1 The Gettier Problem and Fallibilism

Charity Anderson

1.1 Introduction

In his seminal article, Gettier (1963) established that justified true belief (JTB) is not sufficient for knowledge. He did so in the context of two conditions: first, that a belief can be justified and false; second, that justification is closed under deductive entailment. This chapter focuses on the first of these two commitments: that justification and truth can come apart. There is widespread support for the idea that it is possible for one to have justification for a proposition – even justification of knowledge-level strength – and yet believe falsely. Call this thesis JF. JF is commonly taken to be a presupposition of the Gettier problem. Some further insist that JF is a constraint on a satisfactory solution to the Gettier problem. Putative responses to the problem that do not respect this constraint, it is said, run the risk of *avoiding* rather than solving the problem.[1]

The relationship between justification and truth presented in JF has been thought to be a commitment of *fallibilism*. Thus, JF is characterized as a fallibilist constraint. In this chapter, I aim to shed light on the role of fallibilism within the Gettier problem and potential solutions. I investigate the kind of constraints that could reasonably be placed on an acceptable solution and, in particular, the extent to which a fallibilist constraint ought to be placed on a satisfactory solution. Finally, I reconsider the narrative concerning the relationship between JF and fallibilism about knowledge.

The chapter is divided into two sections. Section 1.2 considers the scope of the fallibilist constraint. As it appears in Gettier's article, the fallibilist constraint applies specifically to the justification condition on

I am grateful to John Hawthorne, Alexander Pruss, and Jeffrey Russell for helpful conversation on the issues in this paper, and to Andrew Moon for comments on a draft.

[1] See Zagzebski (1994) and Hetherington (2012; 2016a).

knowledge – where justification is thought of as one among three necessary and jointly sufficient conditions on knowledge (together with truth and belief). Gettier's counterexample revealed the need to add a fourth condition to the traditional analysis of knowledge.[2] It is an open question whether we ought to extend a fallibilist constraint to the elusive fourth condition. That is, it is an open question whether the fourth condition must be such that, when it obtains, the belief can be false. (I'll follow standard usage in thinking of the fourth condition as the "gap filler": that which, when added to JTB, results in a combination that suffices for knowledge.)[3] I examine two ways to think about the scope of the fallibilist constraint. The first extends fallibilism to the fourth condition. Together with extension of a closure principle, these constraints deliver the result that the Gettier problem is unsolvable. A second, more promising option limits the scope of the fallibilist constraint to the justification condition, without extending it to the fourth condition. This option moves us in the direction of a solution, but one which, ultimately, some may find unsatisfying – because it is thought to be "infallibilist" in some objectionable sense. In the course of discussion, I will examine a test case arising in recent literature, with the goal of clarifying the source of this recent dispute.

Section 1.3 examines attempts to make the option of a factive fourth condition acceptable for fallibilists. I argue that such a picture can fall within the family of fallibilist views of knowledge. While JF is rightly considered a "fallibilist" assumption within the Gettier problem, fallibilists about knowledge are not committed to adding only nonfactive conditions to the JTB account. This results in further options for fallibilist responses to the Gettier problem.[4] It is not the case that only an infallibilist picture of knowledge eliminates the Gettier problem.

[2] Either that or we must reconceive of the justification component; hereafter I will focus on the option of adding a fourth condition, although I acknowledge that reinventing justification is also an option that some have taken. I expect that reconceiving of justification is not among the options that my intended interlocutors would find appealing, as this option will more likely be taken to be "infallibilist."

[3] A relevant body of literature concerns whether "warrant" entails truth. See especially Merricks (1995), Ryan (1996), and Moon (2013). "Warrant" standardly names the "gap filler" condition(s), and, according to some, may be factorized into multiple conditions. My discussion here will investigate the components as factorized. One advantage of this approach is that it will be easier to see the properties of the various components and to determine which are fallible and which factive. Here I will be specifically interested in the status of the fourth condition, as distinct from a justification component.

[4] See Howard-Snyder et al. (2003) for an alternative fallibilist response.

