

## Introduction: Meet the Gettier Problem

---

Stephen Hetherington

### I.1 Introduction to an Introduction

How do philosophical movements or eras begin? How do they end? All too often, they fade away with a whisper or two – dying quietly, unnoticed, supplanted through distraction by new ideas or issues or fashions. They might enter the world unobtrusively, too – barely acknowledged, recognized only with hindsight. Not always, though: sometimes, a philosophical movement's or era's end or beginning can be dramatic, even explosive.

That is how post-Gettier epistemology began – with a deafening “kapow!” Epistemology had been proceeding along familiar paths: business as normal. In that spirit, 1962 came and went, as 1961 had done, like 1960 before it, following a calm 1959 . . . .<sup>1</sup> Then 1963 arrived. Fresh issues of that year's array of academic philosophy journals appeared. And then, a few months into the year, without fanfare, one of those journals – *Analysis*, an excellent journal dedicated to publishing concise philosophy papers – gave us just such a paper by Edmund Gettier, a young American philosopher. Epistemology has looked quite different since that June 1963 issue of *Analysis*. Gettier's paper was called “Is justified true belief knowledge?” Having asked the question, he answered it: “No.” And decisively so, in the view of epistemologists *en masse* – then and since. The year 1963 was soon recognized by epistemologists as having been a time of transition for them. Pre-Gettier epistemology had ended; post-Gettier epistemology was under way.

It still is. We are the epistemological heirs of that exciting moment. We are post-Gettier epistemologists. Indeed, we are post-Gettier philosophers, given epistemology's historical importance within philosophy – and given how Gettier's argument struck so many as a clearly successful instance

<sup>1</sup> For a survey of that stage of epistemology, a survey untouched by the torrent that was about to descend, see Hill (1961).

of philosophical reasoning. Gettier initiated what swiftly became a powerfully influential and wide-reaching wave of epistemological inquiry. That wave is still surging – sometimes gently, at other times looming. Not all epistemologists have ridden it, and some fight it even while riding it; many others, though, embrace its flow and power.<sup>2</sup>

More prosaically, Gettier and the impact made by his paper even changed philosophy's everyday lexicon. When someone's name has become as lexically embedded as has happened for Gettier's, this is clearly a mark of professional respect. We now have such philosophically everyday words as "Gettiered," "unGettiered," "Gettierized," "Gettier-proof," and "Gettier cases." Respect, indeed.

## 1.2 The Gettier Argument

In his paper's two-and-a-half pages, Gettier presented a single succinct argument. It was built around two imagined stories – two counterexamples. These were directed by Gettier against what, he also argued, was a discernible form of thesis that philosophers had adopted when trying to understand the nature of knowledge. This form of thesis was, according to Gettier, a form of definition. He claimed that it was a general form instantiated by various specific attempts to understand what it is for someone to have some knowledge – attempts that had begun in ancient Greek philosophy and continued into the 1960s, when Gettier was writing.<sup>3</sup>

This section will briefly describe the key elements of that picture from Gettier – the kind of knowledge that philosophers had been attempting to define, the general form of definition that they had used, and how the counterexamples were to be brought to bear against that general form of definition.

<sup>2</sup> This is a personal note. I was a second-year undergraduate when I first met Gettier's paper. I was introduced to it by William Lycan, then visiting at the University of Sydney from Ohio State University, where his colleagues included George Pappas and Marshall Swain, the editors of a terrific and just-then-appeared collection of papers (1978) on the influence of Gettier's paper within epistemology. The book, like the course, like the issues, felt fresh and energetic. It was one of the two textbooks – the other was Roth and Galis (1970), also excellent, which included Gettier's paper – for the course. This was my most personally exciting undergraduate course. Lycan conveyed so well the sense that these were issues on which real philosophical progress might be made, with new ideas swirling around us, close enough to be touched, and with much precise thinking on display. Gettier's challenge felt very real, very urgent.

