

4 The Gettier Problem and Externalism

Rodrigo Borges

4.1 Introduction

The problem that Edmund Gettier (1963) posed in his classic paper challenges anyone who is trying to understand knowledge. Things are not different for externalists and their theories of justification and knowledge. In what follows, I look at the Gettier problem, and at three externalist solutions to this problem that have been proposed by reliabilists, by truth-tracking theorists, and by virtue epistemologists. Along the way, I identify apparent limitations in those proposed solutions.

4.2 The Gettier Problem

“The Gettier problem” traditionally refers to the problem in conceptual analysis posed by Edmund Gettier in his 1963 classic “Is justified true belief knowledge?”¹ Gettier allegedly produced two counterexamples to the claim that, necessarily, someone, S, falls under the concept KNOWLEDGE if and only if S has a justified true belief.² According to Gettier, Plato (Meno 98), Chisholm (1957, p. 16), and Ayer (1956, p. 34) all accepted this view.³ Although the cases that Gettier presented did not challenge the claim that justification, truth, and belief are *necessary conditions* on KNOWLEDGE, those cases did challenge the claim that justification, truth, and belief form a *sufficient* set of conditions for KNOWLEDGE. So far, there is no consensus in the literature as to whether anyone has succeeded in solving the Gettier problem.

¹ It is not clear that Gettier was the first to raise that *type* of problem for accounts of knowledge. There is good evidence that medieval and ancient Indian philosophers, for example, had considered the same type of problem. See Hilpinen (2017) for discussion of the Gettier problem before Gettier.

² I will use capital letters when referring to concepts.

³ It is not entirely clear whether all of them accepted this analysis of KNOWLEDGE. See Klein (2017) for some doubts concerning Gettier’s exegesis here.

In his paper, Gettier argued that a subject can be shown to be ignorant of a true proposition, p , that she justifiably believes, if we accept two general principles about epistemic justification:

Fallibility. It is possible for a person to be justified in believing a proposition that is in fact false.

Justification Closure. For any p , if S is justified in believing p , and p entails q , and S deduces q from p and accepts q as a result of this deduction, then S is justified in believing q .

Even though Gettier did not tell us why he thought that Fallibility and Justification Closure are true, we may think that those principles are acceptable on the following grounds. Consider Fallibility: why think that it is possible for one to believe a falsehood with justification? Supposedly, because of something like the experience that all of us have undergone at some point in our lives, of being misled into believing something false on the basis of what (at the time) seemed like good evidence (optical illusions and conniving liars do just that). As for Justification Closure, the principle seems to enshrine the centuries-old idea that deduction is the safest method that we can deploy to expand our body of justified beliefs, since the information in the deduced proposition “is already contained” in the premise of our reasoning.

Once we accept Fallibility and Justification Closure, Gettier’s original cases seem to show quite conclusively that Smith, the protagonist of those cases, has a justified true belief but no knowledge. In each of Case I and Case II, Smith has “strong evidence” for his belief in the false proposition in 1, and he accepts 2 as a result of a deduction from 1:

| Case I | Case II |
|--|--|
| 1. Jones is the man who will get the job, and Jones has ten coins in his pocket. | 1. Jones owns a Ford. |
| 2. The man who will get the job has ten coins in his pocket. | 2. Either Jones owns a Ford, or Brown is in Barcelona. |

Given the details that Gettier provides for Case I and for Case II, it seems clear that Smith satisfies Fallibility and Justification Closure in both cases. It seems equally clear that he also fails to know in each case. In Case I, 2 is true in virtue of the number of coins in Smith’s own pocket (a fact of which Smith knows nothing), while Smith’s belief in 2 is based on his counting the coins in