1.2 Nonfactive Justification

As did many of his contemporaries, Gettier took justification to be a necessary component of knowledge, and was writing in a context where justification was thought to be sufficient, in combination with truth and belief, for knowledge.[5] Gettier (1963, p. 121) expected that his counterexample to the JTB account of knowledge could be applied without loss to accounts which substitute either "having adequate evidence for p" or "having the right to be sure that p is true" for "justification" in the analysis of knowledge. Both the idea that justification is a necessary component of knowledge and the idea that one can have justification for a proposition that is false are strongly plausible theses. (To a large extent, both are still widely accepted, though some have moved away from talking about *justification* altogether, and thus are silent on JF.)[6]

JF is widely accepted for good reason: it is strongly intuitive that we can be justified in believing a proposition that, unbeknownst to us, is false. We sometimes have strong evidence for propositions that are false. For example, suppose that a trustworthy friend tells me that she will be at the party tonight and I believe her. On the basis of her testimony, I am justified in believing that she will be at the party, even if she suddenly becomes ill and remains at home. Similarly, I might have justification to believe something that is highly probable, such as that my ticket will lose the lottery, or that my car has not been stolen in the past half hour, and yet my belief can turn out to be false. The experience of discovering that a belief which one holds for good reason is false is commonplace. Assuming a connection between justification and evidence or between justification and reasons, these kinds of examples underlie motivation for JF.

It is seldom suggested that the best resolution to the Gettier problem is to deny JF. One reason for this is that if JF is false, and we cannot have justification for false beliefs, the number of our beliefs that are justified is radically more limited than we typically suppose. If one also holds, as many do, the idea that there is sameness of justification across good and bad cases, the result is that our knowledge will likewise be more limited. To illustrate this, I'll rely on an evidentialist conception of justification (though other pictures

[5] Gettier cites Ayer (1956) and Chisholm (1957) as among those who advocate for the "traditional" JTB view. Some have thought that the JTB account of knowledge was the standard view of knowledge for hundreds of years prior to Gettier – citing support even within Plato's writings. See Dutant (2015) for an alternative to this traditionally accepted narrative.

[6] The notion of *warrant*, in particular, has rivaled *justification* as a replacement concept. Insofar as *justification* is given an internalist reading, it has fallen somewhat out of favor with externalists. But, with the exception of a few extreme externalist pictures (advocates of E = K, the view that one's evidence consists of all and only the propositions one knows, in particular), many still find the notion central to theorizing about knowledge.

could be substituted). The evidence we have for many of the propositions that we take ourselves to know is such that the evidence does not guarantee the truth of the propositions. Consider two agents with the same strong evidence for a proposition – one with an unlucky false belief and the other a true belief. (We need not invoke a BIV counterpart to get an example of this kind off the ground; a parked car case will do. Imagine that the relevant proposition is: *my car has not been stolen today.* Imagine also that the first agent's car was stolen five minutes ago, and that the second agent's car remains in the lot where it was parked, despite the relevant circumstances being otherwise similar.) Many find compelling the thought that, if the first agent is not justified in believing *p*, neither is the second. Pairs of cases of this sort are easily generated, and, together with the intuitive judgment about the sameness of justification across cases, they threaten the idea that we have an abundance of widespread ordinary knowledge. On pain of skepticism, we must affirm JF. In this way, JF is thought to be essential to an antiskeptical picture.[7]

A "gap" between justification and truth is also essential to the construction of Gettier cases using Linda Zagzebski's well-known "double-luck" recipe, where the introduction of bad luck is followed by a dose of good luck resulting in a justified true belief, but one which we intuitively judge is not knowledge. (Without JF we would be unable to follow the first step of the "recipe" – that is, imagining an agent with a justified false belief.) Zagzebski articulates the centrality of JF to the Gettier problem as follows: "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" (1994, p. 69). Zagzebski claims that the lesson of Gettier is that we must either add a fourth component to the JTB analysis or we must reconceive the notion of justification. As long as we do not add a fourth condition that guarantees truth (or reinvent justification such that it is factive), application of the double-luck recipe that Zagzebski articulates will succeed in producing Gettier-style counterexamples.[8]

[7] See Hoffmann (1975) for a defense of this point specifically in the context of the Gettier problem. Hoffmann draws on parity reasoning between the "good" and "bad" cases (with respect to sameness of evidence) to argue against Almeder's thesis that evidence entails truth. Hoffman contends that, if Almeder were correct, skepticism would ensue (Almeder 1974).

[8] See Zagzebski (1994, p. 72): "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." It is not clear whether Zagzebski intends to claim that JF ought to be thought of as extended to include any and all conditions added to true belief, or whether she merely is making the observation that, if we extend JF in this way, we will not be able to solve the problem.

