1 *Introducing Gettierism*

1.1 The Year of Gettier

Within epistemology, 1963 was perhaps most strikingly "The Year of Gettier" – Edmund Gettier. It was the year neither of his birth nor of his death. It was the year when he *changed* epistemology, from then until at least now. It was the year that he published a single paper potentially of much epistemological consequence and actually of much epistemological moment. It was the year of "That Paper."

Gettier's paper asked how we should define knowledge – the kind of knowledge called "knowledge-that" or "propositional knowledge" by philosophers. His paper caused epistemologists speedily to accept that knowledge could not, after all, be defined quite as they had presumed or claimed it should be defined. Unworriedly, they had been talking of knowledge along these lines:

> Any instance of knowledge is a belief. But not just any belief is knowledge: a belief's being a belief is not enough to make it knowledge. What is also needed is the belief's being true, along with its being well supported or justified in some way that bears favourably and purely upon the belief's being true.

This account was sometimes formulated more succinctly, as a conceptually reductive definition. The definition was to be conceptually reductive, in virtue of its *definiens* using only concepts other than that of propositional knowledge, in order to describe fully and only what it is to be an instance of such knowledge and thereby to instantiate the concept of that sort of knowledge. Thus, it would look like this (where "*p*" stands for any given proposition):

> To know that *p* is, by definition, to have a well supported or justified true belief that *p*. That is *all* there is to knowing.

But Gettier in 1963 – epistemologists at the time believed, and still believe, that he – revealed the falsity of that otherwise highly tempting conception of knowledge. His brief paper posed a question – "Is justified true belief knowledge?" – before answering, seemingly incontrovertibly, with "No." By this, he meant that "justified true belief" is not all that there is to a *definition* of knowledge. Epistemologists then and since concurred with him. Post-Gettier epistemology was upon us – and remains so.

The term "post-Gettier epistemology" could be misleading. We do now live in a post-Gettier epistemological time. Yet not all current epistemology *engages* either directly or indirectly with details or even the spirit of Gettier's paper. Nor does all current epistemology reflect upon either the phenomenon or the concept of knowledge. Nonetheless, much current epistemology does do so by engaging directly with aspects and implications of Gettier on knowledge. And, more to the immediate point, almost all current epistemology presumes, accepts, or seeks to build upon the *correctness* of Gettier's answer to his guiding question.

I cannot put the latter point more bluntly than in these terms:

> Gettier is standardly credited with having established a genuine epistemological *result*. Accordingly, the role of Gettier-aware epistemology, when working in this conceptual neighbourhood, is standardly taken to at least presuppose that result – and certainly not to *question* it or to inquire seriously into *whether* it is correct.

David Lewis put that same point more forcefully (1983: x):

> Philosophical theories are never refuted conclusively. (Or hardly ever. Gödel and Gettier may have done it.)

Lewis spoke for many, almost all, epistemologists when attributing to Gettier that achievement – yes indeed, a genuine achievement, nothing less than the incontestable overthrow of what had been a long-standing and epistemologically central thesis or presumption about knowledge. And the word has long since spread beyond epistemology: philosophers in general accord Gettier that achievement. For example, introducing his book on well-being and death, Ben Bradley (2009: xiv) says that

> if a philosopher is asked to point our examples of truths that philosophers have conclusively established, the first would probably be [courtesy of Gettier] that justified true belief is insufficient for knowledge.

Hans-Johann Glock, when trying to ascertain what analytic philosophy is, contemplates the widespread "waning" of the "ambition" to find "analytic definitions" of concepts. He reminds us of "the most spectacular case [of this waning], following Gettier's classic criticism" (2008: 158). And David Chalmers (2012: 13–15), in his quest to resurrect or at least strengthen some elements of philosophy's analytic enterprise, cites Gettier's contribution as an exemplar of what he calls a priori scrutability. The latter is a pivotal component in Chalmers's ambitious "metaphysical epistemology" – his "epistemology in service of a global picture of the world and of our conception thereof" (2012: xx). He trusts *that* strongly in the usual reaction to Gettier.

