THE UNIVERSITY of EDINBURGH

Edinburgh Research Explorer

Analysing the concept of knowledge

OPEN ACCESS

THE ANALYSIS OF KNOWLEDGE

DUNCAN PRITCHARD

University of Edinburgh

ABSTRACT. The recent history regarding the analysis of knowledge is examined, in the context of the so-called 'Gettier problem'. The issues that are explored include: (i) the nature of the Gettier problem itself, and the particular conception of the 'traditional' analysis of knowledge that it presupposes; (ii) why the Gettier problem is not amenable to a straightforward solution; (iii) some of the core motivations for rejecting the 'analytical project' of trying to offering an analysis of knowledge; and (iv) some recent virtue-theoretic solutions that have been offered to the analytical project.

KEYWORDS: Epistemic Luck; Epistemology; Knowledge; Knowledge-First; Virtue Epistemology.

1. THE ANALYTICAL PROJECT IN EPISTEMOLOGY

For much of the last fifty years, the project of offering an analysis of propositional knowledge was thought to be central to the epistemological enterprise. Call this the *analytical project*, for short. The reason why this project came to the fore of epistemological discussion around fifty years ago is of course due to the famous article by Edmund Gettier (1963) which demonstrated that the traditional tripartite account of knowledge in terms of justified true belief was unsustainable. In countering the tripartite account, Gettier also cemented the idea that there was indeed a traditional account of knowledge in the literature along just the lines that he sets out in his critique of it. This might well be contentious, so let us first pause to explore a little what is supposedly involved in the 'traditional' tripartite account.

In what follows, I will take it as entirely uncontentious that knowledge entails truth, in the sense that it is *factive*—if one knows that p, then p is true. It is also reasonably uncontentious that knowledge entails belief—or at least some sort of belief-like propositional attitude—in the sense that if one knows that p, then one believes that p.[1] We should thus focus our attention on the third condition in the tripartite analysis, the justification condition. Gettier himself says very little about this condition. He attributes it to Roderick Chisholm (1957, 16) and to A. J. Ayer (1956), though in both cases he grants that they do not quite use the terminology of 'justification' themselves, writing instead (respectively) in terms of the knowing subject having 'adequate evidence' for their belief in the target proposition or 'having the right to be sure' that the target proposition is true.[2] What seems to be the common thread here as far as Gettier is concerned is that knowledge requires true belief that is supported by good reasons that the subject can offer in favour of the truth of her belief.

Crucially, however, it is also vital to the counterexamples that Gettier offers to the tripartite account, as he himself notes, that justification cannot itself be factive (i.e., one can have a justified belief that p and yet p be false). This is important because so-called Gettier-style cases essentially involve a subject justifiably believing a proposition in a manner that would have ordinarily resulted in a false belief, but where the subject ends up with a true belief regardless, albeit one that is completely disconnected from the justification.

To take one of Gettier's own examples, suppose that one has a good justification for believing that p is true, when it fact it is false. Now imagine that one infers from p the (inclusive) disjunction, either p or q, where q is a proposition which one has no reason for thinking is true. Normally, inferring a disjunction in this way would lead to a false belief (i.e., where one's original belief is in fact false, and where one is adding a disjunct that one has no reason to think is true). Nonetheless, one's belief ought to be justified, at least if the original belief that p is justified, as if one has good reasons for thinking that p is true, then one also has good reasons for thinking that p or q is true. All we now need to do is stipulate that the inferred disjunction is true, on account of q being true (unbeknownst to our subject). Hey presto, we now have a justified true belief, where the truth of the belief is completely disconnected from the original justification (which, recall, was a justification for thinking that p, and hence that p or q, is true).

We will examine the structure of Gettier-style cases a little more thoroughly in a moment. First, I want to reflect on whether Gettier is right to suppose that it was widely accepted, prior to his article, that knowledge has been analysed along these particular

tripartite lines. In fact, there is reason to be sceptical of this claim, as it seems far from straightforward that pre-Gettier epistemologists would have been happy to accept a conception of justification that was compatible with the target belief being false. Ayer, for example, sets out his position in opposition to what he refers to as the 'quest for certainty', which would suggest that he took the standard view to be much more demanding than his own. Moreover, there is a clear philosophical tradition, at least in the UK, of treating knowledge as involving an *infallible* grasp of the truth.[3] This is a tradition that goes back at least to the writings of John Cook Wilson (which have been tremendously influential on twentieth century British philosophy) and those directly influenced by him, such as H. A. Prichard.[4] Accordingly, it might plausibly be argued that the tripartite account of knowledge that Gettier attacked with his counterexamples was in fact not the standard view at all. Moreover, notice that the idea that knowledge involves an infallible grasp of the truth appears to go hand-in-hand with a conception of justification which is similarly infallible, and which would thus be completely unaffected by Gettier-style counterexamples. On this view, after all, the agents in Gettier-style cases are not justified in what they believe, precisely because their justification is compatible with the falsity of what they believe.[5]

