

*In Gettier's wake**

JOHN TURRI

john.turri@gmail.com

1. *Introduction*

One main goal of epistemology is to define knowledge. Legend has it that the 'traditional' or 'standard' view of knowledge is justified true belief (K=JTB) and that this traditional view reigned supreme for decades, centuries even. As one leading epistemology textbook puts it,

It is reasonable to say [that] some version or other of the traditional conception of knowledge was taken for granted . . . by virtually all philosophers seriously concerned with knowledge in the period from the time of Descartes until the middle of the twentieth century. (BonJour 2001: 43)

But that all changed in 1963 when an unheralded young philosopher at Wayne State University in Detroit, Edmund Gettier, published a paper as short as it has been influential: 'Is Justified True Belief Knowledge?'. Gettier's paper has since engendered a half-century worth of responses. If you added up the number of times that this article has been discussed or cited in the literature (thou-

* This is the penultimate version of a paper forthcoming in *Epistemology: The Key Thinkers* (Continuum), ed. Stephen Hetherington. Please cite the final, published version if possible.

sands of times), and divided that by the number of words in the article (approximately 930), the resulting quotient would be larger than the quotient for any other work of philosophy ever published. Now *if* we call this the ‘**citation per** word formula’ for calculating a publication’s influence, then it’s safe to say that Gettier’s article is the greatest philosophical caper of all time.

2. *Gettier cases and their structure*

Chapter 10 of this volume already introduced us to Gettier’s discussion. But let’s expand on what was said there.

Gettier presented two cases that he thought were clear counterexamples to the JTB theory. In particular, Gettier contended that his cases showed that having a justified true belief was *insufficient* for knowledge, from which it follows trivially that $K \neq \text{JTB}$. A case of this sort is called a *Gettier case*. The *Gettier problem* is the problem of identifying why the subject in a Gettier case lacks knowledge. It is widely assumed that unless we solve the Gettier problem, we’ll be unable to adequately define knowledge. Interestingly, Gettier wasn’t the first to come up with what we now call Gettier cases. According to Bimal Matilal (1986: 135–7), the classical Indian philosopher Sriharsa constructed similar examples in the 1100s to confound his opponents, and Roderick Chisholm (1989) reminds us that Bertrand Russell and Alexius Meinong also constructed such cases decades earlier than Gettier. But still it is customary to call them ‘Gettier cases’.

Gettier cases are easy to construct, once you get the feel for them. Here are two prototypical examples (not Gettier's originals).

(LAMB) 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 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. (Vaguely modeled after Lehrer 1965: 169–70)

Most philosophers who consider this case say that (a) Dr. Lamb does *not* know that at least one of her students owns a Lamborghini, even though (b) she has a justified true belief that at least one of her students owns a Lamborghini.

(SHEEP) Shep is trekking through a pasture. He gazes down across the field and notices an animal. Viewing conditions are optimal. It appears to be an unremarkable sheep, so he believes, 'That's a sheep in this field,' from which he con-

cludes, 'There's at least one sheep in this field.' As it turns out, Shep isn't looking at a sheep, but rather a dog dressed up to look just like a sheep — a very cleverly disguised dog! Shep has no evidence of this deception. The thought that it was a disguised dog never even occurs to him. And yet it's still true that there is a sheep in the field. It's directly behind the dog, hidden from Shep's view. (Adapted from Chisholm 1989: 93)

Most philosophers who consider this case say that (a) Shep does not know that there is a sheep in the field, even though (b) he has a justified true belief that there is a sheep in the field.

Gettier claimed that the success of these cases as counterexamples to the JTB theory depends on two principles. First, that justification isn't factive. This means that it's possible to have a false justified belief. Second, that justification is closed under deductive entailment. This means that if you're justified in believing some proposition P, and P entails some other proposition Q, and you deduce Q from P, and believe Q based on that deduction, then you're justified in believing Q too.

In light of these two points, here is one way to understand the 'recipe' for generating Gettier cases (Feldman 2003: 28). Begin with a justified false belief that P (which is possible, if justification isn't factive). Then have the protagonist deduce a true consequence, Q, of the justified belief that P, and have the protagonist believe Q on the basis of this deduction (surely this is possible). The resulting belief in Q will be justified (by the assumption that justification is

closed under deductive entailment). And the overall result will be a justified true belief that Q, without knowledge that Q.

