Text recovered from draft of November 26, 2025
June 18, 2026

Chapter 16

Empiricism, Epistemology and the Gettier Problem

1. Introduction

The preceding chapter investigated the version of fallibilism that arises in small-e empiricism. It applies specifically to relations of inductive support. They cannot provide an absolute assurance of the truth of contingent propositions in science. This fallibility, it was argued, does not compromise the authority of inductively well supported science and even has a visible role in the dynamics of scientific progress.

Fallibilism also arises in another literature, in epistemology in general philosophy. Its formulation and goals differ from those of small-e empiricism. Where fallibilism in small-e empiricism primarily concerns abstract logical relations of inductive support, fallibilism in epistemology primarily concerns knowledge and beliefs, that is, mental states, and allows for many sources of error beyond the fallibility of inductive inference specifically. Fallibilism in some form is widely accepted as an abstract principle in the epistemology literature, but, we shall see, its consequences are resisted.

The difference is clearest in their approaches to defining knowledge. Inductive fallibilism in small-e empiricism led, as a secondary result, to a deflation of the concept of scientific knowledge. A result belongs within scientific knowledge, if it is inductively well supported by empirical evidence. It does not need to be true.

In contrast, a controlling problem in epistemology is to determine how to preserve the requirement that beliefs must be of truths, if they are to count as knowledge. Trouble arises when, in addition to stipulating that knowledge requires truth, it is expected that satisfaction of the other parts of the definition will assure that the belief is of a truth. For then the problem becomes insoluble. Even the strongest inductive justification of a contingent proposition does not assure its truth; and we have no other means to assure the truth of a contingent proposition. The clearest manifestation that this literature seeks to solve the insoluble is the persistence of

inconsequential debates over Gettier cases. The assurance that a belief is of a truth is secured by the unrealizable fiction of an omniscient narrator who just declares the belief in of a truth.

In this chapter, Section 2 reviews some of the versions of fallibilism in the epistemology literature. Truth of the proposition believed is still almost universally demanded of a definition of knowledge, even though the epistemological literature on knowledge has provided no means through which we can be assured that truth has been attained. Section 3 reports some minority attempts in that literature to deflate the concept of knowledge in a way similar to that of small-e empiricism. Section 4 reviews the "KK Principle," that to know something requires that we know that we know. Its failure follows from the opacity of truth for contingent propositions. Section 5 considers two ways that inductive fallibilism can be misconstrued. The pessimistic response is inductive skepticism and the idea that we cannot eliminate rationally fantasy scenarios of deceiving demons. It is mostly resisted in the epistemology literature. The optimistic response is that fallibility can be circumvented if only we try harder. It has led to sustained but futile attempts to evade fallibility.

Sections 6 and 7 report the most prominent failure in the epistemology literature to accept the consequences of fallibilism and opacity of truth. It is the immense literature on Gettier cases. In them, the justification for a true belief succeeds not because the justification operated as intended, but success came by luck or happenstance. They are taken to threaten the adequacy of the definition of knowledge as justified true belief. For they reveal what fallibilists should know, even the strongest of inductive justifications do not assure us of truth. Repeated efforts to protect the definition seek a way to distinguish cases in which strong justification does properly indicate truth and those in which it does not. The inductive opacity of truth entails that all these efforts fail and must fail. That is already evident in the Gettier case narratives, for the truth of a belief is recovered not from further inductive investigations, but from the declaration of a disembodied, omniscient narrator.

Section 9 argues that the enduring problems of this Gettier case literature are resolved if we do what inductive fallibility compels. We reduce the definition of knowledge to well-justified belief. Then Gettier cases cease to be troublesome and are the less worrisome cases. The fallibility of inductive inference is most troublesome when there is strong inductive support for a falsehood. The better case arises when there is some failing in the relations of inductive support,

but the result nonetheless turns out to be true. They are the Gettier cases. Section 10 offers a brief summary conclusion.

2. Fallibilisms in the Epistemology Literature

The doctrine of fallibilism seems to be widely accepted by epistemologists in an enormous literature. Or so we are frequently told. The opening line of the abstract to Cohen (1988) is frequently cited and quoted with approval. It asserts: "The acceptance of fallibilism in epistemology is virtually universal." Brown (2018, p. ix) has found, however, that fallibilism is fading with a rise of infallibilist thinking. That there is or, perhaps, was such widespread agreement in as fractured a community as that of the epistemologists is possible because there are so many versions of fallibilism. Here are two forms of it.

2.1 Human Frailty

The most anodyne version of fallibilism is a simple concession that we humans are fallible and that this fallibility must affect our claims to knowledge. Leite (2010, p. 370) expresses it as:

> Fallibilism. In the broadest sense of the term, an anti-dogmatic intellectual stance or attitude: an openness to the possibility that one has made an error and an accompanying willingness to give a fair hearing to arguments that one's belief is incorrect (no matter what that belief happens to be about).

and

> In some recent discussions, the term "fallibilism" stands for the thesis that human beings are fallible about everything (or just about everything) they believe …

Hetherington (n.d., Section 5) leaves little doubt as to the expansive scope of this sense of fallibilism. In a list whose incompleteness he recognizes, he provides seven possible sources of error in belief formation. They are "misusing evidence," "unreliable senses," "unreliable memory," "reasoning fallaciously," "intelligence limitations," "representational limitations" and "situational limitations."

2.2 Incomplete Evidential Support

This attribution of the origin of fallibility to human frailty is remote from the narrower sense of inductive fallibility developed in this small-e empiricism. This latter sense is closer to a more precise version of fallibility commonly found in epistemological writing. Here are a few

versions of it drawn randomly from the literature over many decades. Although the fallibility arises through a frailty in the means of justification, the core focus is still knowledge. Does this frailty still allow us say that we know something?