<sup>3</sup> He first cited Plato, from the *Meno* and the *Theaetetus*, before mentioning Ayer (1956) and Chisholm (1957).

*The kind of knowledge.* Gettier was discussing what philosophers refer to as propositional knowledge. This is knowledge that we report or describe by according it a propositional content. For instance, he knows (propositionally) that he is standing in his office: “he is standing in his office” is a proposition with an indicative form, one that claims to report or describe a state of affairs (namely, his standing in his office). When philosophers say that they are telling us about the nature of knowledge, propositional knowledge is usually – for better or for worse – their intended prey.<sup>4</sup>

*The form of definition.* Philosophers seek understanding, and one of the phenomena of which epistemologists, in particular, seek understanding is knowledge. But not all philosophical understanding, we might feel, needs to be as strict and complete as is conveyed by a *definition*: “virtue = df . . .,” “evil = df . . .,” etc. Given philosophy’s difficulty, in practice we might well be pleased to attain even some philosophical understanding that is only partial and suggestive, still exploratory and allusive: even that could be a job well done. Nonetheless, if we do find an accurate definition of something that matters to us, we might consider this to be a job *very* well done. The nature of knowledge, like that of virtue, or of evil, say, probably matters to many of us. So, can we understand it insightfully and fully? Can we understand it by defining – accurately – what constitutes having some knowledge?

Gettier claimed at the outset of his paper that he was engaging with a philosophical tradition where, apparently, some of its most prominent practitioners were indeed confident of having achieved a definitional level of understanding of knowledge’s nature. Although those philosophers might disagree over details – and Gettier described those differing details for three such philosophers (Plato, A. J. Ayer, and R. M. Chisholm) – they could still be regarded as giving voice, even across the centuries, to a shared underlying form of definition.

What form were those definitions taking? Since Gettier highlighted it – in preparation for launching his counterexamples against it – this form of definition has generally been called the *justified-true-belief* definition of knowledge. It is a general picture of what knowing is (this is why I use the word “form”). Here is its basic idea:

<sup>4</sup> This is not universally so. In recent years, especially, there has been much epistemological focus on knowledge-how – knowing how to perform some (kind of) action. See, for example, Ryle (1946; 1949), Stanley and Williamson (2000), and Hetherington (2011a).

One's having some instance of knowledge *is* (by definition) one's having a belief that is true (accurate) and justified (somehow rationally well-supported).

This formula is intended as a definition of its left-hand-side – that is, of *what it is* to have an instance of knowledge. We are being told that – by definition (albeit in general terms) – any instance of knowledge is nothing more and nothing less than a well-justified true belief. Different epistemologists might then offer differing suggestions as to how to make this generic picture more detailed; the general picture remains, though.

*The counterexamples.* I mentioned that epistemology now includes the concept of a *Gettier case*. That concept is traceable to Gettier's two counterexamples.<sup>5</sup> I also mentioned that he used them in order to question the justified-true-belief definition. But he was not questioning every aspect of that definition. The definition can be treated as a conjunction of two claims – only one of which Gettier was questioning. Here are the two claims:

*JTB-Necessity.* Being a justified true belief is always needed for being knowledge: each of those three elements – justification, truth, and belief – is always required if an instance of knowledge is to be constituted.

*JTB-Sufficiency.* Being a justified true belief is always sufficient for being knowledge: the combination of those three elements – justification, truth, and belief – is always enough to constitute an instance of knowledge.

Gettier's counterexamples were intended to reveal the falsity only of the latter thesis – JTB-Sufficiency. Each of his two imagined situations tells a story about a particular person, Smith. Setting aside the colorful details of each case, the following combination of circumstances obtains in each: Smith forms a belief that is true, is justified, yet is not knowledge. At any rate, that is how Gettier interpreted these situations. Nor was he alone in having that reaction. The vast majority of epistemologists – from 1963 to now – have agreed with that sort of interpretation of these and like situations. Epistemology thus formed the concept of a Gettier case – and soon regarded it as a powerful concept. Each such case (we are now routinely told) is a successful counterexample to “the sufficiency half” of the justified-true-belief definition of knowledge.