Here is another way of understanding the Gettier recipe (Zagzebski 1996). Start with a belief sufficiently justified to meet the justification requirement for knowledge. Then add an element of bad luck that would normally prevent the justified belief from being true. Lastly add a dose of good luck that ‘cancels out the bad,’ so the belief ends up true anyhow. The justification of the justified true belief appears oddly disconnected from the truth, and the overall result will be a justified true belief, one which doesn’t amount to knowledge.

3. Some proposed solutions to the Gettier problem

Attempted solutions to the Gettier problem are legion. Some responses are conservative, in that they hew closely to the original JTB theory, introducing as little change as possible to handle the cases. Some responses are radical, in that they break decisively with the spirit of the JTB account, either dramatically refashioning the justification requirement, or even eliminating it entirely. Here I’ll review some of the most influential and interesting responses to the Gettier problem (see Shope 1983 and Hetherington 2009 for detail on other approaches).

Some philosophers looked at Gettier cases and thought that the problem amounts to this: the subject has a justified true belief, but the belief is essentially based on a false premise. In LAMB the false

premise is Dr. Lamb's belief, 'My student Linus owns a Lamborghini,' and in SHEEP it is Shep's belief, 'This animal in the field is a sheep' (where 'this' refers to the cleverly disguised dog). This suggests the *no essential false basis theory of knowledge*: (NFB) You know that P just in case (i) P is true, (ii) you believe that P, (iii) your belief that P is justified, and (iv) your belief that P isn't essentially based on any falsehood. (For examples of this idea, see Harman 1973 and Clark 1963).

One problem with NFB is that it can't handle simple variants of Gettier cases. Consider this variant of LAMB.

(LUCKY LAMB): The case is the same as LAMB, except that unbeknownst to Linus he has just inherited a Lamborghini. His cousin died and left it to him.

In this case Dr. Lamb's belief 'My student Linus owns a Lamborghini' is true, so NFB can't handle LUCKY LAMB, because in this case it's true that Linus owns the Lamborghini in question. Another problem with NFB is that it appears to give the wrong verdict in cases like this:

(BLUE DRESS) Bill awaits Monica's arrival. He wonders whether she'll wear a scarlet dress. He hears a step on the staircase and swings around to see Monica enter the room. 'What a dazzling indigo dress!' he thinks, and concludes, 'Monica's dress isn't scarlet.' And he's right: her dress isn't scarlet. But it isn't indigo either. It's ultramarine.

Intuitively Bill knows that Monica's dress isn't scarlet. But his belief

is based on the falsehood that her dress is indigo, so NFB rules that he doesn't know that Monica's dress isn't scarlet. (For examples of this line of thought, see Warfield 2005, Saunders and Champawat 1964).

Other philosophers looked at Gettier cases and thought that the problem is this: the subject has a justified true belief, but the justification is defeated (see, for example, Lehrer and Paxson 1969, Klein 1976). In LAMB Dr. Lamb's justification is defeated by the fact that Linus is deceiving her (or, in LUCKY LAMB, that LINUS is trying hard to deceive her). In SHEEP Shep's justification is defeated by the fact that he's being deceived by a cleverly disguised dog. This suggests the *simple defeasibility theory of knowledge*: (SDT) you know that P just in case (i) P is true, (ii) you believe that P, (iii) your belief that P is justified, and (iv) your justification for believing P is undefeated. Some fact F defeats your justification for believing P just in case (i) you believe P based on evidence E, (ii) E justifies belief in P, but (iii) the combination (E+F) *fails* to justify belief in P. In LAMB the defeater is the fact that Linus is deceiving Dr. Lamb about owning a Lamborghini (or, in LUCKY LAMB, is earnestly trying to deceive Dr. Lamb). In SHEEP the defeater is the fact that Shep is looking at a cleverly disguised dog.

One problem with SDT is that it seems to rule out too much. Consider:

(INSANE) You were just tenured! Excitedly you phone to tell your best friend Sophia the wonderful news. Naturally Sophia believes and congratulates you. However, unbe-

knownst to either of you, your dean just went insane — succumbed to the pressures of profit-driven university governance — and is absolutely certain you were not tenured.