> A better way to understand fallibilism is as the claim that:
> (F1) It is possible for *S* to know that *p* even if *S* does not have logically conclusive evidence to justify believing that *P*. (Feldman, 1981, p. 266)

> … a fallibilist theory allows that S can know q on the basis of r where r only makes q probable." (Cohen, 1988, p. 91)
> … the weaker principle which permits S to know q on the basis of r provided r makes q sufficiently probable. (Cohen, 1988, p. 92)

> More commonly, however, "fallibilism" is used as a name for a thesis about knowledge and justification: that we can have fallible justifications for our beliefs, and that it is possible to know that something is the case even if one has only a fallible justification for believing it. (Leite, 2010, p. 370)

> [Most contemporary epistemologists] claim that one can know even if one lacks evidence which guarantees the truth of what's known. They embrace 'fallibilism': one can know that p even though one's evidence does not guarantee the truth of p. (Brown, 2018, p.2)

These formulations of fallibilism typically omit mention of whether fallibilism allows us to attribute knowledge to belief in evidentially well-supported falsehoods. In the broader literature, truth is taken to be an uncontroversially necessary condition for knowledge.[1] My sense is that most fallibilists agree. Reed (2012, p. 585, his emphasis) makes this condition explicit in his introductory formulation of fallibilism:

[1] Ichikawa and Steup (2017, Section 1.1) "Most epistemologists have found it overwhelmingly plausible that what is false cannot be known."

> Roughly stated, the basic idea is that the subject can know something even though it *could have been false*. This is not the same as saying that the subject can know something that is false–it is very widely accepted by philosophers that, if a belief counts as knowledge, it is true.

2.3 Comparison

We can now summarize how inductive fallibilism within small-e empiricism differs from the formulations of fallibilism in the epistemology literature.

Inductive fallibilism does not arise from any human frailty. It inheres in the non-demonstrative nature of inductive logic itself.

The general literature in epistemology formulates fallibilism in human terms of beliefs and, especially, as a problem for the special human belief state of knowledge. The narrower formulation of inductive fallibilism in small-e empiricism asserts fallibility as a property of the logical relation of inductive support. A connection with knowledge only arises derivatively if we proceed to consider the import of inductive fallibility for beliefs.

The general literature on fallibilism is concerned to preserve the necessity of truth for knowledge, in light of the challenge of fallibilism. The import of inductive fallibility for small-e empiricism is that this requirement of the necessity of truth should be dropped. It is an idealized requirement whose satisfaction will always remain opaque to us.

Considerable attention is given in the fallibilism literature to the question of whether necessary truths, such as those of mathematics, can be known, when fallibilism is adopted. Inductive fallibilism is limited to considerations of the inductive support of contingent propositions.

3. Knowledge Demoted

The mainstream view in epistemology seems to be that knowledge should retain its centrality. However, there have been repeated suggestions, to me welcome, that this centrality is a mistake. I do not have a thorough survey of these proposals. Here are a few of them, collected during my reading.

Jeffrey (1968), in his Bayesian analysis of belief, "Probable Knowledge," writes: (p. 166)

> The obvious move is to deny that the notion of knowledge has the importance generally attributed to it, and to try to make the concept of belief do the work that

> philosophers have generally assigned the grander concept. I shall argue that this is the right move.

Kaplan (1985) captures the essential idea that the transition from strong evidential support to knowledge adds nothing accessible to our investigations. He wrote: (p. 355, his emphasis)

> Suppose again that you have carried out inquiry, come to believe that *P* on good evidence, and that you now ask, with this new analysis of knowledge in mind, "But do I *know* that *P*?" Once again, there is nothing to find out. Insofar as you are satisfied that your belief in *P* is well founded, you will ipso facto be satisfied that you have not inferred *P* from a false premise–otherwise you would not think you had good reason for concluding that *P*.

Perhaps the most colorful expression is Papineau's (2021) "The Disvalue of Knowledge."[2] Mimicking Russell's celebrated riposte on causation, he writes in his abstract: (p. 5311)

> I argue that the concept of knowledge is a relic of a bygone age, erroneously supposed to do no harm. I illustrate this claim by showing how a concern with knowledge distorts the use of statistical evidence in criminal courts, and then generalize the point to show that this concern hampers our enterprises across the board and not only in legal contexts.

4. The KK Principle

That inductive fallibility precludes us from affirming absolutely that we know something has some affinity with what is called the "KK principle" in epistemology. The principle asserts that if we know something, it must also be the case that we know that we know it. The principle was introduced in a formal analysis of a knowledge operator "*K*" in Hintikka (1962, p. 104). It asks if *Kp*, knowing proposition *p*, is equivalent to *KKp*, knowing that one knows *p*. The principle has been the subject of extensive debate. For a short survey, see Hemp (n.d.).

The assessment of the principle for present purposes is quite simple. If a belief must be of a truth to count as knowledge, then the KK principle fails for the reasons already given. That our belief in a contingent proposition is of a truth is not something that we can ascertain, beyond an assurance of very-strong inductive support. To require that we can know that we know in this

[2] I thank Joachim Horvath for directing me to this paper.

case just repeats the same problem. To know that we know a contingent proposition requires again that we can ascertain the truth of the proposition, which is precisely what inductive fallibilism in small-e empiricism precludes.

If we adopt the weaker alternative that knowledge just requires strong inductive support from experience for the contingent proposition believed, then the KK principle raises no special questions of foundational importance. That a believer knows that they know some proposition requires the believer to determine that the proposition is well supported inductively. That determination requires some facility with inductive logic that many do not have. According to it, many people may believe the earth orbits the sun, but they do not know it since they are unable to reconstruct the evidential relations that support the belief. An astronomer however can know the result since an astronomer has the expertise to explore these evidential relations. The distinction between the two cases is clear without any mention of "knowledge."