Jones's pocket, whom Smith mistakenly believe to be the man who will get the job. In Case II, Smith's belief in 1 is false but justified because Jones goes around pretending a rented Ford is his car, and Smith's belief in 2 is true only because Brown, whose whereabouts are unknown to Smith, is in fact in Barcelona. The prevalent intuitive judgment about Cases I and II is that Smith has a justified true belief in 2, but fails to know it. Call this *the Gettier intuition*. Epistemologists usually take the Gettier intuition to support strongly the belief that one does not necessarily fall under the concept KNOWLEDGE if one has a justified true belief.⁴

As a problem in conceptual analysis, one can solve the Gettier problem by either (a) adding more conditions to the existing set of necessary conditions (that is, justification, truth, and belief) until one gets a set of conditions that are jointly sufficient for the analysis of KNOWLEDGE, or (b) by challenging Gettier's characterizations of one or more of the existing necessary conditions. As we will see in the next section, *internalists* have, for the most part, favored the first strategy, while *externalists* have, for the most part, favored the second strategy.

Before we look at externalist solutions to the Gettier problem, it might be a good idea for us to address a particular worry. Recently, some have voiced skepticism concerning the Gettier problem. Timothy Williamson (2000, p. 31), for example, has argued that the pursuit of conceptual analyses in general (and the pursuit of an analysis of KNOWLEDGE in particular) is a "degenerating research programme" and should be abandoned. He goes on to say that KNOWLEDGE is in fact unanalyzable. If Williamson is right, the Gettier problem as Gettier presented it is a waste of time. I will not try to decide here whether Williamson is in fact right about this.⁵ Instead, I would like to suggest that, even if he is right and KNOWLEDGE is unanalyzable, the Gettier problem is not a waste of time. Although the Gettier problem is traditionally a problem in conceptual analysis, this is not the only possible interpretation of the epistemological significance of the Gettier intuition. Regardless of whether we can give a noncircular analysis of KNOWLEDGE, we might still be interested in *explaining why Smith fails to know in Case I and in Case II*. Call this latter problem the *Explanatory Gettier Problem*. This problem contrasts with the related, but, strictly speaking, independent

⁴ Of course, for all that Gettier has said, one sometimes falls under KNOWLEDGE even though one has nothing more than a justified true belief. After all, he showed at most that not *all* justified true beliefs are in the extension of KNOWLEDGE; he did not show that *no* justified true belief is.

⁵ But, for discussion of this and other aspects of Williamson's views on knowledge, see Greenough and Pritchard (2009).

problem in conceptual analysis that we have been discussing, which may now be labeled the *Conceptual Gettier Problem*. Even if Williamson is right about the Conceptual Gettier Problem being a waste of time, the Explanatory Gettier Problem might still be interesting. This is so in part because Smith's situation poses a threat to the proposed analysis of KNOWLEDGE only if his justification is, other things being equal, *good enough for knowledge*. If the justification that his belief enjoys is not good enough for knowledge, then Smith does not satisfy one of the conditions on KNOWLEDGE, and Case I and Case II provide no counterexample to the analysis. If the justification that Smith's belief enjoys is, on the other hand, good enough for knowledge, then the Explanatory Gettier Problem (that is, the question of why Smith fails to know) becomes particularly salient. After all, Smith's belief in 2 has much going for it, epistemically speaking (it is true and it enjoys knowledge-grade justification). How could someone in such an otherwise strong epistemological position fail to know? This latter question should interest us regardless of whether or not KNOWLEDGE is analyzable. Or so I think.

4.3 Externalist Solutions to the Gettier Problem

Externalist views are often described as the rejection of one or more forms of internalism. The dispute between those two camps is over whether the factors that determine epistemic justification are purely internal to one's mind. Externalists say that factors *external* to one's mind may help determine whether (and to what extent) one is justified in believing a particular proposition. Internalists argue that only factors internal to one's mind can determine whether (and to what extent) one is justified in believing a particular proposition. I will expand below on the distinction between internalism and externalism views, but the bulk of this section will be devoted to three influential externalist views – reliabilism, truth-tracking views, and virtue theories – and to what they have to say about the Gettier problem.⁶

