Opened, October 29, 2025

First draft, November 26, 2025

Chapter 15

Inductive Fallibilism in small-e empiricism

1. Introduction

Inductive fallibilism asserts that inductive inferences are always fallible. No matter how strong the inductive support for some contingent proposition, there is always the possibility that the proposition is false, even if its falsity is extremely unlikely. According to small-e empiricism, the only way we can learn the truth of contingent propositions is by their inductive support from experience. If we combine inductive fallibilism with small-e empiricism, it follows that we can never secure inerrantly the truth of a contingent proposition. In the best case, we can be assured only that its falsity is very unlikely.[1]

The import of this simple result has been repeatedly neglected or misread, even by literatures that generally endorse some form of fallibilism and empiricism. The goal of this chapter is to articulate inductive fallibilism in the context of small-e empiricism, its consequences and how they have been neglected or misread.

Section 2 will review how inductive fallibilism arises from the nature of inductive inference and how its import in small-e empiricism is that we can never be assured absolutely of the truth. Section 3 will review how inductive fallibilism in empiricism opens the possibility of the dynamics for scientific progress actually seen in the history of science. In it, a science evolves through a sequence of stages. While results in earlier stages enjoy strong inductive support from experience, some of the results are false. Later stages correct the falsehoods. These later stages are themselves subjected to further corrections by still later stages; and so the dynamics of self-correction continues.

Section 4 explores the import of inductive fallibilism for beliefs when those beliefs are formed by means of Hume's maxim "A wise man, therefore, proportions his belief to the evidence," but does not go beyond it. The principal conclusion is the opacity of truth. We may

[1] For a brief recounting of some of the claims of this chapter, see Norton (manuscript).

have a belief in a true, contingent proposition, but the fallibility of our only means of ascertaining its truth precludes our certainty of its truth. It follows that Big-K Knowledge, whose definition requires the truth of the contingent proposition believed, is a laudable idealization whose realization can never be affirmed. A parable illustrates the difficulty. The imaginary "Library of Truth" contains all contingent truths of which we are certain. Its shelves are bare.

A condition of adequacy for any definition of knowledge is that it must be able to attribute knowledge to the results of our best science. It is not enough to say that some of the results of our best science are knowledge, we know not which. The definition must identify which they are. The outcome is that a demand for truth must be dropped from the definition of knowledge. We arrive at a definition satisfying the condition, if we designate as scientific knowledge, those contingent proposition enjoying strong inductive support from experience.

Discarding the *extra* condition that the propositions must *also* be true is no loss. We have never been able to satisfy this extra condition in our accounts of scientific knowledge. In the past they have included propositions we now find to be false; and propositions we now include in scientific knowledge will, we should expect, be found in the future to be false. That is an unavoidable consequence of the fallibility of inductive inference. Adding further conditions to our definition of knowledge cannot escape this consequence. We should, I suggest, follow a proposal already appearing in the literature. The centrality of knowledge in the epistemology of science should be replaced by a central concern for what enjoys strong inductive support from experience.

Section 5 reviews briefly the treatment of fallibilism in the epistemology literature. The main difference between it and the treatment here lies in the locus of the definition. In the epistemology literature, fallibilism is defined in terms of beliefs directly. Here fallibilism is limited to the logical relation of inductive support. The fallibility of belief arises as a secondary consequence of it. For me, the most sobering aspect of the widespread acceptance of fallibilism in epistemology is that the inevitable consequence of the opacity of truth is neglected. Truth 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 6 identifies two ways that the import of the fallibility of inductive inference has been misconstrued in the literature. A pessimistic misconstrual denies that inductive inference is cogent. It manifests in various forms of inductive skepticism, such as acceptance of the tenability of fables that the world we experience is fabricated by versions of deceiving demons. The correct response is that we have massive inductive evidence that the world of ordinary experience is just as it seems and we have no evidence at all that it is fabricated, for example, for electrically stimulated brains in vats. The optimistic misconstrual takes fallibility not as a result but as a challenge. If inductive inferences cannot assure us of the truth of contingent propositions, then we should just try harder. What follows are elaborate proposals to circumvent the fallibility, such as the eight rules in David Lewis' "Elusive Knowledge." The failure of all such proposals is inevitable. Whether they acknowledge it or not, they are elaborations of an inductive logic, all of which are fallible.

Section 7 reviews the Gettier problem, which has long dominated writing in epistemology. Section 8 assesses the limits of what the enormous range of escapes to Gettier problems can achieve. The principal difficulty is that none of them resolve the most important consequence for knowledge of the fallibility of inductive inference, the opacity of truth. This is already evident in the presentation of Gettier cases. The truth is always announced by an omniscient narrator. As long as we remain within the narrative, we can gain further evidence that would bear on the truth of result in question. But that bearing will always be inductive and thus fallible.

Section 9 argues that Gettier cases are benign. Once the demand for truth in Big-K Knowledge is discarded, Gettier problems cease to be troublesome. The fallibility of inductive inference is 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.

The literature on Gettier cases is populated with fanciful examples contrived to make a point. They are all by construction extremely rare, since they all require two very unlikely occurrences to fit together in just the way needed. Section 10 investigates whether there are Gettier cases or just Gettier-like cases in the history of science. Two aspects of Gettier cases do not readily translate into real scientific cases. Surrogates for them must be found. First, there is no omniscient narrator who declares the truth. In place of the narrator are the judgments of a later stage of the science as to which of the results of the earlier stage are true or false. Second, the distinctive feature of Gettier cases is commonly taken to be that an agent arrives at the truth by unanticipated luck. There is no simple analog of this sort of luck in the sciences. In its place are considerations of method. We can identify that a scientist is using unreliable methods by reviewing the performance of the methods in the history of science. If a scientist used methods that are, as a matter of historical record, unreliable, but nonetheless arrived at a result judged true by later stages, then we have a surrogate for luck.

This last discussion indicated how Gettier-like cases would manifest in the history of science, but not that they must manifest. Section 11 reports that such Gettier-like cases do not just arise occasionally in the history. Rather they are pervasive in the history. I had no trouble finding many just in those parts of the history of science where I have some comfort. The examples recounted are: Bradley's discovery of stellar aberration, Dalton's discovery of atomic theory, Carnot's discovery of the founding principles of thermodynamics, Thomson's discovery of the electron, Bohr's discovery of the quantum mechanics of atoms, Einstein's discovery of the general theory of relativity and Hubble's determination of the Hubble law. In each case, the Gettier-like result is not just preserved as the science progresses. They are routinely celebrated as major achievements in foundations of a science.