To designate only the astronomer's belief as knowledge, as the KK principle would require, is an arbitrary stipulation. All that is at issue is the assigning of an honorific term "knowledge." The important foundational concept is that of strong inductive support from experience. It determines how well we are informed of contingent propositions. A believer's awareness of the details of the evidential case is immaterial. What matters is that the evidential case has been made.

5. Misconstruals of Inductive Fallibilism

The most significant and enduring failure of the epistemology literature has been to insist on a definition of knowledge that requires the truth of its subject and to make that definition central to epistemology. There are two further misconstruals of inductive fallibilism. One is to adopt a far-reaching pessimism that arises most often in popular work: if we cannot ascertain truth inerrantly, then we have no basis to affirm anything. The other is a far-reaching optimism: that inductive fallibilism compromises our capacity to affirm truth inerrantly is merely a temporary obstacle that more careful analysis will surmount. Both of these misconstruals, I am pleased to report, have been identified and disputed by fallibilists in the epistemology literature.

Both misconstruals appear elsewhere. In a later chapter, I argue that they underlie both sides of the realism debate in philosophy of science.

5.1 Fallibilist Pessimism

Fallibilists in epistemology routinely seek to refute the idea that fallibilism leads to skepticism about knowledge. Cohen (1988) bases the refutation on the notion of "relevant alternatives." We judge that we know something by considering only relevant alternatives. Presumably the possibility of an artful deceiving demon with extraordinary powers is not among these alternatives. The difficulty Cohen then faces is to articulate a sufficiently precise sense of "relevant." This proposal is one among many. See Carrier (1993) for another account and Hetherington (n.d., Sections 8-10) for a survey of various forms of skepticism that may arise within fallibilism.[3]

My view is that a simple dictate of rationality precludes us taking deceiving demon scenarios seriously. We should follow what the evidence supports. Deceiving demon scenarios are, by construction, ones in which the evidence strongly supports the prosaic conclusion that world is just as it appears. They provide us no evidence at all for a fantastic, deceiving demon.

5.2 Infallibilist Optimism

To some and perhaps many, that we can never be assured of the truth of a contingent proposition is intolerable. They harbor the view that the fallibilism that supports it must be defective and may even result from some derangement in the thinking of fallibilists. That at least is the import of the metaphor used by David Lewis (1996, p. 550) in his defense of infallibilism if we are to take his light-hearted remarks seriously:

> We are caught between the rock of fallibilism and the whirlpool of scepticism. Both are mad! Yet fallibilism is the less intrusive madness.

After allowing that "people can get used to the most crazy philosophical sayings imaginable," he proceeds to the plea (his emphasis):

> If you are a contented fallibilist, I implore you to be honest, be naïve, hear it afresh. 'He knows, yet he has not eliminated all possible sources of error.' Even if you have numbed your ears, doesn't this overt, explicit fallibilism *still* sound wrong?

[3] Hetherington's (n.d., Section 1) summarizes: "Almost all contemporary epistemologists will say that they are fallibilists. Yet the vast majority of them also wish not to be skeptics."

Over a century after the revolutions of relativity and quantum theory shook our faith in absolute security of our best-established science, Lewis' visceral aversion to the ineliminability of error sounds quaint and nineteenth century.

Lewis is not alone in these sensibilities. Views such as his have promoted a spirited debate over fallibilism and infallibilism in epistemology. Brown (2018) is a monograph length repudiation of the rise of infallibilism and a defense of fallibilism.

Proposals for infallible knowledge of contingent propositions cannot work for all the reasons that support inductive fallibilism in small-e empiricism. I write with some weariness of the inevitable failure of attempts to formulate a cogent infallibilism. They cannot escape the fallibility of the only way we can learn of contingent propositions, inductive support. Of course, it is always possible to contrive convoluted narratives that mask the inevitable failure. Since the proposal must depend on unsupported contingent assumptions or covert, fallible inductive inferences, they will be found and the failure revealed. It is inevitable and requires only that we take the trouble to look for them. I think of the belief that this inevitability can be escaped if only we put more effort into it, the "try-harder fallacy."

Lewis' (1996) proposal, "Elusive Knowledge," is as good an illustration of this mode of failure as any. Lewis makes it easy in so far as he explicitly repudiates the essential role of justification in knowledge. It is, he argues (p. 551), neither necessary nor sufficient. He concludes:

> The link between knowledge and justification must be broken.

If ever we suspected that Lewis was no empiricist, this pronouncement leaves no doubt. To an empiricist, it has an air of magical thinking. We can somehow, it seems to say, know deep truths of the world without experiences of them. To support his infallibilism, Lewis (p. 551, his emphasis) offers the following definition:

> Subject *S knows* proposition *P* iff *P* holds in every possibility left uneliminated by *S*'s evidence; equivalently, iff *S*'s evidence eliminates every possibility in which not-*P*.

Prima facie, the definition sets an unachievable standard for knowledge, since the assured elimination of every possibility other than *P* is breathtaking in scope. Inductive inference cannot achieve it. What saves the definition, temporarily, is a further qualification that some possibilities are to be ignored. Lewis writes (p. 554, his emphasis):

> Our definition of knowledge requires a *sotto voce* proviso. *S knows* that *P* iff *S*'s evidence eliminates every possibility in which not-*P*–Psst!–except for those possibilities that we are properly ignoring.

This addition allows Lewis to escape the artifices of Cartesian skepticism. That a deceiving demon has made his cat invisible, he asserts (p. 562), can be properly ignored.