Even so, one can see why those who initially responded to Gettier's article in the mid-1960s were willing to grant that justification must be understood along fallibilist lines. Wouldn't such an infallibilist conception of justification lead directly to radical scepticism, and hence be unsustainable?[6] In that case, the stark choice facing those responding to Gettier's counterexamples seems to either be to endorse an infallibilist conception of justification, and so deal with the very real threat of radical scepticism, or else embrace a fallibilist conception of justification and inherit the problem of offering an adequate analysis of knowledge that Gettier-style cases expose.

In any case, what is clear from the post-Gettier literature is that commentators very quickly took on board the idea that justification must be understood along fallibilist lines, and hence that the 'traditional' tripartite conception of knowledge was as Gettier described. For example, if one were to open an epistemology textbook written from a post-Getter perspective, then one would expect to see the tripartite account of knowledge being described just as Gettier described it. Moreover, we see this point in action in terms of the initial wave of responses to Gettier counterexamples, which involve not a renewed defence of an infallibilist conception of justification, but rather offer ways of effectively bolstering the tripartite account that has a fallibilist conception of justification at its core.[7] With this point in mind, let us look a little closer at the problem that Gettier identifies.

2. THE GETTIER PROBLEM

As we saw in the last section, Gettier's own cases essentially involve an inference, and so one natural response to these cases is to suppose that this is where one should focus in trying to fix the tripartite account.[8] In particular, in the case that we looked at earlier, the subject forms her justified true belief in the disjunction by inferring it from a false premise (i.e., her justified belief that p). Thus one might suppose that all we need to do to get around the problem is to insist that knowledge is justified true belief—where justification is still understood along the non-factive lines that Gettier attributed to the tripartite account—that is not inferred from any false premises. This general style of approach to the Gettier problem is known as the *no-false lemmas* approach.[9] Although initially appealing, it quickly runs into difficulties, and hence acts as a good 'case-study' when it comes to appreciating why solving the Gettier problem is trickier than it looks.

The first thing to note is that while Gettier's own two counterexamples to the tripartite account of knowledge essentially involve an inference from a false premise, this seems incidental to the general structure of a Gettier-style example. Consider, for example, the 'sheep' case that Chisholm (1977, 105) offers. Imagine a farmer—who we'll call 'Roddy'—who is looking into a field and sees what looks to be a sheep. On this basis, Roddy comes to justifiably believe that there is a sheep in the field. As it happens, though, what he is looking at is not a sheep at all but merely a big hairy dog that looks like a sheep. Ordinarily, then, if one were to form this belief on this basis one would end up with a false belief. Nonetheless, Roddy has excellent (although fallible) grounds in favour of his belief—the big hairy dog does, after all, look just like a sheep, and he has no reason to doubt what he sees—which is why his belief is justified. Moreover, as it happens, Roddy's belief is true since there is a sheep in the field hidden from view behind the big hairy dog. Roddy thus has a justified true belief that does not count as knowledge (since one can't come to know that there is a sheep in the field simply by looking at a big hairy dog).

Note that there is no essential inference involved in this case—Roddy just looks into the field and directly forms the belief that there is a sheep in the field. Moreover, notice that this case has just the kind of structure of a standard Gettier-style case, in that we effectively have good epistemic luck cancelling out bad epistemic luck, such that the target belief is both true and justified, but where the latter is completely disconnected from the former. That is, the bad epistemic luck that Roddy happens to be looking at a sheep-shaped object rather

than an actual sheep is cancelled out by the good epistemic luck that there happens to be a genuine sheep in the field. This is why his justification for his belief is disconnected from the truth of the belief.[10]

The existence of such cases entails that the no-false lemmas approach will need to be much more subtle in order to be effective.[11] In particular, it must allow for false assumptions that are in some sense operative in the background, even if the subject doesn't actively employ them in an inference. In the case just described, for example, Roddy is falsely supposing that what he is looking at is a genuine sheep, even if this is not a premise in his reasoning towards his belief that there is a sheep in the field.