It is a fact that your dean is absolutely certain that you were not tenured, and that fact combined with your testimony fails to support Sophia's belief that you were tenured, technically defeating her justification. Thus SDT rules that Sophia doesn't. But intuitively she does know.

In response it has been suggested that knowledge is *ultimately* undefeated justified true belief. Call this the *modified defeasibility theory*: (MDT) you know that P just in case (i) P is true, (ii) you believe P based on evidence E, (iii) E justifies belief in P, and (iv) E is ultimately undefeated. E is ultimately undefeated just in case there is no fact F such that (E+F) fails to justify belief in P; or if there is such a fact, then there is some further fact F* such that (E+F+F*) does justify belief in P.¹ In such a case F* is a defeater defeater. In INSANE, F* is the fact that your dean's conviction is borne of insanity. The main problem with MDT, however, is that the very device it introduces to give the intuitively correct verdict in INSANITY also deprives it of the ability to handle the original Gettier cases. Consider LAMB. The fact that Linus is deceiving Dr. Lamb is a defeater (=F). But the fact that the modest female student does own a Lamborghini is a defeater defeater (=F*). This last fact is a defeater de-

¹ This is an oversimplification, since there might be defeaters for a defeater defeater. But still, the point is clear enough: whenever justification is defeated (or a defeater defeater is defeated) there's always at least one other fact to defeat the defeater (or to defeat the defeater defeater).

feater because this combination:

E: My student Linus has possession of this Lamborghini, drives it frequently, and has a title to the Lamborghini with his name and birthdate on it.

F: My student Linus does not own this Lamborghini.

F*: That young female student of mine owns a Lamborghini.

justifies Dr. Lamb's belief that at least one of her students owns a Lamborghini. It does this because F* obviously entails that at least one of her students owns a Lamborghini. And it would do so, no matter how many of Dr. Lamb's other students don't own a Lamborghini.

The responses we've looked at so far have been conservative. They respond to Gettier cases by adding a fourth condition to the three conditions featured in the traditional JTB account. But there were more radical responses. Some philosophers looked at Gettier cases and thought that the problem amounts to this: the fact that P doesn't cause the subject to believe that P (Goldman 1967). In LAMB it's a fact that one of Dr. Lamb's students owns a Lamborghini, but it isn't this fact (namely, the fact that the female student owns one) that causes Dr. Lamb to believe that one of her students owns a Lamborghini. In SHEEP it's a fact that there is a sheep in the field, but it isn't this fact that causes Shep to believe it. This observation led to *the causal theory of knowledge*: (CTK) you know that P just in case (i) P is true, (ii) you believe that P, and (iii) the fact that P is true causes you to believe that P. CTK gives up on jus-

tification entirely as a condition on knowledge. The problem with CTK is that it is easy to introduce deviant causal chains into the description of any Gettier case, which would make it true that the relevant fact causes the subject to believe that P. For example, suppose that we add the following background to SHEEP. A clever farmer dressed up the cleverly disguised dog to fool Shep. The farmer was caused to do this, oddly enough, by the fact that there is at least one sheep in the field as Shep treks by. So the fact that there is at least one sheep in this field (=P) caused the farmer to dress up the dog, which caused Shep to believe that there is at least one sheep in this field. So the fact that P caused Shep to believe that P. CTK thus rules that Shep knows that P. But intuitively this is the wrong verdict.

Another problem with CTK is that if we impose a causal requirement on knowledge, then it becomes difficult to avoid skeptical consequences for beliefs about abstract matters, such as mathematical and logical truths, because it isn't clear that, say, the fact that $2+2=4$ can cause anything. It also becomes difficult to explain how we know things about the future, because it doesn't seem possible for future facts to cause our beliefs now. Strategies for overcoming these problems have been proposed, but not to the satisfaction of many.