Which possibilities can be ignored? Lewis lays out an elaborate account in eight rules. The first three specify possibilities that cannot be ignored: "rule of actuality," "rule of belief" and "rule of resemblance."[4] They are followed by five rules about what may be ignored: "rule of reliability," two "rules of method," "rule of conservatism" and "rule of attention."

This totality of eight rules presents an inviting challenge to respondents. There is considerable scope for criticism of each rule individually and, no doubt, still further scope for concerns about how all eight are to be combined consistently. It is tempting to engage in this criticism. The weaknesses of the proposals are low hanging fruits, ripe for the picking. The temptation should be resisted. There is sufficient ambiguity throughout the proposal for it to sustain endless, inconclusive debate. This is how dubious work can generate an apparently energetic literature that can never yield anything definite.

Rather, the failure of the analysis can be seen without engaging in any of the details of these rules, which is why I have not given their formulations. Lewis' project is to provide an account of infallible knowledge. Lewis notes (p. 554) that the truth is not included in his definition of knowing. However, it is secured by the first rule of actuality: "It follows [from the rule] that only what is true is known." That is, following Lewis' rules leads to truths and it does so *infallibly*.

How can the rules achieve this?[5] One possibility is that they somehow tap into a resource that enables us to divine truths of the world without a basis in experience. That Lewis has found

---

[4] This last rule is Lewis' response to Gettier problems. We do not know the Gettier proposition since, Lewis argues by considering several Gettier cases, we cannot ignore the possibilities in which the Gettier proposition is false. (pp. 557-58)

[5] I ignore here the escape that Lewis's system operates vacuous in delivering no truths or delivering only very few in highly contrived truths. It is just another way in which the system would fail.

such a resource is not evident in his analysis. In Chapter 12, in arguing for small-e empiricism, I made the case that there is no such resource.

From the examples in Lewis' text, it is evident that he expects us to secure knowledge through experience and thus, inevitably, his rules take us to truths beyond experience. To recall the examples in Lewis' first paragraph, how else are to know what food penguins eat and that the Australian rules football club, Essendon, won the 1993 Grand Final. That is, Lewis' rules are a contribution to inductive logic and a remarkably convoluted one. Since any inductive logic is non-demonstrative, it is necessarily fallible. The artifice of precluding some possibilities as ignorable does not save it. For inferring that they are ignorable is itself an inductive inference and thus fallible.

I will forgo the exercise of identifying precisely where and how the rules fail to be infallible guides to truth. For identifying any specific failure invites responses that trigger the interminable and inconsequential debates that convoluted proposals like Lewis' engender. The details do not matter. The failure is inevitable.

It seems to me that Lewis' analysis has already shown clear signs of decay. In my analysis of the material theory of induction (2021), I found a common mode of failure afflicting rule-based accounts of inductive inference.[6] Each starts with a simple rule. Counterexamples threaten it. They are parried by adding further conditions to the original, simple rule. Then counterexamples arise for these further conditions, so that still further conditions are needed. The cycles continues without a definitive termination. The rule-based account is in a death spiral. That Lewis' account requires *eight* rules suggests that it is already well-advanced in the death spiral.

[6] See for example the analysis of analogical inference in (Norton, 2021, Chapter 4).

6. The Gettier Problem

6.1 Its Persistence

It would be comforting to imagine that the prevalence of fallibilism in epistemology has led to a relaxed attitude to the requirement of truth for knowledge.[7] We all surely agree that truth is our goal. Might we be prepared to adopt a conception of knowledge that allows us to say that we know something without stipulating that its truth must be included in the definition of knowledge? Might epistemologists shift their emphasis to the strength of a justification in a belief and accept that truth will follow, mostly but not invariably?

The longevity of writing on the Gettier problem is a strong indication that this comfort is, alas, just my wishful thinking. In that literature, as we shall see, there is a near universal agreement that the truth of its subject is part of the *definition* of when a belief is knowledge. Before outlining the details of the Gettier problem, it is worth assessing just how massive is the attention it has attracted. Gettier's paper was published in 1963. In an early 21st century review,[8] Hetherington (n.d., a) noted:

> There is no consensus, however, that any one of the attempts to solve the Gettier challenge has succeeded in fully defining what it is to have knowledge of a truth or fact. So, the force of that challenge continues to be felt in various ways, and to various extents, within epistemology.

The situation has persisted over 50 years after Gettier's paper. Hetherington (2016, p. ix) noted that the Gettier problem literature was characterized by "widespread frustration" that…

> … epistemologists as a group, it seems, remain as far as they have ever been from agreeing on why the Gettier Moral[9] is true.

---

[7] My thanks to the many who have indulged my earlier attempts to understand the Gettier problem, including Joachim Horvath, Kareem Khalifa, James Norton and Timothy Williamson. My apologies to those among them who tried to save me from my Gettier heterodoxy.

[8] The latest reference in this undated article is from 2006.

[9] [of the inadequacy of defining knowledge as justified, true belief]

General search results affirm the extraordinary attention the problem has attracted. Gettier's (1963) paper had been cited over 6500 times[10] by November 2025, according to Google Scholar. An ordinary Google search on November 7, 2025, on the term "Gettier problem" returned 566,000 hits; and on "Gettier case" returned 511,000 hits. These are extraordinarily large numbers, given that the issue is highly specialized and likely only to attract attention from epistemologists.[11]

One might imagine that a provocative paper like Gettier's would initially attract attention after publication, that in the subsequent few decades the literature would find a generally accepted accommodation and then it would fade from the literature. There are indications that something like the reverse has happened. A Google n-gram plot of the presence of the terms "Gettier problem" and "Gettier case" in Google Books' digitized library for the period 1960 to 2022 (Figure 1.) indicates a marked rise in work discussing the terms in the decade 2010-2020.

