

THE INESCAPABILITY OF GETTIER PROBLEMS

BY LINDA ZAGZEBSKI

Gettier problems arise in the theory of knowledge when it is only by chance that a justified true belief is true. Since the belief might easily have been false in these cases, it is normally concluded that they are not instances of knowledge.¹ The moral drawn in the thirty years since Gettier published his famous paper is that either justified true belief (JTB) is not sufficient for knowledge, in which case knowledge must have an 'extra' component in addition to JTB, or else justification must be reconceived to *make it* sufficient for knowledge. I shall argue that given the common and reasonable assumption that the relation between justification and truth is close but not inviolable, it is not possible for either move to avoid Gettier counter-examples. What is more, it makes no difference if the component of knowledge in addition to true belief is identified as something other than justification, e.g., warrant or well-foundedness. I conclude that Gettier problems are inescapable for virtually every analysis of knowledge which at least maintains that knowledge is true belief plus something else.

Notice first that Gettier problems arise for both internalist and externalist notions of justification. On internalist theories the grounds for justification are accessible to the consciousness of the believer, and Gettier problems arise when there is nothing wrong with the internally accessible aspects of the cognitive situation, but there is a mishap in something inaccessible to the believer. Since justification does not guarantee truth, it is possible for there to be a break in the connection between justification and truth, but for that connection to be regained by chance.

The original 'Smith owns a Ford or Brown is in Barcelona' case is an example of this sort. Here we are to imagine that Smith comes to you bragging about his new Ford, shows you the car and the bill of sale, and generally gives you lots of evidence that he owns a Ford. Basing what you think on the evidence, you believe the proposition 'Smith owns a Ford', and from that you infer its disjunction with 'Brown is in Barcelona', where Brown is an acquaintance and you have no reason at all to think he is in Barcelona. It turns out that Smith is lying and owns no Ford, but Brown is

¹ 'Is Justified True Belief Knowledge?', *Analysis* 23 (1963), pp. 121–3.

by chance in Barcelona. Your belief 'Smith owns a Ford or Brown is in Barcelona' is true and justified, but it is hardly the case that you know it.

In this case the problem arises because in spite of the fact that you have done everything to reach the truth from your point of view and everything that anyone could expect of you, your efforts do not lead you to the truth. It is mere bad luck that you are the unwitting victim of Smith's lies, and only an accident that a procedure that usually leads you to the truth leads you to believe the falsehood 'Smith owns a Ford'. The fact that you end up with a true belief anyway is due to a second accidental feature of the situation – a feature that has nothing to do with your cognitive activity. What generates the problem for JTB, then, is that an accident of bad luck is cancelled out by an accident of good luck. The right goal is reached, but only by chance.

Internalist theories are not the only ones afflicted with Gettier problems, contrary to a recent claim made by Alvin Plantinga.² Consider how the problem arises for reliabilism. In this group of theories believers are justified when their beliefs are formed in a reliable, or truth-conducive, manner. On this account also there is no guarantee that justified beliefs are true, and a breakdown in the connection between a reliable belief-forming process and the truth is possible. When that happens, even if you manage to hit on the truth anyway, you do not have knowledge.

The well-known fake barn case can be described as an example of this sort. Here we are to imagine that you are driving through a region in which, unknown to you, the inhabitants have erected three barn façades for each real barn in an effort to make themselves look more prosperous. Your eyesight is normal and reliable enough in ordinary circumstances to spot a barn from the road. But in this case the fake barns are indistinguishable from the real barns at such a distance. As you look at a real barn you form the belief 'That's a fine barn'. The belief is true and justified, but is not knowledge.

As in the first case, the problem arises because of the combination of two accidental features of the cognitive situation. It is only an accident that visual faculties normally reliable in this sort of situation are not reliable in this particular situation; and it is another accident that you happened to be looking at a real barn and hit on the truth anyway. Again the problem arises because an accident of bad luck is cancelled out by an accident of good luck.

Gettier problems cannot be avoided by Alvin Plantinga's new theory either. Plantinga calls the property that in sufficient quantity converts true belief into knowledge 'warrant' rather than 'justification'. On his proposal warrant is the property a belief *B* has for believer *S* when *B* is produced in *S*

² Alvin Plantinga, *Warrant and Proper Function*, (Oxford UP, 1993), p. 48.

by *S*'s faculties working properly in the appropriate environment, according to a design plan successfully aimed at truth.³ But Plantinga does not maintain that every warranted belief is true any more than reliabilists maintain that every reliably formed belief is true or internalists maintain that every internally justified belief is true. Let us see if we can form a Gettier case for Plantinga's theory parallel to the other two cases we have considered. To do so we need to look for a situation in which *S*'s faculties are working the way they were designed to in the appropriate environment, but *S* unluckily has a false belief. We can then add a second accident which makes the belief true after all.