A descendant of CTK is reliabilism about justification and knowledge.² Rather than give up entirely on justification as a condi-

² Both reliabilism and the CTK were discussed also in Chapter 10 of this volume.

tion on knowledge, some argued that we can understand justified belief as belief produced by a reliable cognitive process (Goldman 1979), and then understand knowledge as roughly justified true belief, where ‘justification’ is given the relevant reliabilist reading. The result is *process reliabilism*: (PR) you know that P just in case (i) P is true, (ii) you believe that P, and (iii) your belief that P is produced by a reliable cognitive process (i.e. your belief is ‘justified’). PR might handle Gettier cases by pointing out that in LAMB, for example, Dr. Lamb’s belief is produced by making deductions based on the testimony of someone who is trying to deceive her, which plausibly isn’t a reliable process. And in SHEEP Shep’s belief is produced by making deductions based on mistaken appearances, which plausibly isn’t reliable either.

The main criticism of PR is that it has no principled way of *individuating* cognitive processes, and so no principled way of deciding whether any given true belief amounts to knowledge (Conee and Feldman 1998). For example, why say that Shep is basing his deductions on misleading appearances, rather than on perceptual experience? Perceptual experience is reliable. But then why doesn’t Shep know? Another problem with PR is that it can’t handle simple variants of Gettier cases. Consider this variant of SHEEP.

(SPECIAL DOG) The case is the same as SHEEP, except that the disguised dog is very special. It tracks Shep and appears to him *only* when at least one sheep is nearby. It wouldn’t appear to him unless there were a nearby sheep. It also prevents him from encountering any other non-sheep

sheep-lookalike that would mislead him into concluding that there is at least one sheep nearby.

In SPECIAL DOG the following method seems perfectly reliable for Shep: from the fact that something looks like a sheep nearby, conclude that there is at least one sheep nearby. So PR rules that in SPECIAL DOG Shep knows that there is at least one sheep nearby. But it would be very surprising if Shep knew in SPECIAL DOG but not in SHEEP.

Other philosophers looked at Gettier cases and thought that the problem is this: it's just an accident that the subject's belief is true (e.g. Unger 1968). In LAMB it's just an accident that Dr. Lamb ended up being right that at least one of her students owned a Lamborghini. And in SHEEP it's just an accident that Shep ended up being right that there's at least one sheep in the field. This suggests the *no-accident theory of knowledge*: (NAT) you know that P just in case (i) P is true, (ii) you believe that P, and (iii) it is not at all an accident that your belief that P is true. NAT omits justification from its definition of knowledge, which leaves open the possibility that there can be 'unreasonable' knowledge, that is, knowledge which the subject is unjustified in believing is true (Unger 1968: 164). For example, if an epistemic guardian angel watched over you and ensured that your every wish came true, then wishful thinking would be a way for you to gain knowledge, since it would be no accident that your wishful beliefs turned out to be true. Many judge this to be an absurd consequence of the view — surely believing something because you want it to be true isn't a way of gaining knowledge! An-

other potential problem with NAT is that it's very difficult to explain what clause (iii) amounts to.

A related family of views propose a *safety condition* on knowledge (Sosa 1999, Pritchard 2005), which is intended to give content to the idea that knowledge can't be "accidental" or "lucky." The most conservative version of a safety-based view simply appends a safety condition to the traditional analysis, yielding the *safe justified true belief theory of knowledge*: (SJTB) you know that P just in case (i) P is true, (ii) you believe that P, (iii) your belief that P is justified, and (iv) your belief that P is safe. A true belief is safe just in case it wouldn't easily have been false. What does it mean to say that a true belief 'wouldn't easily have been false'? There's no precise way to define this, but the intuitive idea is that *something significant* would have had to change in order to have made the belief false. One problem with this view is that it fails to handle simple variants of Gettier cases, such as SPECIAL DOG, because the Gettiered belief is not only justified and true, but also safe. To see why, recall that in SPECIAL DOG the cleverly disguised special dog wouldn't appear to Shep unless there were a nearby sheep, and also prevents Shep from encountering any other non-sheep sheep-lookalike that would mislead him into concluding that there is at least one sheep nearby. In effect, the special dog acts as a sort of epistemic guardian angel for Shep on such matters, which ensures that his beliefs about nearby sheep are not only true but also safely formed.