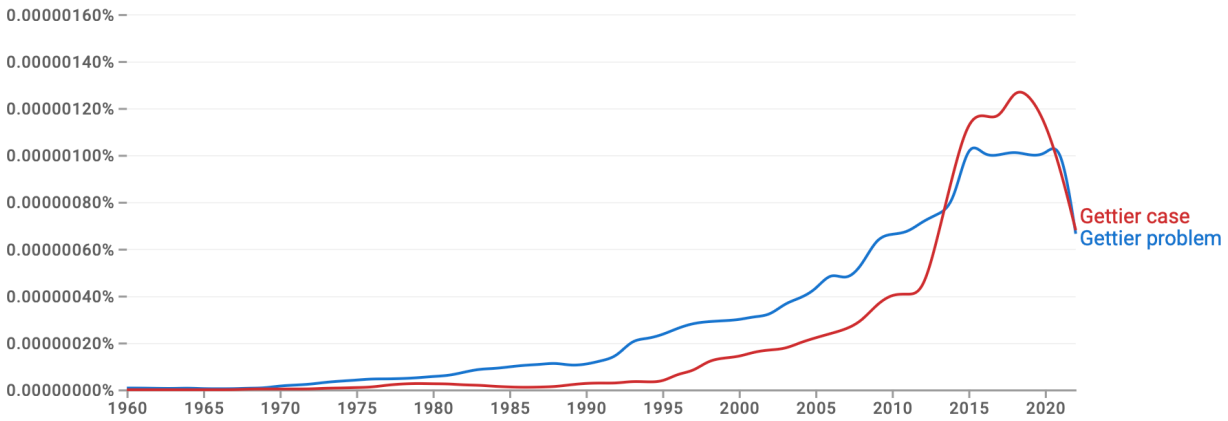

Figure 1. Distribution of the terms "Gettier problem" and "Gettier case."

Whether it is deserved or not, this level of attention and fragmentation of responses shows that the Gettier problem has posed an enduring and vexing problem for epistemology.

[10] Specifically, 6594 citations on November 9, 2025.

[11] For this reason, I despair of a literature search that does justice to the breadth of writing on the topic and apologize to readers whose favorite Gettier problem paper has been overlooked here.

6.2. The Problem

Gettier's (1963) paper argues for the insufficiency of a definition of knowledge that, it is generally allowed, has been recognized since antiquity. It is that a believer *S* knows a proposition *P* if and only if three conditions are met: *P* is true; *S* is justified in believing *P*; and *S* believes *P*. In a short, familiar slogan, knowledge is justified, true belief. The case for its insufficiency is made by displaying counterexamples, in which these three conditions are met, but which we are supposed to see intuitively do not count as knowledge. In them, *S* is justified in believing *P*, but the mechanism through which the justification operates fails. Nonetheless some other circumstance intervenes and *P* turns out to be true after all. Or so an omniscient narrator assures us.

The repertoire of examples has grown large over the decades, as discussion of the problem persisted. An example due to Bertrand Russell, presented prior to Gettier's paper, is typical. In my version of it, in the town square, I notice that the clock on the tower reads noon. Its reading justifies my belief that it is noon. What I do not know is that the clock had stopped a while ago and was left fixed at the noon reading. As it happens, it is noon, so my belief is of a truth after all. Russell (1923, p. 170-71) concludes of his version of the example, in line with the general view of Gettier cases in the literature, that:

> … this man acquires a true belief as to the time of day, but cannot be said to have knowledge. … Such instances can be multiplied indefinitely, and show that you cannot claim to have known merely because you turned out to be right.

The repertoire of Gettier cases in the literature consists, as far as I can see, of situations in ordinary life. The precise notion of justification employed is commonly left open. In the event, the notion is inductive, which keeps Gettier cases within the compass of the present discussion. Most commonly, the justification for a belief is an induction from the evidence of human sense experience. In the first of Gettier's (1963) cases, Smith comes to believe contingent facts about who secured a position. The evidence is an assurance by a company president that Jones will be selected for a position and that Smith has counted ten coins in Jones' pocket. In the other of Gettier's examples, Smith believes Jones owns a Ford since he has always seen Jones drive a Ford. In other examples, recalled in Hetherington (n.d. a, Section 4), we see what looks like a sheep but is in fact a shaggy dog; we strike matches we believe on past experience always ignite, but do not; we see what looks like a barn in a field, but is in fact merely a barn façade. Ichikawa

and Steup (2017, Section 3) offer the case of someone misidentifying a mirage as water, where fortuitously there is water. It is essential to the examples that there is an omniscient narrator who declares the truth.

6.3. Escapes

Over the decades, many solutions to the Gettier problem have been proposed. What they generally share is a retention of the requirement that knowledge requires truth. Ichikawa and Steup (2017, Section 3, their emphasis) note:

> Few contemporary epistemologists accept the adequacy of the JTB analysis. Although most agree that each element of the tripartite theory is *necessary* for knowledge, they do not seem collectively to be *sufficient*.

For over half a century since Gettier's paper, epistemologists have sought to repair the characterization of knowledge as justified true belief, almost invariably by adding further conditions. The range of proposals is so great that no synopsis is possible here. We can however get some sense of its enormity by recalling the categories of responses in surveys of the Gettier problem. Hetherington (n.d. a) lists as "attempted solutions": "Eliminating Luck," "Eliminating False Evidence," "Eliminating Defeat," and "Eliminating Inappropriate Causality"; and as "Attempted Dissolutions": "Competing Intuitions" and "Knowing Luckily." Ichikawa and Steup (2017) list as possible fourth conditions "No False Lemmas," "Sensitivity," "Safety," "Relevant Alternatives," "Reliabilist Theories of Knowledge," "Causal Theories of Knowledge," "Epistemic Luck," "Methodological Options," "Virtue-Theoretic Approaches," "Knowledge First," "Pragmatic Encroachment" and "Contextualism."