Presented in this way, fallibilism appears to be the root cause of the Gettier problem. Holding fallibilism fixed, a solution to the Gettier problem requires that we add a component to the JTB analysis of knowledge. Whether this component must also be nonfactive is an open question. Section 1.3 will discuss the option of a factive fourth condition.

1.2.1 Constraints on a Solution

Gettier's article initiated a search – conceived of as a challenge – to find the elusive fourth condition. It is worth noting, though, that Gettier himself never explicitly states a "challenge." Instead, he makes an observation: justified true belief is insufficient for knowledge. This observation is now nearly unanimously accepted and points to the need for (and marked the start of endless attempts to provide) a further condition on the traditional analysis of knowledge. Of course, some – notably, Timothy Williamson – reject this project of analysis as unfruitful and misguided, but this kind of response is for the most part an outlier position: many have taken the "challenge" over the years.[9] In this section, I will examine what parameters we ought to place on an adequate solution.

Gettier did not himself set down constraints on what could serve as a satisfactory fourth component. Granted, one might suppose that, *in the spirit of* Gettier, we ought to think that his fallibilist constraint on justification – that a belief can be justified and false – should be extended as a constraint on the fourth condition (see Hetherington 2012). But this expectation has unwelcome results. Before we pursue this line of thought, let me address a related matter.

One issue that is somewhat neglected in the literature concerns whether Gettier's other assumption – a closure principle for justification – ought also to be extended so as to apply to the fourth condition. That is, little has been said concerning whether the fourth condition is closed under competent deduction.[10] I will make two points concerning this issue. First, note that the motivation for extending the fallibilist constraint applies equally to extension of a closure principle. At least, if the reason for extending fallibilism is to provide a solution that treats the fourth condition in the same way as Gettier treats the justification condition, this reasoning equally motivates extension of a closure principle.

[9] See Shope (1983) and Lycan (2006) for summaries of those attempts.
[10] Merricks (1995) is a noteworthy exception. See also Ryan (1996) and Howard-Snyder et al. (2003).

Second, as Merricks (1995) notes, there is some independent pressure to think that the fourth condition is closed.[11] Suppose that it is not closed: in settings where the fourth condition fails to obtain for the proposition deduced, knowledge will also fail to obtain. In this way, closure failure for the fourth condition will result in closure failure for knowledge.[12] (One might try a gerrymandered approach, where the fourth condition is closed when the initial belief is true, but fails to be closed when the initial belief is false. Or, alternatively, where the fourth condition is closed only when the initial belief is justified. But, as these strategies are *ad hoc*, they are unlikely to be appealing.) Thus, insofar as one finds closure for knowledge attractive, one ought to expect the fourth condition to be closed.[13]

We can now proceed to the central point. If the fourth condition is closed and the fourth condition is fallibilist – that is, the fourth condition can obtain when a belief is false – then the fourth condition will not be able to play the role of "gap filler" in an analysis of knowledge. Any proposed analysis of knowledge whereby knowledge consists of JTB plus a fourth condition, where the justification condition and the fourth condition each respect both constraints under consideration, will be open to Gettierization.[14] Application of the structure of one of Gettier's original counterexamples will make the point evident. In Gettier's second counterexample, an agent forms a justified true belief, *p or q*, that fails to be knowledge because the belief is the result of the agent's reasoning via disjunction introduction that began with a justified false belief. Supposing that the fourth condition is fallibilist and closed under entailment, this process can be replicated by using a belief that meets the fourth condition.[15] Imagine an agent who has a justified false belief that satisfies the fourth condition. We can then construct cases where the agent,

[11] The relevant closure principle will inevitably be more nuanced than the idea gestured at here. Neither knowledge nor justification is, strictly speaking, closed under deduction. Qualifications must be added: for example, that the deduction is competent, that the proposition is deduced from known (or justified) premises, that the premises are not forgotten throughout the process, and so on. See Hawthorne (2004, ch. 1) for helpful discussion. A comprehensive treatment of the nuances is beyond the scope of this paper, but the reader may supply the needed qualifications throughout.

[12] Some are happy to accept this result. Dretske (2005), for example, adopts this strategy: his sensitivity condition is not closed, but he welcomes the result that knowledge is not closed.