Another family of views proposes a *sensitivity condition* on

knowledge (Dretske 1970, 2005, Nozick 1981). A conservative sensitivity-based view might simply append a sensitivity condition to the JTB analysis, but sensitivity theorists typically dispense with justification altogether. Your belief that P is sensitive just in case the following conditional is true: if P were false, then you wouldn't believe that P. A sensitivity condition on knowledge handles standard Gettier cases. In LAMB if it were false that at least one of the students owned a Lamborghini, then Dr. Lamb would still believe that at least one student did (because Linus would still have deceived her). In SHEEP if it were false that there was at least one sheep in the field, then Shep would still believe that there was (because the cleverly disguised dog would still have tricked him). One problem facing this diagnosis is that it can't handle simple variants of the cases. For example, it can't handle SPECIAL DOG because the special cleverly disguised dog wouldn't have tricked Shep if there were no sheep nearby; and the special cleverly disguised dog would prevent anything else from tricking Shep; so if there weren't a nearby sheep, Shep wouldn't believe that there was one. Another serious problem facing this view is that it implies that knowledge isn't closed under some trivial, known deductive entailments. (We'll soon return to closure and how counterintuitive it can be to deny it, in the next paragraph and again in section 3 below. See also Vogel 1990 and Hawthorne 2005 for a defense of closure.)

Related to the sensitivity-based account of knowledge, Fred Dretske (2005) has also argued for a sensitivity-based account of reasons or justification. On this view, justification isn't closed under

deductive entailment, and even fails to transmit across some simple, known deductive entailments. Generally speaking, a justification to believe P is an indication that P is true. Indications carry information. Information comes from sources. Consider a thermometer, which is a source of information. The thermometer indicates the ambient temperature in the room. Its readout provides a reason for believing that it is twenty-one degrees in here. The readout carries information about the ambient temperature, in this case that it is twenty-one degrees. ‘That it is twenty-one degrees’ entails ‘that it is not eighteen degrees being misrepresented as twenty-one degrees’. But the latter claim is *not* part of the readout’s content — it doesn’t carry that information. And yet the readout’s content entails it. So your reason for believing that it’s twenty-one degrees needn’t also be a reason for you to believe the obvious deductive consequences of the claim that it’s twenty-one degrees. It might be easier to grasp how potentially counterintuitive this is by considering the matter more schematically. According to Dretske, the following is possible: reason R justifies you in believing P, and you *know* that the truth of P *guarantees* the truth of Q, but still, R does *not* justify you in believing Q. If Dretske is right about this, then one of the key assumptions of Gettier’s original discussion — namely, that justification is closed under deductive entailment — is thrown into doubt.³

³ The denial of closure has significant implications for epistemology, beyond the Gettier problem. It would also enable a direct and powerful response to many influential skeptical arguments. In fact, this is precisely how the idea of denying closure entered the contemporary discussion. See Dretske 1970, Nozick 1981, and also Pritchard 2008 for a helpful overview.

A more recent approach to the Gettier problem is to argue that knowledge can be defined as true belief for which the subject *earns credit* for believing the truth, but a Gettier subject doesn't earn credit for believing the truth, which explains why she doesn't know (e.g. Greco 2003, Zagzebski 2009). For example, in SHEEP Shep doesn't earn credit for believing the truth about whether there's a sheep in the field. Rather, we would credit a confluence of odd circumstances for the fact that Shep ends up believing correctly. It's an open question whether the operative notion of *credit* can ultimately sustain this treatment of the Gettier problem. A related view defines knowledge as follows: you know that P just in case you have a true belief that P *because* you believed competently; however, it is argued, although the Gettier subject has a true belief and believes competently, he doesn't have a true belief because he believes competently, which explains why he doesn't know (Sosa 2007). Another related view defines knowledge as follows: you know that P just in case the fact that you have a true belief that P is a manifestation of your cognitive powers; however, it is argued, although the Gettier subject has a true belief and exercises her cognitive powers, the fact that she has a true belief isn't a manifestation of her cognitive powers, which explains why she doesn't know (Turri 2011, forthcoming). The jury is still out on this family of approaches.