But the fundamental problem facing any development of the no false lemmas view is to offer a principled way of understanding what constitutes a false lemma in this regard which is neither too broad as to exclude genuine cases of knowledge, nor too thin as to be unable to deal with Gettier-style cases. In order to see the force of the latter horn, notice that one can construct Gettier-style cases where the subject is not misled at all, at least in a direct fashion. Consider, for example, the famous *barn-façade case*.[12] Here we are to imagine a subject—let's call him Barney—who looks at a genuine barn in good cognitive conditions, and thereby forms the justified true belief that what he is looking at is a barn. The twist in the tail is that although this barn is genuine, it is in fact the only genuine barn in the vicinity, all of the others apparent barns in fact being mere barn-façades. Had Barney happened to look at any of the other apparent barns, then he would have formed a false belief.

This case has just the same feature of good epistemic luck cancelling out bad epistemic luck that we find in other Gettier-style cases. Barney has the bad epistemic luck to be forming a belief that there is a barn before him by looking at the front of the object in an environment where this is not a reliable guide as to whether there is a genuine barn there. But this bad epistemic luck is cancelled out by the good epistemic luck that he happens to be looking at the one genuine barn in the vicinity. Accordingly, his justification for what he believes is disconnected from the truth of his belief.

But what is the false assumption in play in the barn façade case supposed to be on the no false lemmas view? Is Barney really assuming that he is not in barn façade county? Moreover, if we do attribute this stance to the proponent of the no false lemmas thesis, then the other horn of our dilemma becomes apparent. For if the notion of a false lemma is this broad, then isn't it often the case that when we genuinely come to know propositions there will be *some* false assumptions in play, however tacit? The worry is thus that a lot of our *bona fide* knowledge is now under threat.[13]

The point of the foregoing is not to suggest that the no false lemmas approach is completely unsustainable, but merely to highlight the difficulties involved in dealing with the Gettier problem.[14] Indeed, I think we are now in a better position to say a bit more about what the Gettier problem amounts to. In particular, while in the first instance it was treated as the challenge of explaining how we go about fixing the traditional tripartite account of knowledge, once we recognise that there is no easy fix in this regard a deeper construal of the Gettier problem emerges. This concerns our very notion of knowledge. Gettier cases highlight the point that when one knows one's cognitive success—i.e., one's true belief—ought not to be just down to luck in any significant way. Instead, it should be arising out of our cognitive agency.

This is exactly what goes wrong in Gettier-style cases, and thus why the tripartite account of knowledge is unsustainable as it stands. That is, we might have initially thought that having a justification, even a fallible one, would ensure that our beliefs are due to our cognitive agency in such a way that they are not simply down to luck. And yet the good-luck-bad-luck structure of the Gettier-style cases exposes how one can have a justified true belief and yet the truth of one's belief is entirely down to luck and nothing to do with one's manifestation of cognitive agency. Roddy is manifesting his cognitive agency by using his reliable perceptual faculties to pick out (what he takes to be) the sheep before him. And he ends up with a true belief, but the truth of his belief is just down to the good epistemic luck that there happens to be a sheep in the field hidden from view behind the sheep-shaped object that he is looking at. Similarly, Barney is manifesting his cognitive agency by using his reliable perceptual faculties to pick out the barn before him. And he ends up with a true belief, but the truth of his belief is just down to the good epistemic luck that he happens to be looking at the one genuine barn in the vicinity. And so on.

If this is the right way of characterising the Gettier problem, then it suggests that our natural way of thinking about knowledge is in terms of a cognitive success that is not down to luck, but is rather in some substantive way due to our manifestation of cognitive ability. This raises some interesting questions about why we think about knowledge in this fashion. Moreover, given that the moral of the Gettier problem is that (fallible) justification cannot fit the bill in this regard, then the hunt is on to find a way of thinking about knowledge that is adequate.[15]

3. SCEPTICISM ABOUT THE ANALYTICAL PROJECT

Interestingly, however, while there has been a thriving philosophical industry devoted to resolving the Gettier problem, there is also a growing mood of scepticism about the very possibility—and even desirability—of offering such a solution. There are various sources of this scepticism. One source of scepticism is not specific to the analytical project in epistemology, but rather concerns all analytical projects in philosophy. Perhaps they are all wrong-headed, in that no philosophically interesting notion can be made subject to such an account? One might, for example, advance broadly Wittgensteinian reasons in support of this claim.[16]

Whatever the rationale for this sceptical line, I think we can fairly set it to one side for our purposes, given that it is not specific to the analytical project in epistemology. Moreover, it is also notable that this way of critiquing the analytical project tends to take it as given that what it would take to resolve this project is a non-circular set of necessary and sufficient conditions. But it is far from clear that our ambitions should be nearly so ambitious. An analysis of knowledge might well be informative even though it inter-defined several notions, and hence was not ultimately non-circular. Nonetheless, such an account might cast a lot of light onto the nature of these notions and how they relate to one another. Conversely, it is also worth noting that non-circular analyses of knowledge can sometimes be *un*informative, as when they are *ad hoc*, so it is not as if eliminating circularity is a guarantee of a good analysis anyway.[17]