Section 12 offers brief concluding remarks.

2. Inductive Fallibilism in small-e empiricism

This chapter will develop the consequences of combining two results. The first arises from the nature of inductive inference. The second in a basic supposition of small-e empiricism.

First, that inductive inference is fallible follows directly from how inferences are divided into the deductive and the inductive, that is, the demonstrative and the non-demonstrative. Deductive or demonstrative inferences are by definition truth preserving. If the premises are true, then the conclusion must also be true. Inductive inferences are non-demonstrative, by definition. The truth of the premises does not assure the truth the conclusion. The likelihood that the conclusion is true may merely be enhanced, even if very slightly. In the better cases, the truth of the premises may render the truth of the conclusion very likely. If the induction is expressed in relations of inductive support, such as conditional probabilities, the strongest relation of support must still leave open the possibility that the result supported is false. If the relation were to be so strong as to leave no possibility of falsity, then it would be demonstrative.

It follows that even the strongest of inductive supports leaves open the possibility of falsity, even if the falsity is unlikely. We can never be assured absolutely of the truth of a result on the basis of inductive support.

Second, the core doctrine of small-e empiricism, as developed in Chapter 12, asserts that the inductive support of experience is the *only* means of providing support for the truth of the contingent propositions of science. If we allow that this restriction applies to contingent truths in general, it now follows that we can never be assured absolutely of the truth of a contingent proposition.

We should not overestimate the peril that comes with this lack of absolute assurance. In familiar, established sciences, we have levels of inductive support that are so strong that, for all practical purposes, they can and should be taken as assuring us of truth. However, a careful and proper statement of the level of security of even the best inductively supported results is that there is always a possibility of error, even if extremely unlikely in any particular case.

3. Fallibilism and the Dynamics of Scientific Progress

The dynamics of progress in science reflects the fallibilism of inductive inference. It is possible that our inductive inferences from experience in science never lead to strong support for falsehoods. Since that would require the epistemic dice to fall in our favor every time we make an inductive gamble, it can only be realized in an extremely unlikely, fantasy world. In the real world, in our real history, the sciences advance through inductive successes and inductive failures.

At each stage of a science as it develops, some of its results will be correct and some not. We cannot then be sure of which are correct. At the time, all we have is strong inductive support for the results and a lingering sense that this support is fallible. Then the science moves on, commonly through the opening of new experiences. Through them, we learn that some of the results of the past stage of science were incorrect, in spite of their enjoying what we then thought[2] was strong inductive support. The process repeats. Our science grows. We correct past

[2] Recall that, according to the material theory of induction, judgments of inductive support derive from facts in the relevant domain. If we adjust our appraisal of which those facts are, we must make corresponding adjustments to our appraisal of relations of inductive support.

errors and replace them with newer, well-supported propositions. At no stage can we be sure that we have secured the final truth. The best we can say is that the inductive evidence for our most secure results is so strong that their falsity is extremely unlikely.

We will return to this dynamic below in the discussion of the Gettier problem, where it will be illustrated by examples drawn from the history of science. I will argue that it is routine for successes of the earlier stages to be Gettier-like cases, when judged by the results of the later stages, and that these Gettier-like case results are carried forward to the new stage. Here, Gettier cases are not a dire epistemic threat, but a familiar and celebrated part of the progress of science.

4. Inductive Fallibilism applied to Epistemology

4.1 Beliefs

In formulating small-e empiricism, I have sought as much as possible to avoid detailed engagement with beliefs derived empirically. However, a proper analysis of the import of fallibilism for the present philosophical literature cannot avoid such engagement completely. In the following, I will consider beliefs formed in accord with Hume's maxim, as recounted in Chapter 12: "A wise man, therefore, proportions his belief to the evidence." Under this maxim, we should be very confident in our belief in contingent results that enjoy very strong inductive support. That confidence must always be moderated by a recognition that the belief may err, even if the possibility is very slight.

I shall avoid assertions on relations among these confidences in our beliefs, that go beyond those assertions that are supplied directly by Hume's maxim. The prudence of this course is illustrated by Kyburg's (1961, p. 197) lottery paradox. If we consider any particular ticket in a lottery, a full recounting of the operation of a lottery lends strong support to the proposition that the ticket will lose. We may be tempted to assume a conjunctive principle for beliefs. If we should have strong confidence in each of a set of propositions, then we should have confidence in the conjunction. We should be confident that any particular ticket in the lottery will lose. Therefore, under this principle, we should be confident in the conjunction, that all tickets lose. This result contradicts the certainty of lottery design, that one ticket will win.

The paradox is avoided if we adhere to Hume's maxim. It rejects the conjunctive principle, since the conjunction is refuted by the evidence. We have strong support that any nominated ticket will lose, but certainty that one will win. Similar problems will also be avoided by adherence to Hume's maxim, as long as the inductive logic employed is internally consistent. Although the lottery paradox was formulated in terms of probabilistic relations of support, this more general moral applies to all well-formulated relations of inductive support.

Work in epistemology may seek rules relating confidences in our beliefs. As far as small-e empiricism is concerned, such efforts cannot succeed if they go beyond what is authorized by Hume's maxim. For, according to small-e fNecessempiricism, the only basis for confidence in a contingent proposition is inductive support from experience.

It may still happen that conforming our beliefs to what is well-supported inductively leads to internal contradictions. The source of the problem is then identifiable as some defect in the account of inductive support employed. If we conceive of inductive inference as governed by universal rules, we face a difficult foundational problem. One troublesome case will require a global revision of universal rules. The material theory of induction can handle these cases with much less trauma. Since it treats all inductive inference as local, the repair needed will be local. Since background facts replace universal rules, the repair will merely require a revised assessment of the background facts in the domain considered, such as is routine in corrective work in science.

4.2 Knowledge: The Opacity of Truth

The literature in epistemology has expended massive efforts in the search for an unproblematic account of just when our beliefs rise to the level of knowledge. From the perspective of small-e empiricism and its inductive fallibilism, these efforts are misplaced. In science especially and even in ordinary life, what matters is whether some contingent proposition is well-supported inductively by experience. We can ascertain whether such support has been secured. The essential, overriding fact is that this is *all* we can ascertain.