Suppose that Mary has very good eyesight, but it is not perfect. It is good enough to allow her to identify her husband sitting in his usual chair in the living room from a distance of fifteen feet in somewhat dim light (the degree of dimness can easily be specified). She has made such an identification in these circumstances many times. Each time her faculties have been working properly and the environment has been appropriate for the faculties. There is nothing at all unusual about either her faculties or the environment in these cases. Her faculties may not be functioning perfectly, but they are functioning well enough, so that if she goes on to form the belief 'My husband is sitting in the living room', that belief has enough warrant to constitute knowledge when true and we can assume that it is almost always true.

The belief is *almost* always true, we say. That is because warrant in the degree necessary for knowledge does not guarantee truth, according to Plantinga. If it *did* guarantee truth, of course, the component of truth in the analysis of knowledge would be superfluous. Knowledge would simply be warranted belief. So it is possible for Mary to make a mistake even though her faculties are functioning properly enough for knowledge and the environment is normal for the faculties. Let us look at one such case.

Suppose Mary simply misidentifies the chair-sitter who is, let us suppose, her husband's brother. Her faculties may be working as well as they normally do when the belief is true and when we do not hesitate to say it is warranted in a degree sufficient for knowledge. It is not a question of their suddenly becoming defective, or at any rate, more defective than usual, nor is there a mismatch between her faculties and the environment. No one is dressing up as her husband to fool her, or anything like that, so the environment is not abnormal as the fake barn case is abnormal. Her degree of warrant is as high as it usually is when she correctly identifies her

³ The wording I have used can be found in Plantinga's book *Warrant and Proper Function*. A very similar wording can be found in 'Positive Epistemic Status and Proper Function', in J.E. Tomberlin (ed.), *Philosophical Perspectives 2: Epistemology* (Atascadero: Ridgeview, 1988), pp. 1–50. In that paper he calls 'positive epistemic status' what he now calls 'warrant'.

husband since even in those cases it is true that she *might* have misidentified the chair-sitter if it had been her husband's brother instead. Of course, she usually has no reason to suspect that it *is* her husband's brother and we can imagine that she has no reason to suspect so in this case either. Maybe she knows that her husband's brother looks a lot like him, but she has no reason to believe that he is in the vicinity, and, in fact, has strong reason to believe he has gone to Australia. So in the case we are considering, when Mary forms the false belief, her belief is as warranted as her beliefs normally are in these circumstances. In spite of well-functioning faculties and a benign environment, she just makes a mistake.

Now, of course, *something* has gone wrong here, and that something is probably in Mary rather than in the environment. It may even be correct to say that there is a minor defect in her faculties; perhaps she is not perfectly attentive or she is a little too hasty in forming her belief. But she is no less attentive and no more hasty than she usually is in such cases and usually it does not matter. People do not have to be perfectly attentive and perfectly cautious and have perfect vision to have beliefs sufficiently warranted for knowledge on Plantinga's theory. And this is not a *mistake* in Plantinga's theory. It would surely be unreasonable of him to expect perfectly functioning faculties in a perfectly attuned environment as his criteria for the warrant needed for knowledge. So Mary's defect need not be sufficient to bring her degree of warrant down below that needed for knowledge on Plantinga's account.

We can now easily emend the case as a Gettier example. Mary's husband could be sitting on the other side of the room, unseen by her. In that case her belief 'My husband is sitting in the living room' is true and has sufficient warrant for knowledge on Plantinga's account, but she does not have knowledge.

In discussing Gettier problems Plantinga concludes: 'What is essential to Gettier situations is the production of a true belief despite a relatively minor failure of the cognitive situation to match its design'.⁴ But this comment is problematic on his own account. As we have seen, Plantinga considers warrant a property that admits of degree, but it is clear that the degree of warrant sufficient for knowledge does not require faculties to be working perfectly in an environment perfectly matched to them. In Gettier-style cases such as the case of Mary, either the degree of warrant is sufficient for knowledge or it is not. If it is not, then a multitude of beliefs we normally think are warranted are not, and there is much less knowledge in the world than Plantinga's numerous examples suggest. On the other hand, if the degree of warrant *is* sufficient for knowledge, then Plantinga's theory faces Gettier problems structurally identical to those of the other

⁴ 'Positive Epistemic Status and Proper Function', p. 43.

theories. Furthermore, even if some aspect of the Mary example makes it unpersuasive, there must still be cases of warranted false belief on Plantinga's theory if the component of truth in knowledge is not redundant. With such a case in hand a Gettier example can be constructed by adding a feature extraneous to the warrant of the believer which makes the belief true after all. In such a case the degree of warrant is unchanged, but it is not knowledge since it might just as well have been false.

