

For *Central Issues of Philosophy*, (ed.) J. Shand, (Blackwell)

KNOWLEDGE

DUNCAN PRITCHARD

University of Edinburgh

1. TWO PLATITUDES ABOUT KNOWLEDGE

There are many kinds of knowledge. I may know that Paris is the capital of France, or know how to bake a cake, or know where my keys are, or know who was the inventor of the zip fastener, and so on. To keep matters simple, we will focus on a particular kind of knowledge which is of central importance, what is known as *propositional knowledge*. Propositional knowledge, as the name suggests, is knowledge of a proposition. A proposition is, roughly, what is expressed by a sentence which says that something is the case—e.g., that Paris is the capital of France, or that the earth is flat. In focussing on propositional knowledge, then, we are focussing on *knowledge that* such-and-such is the case, rather than, say, on *knowing how* to do such-and-such, or *knowing where* such-and-such is, and so on.

Everyone agrees that knowledge entails true belief, in the sense that if one knows a proposition, p , then one believes p and p is true. (Of course, one might think that one knew a certain proposition which turned out to be false, but in such a case one would thereby discover that one did not really know it after all). Everyone also agrees that there is a lot more to knowledge than merely true belief.

It is easy to formulate cases of true belief that are not also cases of knowledge. For example, imagine a gambler—let's call him 'Lucky'—who believes that the horse that will win the next race is Lucky Lass, where this belief is formed simply on the basis that he likes the name. Suppose further that Lucky Lass does indeed win the next race. Would we say that Lucky *knew* that Lucky Lass would win? Surely not, since his belief is simply the result of guesswork and guesswork is by itself no route to knowledge. And yet he does have a true belief in this proposition.

The tricky task for those working in the theory of knowledge (otherwise known as *epistemology*) is to explain what else is required for knowledge over and above true belief.¹

There are two very natural ways of explaining why Lucky's true belief does not qualify as knowledge. The first is to note that Lucky's belief is only true as a matter of luck. That is, given how he formed his belief, that belief could very easily have been wrong. Compare Lucky's belief in this regard with the belief held by someone who we would regard as having known that Lucky Lass would win. Let us suppose that unbeknownst to Lucky, the race was fixed and the fixer was a local gangster who we'll call 'Mr. Big'. Since Mr. Big knows that the race is fixed in Lucky Lass's favour, we would naturally regard him as knowing that Lucky Lass will win. Notice, however, that given how Mr. Big has formed his true belief, it is not a matter of luck that his belief is true. That is, his true belief could not have very easily been wrong. There is thus a lot to be said for the idea that a pre-condition on knowledge is that it is not a matter of luck that one's belief in the target proposition is true. Call this intuition about knowledge, the *anti-luck platitude*.

A second way in which one might naturally explain Lucky's lack of knowledge is in terms of the fact that his true belief was not in any way the product of ability, but rather simply due to a lucky guess. In contrast, Mr. Big's true belief *was* formed through ability. After all, he saw for himself that all the other horses in the race were drugged and hence, given what he knows about the performance of drugged horses, he thereby knows that Lucky Lass will win. One way of putting this point is by saying that when one knows it is of some credit to one that one has a true belief. In the case of Lucky, however, it is of no credit to him at all that he formed a true belief, since his belief is only true by luck. There is thus a lot to be said for the idea that a pre-condition on knowledge is that one's true belief in the target proposition is gained via ability. Call this intuition about knowledge, the *ability platitude*.

These two platitudes have been extremely influential on contemporary theorising about knowledge. Interestingly, one might be tempted to suppose that they are just two ways of putting the same point, such that there is in effect just one 'super' platitude in play here. After all, one might naturally suppose that any true belief that was gained via ability would not be true as a matter of luck, and that any non-luckily true belief must have been gained via ability. If this were right, then we would be well on the way to understanding what knowledge is, since we would just need to say more about what satisfying these two platitudes would involve. As we will see, however, matters are not quite so straightforward. Indeed, we will see that these two platitudes in fact impose two independent constraints on knowledge.