A more interesting sceptical line regarding the analytical project for our purposes—and which is specific to knowledge—arises out of the post-Gettier literature itself. It involves a kind of pessimistic induction, whereby the continued failure of epistemologists engaged in this literature to offer a counterexample-proof analysis of knowledge is meant to provide inductive support for the idea that there is something inherently wrong-headed about this approach.[18] In effect, the idea is that offering an analysis of knowledge is akin to being involved in a degenerating research programme.[19] Interestingly, this complaint often goes hand-in-hand with a distinct sceptical line to the effect that the kinds of analyses of knowledge that have been offered in post-Getter epistemology tend to be so convoluted and plain ugly that one struggles to see how they could ever be plausible, much less account for why knowledge is such a fundamentally important notion.

Although these two complaints tend to be run together, I think we would be wise to keep them apart. The latter complaint is the weaker of the two. Why can't an analysis of

knowledge be entirely adequate but also ugly? Indeed, why can't it be simultaneously true that knowledge is the kind of thing that we value a great deal, and which also has a complex, and thus ugly, analysis?[20] What makes this complaint even less compelling is that although some solutions that have been offered to the Gettier problem have been complex, there are in fact some elegant solutions available too, as we will see in a moment.

This last point is also relevant to the pessimistic induction complaint. This is because whether such an inductive inference is sensible very much depends on whether we think that progress has been made in the post-Gettier literature. That a research programme has failed does not itself entail that it is a degenerating research programme, after all, since this could be a failure that has been inching towards success. Again, that there are some elegant solutions to the problem available in the literature is thus a good ground to be suspicious of this sceptical line regarding the analytical project.

A final source of scepticism about the analytical project in epistemology arises out of the *knowledge-first* research programme as defended by Timothy Williamson (2000). Unlike the other critiques of the analytical project, this is focused on offering a positive defence of an alternative approach. Williamson argues that once we recognize the conceptual primacy of knowledge with regard to related concepts like belief and evidence, such that we attempt to define the latter in terms of the former rather than *vice versa*, then we will be able to make significant headway on the key epistemological problems. This doesn't just include resolving the Gettier problem—which effectively disappears on this view, at least insofar as we treat that problem as essentially trading on the idea that we should analyse knowledge in terms true belief plus some further epistemic conditions—but also extends to other difficulties such as the problem of radical scepticism.

It would take me too far afield to critically engage with knowledge-first epistemology here, particularly since it is a proposal that has so many important moving parts.[21] Instead, I am going to take a slightly different tack, and briefly outline two recent responses to the analytical project that I think are plausible. As we will see, in doing so I am not necessarily objecting to knowledge-first epistemology, since both of these proposals could be recast along knowledge-first lines (even while—arguably at any rate—remaining potential resolutions of the analytical problem).[22]

4. RESOLVING THE ANALYTICAL PROJECT:
ROBUST AND ANTI-LUCK VIRTUE EPISTEMOLOGY

We noted above that the moral of the Gettier problem is that one can appropriately manifest one's cognitive agency (such that one's belief is justified), and in the process be cognitively successful (i.e., form a true belief), even though there is a complete disconnect between one's cognitive success and one's manifestation of cognitive agency. In particular, even despite one's appropriate manifestation of cognitive agency, one's cognitive success can nonetheless be due to luck. If that's right, then this suggests that the way to respond to the Gettier problem is to insist that the cognitive success be appropriately related to the manifestation of cognitive agency.

How might one do this? One popular proposal—advanced, in various guises, by such figures as Ernest Sosa (1991; 2007; 2009; 2016), Linda Zagzebski (1996; 1999), and John Greco (2009)—is to contend that knowledge should be understood as cognitive success that is *because of* the subject's manifestation of relevant cognitive agency. Take, for example, the Roddy case described above. Although Roddy is cognitively successful and appropriately manifests cognitive agency (such that his belief is justified), his cognitive success is not in any sense because of his manifestation of cognitive agency. Rather, it seems to be because of the happenstance that there is a sheep hidden from view in the field, which is nothing to do with his cognitive agency.