Some have argued that Gettier's intuition about his cases was wrong: Gettier cases *are* cases of knowledge. Stephen Hetherington (1998, 1999, 2011) argues that a Gettier subject knows despite coming perilously close to not knowing, and supplements this by dia-

gnosing intuitions to the contrary. Whereas safety theorists would claim that the unsafety of a Gettier subject's belief disqualifies it as a case of knowledge, Hetherington contends that its unsafety misleads us into thinking that the Gettier subject doesn't know. The Gettier subject's belief might very easily have been false, and we mistake this *near failure* for an actual failure. Gettier subjects straddle the divide between *just barely knowing* and *not knowing*. Although ingenious, Hetherington's view remains a minority position (Lycan 2006, Turri forthcoming).

Many philosophers have taken Gettier cases to show that justified true belief isn't sufficient for knowledge, even though it still is necessary. And as we've already seen, some have tried to replace justification with something else entirely, such as an appropriate causal relation, a safety condition, or a sensitivity condition. But at least one philosopher has argued that they're all wrong because knowledge is simpler than any of them had imagined: knowledge is mere true belief (Sartwell 1991, 1992). Crispin Sartwell's argument for this position is simple: knowledge is the goal of inquiry; the goal of inquiry is true belief; so knowledge is true belief. Inquiry just is the procedure of generating beliefs about particular propositions, and when we ask whether some claim is true, what we want is to *know* whether it's true. In other words, knowledge is the goal of inquiry. But most philosophers will object that we also want our true beliefs to be *justified*, well supported by evidence, so Sartwell has left out an important aspect of our goal. Sartwell accepts that we want justified beliefs, but argues that this is only because justifica-

tion is a good sign that we've got what we really want, namely, true belief. Justification is instrumentally good because it's a good sign that we do know, but isn't an essential part of knowledge.

4. The Scylla and Charybdis of post-Gettier epistemology: Or, teetering between fallibilism and skepticism

Nearly all epistemologists think that Sartwell is wrong, and that knowledge requires something more than true belief. Most epistemologists still think that justification is a necessary condition on knowledge, even if justified true belief isn't sufficient for knowledge. And most epistemologists still agree with Gettier that justification isn't factive (Sutton 2007 dissents). Having a justified belief that P doesn't guarantee that P is true: you could be justified in believing P even though P is false. Moreover, it is widely held that the minimum level of justification required for knowledge is also non-factive: having knowledge-grade justification for believing that P doesn't guarantee that P is true. To put it differently, the conventional wisdom in contemporary epistemology is that knowledge-grade justification is *fallible*: you could be wrong even though you have it.

But fallibilism has struck many as deeply problematic. What follows is one way of explaining why fallibilism can seem both attractive and deeply puzzling (BonJour 2001).

Suppose you have a true belief. In order for it to be knowledge, how much justification must be added to it? Think of justification for a belief as measured by *how probable* the belief is given the

reasons or evidence you have. We can measure probability any way we like, but one convenient way to measure it is to use the decimals in the interval $[0, 1]$ on the number line. A probability of 0 means that the claim is guaranteed to be false. A probability of 1 means that the claim is guaranteed to be true. A probability of .5 means that the claim is just as likely to be true as it is to be false. The question then becomes: how probable, relative to your evidence, must your belief be for it to be knowledge?

Obviously it must be greater than .5 — after all, if it were less than .5, then it would be more probable that your belief was false, given your evidence! But how much greater than .5? Suppose we say that knowledge requires a probability of 1 — that is, knowledge requires justification that *guarantees* the truth of the belief. Call this *infallible justification*. *The infallibilist conception of knowledge* says that knowledge requires infallible justification. We can motivate the infallibilist conception as follows. If the aim of belief is truth, then it makes sense that knowledge would require infallible justification, because it guarantees that belief's aim is achieved. Clearly it's a good thing to have such a guarantee.

But all is not well with the infallibilist conception. It seems to entail that we know nothing at all about the material world outside of our own minds, or about the (contingent) future, or about the (contingent) past. For it seems that we could have had the same justification that we do in fact have, even if the world around us (or the past, or the future) had been radically different. Our justification doesn't guarantee that a material world exists. (Think of Descartes's

evil genius.) Neither does it guarantee that there is a past or future. This dramatic skeptical consequence conflicts with commonsense and counts against the infallibilist conception of knowledge. This is presumably part of the motivation for the widespread agreement that justification isn't factive.