And those few remarks are entirely representative of so many more. The philosophical verdict, it seems, is in: Gettier won; an apparently traditional definition of knowledge lost; and epistemology is thereby wiser. Real conceptual progress has been made, it is widely assumed.

Nonetheless, this book will be questioning part of that familiar picture. Just how much wiser *is* epistemology, courtesy of Gettier's supposed result? I will argue that, when post-Gettier epistemologists have claimed – *even* in a preliminary or merely suggestive or illuminative way – to explain or understand what underlies his putative result, there is a fundamental and systemic reason why those claims are false. My point will not be that pre-Gettier epistemologists, in contrast, understood knowledge fully. Nor will I be saying that Gettier was mistaken in all aspects of his claimed result. But this book will show how, even if Gettier was correct in his core conclusion, epistemologists have not understood at all *how* he was correct. And again, this is so even when that "how" is allowed to remain broad or suggestive, intending only to convey "the basic idea" as to why Gettier was correct.

As a result, we will see why epistemology need never have come to include what is sometimes called "Gettierology" – a Gettier-inspired complexity catering to the standard assessment of what Gettier's paper established and of how epistemologists should consequently approach the task of explaining, understanding, or defining knowledge. Perhaps – we will also see – there is an *alternative* way for an epistemology of knowledge to accommodate what, if anything, was correct in Gettier's challenge.

1.2 Gettierism Introduced

The epistemological orthodoxy against which this book will argue deserves a fresh descriptive name. Although the term "Gettierology" designates Gettier-inspired study – a particular range of epistemological content – it does not capture the unquestioning *acceptance* by epistemologists of the core of that content: specifically, the main premise guiding such study. And that acceptance – the *epistemological embrace* of Gettier's putative result, an embrace routinely claimed to be epistemically secure and perhaps even theoretically explicable – will be receiving this book's critical attention. So I offer a further name for my topic. I do not want that name to be at all disrespectful (a taint sometimes affixed to the term "Gettierology"), yet I will be presenting arguments which might collectively *seem* to reflect a lack of respect for what Gettier contributed to epistemology. Accordingly, before immersing myself in those arguments, let me record that one of my most prized philosophy possessions is a copy of the issue of the journal *Analysis* – volume 23, number 6, June 1963 – in which Gettier's paper was published, and let me make a suggestion that reflects what is actually my contrary opinion of his epistemological contribution – which is that his paper is stimulating, clever, and beguiling. I suggest the title of "Gettierism" for the associated epistemological orthodoxy – that is, for what will be my focus in this book. This is the orthodoxy's title, and becoming clearer on its content will occupy us for much of this chapter. Gettierism springs forth from – but is far from exhausted in content by – this central tenet: *Gettier was right*.

About *what* was he right, though, according to Gettierism? The following: he revealed the falsity of a definition – a somewhat generic justified-true-belief definition of knowledge. Gettier's immediate conclusion was that a *sufficient* condition for a belief's being knowledge had not been adequately delineated by that epistemologically favoured form of definition, for a belief could be true and well justified without being knowledge. How did Gettier claim to establish this shortcoming in that form of definition? Astute counterexampling was his method. He described two possible situations that functioned as decisive counterexamples to the justified-true-belief definition of knowledge.

Phrases along the lines of "Gettier showed that ..." or "Gettier proved that ..." are now *de rigeur* for epistemology graduate

students to learn in this setting. For most, it is apparently an effortlessly gained skill – an *easy* epistemological lesson. What has been an orthodoxy for quite a while – Gettierism – is becoming only more so. And it has long since moved beyond talking just of a form of *definition* of knowledge. These days, many epistemologists would report Gettierism's result also in these broader terms:

> Gettier showed that if we are to understand the nature of knowledge epistemologically, then we need to *conceive* of it in terms other than simply of justification, truth, and belief.