It is not enough, then, to say that Gettier problems arise because of a minor mismatch between faculties and environment. What Plantinga should have said is that the problem is due to a relatively minor failure of the cognitive situation to connect to the truth. As long as the property that putatively converts true belief into knowledge is analysed in such a way that it is strongly linked with the truth, but does not guarantee it, it will always be possible to devise cases in which the link between such a property and the truth is broken but regained by accident. Such is the nature of Gettier cases.

The three examples we have considered suggest a general rule for the generation of Gettier cases. It really does not matter how the *particular* element of knowledge in addition to true belief is analysed. As long as there is a small degree of independence between this other element and the truth, we can construct Gettier cases by using the following procedure: 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. The falsity of the belief will not be due to any systematically describable element in the situation, for if it were, such a feature could be used in the analysis of the components of knowledge other than true belief, and then truth would be entailed by the other components of knowledge, contrary to the hypothesis. The falsity of the belief is therefore 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. The second element must be independent of the element of warrant so that the degree of warrant is unchanged. The situation might be described as one element of luck counteracting another. We now have a case in which the belief is justified (warranted) in a sense strong enough for knowledge, the belief is true, but it is not knowledge. The conclusion is that as long as the concept of knowledge closely connects the justification component and the truth component, but permits *some* degree of independence between them, justified true belief will never be sufficient for knowledge.

It is often observed that in typical Gettier cases the justified belief depends upon or otherwise 'goes through' a false belief, so a way to handle these cases is to add what are commonly called 'defeasibility conditions' to

the analysis of knowledge. This move was especially popular during the sixties and seventies. It adds to the requirement that knowledge be justified true belief the restriction that the belief in question must also be justified in certain counterfactual situations. One way to define these conditions is in terms of the psychological effect on the subject, as in Steven Levy's definition of a defeasibility condition as 'a requirement to the effect that for S to know that p there must be no other evidence against p strong enough to undermine S 's belief that p , should this evidence come to S 's attention'.⁵

The three cases I have just described do have the feature that there is a false belief in the neighbourhood of the belief in question which is such that, should the subject discover its falsehood, that would undermine the belief in the proposition in question. So your belief that either Smith owns a Ford or Brown is in Barcelona is undermined if you discover that Smith does not own a Ford. Your belief that this is a barn is undermined if you discover that most objects that look like barns in this vicinity are not real barns. Mary's belief that her husband is sitting in the living room is undermined if she discovers that that man sitting over *there* in a particular chair in the living room is not her husband. In each case were S to be advised of the falsity of the underlying belief, S would retract the belief under discussion. The belief would be defeated by such new information.

This move puts a strain on the independence of the justification/defeasibility condition and the truth condition. If S 's belief that p is false, there will obviously be many other propositions which are logically or evidentially connected to p which are false also. Should S become aware of any of these propositions, that may easily undermine S 's belief that p , assuming S is rational. This means that the falsehood of p is incompatible with a strong defeasibility condition, contrary to the hypothesis that the justification and defeasibility components of knowledge do not entail the truth condition. This problem is even more apparent in statements of the defeasibility condition in terms of evidential support rather than a psychological requirement, as in Pappas and Swain's definition: 'the evidence e must be sufficiently complete that no further additions to e would result in a loss of justification and hence a loss of knowledge'.⁶ Obviously, if the belief is *false*, further additions to e will result in a loss of justification, and hence a loss of knowledge.

Strong defeasibility conditions, then, threaten the assumption of independence between the justification (warrant) condition and the truth

⁵ Steven Levy, 'Defeasibility Theories of Knowledge', *Canadian Journal of Philosophy* 7 (1977), p. 115.

⁶ George Pappas and Marshall Swain (eds), *Essays on Knowledge and Justification* (Ithaca and London: Cornell UP, 1978), p. 27.

condition for knowledge. But weaker defeasibility conditions are subject to Gettier-style counter-examples following the pattern described above. In each case we find an example of a false belief which satisfies the justification and defeasibility conditions, and then make the belief true anyway due to features of the situation independent of the satisfaction of those conditions.

Suppose Dr Jones, a physician, 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 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.