<sup>5</sup> This is not to say that no similar cases had ever appeared previously within philosophy. It was the cases plus the systematic use to which Gettier and others put them that have given us this concept. For some discussion of earlier cases with a similar internal structure, see Shope (1983, pp. 19–21) and Hetherington (2016a, p. 5).

And there have been many such cases, each of them inspired by Gettier. The book's cover hints at one of them – the famous sheep-in-the-field case. This one is from Roderick Chisholm (1966, p. 23n22; 1977, p. 105; 1989, p. 93). It imagines someone's gazing at a field, noticing what looks like a sheep, and forming the belief that there is a sheep in the field. The belief is true, although only because there is a sheep far away in this same field, hidden from this person's view. (What is being seen is a dog that is disguised as a sheep.)

### 1.3 The Gettier Argument's Apparent Significance

Suppose that Gettier was right in his answer to his article's own titular question: suppose that the justified-true-belief definition is mistaken, because JTB-Sufficiency is mistaken. This looks like a powerful blow to have struck against the justified-true-belief definition of knowledge. Is this also enough to make Gettier's argument significant? That might depend on how important is it to see the falsity of that definition (if indeed it is false). Might it depend, too, on what might be revealed about epistemology, or even about philosophy more broadly, by seeing the definition's falsity in the particular way in which Gettier revealed its falsity (if indeed he did)?

The justified-true-belief definition did strike epistemologists as having been an important philosophical insight – albeit one that suddenly, just like that, Gettier had refuted. At least part of that importance was historical; or so Gettier apparently thought. As Section 1.2 noted, he began his paper by arguing that instances of this general way of defining knowledge had been with us since Plato – almost the beginning of Western philosophy – and were still with us. Epistemologists in 1963 did not question that view of the historical significance of his argument's target.<sup>6</sup>

There are other potential sources of significance, too. Suppose that Gettier was not reacting against a conception of knowledge that had in fact been guiding all epistemologists for all of those centuries. If so, perhaps he was oversimplifying the history of epistemology in this respect. Is it possible, for instance, that he was discussing only one of *two* forms – broadly characterized – that knowledge might take? Ultimately, philosophy might need to choose between fallibilist and infallibilist conceptions of knowledge; perhaps some of philosophy's historically significant figures adopted

<sup>6</sup> But this has subsequently been questioned. For example, see Kaplan (1985), Dutant (2015), and Le Morvan (2016).

infallibilist conceptions of knowledge. Seemingly, though, Gettier's argument was only about fallibilist conceptions of it. (He introduces his argument by *saying* that the sort of justification that he will be discussing leaves open the possibility that the belief being justified is false, even while being justified. That is a fallibilist way of thinking about knowledge's justification component.) This might undermine the *historical* reach of Gettier's argument.

Still, the argument could remain conceptually important, particularly if independent reasoning tells us that some sort of fallibilist conception of knowledge is what we should be adopting. For then we could regard Gettier as having shown, dramatically and decisively, that we need to say more, if we are to understand fallible knowledge (and thereby knowledge at all), than that a belief is knowledge if and only if it is true and fallibly justified. This could be a spur to further fallibilist thinking about knowledge's nature.

Philosophers also saw significance in Gettier's argument precisely because it seemed to show that real philosophical insight – a real result, even if a critical or destructive one, about how to define knowledge – could be attained; and it could be attained simply yet decisively. This was methodologically inspiring for those who might otherwise feel that philosophy is never a way to gain real results, real insights. Thanks to Gettier, maybe there can be real philosophical progress, achievable purely by thinking!