The two main versions of epistemological internalism are *mentalism* and *accessibilism*. According to mentalism, all justifying factors are internal to one's mind in the sense that the degree to which one is justified in believing that *p* at some particular time supervenes on the totality of one's mental state at that time. According to accessibilism, all justifying factors are internal to

⁶ There is a vast literature on the internalism/externalism debate in epistemology, and I will not rehash it here. My focus is on how likely it is that an externalist view would solve the Gettier problem. For a good overview of the debate, see, among many others, Alston (2005).

one's mind in the sense that the degree to which one is justified in believing that p at some particular time supervenes on one's past, present, or at least potential awareness of one's reasons for believing that p . The idea animating each form of internalism seems to be that one ought to follow one's evidence where it leads, and that one can follow only that which presents itself to conscious awareness. Perhaps the best-known implementation of this idea is in the epistemology of Rene Descartes (2008), which takes all justifying factors to be objects of immediate reflective awareness (*ideas*).

It should be noted that mentalism and accessibilism are, in principle, independent of one another. That is, there is no contradiction involved in accepting only one of those views.

Accessibilism does not necessitate mentalism. For one thing, if the class of mental states on which justification supervenes includes only *pure mental* states (that is, states whose obtaining does not depend on any extramental condition being satisfied), and one's version of accessibilism allows for justification to come from one's access to mental states that are *not* purely mental (such as knowledge),⁷ then one has a form of accessibilism that does not entail mentalism. Perhaps more controversially, if, as some (such as Kripke [1980] and Turri [2010]) have suggested, there is such a thing as a priori knowledge of propositions about the world outside of one's own mind, then it would in principle be possible for one to be an accessibilist without also being a mentalist. For example, Saul Kripke (1980) argued that one can know a priori that a certain stick, s , is exactly one meter long, if one baptizes s as the "meter stick." If he is right, then the person baptizing s as the "meter stick" has direct a priori access to what justifies her in believing (and knowing) that s is exactly one meter long (namely, the fact that she baptized s as the "meter stick"); but, since this justifier is not a mental state of the subject but is instead a state of the world, the view would be a form of accessibilism that does not necessitate mentalism.⁸

Mentalism does not necessitate accessibilism. Richard Feldman and Earl Conee, the main proponents of mentalism, conceive of their version of mentalism as being independent of accessibilism. They argue (Conee and Feldman 2004, p. 73) that, in some cases, a mental state plays a justificatory

⁷ That is, *if* knowledge is a mental state, then it is not a pure mental state, for its obtaining depends on the extramental condition of *truth* being satisfied.

⁸ I do not mean to commit myself to the view that there are cases of a priori knowledge of the external world. My point is just that *if* there are such cases, *then* one could be an accessibilist without also being a mentalist. As I suggested earlier, the issue of a priori knowledge of the external world is contentious. See Turri (2010) for detailed discussion.

role even if that state is not accessible to the subject. They imagine a case in which a huge conjunction stored in memory defeats the subject's justification for believing that p . Due to the sheer size of the conjunction, the subject is psychologically incapable of grasping it, and hence does not have reflective access to it. This shows, according to Feldman and Conee, that the conjunction plays a justificatory role even though it is not accessible to the subject.

None of this implies that accessibilism and mentalism are not compatible. One could also, in principle, accept both views. The result would be a type of internalism that takes all justifiers to be reflectively accessible mental states. Usually, however, internalists accept either one or the other version of internalism.⁹

Externalists usually reject both mentalism and accessibilism.¹⁰ The general externalist idea is that some of the factors that help determine whether (and to what extent) S is justified are neither mental states of S nor something that S can access via reflection. For instance, it seems that evidence that one does not possess can sometimes be an external factor that determines one's degree of justification. Consider this example. I walk into the supermarket looking for bread and milk. I left my phone in the car. As I walk down the aisle looking for milk, my wife is sending text and voice messages to me, saying that we do have enough milk at home (she checked the fridge). The fact that her messages say (truthfully) that we do have enough milk seems to be a factor relevant to how justified I am in believing that we do not have enough milk. That is the case even though I currently do not have access to those messages and they are not the content of any of my mental states. There might also be a more general argument against the claim that all factors that are relevant to justification are internal to subjects in the way mentalists and accessibilists say that those factors are relevant. For any ground, g , that the internalist considers relevant