In this case there is nothing defective in the justification of Dr Jones' belief that Smith has virus *X* and no false belief figures causally or evidentially in her justification, nor is there any false belief in the neighbourhood. Furthermore, she would have believed that Smith has virus *X* in a wide range of counterfactual situations. None the less, let us suppose that the belief is false. Smith's symptoms are due to a distinct and unknown virus *Y* and the fact that he exhibits high antibody levels to virus *X* is 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, but it is both justified and undefeated. Of course, given that the belief is false, there must be *some* evidence against it accessible to her in some counterfactual circumstances, so if defeasibility conditions are strong enough, no false empirical belief passes the test. But as said above, that is to impose an unreasonably strong defeasibility condition, one that makes the justification/defeasibility condition entail truth. The most reasonable conclusion to draw in this case, then, is that Jones' belief is justified and undefeated, but 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. So while the evidence upon which Dr Jones bases her diagnosis does make it highly probable that Smith has *X*, the fact that Smith has *X* has nothing to do with that evidence. In this case, then, Dr Jones' belief that Smith has virus *X* is true, justified and undefeated, but it is not knowledge.

It appears, then, that no account of knowledge as true belief plus something else can withstand Gettier objections as long as there is a small degree of independence between truth and the other conditions of knowledge. What are our alternatives? We have already seen that one way to solve the problem is to give up the independence between the justification condition and the truth condition. Justification would be defined in such a way that no false belief can satisfy it. Since Gettier cases are based on situations in which the belief is true, but it might just as well have been false, all such cases would be excluded from the class of justified (warranted) beliefs. On this approach the element of truth in the account of knowledge is superfluous and knowledge is simply justified (warranted) belief. 'S is justified in believing p ' entails p . Few philosophers have supported this view.⁷

So Gettier problems can be avoided if there is no degree of independence at all between truth and justification. A second way to avoid them is to go to the opposite extreme and to make the justification condition and the truth condition almost completely independent. It could still be the case that justification puts the subject in the best position available for getting the truth, but if the best position is not very good, most justified beliefs will be false. Perhaps most justified scientific hypotheses since the world began have been false. Perhaps Plato, Spinoza, Kant and Hegel were justified in believing their metaphysical theories, but most of their theories (at least) were false. Still, if one of them is true, some theorists might be willing to call it knowledge. On this approach the element of luck permitted in the state of knowledge is so great that alleged counter-examples based on luck do not count against it. From this viewpoint, Gettier cases would simply be accepted as cases of knowledge. After all, if knowledge is mostly luck anyway, there will be nothing bothersome about a case in which the truth is acquired by luck.

Perhaps neither of these alternatives will appeal to most philosophers, who find the idea that there is a small but real degree of independence between justification and the acquisition of truth just too attractive to give up. A third reaction to the problem, then, is to accept the fact that no 'true belief + x' account of knowledge will be sufficient, but that it will always be necessary to add the element of luck to the analysis. So knowledge is true belief + x + luck. This approach recognizes the fact that the concept

⁷ An exception is Robert Almeder, 'Truth and Evidence', *The Philosophical Quarterly*, 24 (1974), pp. 365–8. Almeder's reason for maintaining that 'S is justified in believing p ' entails p is that the *determination* of the fact that a belief p is justified entails the determination of the fact that p is true. I find this implausible, since (1) there are many ways to determine the truth-value of a proposition p independently of the justification of a particular believer in believing p ; and (2) even if the act of determining that a belief is justified included the determination of its truth, it does not follow that the fact that a belief is justified entails its truth.

we substitute for 'x' ought to be one that has a strong general connection with the acquisition of truth, but that an inviolable connection would be unreasonable. On the other hand, it also recognizes the fact that we are much less forgiving with the concept of knowledge itself. The connection between justification or whatever it is we substitute for 'x' and truth must exist in each and every particular case of knowledge. The notion of knowledge requires success, both in reaching the goal of truth, and in reaching it via the right cognitive path. The notion of justification or warrant is less stringent, requiring only that the right path is one that is *usually* successful at getting the truth. It is this difference between the notion of knowledge and the notion of justification that is responsible for Gettier problems.

Almost every contemporary theory of justification or warrant aims only to give the conditions for putting the believer in the best position for getting the truth. The best position is assumed to be very good, but imperfect, for such is life. Properly functioning faculties need not be working perfectly, but only well enough; reliable belief-producing mechanisms need not be perfectly reliable, only reliable enough; evidence for a belief need not support it conclusively, but only well enough; and so on. As long as the truth is never assured by the conditions which make the state justified, there will be situations in which a false belief is justified. I have argued that with this common, in fact, almost universal assumption, Gettier cases will never go away.

Loyola Marymount University