shows the right time. Now suppose that I am justified (by whatever standard is appropriate in the various contexts) in forming the beliefs I do, on both those past occasions and on the present one. This can be so, even if it often requires more than a casual glance. My past beliefs were false, hence I did not satisfy the truth-condition of the traditional analysis. But on the present occasion I do satisfy all three conditions. There is, I claim, no more reason to say that I do not know what time it is than there is to say that Henry does not know that what he is looking at is a barn. Should I be aware, however, of the past unreliability of the clock, the bar for meeting the justification condition would rise dramatically, just as it would for Henry if he had reason to suspect that he is in fake-barn country. Neither situation is one in which it is clear that I satisfy all three conditions laid down by the traditional analysis yet lack knowledge.

V. Let us now consider the claim that that not only is there a genuine Gettier problem but that it is, in Zagzebski's word, inescapable.³¹ She argues that the traditional analysis, even if supplemented by a condition or conditions of the sort

what if Jill lives in North Korea? Seeing the later reports with their stock footage of the dear leader would, far from leading one to doubt the original reports, re-inforce one's confidence in their veracity. (“They must have slipped one by the censors!”)

³⁰ E.g., Pritchard, *Epistemic Luck*, 156.

³¹ Linda Zagzebski, “The Inescapability of Gettier Problems,” *The Philosophical Quarterly* 44, 174 (1994): 65-73.

often proposed to block Gettier-style counterexamples, would still be open to these. She offers what is, in effect, a recipe for constructing Gettier cases:

...start with a case of justified (or warranted) false belief. Make the element of justification (warrant) strong enough for knowledge, but make the belief false... due to some element of luck. Now emend the case by adding another element of luck, only this time an element which makes the belief true after all. We now have a case in which the belief is justified (warranted) in a sense strong enough for knowledge, the belief is true, but is not knowledge.³²

Zagzebski gives us a number of examples in support of this sweeping claim. Here is one, aimed at Plantinga's attempt to solve the Gettier problem by appending to the traditional analysis a condition requiring that the subject's faculties be working properly in an appropriate environment. Imagine that Mary has very good eyesight – good enough for her to see, and thus to come to know, when her husband is sitting in this favourite chair in the living room. But on this occasion, Mary's husband's brother, who looks a lot like the husband, is sitting in the chair. Mary is, of course, entitled to conclude from her false belief that her husband is sitting in his favourite chair that he is in the living room. As luck would have it, he is. Thus Mary's belief that her husband is in the living room is true, but, intuitively, it is not an instance of knowledge.

It should be clear that what must be said about this example is just what I argued must be said about Chisholm's. Mary's belief concerns the person she is seeing, not one she is unaware of. That she thinks that that person is her husband is just like thinking that the bush one is looking at is a sheep. If the example exhibits the general form of Zagzebski's recipe for Gettier-style counter-examples even to a strengthened traditional analysis, a reply along these lines will always be available. This also shows, of course, that such strengthening is not required.³³

Another example Zagzebski gives is that of Dr. Jones, who

...has very good inductive evidence that her patient, Smith, is suffering from virus X. Smith exhibits all of the symptoms of this virus, and a blood test has shown that his antibody levels against virus X are extremely high. In addition, let us suppose that the symptoms are not compatible with any other known virus, all

³² Zagzebski, "Inescapability," 69.

³³ It should be obvious that the example fails the 'namely' test. The proposition that the husband is sitting in his favourite chair entails that he is in the living room. It also entails that he is somewhere in the house. But while his being in the attic makes the latter proposition true, surely, that is not what Mary believes. Suggest to her that what she believes is that her husband is somewhere or other in the house, and see what she says. Explain to her that she is wrong because this follows from what *she* says she believes, namely, that her husband is sitting in his favourite chair, and she will say, "You philosophers!"

of the evidence upon which Jones bases her diagnosis is true, and there is no evidence accessible to her which counts significantly against the conclusion. The proposition that Smith is suffering from virus X really is extremely probable on the evidence. None the less, let us suppose that the belief is false. Smith's symptoms are caused by the presence of a different and unknown virus; the antibody levels are due to idiosyncratic features of his biochemistry which cause him to maintain unusually high antibody levels long after a past infection. In this case Dr. Jones' belief that Smith is presently suffering from virus X is false. ... Now to construct a Gettier-style example we simply add the feature that Smith has very recently contracted virus X, but so recently that he does not yet exhibit symptoms caused by X, nor has there been time for a change in the antibody levels due to this recent infection.³⁴