When, we may still be inclined to ask, does a belief rise to the exalted level of knowledge? Here knowledge of a contingent proposition is usually understood to require strong support *and* its truth. The difficulty is that small-e empiricism precludes our ascertaining the truth of the contingent proposition directly, independently of the strength of inductive support of experience. There is no further condition that would indicate truth.

The first reaction to this problem and, in retrospect, the final response of small-e empiricism is just this: why care? Why not just focus on what can be ascertained? If some contingent proposition enjoys very strong inductive support empirically, is that not enough?

We might nonetheless persist in asking the question. We then face a trilemma in seeking to define just what it takes for a belief in a contingent proposition to count as knowledge.

Either we must add something to the condition of very strong inductive support from experience;

or we do not make an addition;

or we give up trying to formulate a precise definition.

The difficulty for the first horn is finding a workable addition. Small-e empiricism asserts that no such addition can succeed in assuring truth, beyond what inductive support already provides. It keeping with small-e empiricist expectations, no such addition has been found. The popular addition, as we shall see in greater detail below, is just to add directly the condition that the proposition must also be true. This direct addition faces the great difficulty that truth by itself is opaque. We have no way to determine the truth of a contingent proposition beyond the fallible assurances supplied by the inductive support from experience. In the examples that populate the epistemology literature, this difficulty is obscured by an omniscient narrator who just declares which propositions are true. In the real world, there is no omniscient narrator. Our sole guide is the strength of inductive support from experience.

For this reason, this first horn in untenable. We might insist on defining knowledge that a well-supported belief must also be true. Call this familiar definition of knowledge as justified, true belief, "Big-K Knowledge." To insist on it is to adhere to a laudable philosophical purity whose outcome is sterile. We have defined Big-K Knowledge as a belief state that we can never be assured is realized. Consider the collection of the best results of our present science. Presumably at least some of them are true and, we hope, most. However, we can point to no specific result, under this first horn, and say it is Big-K Knowledge, that we *know* it.

In my view, either of the two remaining horns are acceptable. They are, in the end, quite similar in their import.

The second horn is just to accept that a belief rises to the level of knowledge if it enjoys strong inductive support. We could call it "small-k knowledge." This is how the term "knowledge" is applied routinely in scientific discourses. There "scientific knowledge" is routinely identified with the core claims of a stable, well-supported science. The term is applied even though there is always a very slight chance that a proposition proclaimed as known may turn out to be false. Sometimes, the possibility is more than slight, as is the case in fundamental physics. Our best theory of gravity is Einstein's general theory of relativity. That it will be found to fail on very small, quantum scales is widely expected. Our best theory of matter is quantum theory. That it will be found to fail in very intense gravitational fields is expected. These failures only arise when we seek to apply these theories in extreme domains. We can surely be very confident of general relativity for a wide range of gravitational situations; and of quantum theory for common molecular scale processes. However, our confidence cannot be absolute; and it must diminish as we approach domains in which we expect theories to fail.

Under the third horn, we would just renounce efforts to maintain knowledge as a central term of art in our analysis. That does not mean that we must excise the terms "know," "knowledge" and the like from our vocabulary. They would only have an informal import, such as the terms have in ordinary, casual discourse.

The overall import of these last two horns combined is that we should remove knowledge from its central position in epistemology. It is replaced by contingent propositions that are well-supported inductively by experience.

That knowledge should no longer occupy a central position in epistemology has been suggested repeatedly in the literature. 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."[3] 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.

4.3 Condition of Adequacy for Knowledge

The discussion of this last subsection embodies a condition of adequacy for any definition of knowledge. We are free to define knowledge as we wish. A definition is a declaration of how we will use a term. That freedom has been exercised extensively in responses to the Gettier problem. The danger, however, with this freedom is that we wander off into a philosopher's fantasy land of idealized but unrealizable imaginings. If knowledge of contingent propositions is to be found anywhere it is in our best science. We should expect that the definition of knowledge, if one is to be employed, identifies our best science as knowledge. This is a condition of adequacy for any definition of knowledge.

Big-K Knowledge fails to satisfy the condition. What also does not suffice is a definition of knowledge that merely tells us that *some* of the contingent propositions of our best science might be knowledge, but does not enable us to identify just which they are, if any.

A parable illustrates the importance of meeting this condition.

4.4 The Library of Truth

A story I once heard when I was more easily led was of a not too distant land, not too long ago, in which the benevolent ruler assembled the land's greatest savants and commissioned them to catalog all the contingent truths of known science. To house them, a huge library was built—the Library of Truth—and over its marble paved doorway stood a sign:

"Let none but truths enter here."

A guard stood at the door and refused entry to all works, but those whose truth could be assured. After a year had passed, the ruler entered the library, expecting its shelves to be overflowing with the greatest of scientific riches. The shelves were bare. They held not a single work. In rage, the

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

ruler demanded an explanation of the guard, who replied, trembling, "But sire, none could prove their truth."

Unknown to the ruler, the savants had seized the opportunity and had secretly started a hidden repository of the best science. They called it "The Library of the Well-Justified." Its shelves filled rapidly and became an invaluable resource for generations to come.

4.5 The KK Principle

The concerns of this section are raised by 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 intuition seems to be that a belief only rises to the level of knowledge if it is so firmly supported that no future evidence can compromise 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 Hemo (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 given in the last section. 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 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 question of foundational importance. For the requirement of strong inductive support from experience is defined fully in terms of inductive logic. The mental state of the believer is not a part of the definition (other than that the believer believes the proposition). It might be comforting to require that believers are explicitly aware of this strong support before they are authorized to declare it as knowledge.

Alternatively, we might find it prudent not to demand it. Most people know many things, without their being able to articulate the evidential support for them. We all know that the earth orbits the sun and not vice versa; and that matter is made of atoms, whereas most of us could not articulate the evidential case for these propositions. The situation is little different for those with professional competence in a science. Much of what they know is simply read from standard textbooks, which do not provide elaborations of the evidential case made for the results. The function of a textbook is to insulate a new generation of scientists from the need to learn the intricacies of these evidential cases for result found to be reliable when the textbook was written. Do these scientists "know" a foundational and unchallenged result that they learn from a textbook?

Does it matter? 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. Fallibilisms in the Philosophy Literature

The formulation of fallibilism employed here is narrow. It applies specifically to the capacity of inductive inference and inductive support and applies to contingent propositions. Hence, I have called it *inductive* fallibilism. Here I will briefly review, as best I can, the enormous literature on fallibilism in the epistemology literature to delineate how inductive fallibilism relates to it.