[13] Nevertheless, some have objected to this type of argument. See Ryan (1996) and Howard-Snyder et al. (2003) for relevant discussion.

[14] The result can be obtained with a weaker principle than deductive closure. Closure under disjunction introduction suffices to deliver the unwelcome result.

[15] Merricks (1995) makes this structural point concerning "warrant." It is worth considering how the argument will apply to a factorized analysis, as I do here, where we consider the fourth condition as distinct from justification.

relying on disjunction introduction, forms a justified true belief that meets the fourth condition, but where the belief fails to be known.

The lesson is that to insist that any "extra" conditions on knowledge – conditions in addition to justification, truth, and belief – respect both of the constraints that Gettier places on justification (a fallibilist constraint and closure) is to set up the problem in such a way that, in principle, it cannot be solved.[16]

(A similar process delivers the result that the fourth condition must entail justification. Suppose that it does not: then we may construct cases as follows. An agent has a justified false belief that p, and an unjustified true belief that q that satisfies the fourth condition. If both justification and the fourth condition are closed, then the agent may deduce p or q and the belief will be justified, true, and meet the fourth condition, but will be unknown. Hence, if the fourth condition is closed and does not entail justification, it will not play the role of "gap filler.")[17,18]

It is clear that we cannot extend both constraints to the fourth condition (unless, of course, we are willing to rest content with an unsolvable problem). Initial investigation suggests that we have reason to extend closure and to loosen our fallibilist grip – at least for those who find closure for knowledge attractive. For now, I merely note that, as constraints on an adequate solution, together these deliver the unfortunate result of ruling out "solving" the problem. In Section 1.3, I will attempt to make the option of an "infallibilist" fourth condition palatable for fallibilists about knowledge.

1.2.2 "Manifest Failure" and Infallibilism

It will be instructive, at this point, to explore a dispute arising in recent literature. The dispute illustrates the specific issue under discussion – namely,

[16] Hetherington (2012; 2016a) insists that a "solution" to the Gettier problem requires that all additional conditions are fallibilist.

[17] One might think that the most plausible conclusion to draw is that closure must be rejected. But notice that rejection of closure in full generality will not suffice to avoid the result. What is needed is not merely the claim that the fourth condition is not closed in every instance of competent deduction; rather, it is that there is no case of the kind described where the fourth condition transmits via competent disjunction introduction. This will be significantly more difficult to motivate than mere rejection of fully general closure, and is likely to look implausibly gerrymandered.

[18] It is likely to be objected that if the fourth condition entails both truth and justification, it also entails knowledge and thus leads to a failed project – where the aim of the project is to provide an analysis of knowledge in terms that are each necessary but none of which are individually sufficient for knowledge. Although this may strike some as an unsatisfying result, I suggest that such an "analysis" of knowledge would nevertheless provide us with an informative identity.

whether fallibilism ought to be extended to the fourth condition. I will first survey the details of the dispute and then draw on the above remarks to attempt to resolve the disagreement.

Turri (2011) suggests a solution to the Gettier problem. His proposal develops Sosa's general framework (2007a), whereby knowledge is a performance that exhibits an "AAA" structure. According to Sosa, knowledge that *p* involves the following: one's belief that *p* must be accurate (true), adroit (display intellectual competence), and apt (true *because* competent). Turri agrees that knowledge requires these conditions, but maintains that they are together insufficient for knowledge. To demonstrate that the AAA structure is insufficient, Turri presents a case where a subject's belief is true because competent but where the truth of the belief is nevertheless a result of a deviant causal chain. His counterexample is structurally similar to the case of the *epistemic guardian angel*.[19] In short, the case is one where a guardian angel admires a subject for being competent and therefore decides to rearrange the world, as needed, to ensure that the subject's beliefs are true. The relevant feature of this kind of counterexample is that the agent's belief is true *because of* the agent's competence, but we tend to judge that knowledge is absent. To avoid this difficulty, Turri's proposed solution expands on Sosa's account as follows: knowledge requires not merely *apt* belief (true *because* competent) but also *adept* belief – true belief that *manifests* a competence. Not all outcomes that occur because of a competence are outcomes that manifest that competence. To illustrate the difference, Turri draws on the idea of fragility. Suppose that two glasses are falling off a table. We can imagine different outcomes for each glass: the first glass shatters on the floor; the second glass is caught before it hits the floor. The glass that breaks when it hits the floor *manifests* fragility, but the glass that is caught before it hits the floor does not manifest fragility – even though it is caught *because* it is fragile. According to Turri, Gettier subjects do not manifest a competence when they believe the truth. If we add *adeptness* to Sosa's account of knowledge, Turri claims, the Gettier problem is solved.