The trouble here is that the belief Dr. Jones supposedly forms on the basis of Smith's symptoms is under-described by Zagzebski. It is not just that Smith has X but that X is what is causing the symptoms. This is not a belief that is, or could be, made true by the fact that a-symptomatic X is present. Obviously, Dr. Jones would never form the belief that Smith has a-symptomatic X on the basis of his evidence. Thus the fact that Smith does have a-symptomatic X is not relevant to the truth of the belief Dr. Jones does form, a belief that is, by hypothesis, false. Once again, it is only by (perhaps inadvertent) slight of hand that we are maneuvered into agreeing that Dr. Jones has a justified true belief. But if he does not, we do not have a Gettier case.

I have not offered a direct argument against Zagzebski's general claim that no fallibilist account of knowledge can escape the Gettier problem.³⁵ But if my misgivings about her pivotal examples is well founded, we have reason to be skeptical about it.

VI. There remains one more putative Gettier case to comment on: Turri's amended Lamborghini story. In its initial version it goes like this:

One of Dr. Lamb's students, Linus, tells her that he owns a Lamborghini. Linus has the title in hand. Dr. Lamb saw Linus arrive on campus in the Lamborghini each day this week. Linus even gave Dr. Lamb the keys and let her take it for a drive. Dr. Lamb believes that Linus owns a Lamborghini, and as a result concludes, "At least one of my students owns a Lamborghini." As it turns out, Linus doesn't own a Lamborghini. He's borrowing it from his cousin, who happens to have the same name and birthday. Dr. Lamb has no evidence of any of this deception, though. And yet it's still true that at least one of her students

³⁴ Zagzebski, "Inescapability," 71.

³⁵ Heathcote ("Truthmaking, Evidence") takes a good stab at one.

owns a Lamborghini: a modest young woman who sits in the back row owns one. She doesn't like to boast, though, so she doesn't call attention to the fact that she owns a Lamborghini.³⁶

As Turri notes, this story is 'vaguely modeled' after Lehrer's Havit/Nogot story, and, it should be easy to see that it succumbs to the same objection: it fails the 'namely' test. But Turri gives it a twist that clearly blunts that objection. He asks us to imagine that "...unbeknownst to Linus he has just inherited a Lamborghini. His cousin died and left it to him."³⁷ Here there is no questioning that Lamb believes exactly what happens to be true, namely, that *Linus* owns a Lamborghini. However, there is another problem, both with the original and the amended version. In neither is Lamb's belief that Linus owns a Lamborghini justified, at least not in the sense needed for satisfying the traditional account. The fact that Lamb does not realize that she is being deceived may explain and excuse her belief. But that is not the same thing as providing adequate grounds for holding it, which is what, I submit, the traditional account requires. Swallowing lies, even if blamelessly, is not a good reason for believing, and a belief so arrived at is not a justified one. For a belief to be justified in the relevant sense, it must be the case that the evidence the believer has supports the it, where what that means is that the evidence's being what it is makes the proposition believed more likely to be true than it would be otherwise. Someone may be forgiven for believing what he does, even if the evidence on which he bases his belief does not support that belief, as long as it appears to support it and the believer has taken sufficient care in assessing it.³⁸ If we assume that there is no evidence available to Lamb that would suggest that Linus may be deceiving him, he may be seen as blameless. Nonetheless, Linus' lies are not evidence that he owns a Lamborghini. If so, the justification condition of the traditional analysis is not satisfied.³⁹

In fact, deliberate deception is not the only thing that can rob a believer of the kind of justification the traditional account says is needed for knowledge. A simple mistake can do so. Imagine that Lamb has the same evidence of Linus' ownership except for Linus' telling her that the car in question is a Lamborghini.

³⁶ Turri, "In Gettier's," 215.

³⁷ Turri, "In Gettier's," 271.

³⁸ Whether it does depends on whether he has taken as much trouble to make sure that he is not being deceived as the situation demands, which, in turn, depends on what is reasonable for him to believe about his students in general, Linus in particular, and so on.

³⁹ Given Iago's machinations, it is understandable that Othello takes Cassio's possession of Desdemona's handkerchief as evidence of her infidelity. It is because it is not that we, who know that the 'evidence' is planted, say that even if she were in fact unfaithful, he would not know it.