#### 1.4 The Gettier Problem

Even so, philosophers soon viewed Gettier as having given to epistemology what they called the Gettier *problem*. It is useful to distinguish between two ways of interpreting this idea. We can talk of the *in-principle* Gettier problem and of the *in-practice* Gettier problem.

*The in-principle Gettier problem.* This was the intellectual puzzle, the philosophical conundrum, posed by Gettier. Could it be solved? The usual formulation of it was along these questioning lines:

How should knowledge be defined, if not merely as justified true belief? What more is needed, if a belief – presumably a true and justified one – is to be knowledge, given that Gettier cases reveal the insufficiency, for describing something as knowledge, of describing it merely as a belief that is true and justified?

The challenge of answering these questions is the Gettier problem, understood purely as an intellectual challenge.<sup>7</sup>

*The in-practice Gettier problem.* The intellectual challenge posed by the in-principle Gettier problem soon proved to be, as a matter of professional practice among epistemologists, exceedingly difficult to meet. One epistemologist after another would claim to have solved the Gettier problem (the in-principle problem). Some such claims would attract support; never enough support, though, to convince epistemologists as a group. And this was a problem of its own – a professional problem, a collectively lived problem of philosophers devoting much time to trying to solve the in-principle Gettier problem while never agreeing on success being achieved. “Why cannot we agree on how to solve the Gettier problem?” That question was ever present. The result is that, by now, we have long had the in-practice Gettier problem.<sup>8</sup> Individual Gettier cases have often seemed to offer clear indicators of something going awry within them; and hence of what to fix within them if knowledge was to have been present within the particular case’s circumstances; hence, too, of what knowledge even is. The in-practice problem, however, is that all of these proposed solutions have continued being susceptible to Gettier cases – that is, to further instances of the (in-principle) Gettier problem. The (in-practice) Gettier problem thus grew and grew. It is still growing.

## 1.5 Post-Gettier Epistemology

The in-practice Gettier problem is therefore still with us: rampant disagreement persists among epistemologists as to how to solve the in-principle Gettier problem. Is that also a worrying state of affairs? Should epistemologists be perturbed, for example, at the possibility of there being something problematic about the practice of epistemology itself? Is the Gettier problem’s in-practice version an indication of a wider problem – about epistemology as a whole, about epistemology’s potential ever to solve its own (in-principle) problems?

Not all epistemologists are worried. Many are encouraged by the range of independently interesting epistemic concepts and theses – about knowledge, but also about associated phenomena, such as belief and justification – that seem to have emerged from epistemology’s post-Gettier efforts. Post-Gettier

<sup>7</sup> The justified-true-belief definition was sometimes called “JTB” for short. Correlatively, the in-principle Gettier problem was sometimes referred to as the quest to ascertain the correct “JTB+.”

<sup>8</sup> Lycan (2006) refers to it as “the Gettier Problem problem.”

epistemology has struck many as being quite fruitful – in a range of ways, regardless of whether it has reached an in-practice solution to the in-principle Gettier problem.

The independently interesting ideas to emerge from post-Gettier epistemology have also included potential morals about epistemology as a whole – meta-epistemological thoughts. We might wonder about the methodology on display in Gettier’s argument – the use of imagined or intuited counterexamples to a definition that aspires to being a full understanding of some phenomenon such as knowledge. Should we always seek a definition, when seeking philosophical understanding? What alternatives are there? How reliable are intuitions about how to “read” particular cases (either real or imagined)? Are such intuitions decisive? If not, how strongly should they guide our thinking? And so on.

## 1.6 This Book’s Structure

This book is organized to reflect (i) the intellectual structuring of the in-principle Gettier problem – how Gettier’s article introduced its two counterexamples – and (ii) the post-Gettier professional realities that have constituted both the optimistic and the possibly worrying elements of the in-practice Gettier problem. Accordingly, the book has the following three parts (only informally, I should add).