The doctrine of fallibilism seems to be widely accepted by epistemologists. 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.

5.1 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.”

5.2 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 chapter. 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.

> 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.[4] My sense is that most fallibilists agree. Reed (2012, p. 585, his emphasis) makes this condition explicit in his introductory formulation of fallibilism:

> 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.

5.3 Charles Sanders Peirce

Although the term "fallibilism" already had some limited use prior to Peirce's work, the term seems to have entered the philosophical literature in a more precise form through his writings. The formulations of the last subsection resemble at least the verbiage of Peirce's formulation. Here is one synoptic formulation (1940, p. 58):

> … On the whole, then, we cannot in any way reach perfect certitude nor exactitude. We never can be absolutely sure of anything, nor can we with any probability ascertain the exact value of any measure or general ratio.
>
> This is my conclusion, after many years study of the logic of science; and it is the conclusion which others, of very different cast of mind, have come to, likewise.

Peirce explicitly restricts the means through which we can learn in science to some form of enumerative induction, whose scope he finds to be very weak. His assessment begins (p. 56):

> All positive reasoning is of the nature of judging the proportion of something in a whole collection by the proportion found in a sample. Accordingly, there are three things to which we can never hope to attain by reasoning, namely, absolute certainty, absolute exactitude, absolute universality. We cannot be absolutely

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

> certain that our conclusions are even approximately true; for the sample may be utterly unlike the unsampled part of the collection. We cannot pretend to be even probably exact; because the sample consists of but a finite number of instances and only admits special values of the proportion sought. …

It was welcome to find that Peirce continued to assert explicitly a key tenet of small-e empiricism: the *exclusive* power of induction to inform us (p. 56):

> Now if exactitude, certitude, and universality are not to be attained by reasoning, there is certainly no other means by which they can be reached.

My reporting of Peirce's views here is limited to the verbiage of the passages selected from Peirce's writings and what they appear to assert on their face. Whether these conform with Peirce's other writings is a question I must leave open for careful Peirce scholarship. That Peirce restricts his remarks to some form of enumerative induction must be reconciled with his celebrated introduction elsewhere of *abductive* inference.

5.4 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 terms of beliefs and, especially, as a problem for the special belief state of knowledge. The formulation of inductive fallibilism 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.

6. Misconstruals of Inductive Fallibilism

The most significant and enduring failure of the epistemology literature has been to persist in insisting on a definition of knowledge that requires the truth of its subject and of making that definition central to epistemology. There are two further misconstruals of inductive fallibilism. One is to adopt a far-reaching pessimism: 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.

6.1 Fallibilist Pessimism

Inductive skepticism in its most expansive form amounts to a denial of the cogency of inductive inference. Its most vivid expression comes in the form of Cartesian skepticism, which has already been discussion in Chapter 8, section 7.[5] The idea that the apparent world is merely an artfully contrived deception has persisted. In the past century, in the philosophical literature, we have Putnam's example that we are all brains in vats. In the popular media, we have the Matrix series of movies. In popular science, we have the proposal that our world is really a computer simulation. I like to think that the idea of an illusory reality has sufficient entertainment value that otherwise astute thinkers willingly suspend their normal critical standards for the fun of it. However, the comforting thought that the lapses in rationality are willful is challenged by the persistence of the idea that our reality is computer simulation.

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

[5] We may exclude Descartes himself from arguing for inductive skepticism. He sought deductive certainty and was correct to point out the difficulty of excluding demonic deceptions deductively.

"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. [6]

Since small-e empiricism is not committed to defending knowledge as true belief, its response to Cartesian skepticism is easy. An essential component of the skeptical scenarios from Descartes to the present is that matters are so contrived that the deceptive truth never manifests in our experience. On the contrary, all our experiences are so set up that they match perfectly with the reality supposed to be illusory. In that circumstance, a small-e empiricist has massive evidence for that reality and *no evidence whatever* for any of the many deceptive scenarios envisaged. If we follow Hume's maxim, we should assign very strong belief to the reality of the world as it manifests. The strength of this belief is stronger than our belief in even the most secure results of science, since the evidential foundations of those results in experience depend on the corrigible veracity of those experiences. That makes even more dubious any skeptical scenario that tries to draw on results of present science.[7]

6.2 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):

---

[6] 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."

[7] While writing this last paragraph, I found it hard to shake my incredulity that there is still a need in an academic treatise in philosophy to argue that the ordinary world really is just the way it seems.

> 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?

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 and 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, that we can somehow 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."[8] 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*.

[8] 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)

How can the rules achieve this?[9] 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 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.[10] 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.

---

[9] 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.

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

7. The Gettier Problem

7.1 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.[11] 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,[12] 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[13] is true.

---

[11] 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.

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

[13] [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[14] 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.[15]

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.

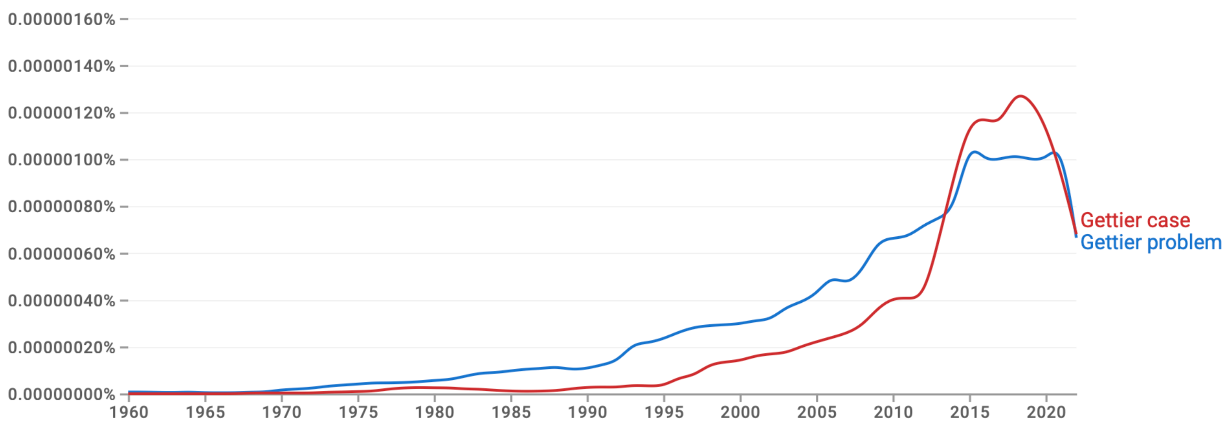

Figure 1. Distribution of the terms "Gettier problem" and "Gettier case."