We seem compelled to conclude that knowledge requires justification that makes the belief very likely true, but needn't guarantee it. This is the *fallibilist conception of knowledge*. But a question about this view immediately arises: what level of justification does it require? Any point short of 1 would seem *arbitrary*. Why should we pick that point exactly? The same could be said for a vague range that includes points short of 1 — why, exactly, should the vague range extend that far but not further? This might not seem so troubling in itself, but as Laurence Bonjour (2001) points out, it suggests an even deeper problem for the weak conception. It brings into doubt the value of knowledge. Can knowledge really be valuable if it is arbitrarily defined? It would count heavily against the fallibilist conception of knowledge if it implied that knowledge wasn't valuable. (Kaplan 1985 raises related worries about knowledge's value in light of the Gettier problem.)

A related problem for the fallibilist conception of knowledge presents itself, which relates to the second of Gettier's assumptions. Suppose for the sake of argument that we settle on .9 as the required level of probability. Suppose further that you believe Q and you believe R, that Q and R are both true, and that you have reached the .9 threshold for each. Thus the fallibilist conception en-

tails that you know Q, and it entails that you know R. Intuitively, if you know Q and you also know R, then you know the conjunction (Q&R), just by simple deduction. But, surprisingly, the weak conception of knowledge can't sustain this judgment! To see why, consider that the probability of the conjunction of two independent claims, such as Q and R, equals the product of their probabilities. (This is the special conjunction rule from probability theory.) In this case, the probability of Q = .9 and the probability of R = .9. So the probability of the conjunction (Q&R) = $.9 \times .9 = .81$, which falls short of the required .9. So the weak conception of knowledge along with a law of probability entail that you *don't know* the conjunction (Q&R), because you aren't well enough justified in believing the conjunction. Can we tolerate this result?

So we are faced with a choice between two views, fallibilism and infallibilism, each of which has seemingly unpalatable consequences. If we accept fallibilism, then we seem poised to surrender the intuitive claim that (knowledge-grade) justification is closed under simple, known deductive entailments, and also the intuitive claim that knowledge is valuable. And if we accept infallibilism, then we seem poised to surrender the intuitive claim that we're in a position to know lots of things about the material world, the past and the future.

Notice how Gettier's two assumptions relate to these unpalatable consequences. In setting up his problem, Gettier assumed that (1) justification isn't factive and (2) justification is closed under deductive entailment. Infallibilism threatens to falsify something in

the ballpark of (1), whereas fallibilism threatens to falsify something in the ballpark of (2). Gettier's lasting legacy might well be to force us to choose between these two claims. Are we forced to choose between them, or can we find some way to have our epistemological cake and eat it too?⁴

⁴ For helpful feedback, I thank Stephen Hetherington and Angelo Turri.

References

- BonJour, Laurence. 2001. *Epistemology: Classical Problems and Contemporary Responses*. Lanham, Maryland: Rowman and Littlefield.
- Chisholm, Roderick M. 1989. *Theory of Knowledge*, 3rd ed. Englewood Cliffs, NJ: Prentice Hall.
- Clark, Michael. 1963. "Knowledge and Grounds: A Comment on Mr. Gettier's Paper." *Analysis* 24.2: 46–48.
- Conee, Earl and Richard Feldman. 1998. "The Generality Problem for Reliabilism." *Philosophical Studies* 89.1: 1–29.
- DePaul, Michael and Linda Zagzebski, eds. 2003. *Intellectual Virtue: Perspectives from Ethics and Epistemology*. Oxford: Oxford University Press.
- Dretske, Fred. 1970. "Epistemic Operators." *The Journal of Philosophy* 67.24: 1007–1023.
- Dretske, Fred. 2005. "The Case Against Closure." In *Contemporary Debates in Epistemology*. Ed. Matthias Steup and Ernest Sosa. Malden, Mass.: Blackwell.
- Feldman, Richard. 2003. *Epistemology*. Upper Saddle River, NJ: Prentice Hall.
- Gettier, Edmund. 1963. "Is Justified True Belief Knowledge?" *Analysis* 23.6: 121–3.
- Goldman, Alvin I. 1979. "What Is Justified Belief?" In *Justification and Knowledge*. Ed. George Pappas. Dordrecht: Reidel.
- Greco, John. 2003. "Knowledge as Credit for True Belief." In DePaul and Zagzebski, eds.
- Harman, Gilbert. 1973. *Thought*. Princeton: Princeton University Press.
- Hawthorne, John. 2005. "The Case for Closure." In *Contemporary Debates in Epistemology*. Ed. Matthias Steup and Ernest Sosa. Malden, Mass.: Blackwell.
- Hetherington, Stephen. 1998. "Actually Knowing." *The Philosophical Quarterly* 48.193: 453–469.