Gettierism thus becomes a meta-epistemological thesis, too. It is the thesis that Gettier was right *and* hence that a constraint upon our conceiving accurately of knowledge is that our conception will incorporate his contribution. Naturally, we may choose to say much *more* than that about knowledge's nature, nor need we engage in *detail* with Gettier's challenge; yet we are at least required not to *reject* what Gettier showed. This is a *minimal* implication of Gettierism. A less minimal implication is Gettierism's being regarded as apt motivation for moving epistemology beyond conceiving of knowledge merely as a justified true belief. And many post-Gettier epistemologists have sought to be helpful in this respect, by conceiving of knowledge in some way that reflects knowledge's being not only a justified true belief.

1.3 Gettier Cases Introduced

It is more than fifty years since Gettier startled epistemologists with his two counterexamples. In fact, it is now widely agreed that there were a few earlier intimations of the possibility of such tales being used to much the same conceptual effect as Gettier used them. Alexius Meinong and Bertrand Russell are usually mentioned at this point (e.g. Shope 1983: 19–21). Perhaps (see Stoltz 2007 and, more fully, Ganeri forthcoming 2017), similar cases were also discussed in earlier Indian and Tibetan epistemology. Even so, 1963 was when the idea behind such cases made its impact within Western epistemology, courtesy of Gettier. By now, we have long had his presumed lesson in our epistemological minds. Those two stories of his are familiar fare – old friends, soothingly instructive. These stories, and kindred others, have become known as "Gettier cases."

I will not attempt a full account of what makes a situation a Gettier case, even though it is epistemologically usual to talk in a general way of *Gettier cases as such*. This category is treated as encompassing Gettier's own two stories plus the sundry sufficiently similar ones that have been called "Gettier cases" when epistemologists have been claiming to explain or illustrate the thesis that it is possible for a belief to be true and well justified without being knowledge. So the category is rough-hewn and flexible. Still, epistemologists are *very* confident in talking of Gettier cases (for overviews of such talk, see Shope 1983; Lycan 2006; Hetherington 2011c; and Turri 2012a) – and in concluding each time that Gettier cases, whenever these do arise, exemplify the falsity of the justified-true-belief conception of knowledge.

Here, for specificity, is Gettier's Case I (1963: 122) – often called the "job/coins case":

> Smith has good evidence that Jones is about to be offered a particular job with a particular company. Smith also gains good evidence for Jones having ten coins in his pocket. Thereupon, Smith infers – with good justification – that the person who will get the job has ten coins in his pocket. But actually it is Smith, not Jones, who will get the job; and this will be a complete surprise to Smith. He is equally ignorant of there being ten coins in his own pocket; yet there are. Accordingly, even though Smith has managed to gain a well justified true belief (that the person who will get the job has ten coins in his pocket), is this belief of his knowledge?

Clearly not, said Gettier. And his view has since been shared by epistemologists in general: Smith does not know.

The same is true, *mutatis mutandis*, of other epistemologically popular Gettier cases. One of these is Roderick Chisholm's sheep-in-the-field case (1966: 23n22; 1977: 105; 1989: 93), presented here in a slightly embellished version and applied to you as the epistemic agent:

> From outside a field, you see within it what looks just like a sheep. You infer that there is a sheep in the field. And you are right because, hidden behind a hill, there is one. Yet what you are seeing is not a sheep; it is a dog, disguised as a sheep. Is your belief – that there is a sheep in the field – therefore not knowledge?

"Indeed so," choruses the standard epistemological reaction: you do not know.

Another interesting Gettier case is Keith Lehrer and Thomas Paxson's (1969) Tom Grabit case:

> When inside the library, you see a man take a book from a shelf, secrete the book inside his coat, then walk directly out of the building – seemingly stealing the book. He looks exactly like someone whom you know – Tom Grabit. Moreover, he *is* Tom Grabit and he *is* stealing that book. So you have a true belief, supported by good visual evidence, that Tom has stolen a book from the library. Yet Tom's mother proceeds to tell people (without your being aware of her doing so) that Tom was nowhere near that library at that time. She claims that his identical twin John was in the library. Is your belief – that Tom Grabit has stolen a book from the library – therefore not knowledge?