This proposal is known in the literature as *robust virtue epistemology*. It is a form of virtue epistemology in that it defines knowledge in terms of the manifestation of cognitive agency (or intellectual virtue, broadly conceived). It is a robust form of the view, in that it only appeals to cognitive agency in its account of knowledge, rather than incorporating a non-virtue-theoretic condition as well (we will be looking at a non-robust, or *modest*, form of virtue epistemology in a moment).[23] There are lots of technicalities regarding how best to understand this proposal—for example, how best to unpack the 'because of' relation which figures in this account—but we will need to restrict ourselves to making a few general points about the view.

First, it is important to note some of the attractive features of the proposal. In particular, it is able to offer a striking diagnostic story as to why knowledge might be structured in this fashion. The idea is that knowledge is simply a sub-class of the broader category of *achievements*, where an achievement is a success that is because of the subject's manifestation of relevant skillful agency. This explains why one can construct analogous

Gettier-style cases outside of epistemology. Take someone firing an arrow at a target. Merely hitting the target would not suffice for an achievement, since perhaps there is no skill involved. But achievements are also undermined when there is both success and the relevant manifestation of skillful agency but where the latter is not because of the former. This is where the Gettier-style luck comes into play. An archer who skillfully fires an arrow at a target, and hits the target, but only hits the target because (say) a dog jumped up mid-flight, snatched the bolt and placed it in the target, would not be credited with an achievement.

A further advantage of this proposal is that it seems able to account for the special value that we place on knowledge. Achievements, after all, are of special value to us—a good life is surely one that is rich in achievements, for example—and hence it is unsurprising that knowledge, *qua* cognitive achievement, should also be prized.[24]

Robust virtue epistemology thus offers a compelling response to the analytical project, and thus the Gettier problem. Note, though, that the analysis of knowledge that it offers may well turn out to be non-reductive. Perhaps, for example, there is no way of defining cognitive agency that doesn't appeal to knowledge. In that case, the proposal would be ultimately circular. Nonetheless, it could still be an informative, for the reasons given above, and hence an adequate response to the analytical project. (Relatedly, notice that robust virtue epistemology and knowledge-first epistemology need not be in conflict, in that a proponent of the latter could consistently endorse a version of the former).

Despite the attractions of robust virtue epistemology, it does face some problems. One issue is that although it fares well with regard to standard Gettier-style cases like the Roddy scenario, it struggles with some non-standard cases, particularly the Barney example. Wouldn't we naturally describe this case as being one where the subject's true belief is indeed because of his manifestation of relevant cognitive agency (he is looking at a real barn after all)? More generally, some have objected that knowledge and achievements come apart in terms of their relationship to luck. While an achievement can be *fragile*—in the sense that although it is a success that is because of one's skillful agency, it nonetheless could have easily been a failure—the same doesn't seem true of knowledge. The archer who skillfully hits the target, but who could have easily missed (had the mischievous dog just described not been held back for example), nonetheless exhibits an achievement. But would we ever treat a cognitive success that could have very easily been a cognitive failure—as in the Barney case—as an instance of knowledge?[25]

This problem concerns whether robust virtue epistemology is too weak as an account of knowledge, in that it treats some agents as having knowledge which intuitively they do not possess. But the account has also been charged with being too strong, in that it seems to deprive some agents of knowledge that they genuinely possess. Often, for example, testimonial knowledge is gained by, to a large extent, trusting the word of another. But would we want to say that in such cases one's cognitive success is because of one's manifestation of cognitive agency? The point is that although one does manifest some degree of relevant cognitive agency in such cases—testimonial knowledge isn't acquired by simply trusting anyone, after all—nonetheless in such scenarios it seems as if one's cognitive success is because of one's *informant's* manifestation of cognitive agency, if anyone's.

As is usual in philosophy, none of these objections, at least taken alone, are fatal, as there are plausible counter moves that the proponent of robust virtue epistemology can make. For example, some commentators have observed that the intuition that knowledge is lacking in the Barney case is far less strong than in more standard Gettier-style cases, like the Roddy scenario. Accordingly, one could reasonably argue that rather than being an objection to the proposal, one should instead embrace the consequence that Barney has knowledge in this case, on account of the fact that he has exhibited a genuine cognitive achievement.[26]

If one is unconvinced by robust virtue epistemology, however, there is another proposal available in the general vicinity. This is a modest form of virtue epistemology, in that it doesn't exclusively define knowledge in terms of a virtue-theoretic condition. According to this proposal, we should take seriously the idea that there are two distinct constraints operative on knowledge. One is that one's cognitive success should be appropriately related to one's cognitive agency. The other is that one's cognitive success shouldn't be due to luck; alternatively, that it shouldn't be a risky success (i.e., a success that could have very easily been a failure). Gettier-style cases demonstrate that one can satisfy the first constraint without thereby satisfying the second constraint. Indeed, Barney-style cases seem to demonstrate that one can exhibit a genuine cognitive achievement even while lacking knowledge on account of the epistemic luck involved. Conversely, it is uncontentious that one's cognitive success can be neither due to luck nor anything to do with one's manifestation of cognitive agency—perhaps there are environmental factors that simply guarantee that one's belief will be true, regardless of its epistemic pedigree.