John Biro

It is, in fact a Maserati, but one has to be quite knowledgeable about Italian cars not to mistake one make for the other. Lamb is not and comes to believe that Linus owns a Lamborghini. As it happens, Linus does own both the Maserati he shows Lamb and a Lamborghini he keeps in his garage. Gettierites will be quick to say that Lamb does not know that Linus owns a Lamborghini, and I agree. But the reason why he does not is not, as they claim, that having a justified true belief is not sufficient for knowing. It is that he is clearly not justified in believing what he does.

A similar twist may be given to Gettier's own first case. Suppose Smith, having been asked by Jones to hold the latter's jacket while he is changing a flat tire, is furtively checking all the pockets to see how much money Jones has. He thinks he has counted ten dimes and thus forms the belief that Jones has a dollar. However, in his hurry he missed one of the pockets but counted one of the dimes in another twice. Even if the missed pocket has a dime in it, Smith is not justified in believing that Jones has a dollar, and *that* is the reason he does not know that the man who will get the job has a dollar in his pocket, even if that man turns out to be Jones.

What happens in such cases is that the justification condition is not really satisfied, even though it appears to be to the subject, because the evidence the subject is relying on is not probative of the proposition he believes. To be probative, the evidence must in fact support the proposition, not just be taken by the believer to do so.⁴⁰ The traditional analysis is not threatened by such cases; it requires that the justification condition be in fact satisfied, not merely that it appear to be satisfied to the believer. (If evil-demon worries remain, they have nothing to do with the problem Gettier is alleged to have posed.)

In discussing Turri's example, I have challenged his claim that in it the justification condition is satisfied (as I suggested it may be the case with the stopped clock). That is a different complaint from my earlier one that in the typical Gettier examples it is the belief condition that is not satisfied. But, as I noted before, the real question is whether in any of these examples *both* conditions are, as they need to be if the cases are to be counterexamples to the traditional analysis.

⁴⁰ To say that the belief must in fact be justified is not to abandon fallibilism about justification, a desperate measure advocated by some (Robert Almeder, "The Invalidity of Gettier-Type Counterexamples," *Philosophia* 13, 1-2 (1983): 67-74; Scott Sturgeon, "The Gettier Problem," *Analysis* 53, 3 (1993): 156-164; Trenton Merricks, "Warrant Entails Truth," *Philosophy and Phenomenological Research* 55, 4 (1997): 841-855.) As stressed by Lewis, doing so quickly leads to skepticism. (David Lewis, "Elusive Knowledge," *Australasian Journal of Philosophy* 74, 4 (1996): 549-567.) Being probative should not be confused with being factive.

Even if one's justified but mistaken belief entails a true proposition, and one recognizes that it does, it does not follow that one must come to believe the entailed proposition. One way to see this is through seeing what is wrong with the following, seemingly plausible, objection: Surely, someone can come to believe Q if she believes P and believes that P entails Q. Why not? Can we not infer new beliefs from old beliefs, recognizing logical connections? Is that not what is called deductive reasoning? Why is it not possible for someone to infer from Fa to ExFx, and to really believe, and not just 'pickwickian believe,' that ExFx?

Of course, it is, in most cases. But not if the inferred belief can be true even if the belief from which it is inferred is false. Suppose Poirot says "someone in this room is the murderer" because he believes that the nephew killed the uncle. On subsequently discovering that the nephew has a cast-iron alibi and it was the butler, also present, who committed the dastardly deed, we would not allow Poirot to get away with saying (not that he would), "I was right all along!" Similarly, we would think it a poor joke if Chisholm's sheep-spotter claimed that he had a true belief all along.

There is no harm in saying that the inferred proposition is believed when the proposition from which it is inferred is false as long as we remember that it is believed only in a manner of speaking. But the acceptability of such a *façon de parler* should not be allowed to mislead us into thinking that the belief condition of the traditional analysis is satisfied with respect to the proposition that turns out to be true.

Something like this is recognized by Gaultier in a recent paper in which he argues that it is impossible for Smith to form the belief that is made true by his getting the job and having ten coins in his pocket:

... when the belief that John owns a Ford has been formed in the way indicated in the description of the Gettier case, this belief cannot lead one to form, in addition, the different belief that someone in the company owns a Ford.⁴¹

The reason, according to Gaultier, is that in general

...one cannot believe about the question whether *p* something weaker, more indefinite or undetermined, than what... [one's evidence] appears to one to support or establish about the question whether *p*.⁴²

⁴¹ Gaultier, "An Argument," 267.