Again, the standard epistemological reaction has been that in such a circumstance the belief in question is not knowledge.

Finally, we must not forget Carl Ginet's famous case, first published by Alvin Goldman (1976), of the fake barns:

> Henry is driving in the countryside with his small son. 'Look, there's a barn,' says Henry, ever the helpful and enlivening father. Does Henry know that he is seeing a barn? In answering this, bear in mind that in fact – without his realizing it – Henry is looking at the only real barn present in that county at that time. Everything else in that county that would look like a barn, to anyone driving past it, is actually a fake barn – a papier-maché barn façade. So, Henry would have been deceived, had he been looking at one of these. As it happens, he is not deceived, because he is seeing the one genuine barn. But does his belief ('I'm seeing a barn') fail to be knowledge anyway?

Most epistemologists still reply "Yes" to that question: "Yes, Henry does not know."

With these and other cases in hand, then, let us attempt an initial and schematic taxonomic characterization of Gettier cases in general. We may say that, be it actual or possible, and be it realistic or not, any Gettier case is a situation where the following generic description is instantiated, usually in quirky detail:

> An epistemic agent forms a belief that is true and well-even-if-not-perfectly justified. Yet there is some kind of marked oddity in the surrounding circumstances – an oddity in how that justified true belief exists, an oddity bearing in a puzzling way upon whether the justified true belief is knowledge.

A helpful piece of further terminology ensues: a *belief* is Gettiered if and only if it is the belief at the heart of a Gettier case.

A cautionary point is already apt. Epistemologists often, even when simply introducing the idea of a Gettier case, prejudge pending interpretive details by saying from the very outset what form that oddity takes within the case. For example, Tim Black (2011: 187) *begins* his survey article by saying that "Gettier's counterexamples ... involve epistemic agents whose beliefs are true simply as a matter of luck." It is certainly tempting to point already to the presence and influence of such a feature within Gettier cases – within specific ones, even within Gettier cases in general. Nonetheless, we must resist the temptation for now. Section 1.6 will mention several ways that epistemologists have thought to be appropriate for conceiving of Gettier oddity. (And, of course, the most unfortunate of all such pre-emptive judgments occurs when a Gettier case is introduced by including this within the case's initial description: "and so the justified true belief fails to be knowledge." If describing the belief as not being knowledge is as manifestly and trivially appropriate as epistemologists are thereby taking it to be, then we should have no difficulty in explaining its being so. Later chapters, however, will suggest that such a project is not so epistemologically simple, to say the least.)

Note, too, that in introducing the generic idea of a Gettier case, I described each case's oddity as "bearing in a puzzling way upon whether the justified true belief is knowledge." This was done in order to maintain further interpretive neutrality at this early stage of my discussion. But (as Section 1.4 will explain) some may regard "bearing in a puzzling way" as too weak to do justice at all to the idea of a Gettier case. Can we strengthen this description without conceding what is yet to be argued, explained, or shown – namely, that the belief within a Gettier case is not knowledge? We could do so, if we wish, by replacing the phrase in question with one like this: "an oddity that has seemed to most epistemologists since 1963 to prevent the justified true belief's being knowledge." Otherwise, let us resist for now any further precisification of this element of the concept of a Gettier case.

It might also be helpful to mention a way of distinguishing between two kinds of Gettier cases (assuming, with most epistemologists, that the fake-barns case is a Gettier case). Thus, each Gettier case is either *helpful* or *dangerous* (Hetherington 1999; 2001a: chap. 3).

Helpful cases include Gettier's own two and the sheep-in-the-field case. The fake-barns case is a dangerous one, as is the Tom Grabit case. The difference between these two categories of Gettier cases may briefly be conveyed as follows.