It will be helpful to consider briefly how Turri's account handles fake-barn cases. Rather than commit to one solution, Turri sketches a few options, all of which are consistent with his account. Turri himself boldly reports that he lacks the requisite intuition regarding fake-barn cases, and so his preferred option is to deny that fake-barn cases are in fact Gettier cases; nevertheless, he provides several options to accommodate the intuition. In the familiar

[19] Greco (1999) presents this case as a difficulty for simple reliabilism.

fake-barn case, an agent – call him Henry – uses his eyesight to form a true belief that *there is a barn*. Unfortunately, his environment is such that he easily could have formed a false belief – due to the many nearby barn façades. To respect the received intuition – that Henry does not know that *there is a barn* – one option for Turri's account is to claim that although Henry's competence is manifested, it could easily have failed to be manifested, and this feature of Henry's epistemic position is relevant in a knowledge-depriving way. Henry would have formed a false belief had he looked at a nearby façade, and in that case his competence would not have been manifested: successful manifestation of a competence occurs only in cases of knowledge. Note that this strategy makes manifestation of a competence insufficient to play the role of the missing Gettier condition; on this picture, knowledge requires *safe* manifestation of a competence. (This is obviously not as elegant a solution as Turri's original proposal alone.)

Another option for explaining the intuition that Henry does not know is to maintain that the manifestation of a competence is environment-sensitive. One might simply deny that Henry has a competence to identify barns when he is in barn façade country. Clearly, if Henry does not possess the requisite competence, he cannot manifest the competence. (Of course, in these kinds of case the agent might think that he is manifesting a competence, even when he is not, but this is merely a result of the fact that the manifestation of a competence fails to be luminous: an agent is not always in a position to know when he manifests a competence.)[20]

Hetherington (2012) objects that Turri's proposal does not qualify as a solution, because it does not maintain the gap between justification and truth that Gettier assumes. Therefore, Hetherington claims, Turri's alleged solution fails to even interact or engage with the problem. Specifically, Hetherington maintains that Turri does not respect the fallibilist constraint: his objection is that Turri's "solution" is "covertly infallibilist" and so is irrelevant to the challenge of the Gettier problem. At the center of his concern is the idea that Gettier's challenge was posed within a fallibilist conception of knowledge, and so should be answered in that spirit.

Let's consider whether Turri in fact violates the methodological parameters of the problem. Turri's fourth condition does indeed eliminate any gap between manifestation of competence and truth: one cannot

[20] See Turri (2012a) for additional possible explanations of the absence of knowledge in this kind of Gettier-style case.

manifest a competence and have a false belief. But his account builds on Sosa's AAA analysis, which includes a justification-like component which is nonfactive – a belief can be adroit and false. Thus, Turri's solution technically respects JF.[21] Consider an account of knowledge as justified true belief plus a fourth condition where the fourth condition is factive and where justification (or whatever plays the justification role) is non-factive. An account of this type respects Gettier's *actual* constraint: insofar as the account maintains a "gap" between justification and truth, JF is maintained.[22]

We have seen that we are faced with a choice: either the fourth condition ensures truth or it does not. If it does not, and the fourth condition is closed under known entailment, the account will be Gettierizable. If it does ensure truth, the potential for a solution is at least a possibility. I suggest that, on these grounds, we are hard-pressed to insist that Turri's solution – or any solution that fails to extend fallibilism to the fourth condition – be ruled out for this reason alone.

As our discussion has noted, it is not obvious that Gettier himself envisioned the extension of fallibilism and closure to the fourth condition as a constraint on a solution. Moreover, it is implausible that Gettier intended his problem to be constrained in such a way that it is in principle unsolvable. And even if that were his intention, it would be foolish to play by those rules once we see the impossibility of succeeding. I suggest that we are better off thinking of Gettier's article as presenting a puzzle with various options for responding, even if no option is fully satisfying. The strategy that I have explored here – an "infallibilist" fourth condition – opens the door to various solutions to the Gettier problem. Perhaps not everyone will be content to call this kind of reply a "solution," but this option strikes me as preferable to an in-principle unsolvable problem. In the next section, I will argue that a factive fourth condition does not, in fact, lead to infallibilism about knowledge.