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.

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

[15] 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.

7.2. 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, 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.

7.3. 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):[16]

> *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.[17]

7.4 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

[16] 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.

[17] 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.

7.5 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:[18]

> 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

---

[18] 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.

8. The Tenacity of the Opacity of Truth

8.1 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 the condition of adequacy described in Section 4.3 above. To recall, since 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 our best science is knowledge.

The core problem for definitions of knowledge addressed in this chapter remains unsolved. Truth is still opaque. That is so, whether we consider a Gettier case or a more normal case of inductive inference. The best that we learn is that the result is well-supported by the evidence. We are never given an absolute assurance of its truth.

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 concerning 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 6.2 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 6.1, 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 a futile project 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 fails. 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 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 a truth. No harm was done.

10. Gettier-Like Cases as Successes in Scientific Progress

This last section argues that Gettier cases are benign in the context of the contrived examples familiar in the Gettier literature, such as my walk in the town square. These cases are so contrived as to be very rare in ordinary life. They need not just one, but two unlikely occurrences in tandem. An otherwise reliable justification must fail and a second unlikely occurrence must restore the truth of the result. Given their rarity, do Gettier cases or perhaps just Gettier-like cases have any presence in science? Only then would they be relevant to an epistemology for science like small-e empiricism.

We will see below that there are Gettier-like cases in science. It will take a little re-analysis to identify them, since the conditions for Gettier cases in the present literature do not map directly onto conditions in science. These Gettier-like cases are, it will be argued, a routine feature of scientific progress. They are not treated as dire threats to our understanding of scientific knowledge. On the contrary, they are routinely celebrated as great scientific successes.

10.1 A Surrogate for the Omniscient Narrator

The first difficulty in identifying Gettier cases in science is that we have no omniscient narrator to tell when such a case has arisen. The dynamics of scientific progress, described in Section 3 above, supplies a suitable surrogate. We will never have an absolute affirmation of the truth of core propositions in science. As science progresses, each stage can identify what it considers to be the errors of the earlier stages and which are to be accepted as its truths. Gettier-like cases will be those in which an earlier stage is, informally, lucky in its determination of a truth, as judged by the next stage.

10.2 Method Replaces Luck

Luck, or more precisely, *unanticipated* luck is the key fact in a Gettier case.[19] The difficulty in scientific examples is to determine when some truth has been achieved by unanticipated luck. In the contrived examples from ordinary life, we have a default expectation on the range of possibilities in the fictional Gettier scenario. If we are willing to distribute chances over them, we can then determine whether some particular instance has a small chance and thus can be lucky. It is plausible from default assumptions about how I might behave that my visit to the town square could reasonably happen over a wide time interval. It does seem unlikely and thus lucky that my visit coincides precisely with noon. It is unanticipated luck since I have no reason to think that the precise timing of my visit has special import.

It is, however, quite unclear that a similar concept of luck can be implemented in the cases discussed below in the history of science. Niels Bohr, for example, used a collection of ideas from 19th century electrodynamics, applied inconsistently, to arrive at some key results of the coming quantum theory of the atom. He also used these same ideas to derive many more results inconsistent with the coming theory. We can loosely assign chances to the different times for my visit to the town square. Over what possibilities do we assign chances so that Bohr's successes are lucky?

A better way to identify Gettier cases in the history of science, is supplied by de Grefte's (2023, p. 533) reliabilist characterization above of Gettier cases. He wrote: "…it is a matter of luck that the method one used to form one's belief produced a true belief." The key idea for me is that luck is defined in terms of what the *method* can properly deliver. We can discern the methods used by a scientist and we can determine retrospectively whether the scientist had any basis for expecting those methods to provide results that are important to the next stage of science.

A Gettier-like case in science in one in which a method that turns out in retrospect to be ill-adapted to the science nonetheless proves to deliver a truth, as judged by the next stage.

---

[19] Is there such a thing as *anticipated* luck? A beachcomber, sifting the sand each day, does not know, on any particular day, whether luck will turn up some valuable. However, the beachcomber anticipates that if the sifting is carried out repeatedly, day after day, luck will eventually come their way.

Among the many parts of the method is one that just happens to work by the retrospective judgment of a later stage of the science. That is, some component of the method used replicates something that works according to the later science. The scientists of the time cannot know of the future validation of this component, since they do not know the future science. Call it an "unanticipated saving factor." It is the distinctive feature of Gettier-like cases in science. They are, informally speaking, the lucky occurrence that makes the case Gettier-like.

10.3 Actual Performance Replaces Possibility

In the Gettier literature, it is routine to make judgments of luck by means of possible world semantics. A success is lucky if the same success would not be realized in some suitably defined nearby possible world. I am loath to use this approach to assess the adaptation of the methods. My concern over judgments concerning possible worlds is their malleability. We need to determine which possible worlds are closer to the actual and which are farther away; and we need to determine what it is for the conditions and methods of the actual world to be realized in them. My suspicion is that there is sufficient elasticity in the complications of these determinations that a creative philosopher could manipulate them to get whichever result happens to suit their purposes.

Here the fact that the examples are real and not fictional stories gives us a better way to proceed. In each case, we can identify the methods actually used by the scientists; and we can determine their prospects. In the Gettier-like cases, we will see that the methods used routinely returned many results deemed false in the retrospective view of later stages. Thus, we have no good basis for expecting these methods to return a result amenable to the later stage. Considerations of luck have been replaced in the analysis by considerations of the actual performance of the method. There is no need to entertain fragile counterfactuals as to how the method would perform in some imaginary possible world. We consult the actual history.

10.4 Gettier-like Cases are Common

These general reflections open the possibility of Gettier-like cases in scientific progress. It does not establish their frequency. When I began looking for them in the history of science, I was astonished at how often they appear. Gettier cases are *pervasive* in the history of science. The following sections recount briefly a selection of them from a range of sciences with whose history I have some comfort. The examples selected are notable for their importance in the history of science and for their subsequent celebration as major discoveries in science. The examples are: Bradley's discovery of stellar aberration, Dalton's discovery of atomic theory, Carnot's discovery of the founding principles of thermodynamics, Thomson's discovery of the electron, Bohr's discovery of the quantum mechanics of atoms, Einstein's discovery of the general theory of relativity and Hubble's determination of the Hubble law.