If we take these observations seriously, then they suggest a more complex analysis of knowledge than that offered by robust virtue epistemology. According to this proposal—what

is known as *anti-luck* (or *anti-risk*) *virtue epistemology*—knowledge is to be understood as non-lucky/non-risky cognitive success that is significantly attributable to one's manifestation of relevant cognitive agency. Notice that this account is both in one aspect more demanding, and in another aspect less demanding, than robust virtue epistemology. It is more demanding, in that it is not just one's cognitive success that needs to be significantly attributable to one's manifestation of relevant cognitive agency, but rather one's non-lucky/non-risky cognitive success. It is less demanding, in that the relation in play here is not the because of relation employed by robust virtue epistemology. Instead, all that is demanded is that one's manifestation of relevant cognitive agency should play some significant explanatory role in one's non-lucky cognitive success.[27]

Anti-luck virtue epistemology is designed to avoid the kinds of problems that face robust virtue epistemology. Barney-style cases are dealt with on account of how they involve lucky success. But one can also account for why cases of testimonial knowledge that involve a high degree of trust are genuine, since although the subject's cognitive agency does not play an overarching explanatory role in her cognitive success, it does play a significant explanatory role.

Anti-luck virtue epistemology is thus able to offer an alterative way of responding to the analytical project, and hence the Gettier problem (although, like robust virtue epistemology, this doesn't mean that it is thereby offering a reductive account of knowledge). Admittedly, however, anti-luck virtue epistemology lacks the elegant simplicity of robust virtue epistemology, and that might be thought to be a substantive strike against it.

5. CONCLUDING REMARKS

That there exist plausible responses to the analytical project that arose in response to the Gettier problem obviously doesn't entail that this project has been resolved, much less that it is a coherent philosophical project to undertake. But it does suggest that we shouldn't be too quick to dismiss this project, or else to take it for granted that it represents a degenerating research programme. It also highlights that the challenge that Gettier posed epistemology in 1963 is still alive today.[28]

REFERENCES

Ayer, A. J. (1956). *The Problem of Knowledge*, (London: Macmillan).

Chisholm, R. M. (1957). *Perceiving: A Philosophical Study*, (Ithaca, NJ: Cornell University Press).

—— (1977). *Theory of Knowledge* (2nd Ed.), (Englewood Cliffs, NJ: Prentice-Hall).

Clark, M. (1963). 'Knowledge and Grounds. A Comment on Mr. Gettier's Paper', *Analysis* 24, 46–48.

Code, L. (1987). *Epistemic Responsibility*, Hanover, NH: University Press of New England.

Cook Wilson, J. (1926). *Statement and Inference*, (Oxford: Clarendon Press).

DePaul, M. (2009). 'Ugly Analyses and Value', *Epistemic Value* (eds.) A. Haddock, A. Millar & D. H. Pritchard, 112-38, (Oxford: Oxford University Press).

Dutant, J. (2015). 'The Legend of the Justified True Belief Analysis', *Philosophical Perspectives* 29, 95-145.

Gettier, E. (1963). 'Is Justified True Belief Knowledge?', *Analysis* 23, 121-23.

Goldman, A. (1976). 'Discrimination and Perceptual Knowledge', *Journal of Philosophy* 73, 771-91.

Greco, J. (2009). *Achieving Knowledge*, (Cambridge: Cambridge University Press).

Hetherington, S. (2013). 'There Can be Lucky Knowledge', *Contemporary Debates in Epistemology* (2nd Ed.), (eds.) M. Steup & J. Turri, §7, Oxford: Blackwell.

Hookway, C. (2003). 'How to Be a Virtue Epistemologist', *Intellectual Virtue: Perspectives from Ethics and Epistemology*, (eds.) M. DePaul & L. Zagzebski, 183-202, Oxford: Oxford University Press.

Kallestrup, J., & Pritchard, D. H. (2014). 'Virtue Epistemology and Epistemic Twin Earth', *European Journal of Philosophy* 22 (2014), 335-57.

Kvanvig, J. (1992). *The Intellectual Virtues and the Life of the Mind*, Savage, MD: Rowman & Littlefield.

Lakatos, I. (1978). *The Methodology of Scientific Research Programmes: Philosophical Papers Volume 1*, (ed.) J. Worral & G. Currie, (Cambridge: Cambridge University Press).