⁹ The distinction between mentalist and accessibilist forms of internalism has its limitations. One of those limitations is brought out by Williamson's knowledge-first epistemology. According to him, knowledge is not only a mental state but is also the only thing that can justify belief. (Experience, although an evidence-provider, is not itself a justifier, on this view.) This suggests that Williamson is a mentalist, for the justification of belief supervenes exclusively on mental states (that is, on knowledge). Moreover, Williamson's arguments against the principle according to which S is in a position to know that S knows that p whenever S knows that p , seem to amount to a rejection of accessibilism p . This package of views complicates the distinction between accessibilism and mentalism, because knowledge is not usually seen as a mental state in its own right, but only as a special form of the mental state of belief.

¹⁰ One possible exception is Williamson's knowledge-first view, which rejects accessibilism but accepts a form of mentalism. See note 9 for discussion. Another possible exception is the synthesis between reliabilism (an externalist view) and mentalism proposed by Alvin Goldman (2011).

to how justified S is in believing that p , there is a further fact (call it a 'supporting fact') having to do with whether g is a good or a bad reason for S to believe that p . Clearly, the obtaining of a supporting fact is relevant to whether S is justified in believing that p . The problem for the internalist is that supporting facts are relevant to one's justification even if they are not reflectively accessible to one or if none of one's mental states are about that fact (cf. Comesaña 2005b).

I will now discuss three externalist views and their proposed solution to the Gettier problem.

4.3.1 Reliabilism

David Armstrong (1973), Alvin Goldman (1979), Fred Dretske (1971), and others articulated the view according to which one's awareness of the reasons why one believes what one does is not a requirement on one's belief justifiedness. All that is required for epistemic justification, they contended, is the *reliability* of the cognitive process that causes the belief in question. For example, according to reliabilism, beliefs caused by vision are justified not because the subject who is holding those beliefs is (or could be) reflectively aware of something, but rather in virtue of vision being a reliable cognitive process, one that tends to produce true beliefs more often than false ones when used in favorable enough circumstances (such as good lighting conditions). The reliability of vision, of other cognitive processes that normally we deem reliable (such as the other sensory modalities), and the justification of the beliefs that they cause contrasts sharply, says the reliabilist, with the unreliability and the consequent lack of justification of beliefs caused by processes such as guesswork, wishful thinking, and motivated reasoning. The point that the reliabilist is making here applies not only to innate cognitive processes such as our sensory modalities, but also to learned methods such as calculus, Newtonian mechanics, and the prediction of future events based on the 'reading' of chicken entrails. Supposedly, only the first two methods are, once mastered, reliable in the desired sense.

We need to make one final distinction before we are ready to look at what the reliabilist has to say about the Gettier problem. This is the distinction between *conditional* and *unconditional* reliability. According to the reliabilist, the reliability of processes (e.g. reasoning) whose inputs involve beliefs should be assessed differently from the reliability of processes (e.g. vision) whose inputs do not involve beliefs. That is because the

former type of process tends to output true beliefs *conditional on the input-beliefs also being true*. Since beliefs are not part of the input to the latter type of cognitive process, their propensity to output true beliefs is independent of the truth of other beliefs.

With this in mind, let us look at what the reliabilist has to say about the Gettier problem. Given the influential version of reliabilism proposed by Goldman (1979), for example, Smith fails to know in Case I and in Case II, because he is using a conditionally *unreliable* method. Deduction is a conditionally reliable method. As such, it can only produce justified beliefs if its input-beliefs (that is, the premises in one's deductive inference) are true. The problem is that Smith's deductive inferences have false input-beliefs, thereby yielding an unjustified belief in the conclusion.