Among many recent proposals, I found noteworthy attempts to strengthen the justification for knowledge. Horvath and Nado (2021) define a threshold of "normality" above which justification must rise before it suffices for knowledge. de Grefte (2023) argues that the original definition of knowledge as justified true belief survives if we require that the justification arises from a reliable method.

Prominent among all these proposals is the idea that Gettier cases do not count as knowledge because their agents arrived at the truth by luck. This proposal seems to capture well

a sense of what goes wrong in Gettier cases and, for that reason, seems to be quite widely accepted. A popular version was defined by Prichard (2005, p. 146):[12]

Veritic epistemic luck

It is a matter of luck that the agent's belief is true.

That knowledge of truths cannot be acquired by luck is an anti-luck principle, which can be formulated in positive terms as a principle of safety (p. 71)

The safety principle

For all agents, φ, if an agent knows a contingent proposition φ, then, in most nearby possible worlds, that agent only believes that φ when φ is true.

It is a version of Sosa's (1999) safety principle, but now expressed in terms of possible world semantics. Further reflection led Prichard (2009, p. 41) to add a second condition:

… an ability condition of some sort—i.e., a condition to the effect that the true belief was gained via the employment of the agent's reliable cognitive abilities.

The notion of an agent's ability leaves the content of the condition open. A more precise version is found in de Grefte (2023, p. 533):

VERITIC LUCK: A belief is veritically lucky if and only if it is a matter of luck that the method one used to form one's belief produced a true belief.

This is the understanding of veritic luck that I will use below.[13]

6.4 What the Escapes Cannot Do

What can these escapes hope to achieve? Here an obvious point bears repetition and emphasis. No matter what condition is added to the definition of knowledge to tame the Gettier cases, the methods of justification available remain inductive. Therefore, they are fallible and can always lead us into belief in a falsity. No escape can eliminate this enduring difficulty of inductive support. The best that all these efforts can achieve is some sort of comfort that the true belief arising in Gettier cases fails to be secured in a way that merits calling it "knowledge." The

[12] The term "veritic" is, as far as I can tell, an adjectival derivative of "verity" meaning truth. It is idiosyncratic and does not appear in the dictionaries I consulted. Its use was rare and sporadic until Prichard's introduction of it into the literature.

[13] A minor point is that the "only if" conjunct is too strong in that it rules out veritic luck that cannot be attributed some aspect of a method.

presumption is that there is a way to distinguish true beliefs secured meritoriously from those not so secured and that this way can be captured in simple formulae or rules. That the debate remains open over 60 years after Gettier's paper suggests that this presumption is false.

6.5 Why the Fuss?

The Gettier problem literature is driven by a sense that Gettier cases pose a profound challenge to the concept of knowledge. As the discussion elsewhere in this chapter indicates, I do not sense that they are troublesome for a fallibilist empiricist. My puzzle has been to understand why my indifference is exceptional, even while epistemologists now generally accept fallibilism and have some sympathies with empiricist ideas. It is not really a question in philosophy, but one in the psychology and sociology of epistemologists. Hetherington (2012, p. 217) has given an answer that I find plausible:[14]

> A methodological mistake is allowing them not to notice how they are simply (and inappropriately) being infallibilists when regarding Gettiered beliefs as failing to be knowledge. There is no Gettier problem that we have not merely created for ourselves by unwittingly being infallibilists about knowledge.

This matches my sense that the acceptance of fallibilism by epistemologists is limited and does not extend to an important consequence of fallibilism, the opacity of truth.

In earlier centuries, it seemed quite possible to attain truth and recognize that it was attained. In the eighteenth century, it was widely felt that Euclid had done just this with the truths of geometry and Newton had done the same with truths of dynamics. The challenge to philosophy was to find an infallibilist epistemology that capture the methods they used. Might they employ a Cartesian deductivism? Euclid deduced his geometry from self-evident truths. Or might we look to some development of Kantian notions of synthetic a priori truths? They assure us of the truth of Euclid's geometry and Newton's mechanics.

The alarm over Gettier cases indicates to me a lingering attachment to this older, infallibilist project. For the Gettier cases suggest a mismatch between our best methods and truth

[14] Hetherington does not adopt my approach of dropping the requirement of truth from the concept of knowledge. Rather he finds that Gettier cases are thought mistakenly to pose a problem for justified true belief through a modal fallacy. It conflates partial and holistic characterizations of Gettier cases through a tacit assumption of infallibilism.

that would trouble infallibilists, but not fallibilists. If we look for it, we can find attachments to infallibilism in at least some responses to the Gettier problem. An early version of the anti-luck escape from Gettier cases is provided by Unger (1968). His formulation is that knowing precludes accidental truths: (p. 161)

> For any sentential value of *p*, a man knows that *p* if and only if it is not an accident that the man is right about its being the case that *p*.

This wording leaves open whether accidents allow the attaining of truth by a means compatible with fallibilism, that is, by a means only assured of success with very high probability. Unger goes to some pains to rule out this fallibilist possibility. He imagines a well-shuffled deck of a billion white cards and just one of another color. He asks if a man can know that he will win a bet if the top card is white. Unger (p. 162) inclines to answer that he does not:

> Still, we may also find ourselves saying that he cannot really know that he has won until the color of the card is actually revealed. Rather, the increase to such high chances of success furthers our readiness to apply our analytic condition, to say that it is not at all accidental that the man is right (assuming of course that he is right). But again, and equally I think, our willingness here is not so complete as it might be.

The one in a billion chance of failure is enough to preclude knowledge. Whether Unger's attachment to infallibilism is shared more widely, even in some deep corners of epistemologists' thinking, goes beyond what I can ascertain.