⁴² Gaultier, "An Argument," 267. (While Gaultier's discussion is confined to Gettier's first example, it is easy to see how it can be extended to the second: "*p* or *q*" clearly expresses a weaker claim than does "*p*." Gaultier thinks that a commitment to the voluntariness of belief is essential to his argument. I am not sure that I see why. But it is not essential to mine. He also

I think Gaultier is on to something; however, through failing to distinguish between serious and pickwickian beliefs, he commits himself to a stronger thesis than is needed to block the supposed counterexample. Consider again the sheep-in-the-field case. While the proposition expressed by “*There* is a sheep” entails the proposition expressed by “There is a sheep (somewhere),” believing the former does not entail seriously believing the latter. If someone who believed the former were asked whether he believed the latter, no doubt he would say that he did – though in a special tone of voice, I think. We can allow that there is a sense – the philosopher’s sense – in which one believes any proposition one thinks is entailed by a proposition one believes. But we need not agree that one seriously believes everything one thinks is entailed by something one believes.

Thus if by ‘belief’ we mean ‘serious belief,’ Gaultier is right that one cannot infer a ‘There’ belief from a ‘*There*’ belief. But this is not just, or primarily, because the former is weaker. He offers as a general principle that

...at *t*, one cannot believe about the question whether *p* something weaker, more indefinite or undetermined, than what, at *t*, *E* appears to one to support or establish about the question whether *p*.⁴³

This seems to me to be too strong a claim. Having just seen my neighbour enter his house, I believe that there is someone in that house; I will bet you that there is if you claim otherwise. In fact, my neighbour has left through the back door but a burglar has snuck in the same way. The reason why this is not a Gettier case is not that believing that there is someone in the house is weaker than believing that my neighbour is in the house – though it is – but because believing that there is someone (i.e., my neighbour) in the house is not the same thing as believing that there is someone or other in the house.

A mark of seriousness in a belief is that it guides action. If I am in the market for a used Ford and believe that Havit owns the one in the parking lot, it would not be rational for me to go around asking who owns it, as it would be if what I seriously believed was that someone or other in the building did. And if I believed that Secretariat was a dead certainty to win the 1973 Belmont, even though I realized that from the proposition that Secretariat will win it follows that either Secretariat or Twice a Prince, the rank outsider, will win, it would not be rational to split my bet between the two. Not only that – taken as serious beliefs,

thinks, as I do not, that perhaps the fake-barn case and certainly the stopped-clock case are counterexamples to the traditional analysis.

⁴³ Gaultier, “An Argument,” 267.

the two in each pair are incompatible, as shown by the fact that I cannot act on both.⁴⁴

VII. To have a counter-example to the traditional analysis, we need a case where it is clear that one satisfies the conditions it lays down as sufficient for knowledge and it is also clear that one lacks knowledge. I have argued that Gettier's original cases and those that follow the same pattern do not succeed as counter-examples because in those cases one of the three conditions is not satisfied. In many, the subject does not seriously hold the belief that turns out to be true. With some (fake-barn, stopped-clock), I am content with a Scotch verdict: it is at least not *clear* that they are counter-examples to the traditional analysis, because it is not clear that the subject lacks knowledge.

I have no proof that a case cannot be described (not even that one has not been) that does not exhibit one of these three patterns and in which it is clear that all three conditions are satisfied and also that the subject does not know. But I do think that seeing the ways in which the well-known cases I have discussed fail strongly suggests that the confident claims that the traditional analysis has been decisively refuted should be treated with caution. There is, perhaps, life in the old analysis yet.

⁴⁴ It gets worse. If I believed, on the basis of recognizing that it follows from my belief that Secretariat would win, that either Secretariat or Twice a Prince would win, I would have to believe that Secretariat or Sham would win and that Secretariat or My Gallant would win and that Secretariat or Private Smiles would win. (There were only five runners.) Ten dollars on Secretariat would have netted me a profit of two dollars; splitting my ten dollars five ways, a loss of seven dollars and eighty cents. In the 1978 Belmont, even though I believed, along with many, that, given the extra furlong, Alydar would finally catch up with Affirmed, I did also believe that either Alydar or Affirmed would win. But that was not by way of deducing the second proposition from the first. It was because there was independent reason to believe the latter, as reflected in the respective odds (6/5 and 3/5). I did not also believe that either Alydar or Darby Creek Road or that either Alydar or Judge Advocate or that either Alydar or Noon Time Spender would win. In fact, believing any of these is incompatible with believing what I did, as is believing any one of Gettier's 'constructed' disjunctions is with believing one of the others.