[21] Another option is that manifestation of competence is not an added fourth condition, but instead a way of "reinventing" justification. If this is how Turri intends it, then he rejects a key assumption on which Gettier relies. In the next section I'll argue that even though this kind of account rejects JF, it is still compatible with fallibilism about knowledge.

[22] One might worry that "manifestation of a competence" or any other fourth condition that entails truth is no more basic than knowing. If our goal is to analyze knowledge, one way in which this strategy might fail is by providing a component whose extension overlaps exactly with knowledge. We might question whether such an analysis is doing the work that we wanted it to do for us. See Kearns (2007) for a defense of the significance of this type of analysis.

1.3 Fallibilism and Factivity

Suppose that you are a committed fallibilist. What, minimally, is required to be included in the family of views called "fallibilism about knowledge"? Here I suggest that views according to which knowledge involves JTB plus a factive condition, but a factive condition which is not luminous, can plausibly be characterized as fallibilist.[23] The case that I present is stipulative: suppose that infallibilism is the view that knowledge requires a condition that is both factive and luminous.[24] (Following Williamson, I will understand "luminous" as follows: a condition C is *luminous* for S if, whenever C obtains, S is in a position to know that it obtains.)[25] Fallibilism, in turn, will be the theory that knowledge does not require such a condition. I'll call these views infallibilism* and fallibilism*, respectively. On this way of thinking, fallibilism* is compatible with a factive fourth condition, so long as the condition is not luminous.

In favor of this stipulative picture, I will show that that fallibilism* is not encumbered by the same ill-favored consequences in which infallibilism is taken to result: most importantly, fallibilism* does not result in skepticism. (Avoidance of skepticism is, again and again, appealed to as the central motivation for fallibilism.)[26] Fallibilism* does not place overly demanding standards on knowing. It is consistent with the fact that we often make mistakes – even in cases where we think that we know. We have already observed that fallibilism* offers us a path to a solution to the Gettier problem. In this section, we will see that it has other advantages.

Here is one further preliminary clarification. As I use the terms here, both fallibilism about knowledge and infallibilism about knowledge are theses about the nature of knowledge – that is, they are views about what knowledge

[23] One might be tempted to think that a view that affirms JF is sufficient for fallibilism about knowledge, but note that, by that reasoning, Turri's account – and any account that affirms JF but involves a factive fourth condition – will clearly fall into the fallibilist camp.

[24] There is not space here to argue for this view, though I have done so elsewhere (Anderson, manuscript). The salient point of my argument is that factivity without luminosity provides us with, at best, a "cheap" kind of infallibilism. Williamson's knowledge-first picture – which has been called "infallibilist" by some – falls into this camp. Many common tests for infallibilism – skepticism, Cartesian standards, and so on – deliver the result that views of this kind fall on the fallibilist side.

[25] Infallible knowledge might additionally require that the absence of the condition be luminous to the agent – that is, it may require that when C is absent, S is in a position to know that C is absent – but luminosity as stated above will suffice for our purposes here.

[26] See Cohen (1988) and Lewis (1996) for a sample. The idea that infallibilism leads to skepticism is so entrenched in the literature that it is almost a datum.

requires. They are not statements about the extent of our knowledge. Thus, the way I am using the terms, it is open to one to be a fallibilist about knowledge and yet maintain that we have some infallible knowledge. This picture is one whereby knowledge does not *require* demanding standards, but leaves open whether we meet those standards for some items of knowledge (such as knowledge that 2 + 2 = 4). By contrast, it is not an option to be an infallibilist about knowledge but maintain that some of our knowledge does not meet those demanding standards.

Recently, fallibilism about knowledge has enjoyed near-unanimous acceptance in contemporary epistemology. But many representative statements of fallibilism are subject to a problem: on the assumption that fallibilism and infallibilism about knowledge are exhaustive options, many statements of fallibilism in the literature offer a corresponding picture of infallibilism that is inadequate. One could have a fantastic epistemic position – entailing evidence, for example – but have no idea that one has this evidence. Thus, there is a clear sense in which one fails to have the best epistemic position that one could have: on the assumption that one sometimes has entailing evidence and sometimes does not, but one cannot tell when one has it and when one fails to have it, one will be subject to all sorts of error, and unable to avoid the uncertainty and inclinations to hedge that often accompany agents who make such errors.[27] While surely it would be advantageous, in one sense, to have entailing evidence, to possess it without being aware that one possesses it will not result in the kind of secure epistemic position that we often associate with infallibilism.