7. The Tenacity of the Opacity of Truth

7.1 The Generality

Each of the various forms of the escapes to the Gettier problem amount to a variant definition of what it is to know. Considerable energy has gone into debates over which is the best definition. In so far as these definitions require truth as a condition for knowledge or that truth follows from the definition, these debates are futile. For they are debates over definitions that fail to meet a condition of adequacy for a definition of scientific knowledge developed in the previous chapter: there must be at least some results in our best science that can be designated as belonging within scientific knowledge. To recall, the difficulty is that even the most prized items of our best science are only supported by fallible inductions. We cannot positively affirm their

truth. The best we can affirm is very strong support such that we deem their truth very likely. It follows that *none* of the definitions of knowledge that require truth enable us to affirm that anything in our best science is knowledge.

8.2 Illustrations

Consider again my walk in the town square, where I notice that the clock on the tower reads noon. I am now well justified in believing that it is noon. That is all I can say. It will be that way if the clock is well functioning and it is noon; if the clock is stopped and it not noon, so my belief is mistaken; or if the clock is stopped and it just happens to be noon, so it is a Gettier case.

There is a temptation to imagine that we can somehow step outside of the case at hand and determine which is the real situation. That is an illusion. We can only assemble more evidence that allows us to be well justified in our belief about the situation. Or we may find evidence that may undo that justification. We will never ascertain the truth with absolute certainty. We will forever be trapped in the narrative of the case.

To see this, continue the above case. After reading the clock on the tower and forming the belief that it is noon, I glance towards a shop window and notice a second clock that also reads noon. I am reassured in my belief that it is noon. Can I now affirm infallibly the truth of my belief? No. I am still trapped in the narrative. It may well be that this clock is functioning well, so I now have further justification that it is noon. Or I may just be in an extended Gettier case. It may turn out, let us say, that all the clocks in the square are synchronized to a single standard clock. All of them are set to noon temporarily as part of regular maintenance of the synchronization system. If it is not noon, my belief would be mistaken. If it is noon, I am now in an extended Gettier case.

This situation can be continued indefinitely. No matter how the story unfolds, we will only ever have fallible support for our belief; or we may end up with evidence that properly shakes our confidence in that belief. That would be the case if the clock in the shop window were to have read 11am. We will never have an infallible affirmation of the truth of the belief. In the real world, there is no omniscient narrator whose authoritative, disembodied voice murmurs through the square for us to hear: "… as it turns out, the clock has stopped at noon, but it just so happens that it is noon."

That we can never step out of the narrative and adopt the position of an omniscient narrator has long been noted in the literature. Using examples similar to those above, Zagzebski (1994, p. 73) concludes:

> As long as the truth is never assured by the conditions which make the state justified, there will be situations in which a false belief is justified. I have argued that with this common, in fact, almost universal assumption, Gettier cases will never go away.

8.3 Cartesian Doubt

This analysis makes clear that definitions of knowledge that require truth do not just fail the condition of adequacy for scientific knowledge. They also fail to be adequate to the most prosaic of examples. Famously, Descartes, in his First Meditation, asked provocatively: (1641, p. 14)

> Indeed, that these hands themselves, and this whole body are mine—what reason could there be for doubting this?

We would like to be able to affirm that we *know* we have hands. However even our seemingly most secure beliefs such as these are only supported inductively and may fail to be true. In the "phantom limb syndrome," amputees continue to have the sensation of the presence of the amputated limb. The sensation is not veridical. There is something along the lines of a reverse syndrome in "body integrity identity disorder." Suffers become convinced that some body part just should not be part of their body. They doubt, in agreement with Descartes' provocation, that the hands they sense are theirs.

As far as I can see, none of the Gettier inspired amendments to the definition of knowledge as justified true belief escape this difficulty. Briefly, amendments that require justification to meet stringent, further conditions fail, since under them justification remains inductive and thus fallible. Can we dismiss the threat of inductive failures by casting them out as requiring possibilities that can properly be ignored, as does Lewis in Section 5.2 above? The suggestion faces the same problem. The decision over what can be properly ignored is itself inductive and thus fallible. It might seem proper to ignore the possibility that we are mistaken in sensing that we have hands. That may be seem so until we learn of phantom limb syndrome. More significant examples arise in science. A nineteenth century physicist might well have felt it

proper to ignore the rather extreme possibilities countenanced by the quantum theory of the following century.

There is, I believe as indicated above in Section 5.1, a quite serviceable response to Cartesian skepticism. It is that our evidence provides overwhelming support for a world that is as it seems and no evidence for the fantasies of a brain-in-a-vat world. It does so without requiring an absolute affirmation of the truth of our beliefs. If we know that for which we have very strong evidence, then we know we are in a world that is just as it seems.

9. Gettier Cases are Benign

The remedy offered in this chapter to the travails of Gettier cases is straightforward. Attempts to augment the definition of knowledge for contingent propositions as justified true belief cannot overcome the opacity of truth. It is futile to try. The fallibility of inductive inference cannot be erased by even the most artful of added conditions.

Precisely because of its opacity, there is no loss to us if we drop truth from the definition of knowledge. In its place, I have suggested above, we should put beliefs that are well supported inductively by the evidence of experience without demanding their truth. Call such beliefs "knowledge" if you like. Or, to avoid confusion, just call them "beliefs well-supported by experience." We lose nothing, since all we ever have had are such well-supported beliefs. We have no means to go beyond what strong inductive support provides. We do not have an omniscient narrator to consult.