Although this solution to the Gettier problem is a direct consequence of the reliabilist account of inferential justification, the latter might be problematic. For one thing, it is not so clear that it yields the right result in cases where the input-belief to one's deductive inference is justified and true but one's belief in the conclusion does not seem to amount to knowledge. Consider, for instance, the following deductive inference:

- (1) My ticket lost.
- (2) If my ticket lost, I cannot buy a Lamborghini.
- (3) Hence, I cannot buy a Lamborghini.

Suppose that I know 2 because I know that I have no way to afford a Lamborghini unless I win the lottery. Suppose further that, although it is true that my ticket lost, I had not checked the lottery result before I performed this deduction. Instead, I based my belief in 1 solely on my statistical knowledge of the odds of any single ticket winning this large and fair lottery. Most philosophers think that I do not know 3 in these circumstances. And they think this in spite of the fact that they also think that I know 2 and that I have (at least) a justified true belief in 1. This is a potential problem for the reliabilist account of inferential justification because my belief in 3 is not a case of knowledge, even though I am not in a Gettier situation, and even though my belief in 3 is true and the output of a conditionally reliable method (hence it is also justified). Given the reliabilist account of conditional reliability and justification, and its treatment of the Gettier problem, cases such as the lottery case should not exist. The view should exclude the possibility of a reliably acquired true belief that is not a case of knowledge. The view seems unable to do that,

however. This case raises doubts about the reliabilist treatment of deduction and the way in which this method transmits justification and knowledge. Since the reliabilist solution to the Gettier problem depends on this view of the reliability of deduction, it is also subject to doubt.

4.3.2 Truth-Tracking

The truth-tracking theory of knowledge proposed by Robert Nozick (1981), Fred Adams and Murray Clarke (2005), among others, radicalizes the externalist position. Even though the reliabilist rejected mentalism and accessibilism, she still saw herself as offering an account of *epistemic justification*. Nozick, on the other hand, dispensed in his analysis of KNOWLEDGE with the notion of justification. Truth-tracking accounts pursue the idea that a belief is an item of knowledge just in case this belief is *sensitive* to changes in the truth-value of the target proposition in circumstances that are similar enough to the actual circumstances in which the belief is formed. If the subject's belief displays this sensitivity, the truth-tracking theorist says that the belief *tracks the truth* and that the subject thereby *knows* the target proposition. According to this view, one knows that p via a certain method, M , if and only if one not only believes truly that p , via M , in one's actual circumstances, but also tracks the truth of p in close enough circumstances where one continues using M . Whether one tracks the truth of p in close enough circumstances is, in turn, a matter of satisfying the following two subjunctive conditionals:

- If p were false and S were to use M to believe that p , then S would not believe, via M , that p .
- If p were true and S were to use M to believe that p , then S would believe, via M , that p .

To the extent that satisfying these conditions is what turns true belief into knowledge, the conditions play a role that is similar to the one played by the justification condition in the analysis of KNOWLEDGE that Gettier criticized. In that sense, we may speak of Nozick's externalist account of justification, since whether S satisfies either of those subjunctive conditionals in a particular circumstance is a fact that is usually neither accessible to, nor the content of, a mental state of S .

As an illustration of the view, consider my knowledge of the fact that there is a computer screen in front of me right now. According to the truth-tracking account, I know this to be the case because I believe truly that there is

a computer screen in front of me right now via normal vision, and if there was no computer screen in front of me right now I would not believe (while still using normal vision) that there was one, and if there was a computer screen in front of me right now, I would believe (if still using normal vision) that there was one.