As this introduction has explained, Gettier’s paper began by situating its target (the justified-true-belief definition of knowledge), and thereby the paper itself, within what it claimed is a clear historical tradition. Gettier then began his argument by highlighting two key epistemic principles, describing each as vital to the counterexamples that he was about to describe. Chapters 1 (Charity Anderson) and 2 (Claudio de Almeida) are therefore about those principles – respectively, *fallibilism about knowledge’s justificatory component* and *justificatory closure* (Gettier 1963, p. 121):

First, in that sense of “justified” in which S’s being justified in believing P is a necessary condition of S’s knowing that P, it is possible for a person to be justified in believing a proposition that is in fact false. Secondly, for any proposition P, if S is justified in believing P, and P entails Q, and S deduces Q from P and accepts Q as a result of this deduction, then S is justified in believing Q.



Each of those pivotal principles has attracted much attention. Partly, this is because, once noticed, they have been deemed to possess independent epistemological interest. But the attention paid to them is also due to their having helped to generate the internal details of Gettier's two cases – and thereby the in-principle Gettier problem.

What I think of as the book's second part is on various epistemic concepts upon which, *especially* because of Gettier's challenge, epistemologists have focused. Like those two key epistemic principles of Gettier's, mentioned a moment ago, these concepts have been widely regarded within epistemology as independently interesting. But they have indeed been part of the in-practice Gettier problem's taking shape since 1963. As Section I.5 mentioned, once epistemologists began trying to understand Gettier cases with the goal of reaching an improved definition of knowledge, a diversity of seemingly associated epistemic concepts were introduced, highlighted, sharpened, sometimes discarded, often revived, etc. So, Chapters 3 through 8 are about what have probably been the most epistemologically discussed and influential of those concepts:

Chapter 3 (Clayton Littlejohn): Gettier and evidence (including disjunctivism);

Chapter 4 (Rodrigo Borges): Gettier and externalism;

Chapter 5 (Delia Belleri and Annalisa Coliva): Gettier and context (including contextualism and contrastivism);

Chapter 6 (Duncan Pritchard): Gettier and luck (including epistemic safety);

Chapter 7 (Kelly Becker): Gettier and epistemic sensitivity;

Chapter 8 (John Greco): Gettier and cognitive virtues;

Chapter 9 (Ernest Sosa): Gettier and epistemic wisdom (a synoptic vision of what we might learn from post-Gettier epistemology, by an epistemologist who was involved in such epistemology from the outset).

Finally, the book's third part discusses some fundamental *methodological* questions. Gettier posed such an apparently powerful challenge to what would otherwise, it seems, have felt like such an unassailable definition of knowledge that questions were bound to arise about what it was about epistemological thinking that had – mistakenly, it suddenly seemed to epistemologists – made that definition appear so impregnable. Thanks to Gettier, should new forms of epistemological thinking be envisaged? In short, some potentially significant methodological questions have been posed within – and about – post-Gettier epistemology:

Chapter 10 (Patrick Rysiew): Gettier and conceptual analysis (including the rise of knowledge-first epistemology, due mainly to Williamson [2000]); Chapter 11 (Elijah Chudnoff): Gettier and intuition (including the rise of experimental philosophy, due especially to Weinberg, Nichols, and Stich [2001]); Chapter 12 (John Turri): Gettier and experimental philosophy, in more detail (given how closely its origins were tied to talk about Gettier cases); Chapter 13 (Stephen Hetherington): Gettier and explicability (whether any Gettier case can be understood so as to support the usual interpretation of its central belief as failing to be knowledge).<sup>9</sup>

It is unquestionable that the Gettier problem has been a philosophically rich field of study, as this book demonstrates. What is the Gettier problem's future? Will it be solved? (Has it already been solved?) Will it continue being so epistemologically influential? What would *post*-post-Gettier epistemology be like? Tantalizing questions, all.

<sup>9</sup> For fuller introductions to the nature and history of the Gettier problem, see Shope (1983), Lycan (2006), and Hetherington (2011b; 2016a, ch. 1; 2016b).