The outcome is that Gettier cases cease to be troublesome. They just become another instance of the familiar problem of dealing with the fallibility of inductive support. And they are a less troublesome case. The most common difficulty is that a result that is well supported by the evidence turns out to be false. Something went awry so that what was judged to be a reliable mode of justification failed. The usually safe assumption that the clock in the square is running well was not safe. If the frozen reading is wrong, I form a false belief about the time and will suffer the consequences.

There is a better case. The inductive support we relied for our belief may not function as we thought, but some other circumstance may intervene so that the belief turns out to be right after all. This is the better case. It is the Gettier case. If we have freed ourselves from the worry that the Gettier case troubles our definition of knowledge, all that remains is relief that things

worked out. Or so we should think, in so far as we are able in the event to determine that ours is a Gettier case. I may be mistaken in assuming that the clock in the square is running well, but the belief that it was noon is of a truth. No harm was done.

10. Conclusion

Fallibilism entered the mainstream of thought in philosophy of science over a century ago, after relativity theory and quantum theory overturned what centuries before it had deemed unassailable certainties. An important outcome of this chapter is that epistemology has still to accommodate this most momentous of changes in science. In failing to do so, it continues to err in giving knowledge a central position in the epistemology of science, where the truth of a belief is deemed necessary for it to be knowledge. It so doing it has given pride of place to belief states whose attainment will forever be opaque for the contingent propositions of a science.

My sense is that the epistemology literature persists in this formulation since it has been unable to discard an earlier optimism in which the truth of contingent propositions could be ascertained absolutely. It avows an adherence to fallibilism but not to its import. The outcome is that epistemology has taken on the impossible task of giving a precise account of knowledge that still allows us to attribute knowledge to the content of our best science. Once fallibilism is adopted, there is little choice but to replace what we know assuredly in science with that for which we have the strong inductive support of experience. To persist with attempts to make sense of scientific knowledge as true belief is to engage in a project that cannot succeed. That unhappy fact, I believe, lies at the heart of the continuing struggles in epistemology over knowledge.

References

Brown, Jessica (2018) *Fallibilism: Evidence and Knowledge.* Oxford University Press.

Carrier, L.S. (1993) "How to define a nonskeptical fallibilism," *Philosophia*, **22**, pp. 361–372.

Cohen, Stephen (1988) "How to be a Fallibilist," *Philosophical Perspectives, Vol. 2, Epistemology*, pp. 91-123.

de Grefte, Job (2023) "Knowledge as Justified True Belief" *Erkenntnis*, **88**, pp. 531–549.

Descartes, René (1641) *Meditations on First Philosophy*. Trans. M. Moriarty. Oxford: Oxford University Press, 2008.

Feldman, Richard (1981) “Fallibilism and Knowing that One Knows,” *The Philosophical Review*, **90**, pp. 266-282.

Gettier, Edmund L. (1963) “Is Justified True Belief Knowledge?” *Analysis*, **23**, pp. 121-123.

Hemp, David (n.d.) “The KK (Knowing that One Knows) Principle,” *The Internet Encyclopedia of Philosophy* https://iep.utm.edu/kk-princ/

Hetherington, Stephen (n.d.) “Fallibilism,” *The Internet Encyclopedia of Philosophy*, https://iep.utm.edu/fallibil/

Hetherington, Stephen (n.d., a) “Gettier Problems,” *The Internet Encyclopedia of Philosophy*, https://iep.utm.edu/gettier/

Hetherington, Stephen (2012) “The Gettier-illusion: Gettier-partialism and infallibilism” *Synthese*, **188**, pp. 217–230.

Hetherington, Stephen (2016), *Knowledge and the Gettier Problem*. Cambridge: Cambridge University Press.

Hintikka, Jaakko (1962) *Knowledge and Belief*. Ithika, NY: Cornell University Press.

Horvath, Joachim and Nado, Jennifer (2021) “Knowledge and Normality,” *Synthese* (2021) **198**, pp. 11673–11694.

Ichikawa, Jonathan and Matthias Steup (2017) “The Analysis of Knowledge,” *The Stanford Encyclopedia of Philosophy* (Fall 2024 Edition), Edward N. Zalta & Uri Nodelman (eds.), https://plato.stanford.edu/archives/fall2024/entries/knowledge-analysis/

Jeffrey, Richard (1968) “Probable Knowledge,” *Studies in Logic and the Foundations of Mathematics.* **51**, pp. 166-190.

Kaplan, Mark (1985) “It's Not What You Know that Counts,” *The Journal of Philosophy*, **82**, pp. 350-363.

Leite, Adam (2010) “Fallibilism,” in J. Dancy, E. Sosa and M. Steup, *A Companion to Epistemology*. 2nd ed., Wiley-Blackwell, pp. 370-75.

Lewis, David (1996) “Elusive Knowledge,” *Australian Journal of Philosophy*, **74**, pp. 549-67.

Norton, John D. (2021) *The Material Theory of Induction*. BSPS*Open*/University of Calgary Press.

Papineau, David (2021) “The Disvalue of Knowledge,” *Synthese.* **198**, pp. 5311–5332.

Prichard, Duncan (2005) *Epistemic Luck*. Oxford Clarendon Press.

Prichard, Duncan (2009) “Safety-based Epistemology,” *Journal of Philosophical Research*, **34**, pp. 33-45

Reed, Baron (2012) “Fallibilism,” *Philosophy Compass* 7/9, pp. 585–596.

Russell, Bertrand (1923) *Human Knowledge: Its Scope and Limits*. London: George Allen and Unwin. Repr. 1948.

Sosa, Ernest (1999) “How to Defeat Opposition to Moore,” **33**, *Supplement: Philosophical Perspectives,* 13, *Epistemology*, pp. 141-153.

Unger, Peter (1968) “An Analysis of Factual Knowledge,” *The Journal of Philosophy*, **65**, pp. 157-170.

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