What does the truth-tracking theory say about the Gettier problem? According to this view, Smith fails to know because he fails to track the truth of the conclusion of his inferences. The view says that one knows something via inference only if one would not believe the premises in one's inference if the conclusion were false, and if one would still believe the conclusion if the premises were true.¹¹ Smith fails the first condition, for he would believe that Jones is the man who will get the job and that Jones has ten coins in his pocket, even if it were not the case that the man who will get the job has ten coins in his pocket. Smith would also believe that Jones owns a Ford even if it were not the case that either Jones owns a Ford or Brown is in Barcelona.

In order for this solution to the Gettier problem to work, the account of inferential knowledge on which it rests should not exclude cases of inferential knowledge. However, there is some reason to worry whether the truth-tracking theory can in fact do that. To see that, consider the following case.¹² After counting fifty-three people in the audience for my talk, I conclude, on that basis, that the 100 handouts that I printed will be enough. However, unbeknownst to me, there are in fact fifty-two people in the audience: I double-counted an audience member who changed seats during my counting of heads. Most people seem to agree that I know that the 100 handouts that I printed will be enough, even though this knowledge is based on a false premise.¹³ The problem is that the truth-tracking theory does not seem able to deliver this result, for my belief that 100 handouts will be enough seems to fail the condition according to which I would still believe that there are fifty-three people in the audience in close-enough circumstances in which my hundred handouts are enough. Arguably, I fail to satisfy this subjunctive conditional in close-enough circumstances where I still believe that my 100 handouts are enough but where I count only fifty-two people in the audience because – in some close-enough scenario – the one person whom I double-counted in the actual scenario does not move. The upshot is that the account of inferential knowledge to which the

¹¹ Cf. Nozick (1981, pp. 233–4). ¹² This case is adapted from Warfield (2005).

¹³ But see Montminy (2014), Schnee (2015), and Borges (2017) for some push-back on this claim.

truth-tracking theory appeals in its proposed solution to the Gettier problem needs to be mended.

4.3.3 Epistemic Virtues

Finally, according to virtue theorists such as Linda Zagzebski (2000) and Ernest Sosa (2007a), a belief is justified in an externalist sense only if this belief is the output of a virtuous cognitive process. This view is a form of externalism because the fact that one's belief is true in virtue of its being the output of a virtuous cognitive process is neither something to which one usually has reflective access, nor is among one's mental states. According to Sosa, for instance, one knows that p if and only if one's belief that p is not only true, but is true *in virtue of* being the output of a virtuous cognitive process. Now, the idea of a belief being true in virtue of being the output of a virtuous cognitive process might sound peculiar at first: one might think that the view confuses what makes the believed proposition true with what causes one to believe the proposition. This is not what the virtue theorist has in mind, however. Rather, he is trying to explain why the subject has a belief that is true. An example might make this distinction clearer. The fact that I cooked dinner might explain why there is dinner, without explaining why dinner is delicious. In fact, given how bad a cook I am, dinner is delicious *in spite of* the fact that I cooked it! Similarly, the fact that my cognitive process, c , caused my belief that p might explain why I have the belief that p , without explaining why my belief that p is true. The thought is that the fact that a certain belief was caused by a *virtuous* cognitive process helps to explain not only why one has the belief, but also why one has a belief that is true: that's what virtuous cognitive processes do – they cause true beliefs a lot more often than they cause false ones. According to Sosa (2007a, pp. 95–6), this distinction explains why Smith fails to know in Case I and in Case II: Smith's belief is true, but it is not true in virtue of being deduced from his premise belief, for this premise is false. However, this proposed solution to the Gettier problem is threatened by cases such as the handout case above, for in those cases the subject seems to have knowledge even though her deduction depends on a false premise.

4.4 Conclusion

Reliabilism, truth-tracking accounts, and virtue theories are externalist views with plenty to say about the Gettier problem. However, it is

not a settled matter whether those views succeed in explaining why Smith fails to know in Case I and in Case II. In particular, the view of inferential knowledge on which those views rely might not be as sound as their solution to the Gettier problem would suggest and require.

PROOF