Lycan, W. G. (2006). 'On the Gettier Problem Problem', *Epistemology Futures*, (ed.) S. Hetherington, 148-68, (Oxford: Oxford University Press).

McDowell, J. (1995). 'Knowledge and the Internal', *Philosophy and Phenomenological Research* 55, 877-93.

McGlynn, A. (2014). *Knowledge First?*, (London: Palgrave Macmillan).

Montmarquet, J. (1993). *Epistemic Virtue and Doxastic Responsibility*, Lanham, MD: Rowman & Littlefield.

Prichard, H. A. (1950). *Knowledge and Perception*, (Oxford: Clarendon Press).

Pritchard, D. H. (2005). *Epistemic Luck*, (Oxford: Oxford University Press).

—— (2007). 'Recent Work on Epistemic Value', *American Philosophical Quarterly* 44, 85-110.

—— (2009). 'Knowledge, Understanding and Epistemic Value', *Epistemology* (*Royal Institute of Philosophy* Lectures), (ed.) A. O'Hear, 19-43, (Cambridge: Cambridge University Press).

—— (2010). 'Achievements, Luck and Value', *Think* 25, 1-12.

—— (2012*a*). 'Anti-Luck Virtue Epistemology', *Journal of Philosophy* 109, 247-79.

—— (2012*b*). *Epistemological Disjunctivism*, (Oxford: Oxford University Press).

—— (2012*c*). 'On Meta-Epistemology', *Harvard Review of Philosophy* 18, 91-108.

—— (2013). 'There Cannot be Lucky Knowledge', *Contemporary Debates in Epistemology* (2nd Ed.), (eds.) M. Steup & J. Turri, §7, Oxford: Blackwell.

—— (2015*a*). 'Anti-Luck Epistemology and the Gettier Problem', *Philosophical Studies* 172, 93-111.

────── (2015*b*). *Epistemic Angst: Radical Skepticism and the Groundlessness of Our Believing*, (Princeton, NJ: Princeton University Press).

────── (*Forthcominga*). 'Knowledge, Luck and Virtue: Resolving the Gettier Problem', *The Gettier Problem*, (eds.) C. Almeida, P. Klein & R. Borges, (Oxford: Oxford University Press).

────── (*Forthcomingb*). 'Epistemic Risk', *Journal of Philosophy*.

Pritchard, D. H., Millar, A., & Haddock, A. (2010). *The Nature and Value of Knowledge: Three Investigations*, (Oxford: Oxford University Press).

Radford, C. (1966). 'Knowledge - By Examples', *Analysis* 27, 1-11.

Roberts R., & Wood, W. J. (2007). *Intellectual Virtues: An Essay in Regulative Epistemology*, Oxford: Oxford University Press.

Shope, R. K. (1983). *The Analysis of Knowing. A Decade of Research*, (Princeton, NJ: Princeton University Press).

Sosa, E. (1991). *Knowledge in Perspective: Selected Essays in Epistemology*, (Cambridge: Cambridge University Press).

────── (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, (Oxford: Oxford University Press).

────── (2009). *Reflective Knowledge: Apt Belief and Reflective Knowledge*, (Oxford: Oxford University Press).

────── (2016). *Judgment and Agency*, (Oxford: Oxford University Press).

Strawson, P. F. (1992). *Analysis and Metaphysics: An Introduction to Philosophy*, (Oxford: Oxford University Press).

Travis, C. (2005). 'A Sense of Occasion', *Philosophical Quarterly* 55, 286-314.

Williamson, T. (2000). *Knowledge and its Limits*, (Oxford: Oxford University Press).

Wittgenstein, L. (1953). *Philosophical Investigations*, (eds.) G. E. M. Anscombe & R. Rhees, (tr.) G. E. M. Anscombe, (Oxford: Blackwell).

Zagzebski, L. (1994). 'The Inescapability of Gettier Problems', *Philosophical Quarterly* 44, 65-73.

────── (1996). *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*, (Cambridge: Cambridge University Press).

────── (1999). 'What is Knowledge?', *Blackwell Guide to Epistemology*, (eds.) J. Greco & E. Sosa, 92-116, Oxford: Blackwell.

NOTES

[1] This condition has been disputed—for example, by Radford (1966)—but is nonetheless widely accepted.

[2] In addition, Gettier suggests that Plato's discussion of knowledge in the *Theatetus* in terms of true opinion, or judgement, with an account as potentially also being an endorsement of the traditional tripartite account of knowledge.