10.5 Variability of Methods

We will find that the methods used in the historical examples change over time. This is an expectation of the material theory of induction. Central to the methods are judgements of inductive support. These judgments are in turn dependent on the facts that prevail in each domain. Thus, the scientists' assessments of these relations of support depends in turn on their assessments of which are the facts of their domains. Since these last assessments are fallible, it will turn out commonly that methods that they deem reliable are not so. That variability is what is responsible for the prevalence of Gettier-like cases in the progress of science.

10.6 There is No Cause for Alarm

The sense of alarm surrounding Gettier cases is misplaced in science. The worst we can say of the Gettier-like cases in the history science is that the evidential case made for their results was, at the time of their discovery, imperfect. Nonetheless, they were clearly important to subsequent developments in their science and for this are celebrated by later scientists.

Here, I do not intend a skeptical moral about discovery in science. Rather, I intend an optimistic moral. These cases show how science can proceed successfully, in spite of the fallibility of its methods. We see scientists proceeding responsibly with the best methods then available, as they should, and that later work is able to identify the errors and correct them. In all the cases considered below, the Gettier-like successes do not result from pure happenstance. Rather, unknown to scientists at the time, there is something right about the methods they are using. What they cannot then distinguish are which are those parts and which are not. That is a retrospective judgment. In the case of Bohr's model of the atom, for example, his key results of the stability and energetic discreteness of his electron states was properly a reflection of the stability and discreteness of the observed atomic spectra. Bohr would have recognized this at the time, but he would not have been able to identify just which of his remaining results were erroneous, for his methods did not allow it.

11. Gettier Cases in the History of Science

In order to make it possible to recall a range of Gettier cases in the history of science, I have given only abbreviated accounts of each.

11.1 Bradley's Discovery of Stellar Aberration

James Bradley (1729) reported his discovery of stellar aberration to the Royal Society. According to it, the apparent angular position of a star is altered slightly as a result of the motion of the Earth. In the simplest case, in which the star is located perpendicular to the Earth's motion, we subtract vectorially[20] the Earth's velocity v from the velocity of the light c and find the angular displacement in the direction of the Earth's motion to be $\alpha = v/c$. We saw in discussion of this result in Chapter 11, Section 10, that the general formula is
$\tan \alpha = \sin \theta (v/c) / [1+(v/c) \cos \theta]$, where θ is the angular elevation of the star in the direction of motion of the Earth. For Bradley's case of $v << c$, this general formula reduces to the one Bradley reported, $\sin \alpha / \sin \theta = v/c$.

Bradley's discovery, as we saw in Chapter 11, was a celebrated result and one of the few experimental results that Einstein reported as decisive for his discovery of special relativity. It is, on our best present understanding, a truth. However, Bradley's method was flawed. For he derived the result by Newtonian addition of velocities. Famously, according to special relativity, this method fails for motions close to or at the speed of light. Under the Newtonian rule, the motion of the Earth has to be added vectorially to the speed of a light signal to produce an altered speed for the light. A postulate of special relativity is that the speed of light remains unchanged, even if measured from a moving Earth. Special relativity offers a different rule for the composition of velocities in place of the Newtonian rule of vector addition that preserves the constancy of the velocity of light.

Bradley's result survived, however, because of an unanticipated saving factor. The relativistic[21] and Newtonian rules for velocity addition give different results for the composed speeds, but the *same* result for the angular deflection under the small v/c conditions of aberration.

---

[20] To use the vector concept, not introduced until the following century.

[21] The relativistic rule yields $\tan \alpha = (\sin \theta (v/c) / [1+(v/c) \cos \theta])\sqrt{1 - v^2/c^2}$. The correction to the Newtonian result is a term second order in (v/c) and thus negligible for stellar aberration.

There is an intermediate stage in the history of astronomy in which the Gettier-like character of Bradley's discovery was explicitly recognized. Bradley's Newtonian addition of velocities applied as long we could conceive of light as the rapid propagation of Newtonian light corpuscles. In the nineteenth century, this Newtonian theory of light was replaced by a wave theory. The Newtonian addition of velocities could no longer be applied in a straightforward way to the direction of wave propagation and it became quite unclear as to how the effect of stellar aberration was to be recovered. Arthur Berry, writing a history of astronomy at the end of the nineteenth century, expressed puzzlement over Bradley's success. His account captured rather perfectly its Gettier-like character with the verbiage I have emphasized below in his remarks (1898, p. 265):

> The curious inference may be drawn that, if the more correct modern notions of the nature of light had prevailed in Bradley's time, it must have been very much more difficult, if not impracticable, for him to have thought of his explanation of the stellar motions which he was studying and *thus an erroneous theory led to a most important discovery*.

The difficulty for the wave theory had been resolved in electrodynamic theory, though apparently not to Berry's satisfaction, by H. A. Lorentz' 1895 theorem of corresponding states. The mathematics of that theorem was then incorporated by Einstein into his special theory of relativity. Stellar aberration is, we now recognize, a relativistic effect, arising from the relativity of simultaneity.

11.2 Dalton's Discovery of Atomic Theory

John Dalton's (1808, 1810) *New System of Chemical Philosophy* is celebrated as the beginning of modern atomic theory in chemistry. Dalton proposed facts that we still believe lie at the foundations of chemistry: that each element is composed of identical atoms of that element; and that the atoms combine to form the familiar compounds of ordinary chemistry, such as water, muriatic acid, ammonia and so on.

Our concern is with the methods that Dalton used to arrive at his results. What is clear is that whatever methods he was using, they were quite unreliable and produced false results throughout his work. Part I of Dalton's *System* (1808) was concerned with heat and the physical properties of bodies formed from atoms. He determined (Chapter II) that the atoms of gases consist of particles enveloped in atmospheres of heat, where the latter was conceived as an element, caloric, following Lavoisier's system of elements. The enveloped cores are spherical and thus he could characterize a gas as arranged as a neat pile of shot. He used this static model to account for the elastic behavior of gases under changes of pressure and temperature, in direct conflict with the kinetic theory of heat that emerged a few decades later.

Dalton's treatment of the chemical properties of material in Part II (1810) consisted largely of extensive reports of chemical experiments and what could be inferred from them. What little he could infer of the details of the atomic constitutions of these chemicals was, again, largely incorrect. He did not recognize that elemental hydrogen, oxygen, nitrogen and so on consist of diatomic molecules. He had no developed theory of chemical bonding to account for how compounds formed. In its absence, he determined that water is a compounding of *one* atom of hydrogen with *one* atom of oxygen. In modern language, that amounts to the formula HO and not the familiar H_2O.