A further and more troubling problem results from the fact that, given some of the prominent glosses of fallibilism, there are "cheap" ways to possess infallible knowledge.[28] For example, E = K – the view that one's evidence consists of all and only the propositions that one knows – results in an agent's trivially possessing entailing evidence for each and every proposition that the agent knows. Although, by some lights, this picture will count as infallibilist, however trivial, since neither knowledge nor evidence is luminous, it will not suffice for infallibilism* (Williamson 2000).

It will be helpful to consider in more detail some of the common motivations for fallibilism. Among the most prominent reasons for

[27] As a matter of fact, in my view, a full picture of infallibilism will involve more than this, but this minimalist picture will suffice for what I want to argue here.

[28] See Anderson (manuscript) for a more comprehensive treatment of this issue and an argument that Cartesian infallibilism is more difficult to attain than common statements of fallibilism suggest.

adopting fallibilism is the widely accepted idea that infallibilism leads to skepticism. Infallibilism places overly demanding conditions on knowledge – conditions that we meet for too few, if any, propositions. Since most contemporary epistemologists are antiskeptics, given this narrative, fallibilism is the only available option.

Another thought that inspires fallibilism is the fact that we make mistakes. Our faculties are reliable enough to deliver knowledge in a wide variety of situations, but they are not perfectly reliable: sometimes we mess up even when our faculties are working properly and we are in a suitable environment for their use. Yet we have knowledge. Fallibilists about knowledge often aim to describe a theory of knowledge that captures, or is in response to, these central observations.

With this in mind, we can see that a factive fourth condition does not lead to an infallibilism about knowledge that fits with these motivations. A factive fourth condition on knowledge does not lead to skepticism. If satisfaction of the fourth condition needed to be internally accessible, then skepticism might threaten. But the account under investigation includes no such stipulation, and in fact prohibits it. A factive fourth condition is also compatible with the observation concerning our making mistakes. Our faculties can be less than perfectly reliable even if there is no gap between the fourth condition and truth. As long as the fourth condition is externalist, so that we are not always in a position to tell when it holds, the fourth condition can guarantee truth without a commitment to infallibilism about knowledge.

A further advantage of the picture of fallibilism* is that it is compatible with a factive fourth condition and thereby results in a fallibilist picture of knowledge that can solve the Gettier problem. A final advantage is that if one's analysis of knowledge is such that knowledge requires a fourth condition that is factive and fails to be luminous, but which is distinct from the justification condition (which is not factive), then fallibilism* is compatible with JF. Thus, fallibilists* need not give up the intuitive thought mentioned at the start of the chapter – that justification is the same across good and bad cases.

These points will be most easily demonstrated by drawing on Turri's account as an example. It is important to recall that, on his view, manifesting a competence is not something accessible to the agent. One cannot tell by mere reflection alone (and one is not always in a position to know) whether one manifests a competence when one believes p or whether one merely appears to manifest a competence. The presence and absence of manifestation is not luminous to the agent. So, even though Turri's fourth condition is

factive, on my view his account falls within the family of fallibilist views since it does not place a condition on knowledge that is both luminous and factive.[29]

It is instructive to consider how a view with a factive fourth condition might diagnose two different types of Gettier-style cases. Consider first a standard Gettier case involving inference from false belief. Suppose that Smith justifiably believes *JONES* (Jones will get the job and Jones has ten coins in his pocket), and infers *POCKET* (the man who will get the job has ten coins in his pocket). If the fourth condition is closed and factive, then the explanation for why Smith does not know *POCKET* is that the first proposition, *JONES*, which is false, is such that the fourth condition does not obtain with respect to it. Of course, Smith won't be able to tell that the fourth condition fails to obtain with respect to his belief that *JONES*. But since, on our current proposal, the fourth condition is factive, it will not obtain when the proposition is false.[30]

Consider next a second type of Gettier-style case, which depends not on inference from a false belief, but rather on an unfriendly environment. Henry is in fake-barn country looking at a barn. His belief BARN, *that's a barn*, is true, but easily could have been false. On the assumption that we aim to respect the standard intuition – that Henry's belief that BARN is unknown – an explanation is needed as to why he fails to know.