- Hetherington, Stephen. 1999. "Knowing Failably." *The Journal of Philosophy* 96.11: 565–587.
- Hetherington, Stephen, ed. 2006. *Epistemology Futures*. Oxford: Oxford University Press.
- Hetherington, Stephen. 2009. "The Gettier Problem." In *The Routledge Companion to Epistemology*. Ed. Sven Bernecker and Duncan Pritchard. New York: Routledge.
- Hetherington, Stephen. 2011. *How to Know: A Practicalist Conception of Knowledge*. Malden, Mass.: Wiley-Blackwell.
- Kaplan, Mark. 1985. "It's Not What You Know that Counts." *The Journal of Philosophy* 82.7: 350–63.
- Klein, Peter. 1976. "Knowledge, Causality, and Defeasibility." *The Journal of Philosophy* 73.20: 792–812.
- Lehrer, Keith. 1965. "Knowledge, Truth and Evidence." *Analysis* 25.5: 168–175.
- Lehrer, Keith and Thomas Paxson Jr. 1969. "Knowledge: Undefeated Justified True Belief." *Journal of Philosophy* 66: 225–237.
- Lycan, William G. 2006. "On the Gettier Problem problem." In *Epistemology Futures*. Ed. Stephen Hetherington. Oxford: Oxford University Press.
- Matilal, Bimal Krishna. 1986. *Perception: An Essay on Classical Indian Theories of Knowledge*. Oxford: Oxford University Press.
- Nozick, Robert. 1981. *Philosophical Explorations*. Cambridge, Mass.: Harvard University Press.
- Pritchard, Duncan. 2005. *Epistemic Luck*. Oxford: Oxford University Press.
- Pritchard, Duncan. 2008. "Sensitivity, Safety, and Anti-Luck Epistemology." In *The Oxford Handbook of Skepticism*. Ed. John Greco. Oxford: Oxford University Press.
- Sartwell, Crispin. 1991. "Knowledge Is Merely True Belief." *American Philosophical Quarterly* 28.2: 157–65.
- Sartwell, Crispin. 1992. "Why Knowledge Is Merely True Belief." *The Journal of Philosophy* 89.4: 167–180.

- Saunders, John and Narayan Champawat. 1964. *Analysis* 25.1: 8–9.
- Shope, Robert. 1983. *The Analysis of Knowing: A Decade of Research*. Princeton, NJ: Princeton University Press.
- Sosa, Ernest. 1999. “How to Defeat Opposition to Moore.” *Philosophical Perspectives* 13: 141–53.
- Sosa, Ernest. 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, v. 1. Oxford: Oxford University Press.
- Sutton, Jonathan. 2007. *Without Justification*. Cambridge, Mass.: MIT Press.
- Turri, John. 2011. “Manifest Failure: The Gettier Problem Solved.” *Philosophers’ Imprint* 11.8: 1–11.
- Turri, John. Forthcoming. “Is Knowledge Justified True Belief?” *Synthese*.
- Unger, Peter. 1968. “An Analysis of Factual Knowledge.” *The Journal of Philosophy* 65.6: 157–70.
- Vogel, Jonathan. 1990. “Are There Counterexamples to the Closure Principle?” In *Doubting*. Ed. M.D. Roth and G. Ross. Dordrecht: Kluwer.
- Warfield, Ted. A. 2005. “Knowledge from Falsehood.” *Philosophical Perspectives* 19: 405–416.
- Zagzebski, Linda. 1996. *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge: Cambridge University Press.
- Zagzebski, Linda. 2009. *On Epistemology*. Belmont, Calif.: Wadsworth.