Methods dependent on this factually erroneous portrayal of materials are quite unreliable. A specific example are the rules he developed in Part I, p. 214, to determine the ratios of atoms in each compound. Those rules were responsible for the determination of water as HO. They enjoin us to use the simplest ratios possible; and lead directly to the one-to-one ratio of hydrogen to oxygen.[22]

With methods so unreliable, there can be no expectation that Dalton could secure the credit for introducing modern atomism to chemistry. My sense is that the key factor lay in his timing. The idea that all matter is atomic has been proposed routinely since antiquity. All such proposals were ill-fated prior to Lavoisier's epochal determination of the elements in his *Elements of Chemistry* (1790). After it, success is assured for the first chemist to propose an atomic constitution for the chemical element, no matter how flawed their other results may be.

11.3 Carnot's Discovery of the Foundations of Thermodynamics

Sadi Carnot's (1824) *Reflections on the Motive Power of Fire*, is, in my view, one of the greatest works of science of all eras. Its pre-eminence lies in it not just laying the foundations of the new science of thermodynamics, but in introducing a novel form of theorizing that is

[22] See Norton (2024, Ch. 11) for an account of the half century of work needed to establish the correct atomic weights and molecular formulae.

distinctive of thermodynamics.[23] He recognized that it is possible to have a general theory of efficiency covering heat engines of all types, not just steam engines, and no matter their operating fluids or other details. The maximum capacity of these heat engines to produce motive power for each unit of heat supplied was determined solely by the temperatures of the hot source from which it drew heat and cold sink to which it discharged heat. That maximum capacity was realized by reversible operations, that is, ones that could run equally in the forward or the reverse directions. The key insights leading to the later formulation of the second law of thermodynamics are routinely and correctly attributed to Carnot's work.

Nonetheless, Carnot's successes are a Gettier-like case. For Carnot's analysis was carried out within the caloric theory of heat, using methods that routinely led to incorrect results. In the later thermodynamic theory, a heat engine *converts* the energy of heat into work energy. In Carnot's analysis, caloric, that is heat, is a conserved substance. There is no conversion of heat to work. It is solely the *transfer* of heat from hot to cold that engenders motive power. It is then inevitable that much of Carnot's analysis produced results that we now think incorrect. Most notably, because he assumed that caloric is conserved, the quantity of heat drawn from the heat source must match the quantity discharged to the heat sink.

How did Carnot achieve his successes using an understanding of heat engines that is so at variance with the modern view? Here are the saving factors he might not have foreseen. The general thermodynamic principles that were carried forward to modern thermodynamics prove to be insensitive to the particular assumptions that Carnot made erroneously about heat and heat engines. A quite specific one of his erroneous assumptions made his successes possible. For, if we assume that heat engines convert heat energy into work, it would be natural to conclude that an engine that converts heat fully into work is the most efficient. There would be no point in including a heat sink for waste heat in the analysis of the idealized, most efficient engine. Modern thermodynamics recognizes that no heat engine can operate without a colder heat sink to which to discharge waste heat. It is a necessity of its operation and, without it, an analysis of the style of Carnot cannot be carried out. Carnot's assumption about the conservation of caloric, however, requires the heat sink and made it a fixture in the analyses of later thermodynamicists.

[23] For an analysis of how he was able to achieve this feat, see Norton (2022).

11.4 Thomson's Discovery of the Electron

J. J. Thomson (1897) reported the results of his experiments on cathode rays. They were, he announced, beams of negatively charged particles. This announcement is now celebrated as the discovery of the electron, the first particle of the many to come in modern particle physics. In his paper, Thomson sought to determine whether these cathode rays are waves propagating in the ether, as Hertz, Lenard and others had supposed; or whether they are streams of charged particles. Thomson found that the rays are deflected by electric and magnetic fields exactly as would a stream of particles, if each particle carried a negative electric charge. The deflections allowed him to determine a unique mass to charge ratio that is characteristic of the particles that would soon come to be called "electrons."

The case against wave propagation specifically was laid out elsewhere in Thomson's writings. (For a more detailed narrative of the evidential arguments deciding between waves and particles, see Norton, 2021, Ch. 9.) The decisive consideration was that cathode rays are deflected by a homogeneous magnetic field. There was no mechanism in the theory of waves in the ether that would allow this. Charged particles, however, are deflected by such a homogenous field by the Lorentz force law.

Thomson's methods were those dictated by late 19th century electrodynamics. Using them, he made his case well. What he could not foresee was that the quantum theory of the early 20th century would find cathode rays to be the propagation of a quantum electron wave. His case for electrons as localized particles and not waves collapsed. A quantum electron wave can be deflected by a homogeneous magnetic field. Thomson's 19th century methods were ill-adapted to the quantum physics of the 20th century.

However, in spite of the failings of his methods, enough of his conclusions for his particles survived. The unanticipated saving factor is that the deflections he attributed to the behavior of particles were recovered in the quantum theory by Ehrenfest's theorem. It showed that these are the motions of a quantum wave, if the quantum wave is sufficiently localized in space. That then allowed the mass and charge Thomson had determined for his particles to become the corresponding parameters of the quantum electron.

11.5 Bohr's Discovery of the Quantum Character of Atoms

Niels Bohr's (1913) "On the Constitution of Atoms and Molecules" depicts hydrogen and other atoms as consisting of positively charged nuclei, surrounded by one or more electrons in stable, discretely spaced energy states. The electrons can transition up and down between these stable energy states by absorbing or emitting the energy difference between the states in electromagnetic waves. These results provided the basis for the quantum theory of atoms that built upon Bohr's findings.

Even though the results just described are foundational in the later quantum theory of atoms, Bohr's methods were quite ill-adapted to the new quantum theory. His analysis employed the methods and results of 19th century electrodynamics. Bohr's electrons were not the quantum waves that surround the nucleus of an atom in diffuse clouds. His were still the localized particles of Thomson and they orbited in circular or elliptical orbits. Bohr's electrons jumped instantaneously between their stable energy states. Their transitions were not governed by the Schrödinger evolution of the yet to be developed quantum electrodynamics. The electromagnetic energy absorbed and emitted by Bohr's electrons were in the form of classical waves, not as the light quanta Einstein had introduced in 1905 that became the foundation of the quantum treatment of electromagnetic systems.