At this point, it will be helpful to bear in mind that I have not argued that a factive fourth condition is sufficient in itself to solve the Gettier problem. I have argued here that a factive fourth condition is necessary to solve the problem, but it remains to be seen whether the fourth condition must involve more than factivity. (Perhaps, as Turri mentions, there are additional constraints, such as safety or an appropriate environment.) My aim here has been to show that views with a factive fourth condition ought to be evaluated in terms of how they fare with respect to solving the problem, rather than being ruled out in advance for being infallibilist.

Return again to Hetherington's critique of Turri. We are now in a better position to assess whether Turri's solution violates the spirit of Gettier's problem by adopting infallibilism. The motivations for infallibilism discussed here suggest that it does not. Suppose that knowledge requires true belief that

[29] This point applies more broadly to any view that shares this feature with Turri's account – including those that have been called "warrant infallibilism." Thanks to Andrew Moon here.

[30] One might worry that this type of account makes the truth condition superfluous. I expect we can be persuaded to take on this consequence once we appreciate that there is no way to avoid the result.

manifests a competence. This does not lead to skepticism, and it is consistent with the fact that we sometimes make mistakes even when we think we know. Turri's solution does not place overly demanding requirements on knowledge. Arguably, his proposal falls into the family of views that I have argued here belong to fallibilism* about knowledge.

The dissatisfaction with, and worries about, infallibilism are not applicable to Turri's solution.[31] This lesson applies more generally: a nonluminous "infallibilist" fourth condition does not land us in skepticism, is compatible with the general imperfection of our faculties, and does not place overly demanding standards on knowledge. Recall that one motivation for JF was that we wanted knowledge-level justification to be the kind of thing that we have often and for many propositions. The type of manifestation of a competence that Turri suggests is required for knowledge is an ordinary, everyday sort of manifestation. It is ordinary in that it does not require a kind of ability beyond what is possessed by an average person, and it is everyday in the sense that it affirms the idea that most of us have a vast amount of knowledge that we gain on a daily basis in normal circumstances. It does not require Cartesian certainty, or a modal success-rate above our usual abilities. Thus, I suggest that, although it includes a factive fourth condition, Turri's solution does not count as infallibilist. While conceding that Turri is not fallibilist by the lights of some statements of fallibilism in the literature, his solution is nevertheless fallibilist in a substantive way. In light of this, I think that we would be misguided to reject his solution – and relevantly similar solutions – simply because they do not align with certain common notions of fallibilism.

Might a solution of this type solve the Gettier problem? Undoubtedly, some will decry such a solution as not playing by the rules. But, as I have tried to argue here, while this kind of strategy does not provide us with an analysis of knowledge according to which all conditions (excluding truth) are nonfactive, it clears the way for a solution – and it does so without landing us in skepticism, the long-feared result of infallibilism. Without adoption of fallibilism*, the possibilities for a fallibilist theory of knowledge to avoid the Gettier problem look grim.

Reluctance to embrace this type of strategy in response to the Gettier problem seems to rest on a misconception concerning the scope of fallibilism.

[31] Of course, we might find Turri's account unsatisfactory for other reasons. Greco (2012), for example, notes that he finds the account unhelpful because it fails to provide a principled explanation for certain contrast cases – a failure that arises precisely because agents are not in a position to know when they manifest a competence and when they fail to do so.

Perhaps we are too committed to the unsolvability of the problem to be open to this route. The Gettier problem has been with us for many years: its presence pervades introductory textbooks, and any account of the recent history of epistemology will inevitably include this formidable puzzle. This alone provides reason to resist what might be thought of as too "easy" a response. But a solution of the type that Turri offers (and it is not the only fallibilist option involving a factive fourth condition) shows us that, as we move away from a fallibilist picture of knowledge to a fallibilist* picture, the Gettier problem's threat weakens its grip; hence, even fallibilists may find solutions.

If what I have argued is correct, Gettier's article is no less significant: his observation – that justified true belief is insufficient for knowledge – was correct. Perhaps, however, we are in danger of overlooking the full significance of his article. Even after more than fifty years, we have yet to fully exhaust the insights to be drawn from his observation.