Helpful Gettier Cases. In each such case, the justified true belief is attained partly *because* of the case's odd circumstance. The oddity helps to bring about that combination of truth, belief, and justification. Without quite that oddness, maybe that justified true belief would not have emerged. And so the person would not have been even *that* close (however close it is) to having knowledge that *p* by having this combination of justification, truth, and belief that *p*.

Dangerous Gettier Cases. Each of these cases likewise includes an odd circumstance. In this sort of case, however, the odd circumstance remains only a *threat* to the existence of the justified true belief (which combination, this time, comes about without the odd circumstance's help). If that circumstance was not present, the belief would be justified and true in the way it is – but with all else being normal, not odd. As it is, the belief remains *close* – due to that odd circumstance – to not being both true and justified in the way that in fact it is.

I should note that some epistemologists (e.g. Lycan 2006: 162–3), even while maintaining a commitment to Gettierism, do not agree that Henry in the fake-barns case lacks knowledge. Hence, they do not regard this *as* a Gettier case. Correlatively, they might set aside – as not really directing us to Gettier cases – the second category of case that I have described just now. I have mentioned the fake-barns case in this section, though, because most epistemologists continue regarding it as a Gettier case, from which knowledge is absent. For simplicity, however, when applying a general argument to a particular Gettier case, I will usually focus upon one of the other cases, those from the first category that I have mentioned – helpful Gettier cases. Almost all epistemologists accept these as being Gettier cases.

1.4 Gettierism Refined

At the core of this chapter is the quest to formulate Gettierism accurately, succinctly, and helpfully. Section 1.2 began that process by introducing Gettierism as, first of all, the thesis that Gettier was right in dismissing the justified-true-belief definition of knowledge, given his imagined (Gettier) cases. I then noted that Gettierism gains

a meta-epistemological element – the conviction that when we are pondering knowledge's nature we are following an epistemologically fruitful lead only once we accept or at least presume Gettier's being correct (in his dismissal of the idea that knowledge may be understood fully in terms of justification, truth, and belief). Let us now add some further precision to this initial formulation of Gettierism.

Did Gettier describe the *only* instances that there are or could be of a well-even-if-not-perfectly justified true belief's failing to be knowledge? No, he pointed to just two possible instances of this. But subsequent epistemologists extended his contribution, by noticing and imagining a multitude of actual and non-actual instances. Epistemologists' reactions to those cases have been constant – continually inferring that, whenever a belief is true and well-even-if-not-perfectly justified within a situation *relevantly like* one of those that was described by Gettier, it is part of a Gettier case *and it is therefore not knowledge*. If a belief is the centrepiece of a Gettier situation, it is not knowledge. More succinctly, if a belief is Gettiered, it is not knowledge.

Moreover, this thesis is treated by epistemologists as conceptual – as necessarily true. Its message aspires not merely to being the contingent truth that, as the world turns, no Gettiered beliefs are knowledge. It claims, more strikingly, that in *principle* no Gettiered beliefs are knowledge. Its message is that, *necessarily*, a belief's being Gettiered precludes the belief's being knowledge. And we should not mean this to be a trivial necessity. Sometimes the term "Gettiered" is used with the presumption that it *obviously means*, in part, "fails to be knowledge." But that implication is not part of the *obvious* meaning as I am using the word. Whether all Gettiered beliefs fail to be knowledge is something to be *discovered* as an implication. At any rate, this is so within this book. Nonetheless, I am happy to allow at this book's outset that *if* it is true that no Gettiered belief is knowledge, then this is a necessary truth.

And we may call this thesis the *standard interpretation* of Gettier cases. It implies that, as a matter of conceptual principle, being true and well-even-if-not-perfectly justified is insufficient for a belief's being knowledge: more succinctly, a belief's being true and justified does not entail its being knowledge. Why should that non-entailment, though, entail the belief's not being knowledge? It tells us that being true and justified is not *definitionally* enough for a belief's being