Worse, Bohr's 19th century methods were employed inconsistently. The classical theory insisted that an orbiting charge could not adopt a stable state, but would spiral into the atomic nucleus. Bohr posited stability. The classical theory allowed that an electron could adopt a continuous set of energy states in its orbiting the atomic nucleus. On the strength of the evidence of discrete atomic spectra, Bohr inferred otherwise.

In sum, Bohr's methods were unreliable when judged by the later quantum theory. Most of his results were incorrect. But within the compendium of false results lay a few that became the foundation of the new quantum theory of atoms. These were the ones that reflected the characteristics of the atomic spectra that formed the empirical basis of his theory: the stability and energetic discreteness of his electron states reflected the stability and discreteness of the observed atomic spectra. The unanticipated saving factor is that just these were the results that would survive into the new science to come.

11.6 Einstein's Discovery of the General Theory of Relativity

For over seven years, from 1907 to 1915, Einstein expended at times exhaustive efforts to generalize his special theory of relativity. The result was his general theory of relativity. It was celebrated then and still today as our most successful and most profound relativistic theory of gravity. Private notes, correspondence and intermediate publications have provided us extensive documentation of the stages of Einstein's thinking and the methods he used. On his later summary, his methods were controlled by three principles: the idea that the relativity of motion was to be extended to acceleration; Mach's principle, that the inertia of a body results from an interaction with the other masses of the universe; and the principle of equivalence, which asserted the identity of gravity and acceleration.

What makes Einstein's discovery Gettier-like is that two of these ideas, so foundational to his discovery, prove to be incompatible with his final theory. His theory does not generalize the relativity of motion to acceleration; and his rewriting of the generalized principle as a principle of general covariance was shown by Kretschmann in 1917 to be vacuous. Mach's principle also failed in the theory. Einstein disowned it for its dependence on an ontology of interacting masses at variance with the notion of unified field theory that dominated his later thinking. In so far as Einstein's methods were controlled by these principles, we could not now expect them to lead him successfully to his final theory. Yet he did arrive at the theory.[24]

In retrospect, I believe his success depended at least in part on his insistence on further conditions that are sufficiently unremarkable as to elude inclusion in celebratory histories, but are, in retrospect, powerful guides toward the final theory. I have explored some of them in Norton (2020). During the crucial period of the discovery of the theory, Einstein could not have anticipated that these were the saving factors and that the principles he prized would fail in the theory.

11.7 Hubble's Discovery of the Hubble Law

Hubble's (1929) celebrated discovery of the recession of the galaxies was recounted in Chapter 10 here. Its main result is a linear relationship between the velocity of galactic recession and the distance to the galaxy. In retrospect, his results were a mix of successes and failures. The linearity of the relation survives and carries his name. Hubble's determination of the constant of proportionality, 500 km/sec/megaparsec, proved to be 6 or 7 times greater than the later value of roughly 70 km/sec/megaparsec.

[24] For my early recounting of Einstein's discovery, based on private calculations in his Zurich notebook, see Norton (1984). For more details of the incompatibility of these principles with his theory, see Norton (1993).

The analysis in Chapter 10 recounts where Hubble's methods failed in the determination of the constant. The difficulty was that he did not know that there is more than one type of Cepheid variable star that formed the basis of his estimates of the distanced to the galaxies. He had presumed the wrong type in estimating distances to nearby galaxies and this compromised all this remaining estimates. The saving factor that Hubble could not anticipate was a result of later investigations. There is only one other type of Cepheid to be considered. Accommodating it into the analysis merely requires a recalibration of Hubble's distance estimates by a constant factor. The correction altered the constant, but not the celebrated linearity of Hubble's law.

12. 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 of 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.

Kyburg, Henry E. (1961), *Probability and the Logic of Rational Belief*. Middletown, Connecticut: Wesleyan University Press.

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.

Norton, John D. (2024) *The Large-Scale Structure of Inductive Inference.* BSPS*Open*/University of Calgary Press.

Norton, John D. (manuscript) "Science Progresses with Gettier-Like Cases."

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

Peirce, Charles Sanders (1940) *Philosophical Writings of Peirce*. Ju. Buchler, ed., Routledge and Kegan Paul Ltd. Repr. New York: Dover, 1950.

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.

References for Gettier Cases in History of Science

Berry, Arthur (1898) *A Short History of Astronomy*. London: John Murray.

Bohr, Niels (1913) "I. On the Constitution of Atoms and Molecules," *Philosophical Magazine* Series 6, **26**, No. 151, pp. 1-25.

Bradley, James (1729). "An account of a new discovered motion of the fixed stars." *Philosophical Transactions of the Royal Society*, **35**, pp. 637–661.

Carnot, Sadi (1824). *Réflexions sur la Puissance Motrice du Feu*. Paris: Bachelier. Trans., *Reflections on the Motive Power of Heat*. Wiley. 1897.

Dalton, John (1808) *A New System of Chemical Philosophy*. Part I. Manchester: R. Bickerstaff.

Dalton, John (1810) *A New System of Chemical Philosophy*. Part II. Manchester: R. Bickerstaff.

Hubble, Edwin (1929) "A Relation between Distance and Radial Velocity among Extra-Galactic Nebulae." *Proceedings of the National Academy of Sciences*, **15**, pp.168–73.

Lavoisier, Antoine-Laurent (1790) *Elements of Chemistry.* Trans. Robert Kerr. Edinburgh: William Creech.

Norton, John D. (1984) "How Einstein Found His Field Equations: 1912-1915," *Historical Studies in the Physical Sciences*, **14,** pp. 253-315.

Norton, John D. (1993) "General Covariance and the Foundations of General Relativity: Eight Decades of Dispute," *Reports on Progress in Physics*, **56**, pp.791-858.

Norton, John D. (2020) "Einstein's Conflicting Heuristics: The Discovery of General Relativity," pp. 17-48 in *Thinking about Space and Time*: *100 Years of Applying and Interpreting General Relativity*. *Einstein Studies*, Volume 15. C. Beisbart, T. Sauer, C. Wüthrich (eds). Cham, Switzerland: Birhäuser/Springer Nature.

Norton, John D. (2022) "How Analogy Helped Create the New Science of Thermodynamics," *Synthese*. **200**, article 269, pp. 1-42.

Thomson, Joseph J. (1897) "Cathode Rays," *Philosophical Magazine,* **44,** pp. 293–316.