[3] And of course there is also an earlier, and far more influential, Cartesian tradition of seeking certain foundations for one's knowledge.

[4] See Cook Wilson (1926) and Prichard (1950). I am grateful to Charles Travis for first alerting me to the work of Cook Wilson and its influence on UK philosophy, particularly via Prichard, but also in the writings of such figures as J. L. Austin, John McDowell, and Timothy Williamson. See especially Travis (2005).

[5] For a detailed defence of the idea that an infallibilist conception of justification was part of the standard pre-Gettier account of knowledge, see Dutant (2015). For a helpful discussion of the role of the notion of fallible justification in the Gettier problem, see Zagzebski (1994).

[6] This is the line that Ayer (1956, 21) took, for example.

[7] See Zagzebski (1994). A good example of this general tendency is the *no false lemmas* approach that we will examine below. Note, however, that in the recent literature there has been a resurgence of interest in a factive conception of justification, largely inspired by McDowell's (e.g., 1995) defence of (a position which is these days called) epistemological disjunctivism. See, for example, Pritchard (2012*b*).

[8] For a detailed account of the first wave of responses to the Gettier counterexamples in the literature, see Shope (1983).

[9] For an early version of this approach, see Clark (1963).

[10] I think the first person to characterise Gettier-style cases in terms of this recipe of good epistemic luck cancelling out bad epistemic luck was Zagzebski (1994). See also Zagzebski (1999). For more on the notion of epistemic luck and its role in Gettier-style cases, see Pritchard (2005; 2015*a*).

[11] Note too that there are lots of cases of this general type in the literature, and hence it would not be a promising strategy to try to side-step this difficulty by claiming that in this particular example the subject is making a hidden inference from a false premise (e.g., such as from *that's a sheep*, to *there is a sheep in the field*).

[12] This was first offered by Goldman (1976), though he credits the example to Carl Ginet.

[13] Note that one response to the barn façade case is to deny that it is a Gettier-style case at all, and thus insist that the agent acquires knowledge. Although this approach was not initially seen as very plausible, it has gained some traction in the recent literature. For a high-profile proponent of this line, see Sosa (2007, 31–32). For a recent critical discussion of this line of argument, see Pritchard (*forthcomingb*). A more radical line—one that applies not just to the barn façade case but also to other Gettier-style cases—is to insist that knowledge is compatible with the agent forming her true belief entirely by luck. For a defence of this view, see Hetherington (2013). For a response, see Pritchard (2013).

[14] For a particularly interesting recent development of the no false lemmas line, see Lycan (2006).

[15] For further discussion about the right way to conceive of the 'Gettier problem', see Pritchard (2015*a*; *forthcominga*).

[16] See, especially, Wittgenstein's (1953) famous remarks on family resemblance concepts.

[17] See Zagzebski (1999) for an insightful discussion of how we should think about the project of analysing knowledge. More generally, see Strawson's (1992) influential account of elucidation in philosophical analysis.

[18] There are many exponents of this complaint, but the most prominent is Williamson (2000, chapter 1).

[19] See Lakatos (1978).

[20] See DePaul (2009) for a very interesting discussion of just this point.

[21] For a very helpful recent book-length critique of knowledge-first epistemology, see McGlynn (2014).

[22] A further source of scepticism about the analytical project can be found in the particular wing of virtue epistemology that explicitly rejects the idea that defining knowledge should be central to the epistemological enterprise. See, for example, Code (1987), Kvanvig (1992), Montmarquet (1993), Hookway (2003), and Roberts & Wood (2007). Sometimes this proposal is allied to one of the complaints regarding the analytical project just noted, but often the claim is simply that focussing on the analytical project in some fundamental way distorts what the epistemological enterprise ought to be about (i.e., epistemically virtuous agents). For further discussion of the methodology of epistemology, and the role of the analytical project within it, see Pritchard (2012*c*).

[23] I introduced the distinction between robust (or 'strong') and modest (or 'weak') forms of virtue epistemology in earlier work. See especially Pritchard (2009) and Pritchard, Millar & Haddock (2010, chs. 1-4).

[24] For further general discussion of the value of knowledge, see Pritchard (2007). For further discussion of the specific topic of the value of achievements, see Pritchard (2010).

[25] For a recent development of this general line of objection in terms of an epistemic twin earth argument, see Kallestrup & Pritchard (2014).

[26] See also endnote 13.

[27] For the main defences of this proposal, see Pritchard (2009; 2012*a*; *forthcominga*; *forthcomingb*) and Pritchard, Millar & Haddock (2010, chs. 1-4).

[28] I am grateful to Stephen Hetherington for helpful comments on an earlier version of this article.