2. GETTIER-STYLE CASES

Traditionally, the way of explaining what knowledge is in a manner that is consistent with both the ability and the anti-luck platitudes is to appeal to a *justification* condition, where satisfying this condition involves the agent being able to cite good grounds in favour of what one believes. Such an account of knowledge is known as the ‘tripartite’ account of knowledge, since it defines knowledge as having three parts: justification, truth and belief.

When it comes to the example just described involving Lucky and Mr. Big this proposal fares pretty well. After all, Lucky is unable to offer any good grounds in favour of what he believes, unlike Mr. Big who can offer excellent grounds in favour of why he thinks that Lucky Lass will win. The proposal does not fare well when it comes to other cases, however. One problem concerns the fact that we often attribute knowledge in cases where the agent concerned is unable to offer any good grounds in favour of what she believes. We will consider one example of this type later on. First, though, we need to look at an even more fundamental problem facing the tripartite account.

Consider the following example. Suppose that our agent—we will call him ‘Edmund’—comes downstairs one morning and forms his belief about what the time is by looking at the grandfather clock in his hall. The belief he forms, let’s say, is that it is 8.20am. Suppose further that this clock has been very reliable in the past and Edmund knows this, and also that Edmund has independent grounds for thinking that the time is roughly 8.20am (e.g., it’s light outside, he usually gets up around this time, and so forth). Finally, let us stipulate that Edmund’s belief is true, it is 8.20am. Edmund thus has a true belief in this proposition, and he is also in a position to offer excellent grounds in favour of his belief—i.e., his belief is justified. According to the tripartite account, then, he must know that the time is 8.20am. Here comes the twist, however. As it happens, and unbeknownst to Edmund, the clock stopped working twenty-four hours previously and is stuck at the time 8.20am. Does Edmund know what the time is? Surely not. After all, one cannot find out what the time is by looking at a stopped clock. The moral of the story is thus that whatever knowledge is, it is not justified true belief.²

These cases are called ‘Gettier-style’ cases, since they were first formulated as an objection to the tripartite account in a famous article by Edmund Gettier (1963).³ There is a recipe for creating such cases. First, you take a belief that is formed in a way which would normally result in a false belief (e.g., in this case, a belief that is formed by looking at a stopped clock). Next, you set up the case such that the agent has good citable grounds in favour of her belief (e.g., in this case, Edmund has excellent

grounds for believing that the time is indeed 8.20am). Finally, you add the further detail that the belief so formed is, as it happens, true anyway.

Here is a second example which further illustrates this recipe for Gettier-style cases. 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 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 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. 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 which 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).⁴

What is interesting about Gettier-style cases is that they demonstrate that merely having a justification in favour of what you truly believe is insufficient to deal with the constraint on knowledge imposed by the anti-luck platitude. For in all such cases what you have is a justified true belief which doesn't count as knowledge because the belief in question is only luckily true—i.e., it could very easily have been wrong. In the case of Edmund, for example, had he come downstairs a minute earlier or a minute later (or if the clock had stopped a minute earlier or a minute later), then he would have formed a false belief by looking at this clock. The same goes for Roddy. Had there not been a sheep hidden from view behind the big hairy dog, then he would have formed a false belief by looking at the big hairy dog. In both cases, then, the belief formed is only luckily true even though it is justified.

A second point to notice about Gettier-style cases is that they *do* satisfy the constraint on knowledge imposed by the ability platitude. After all, both Edmund and Roddy form their respective beliefs through ability—i.e., these cases are not like the case of Lucky who gains a true belief simply by guesswork. Recall that Edmund has every reason to trust what this clock tells him, and Roddy is indeed looking at something that looks very much like a sheep. Nevertheless, merely forming one's belief through ability does not appear to suffice for knowledge. Instead one must in addition form one's belief in a way that is non-lucky. The demands imposed by the ability and anti-luck platitudes are thus distinct.

3. THE LOTTERY PUZZLE

There is another type of example which illustrates that the demands imposed by these platitudes are distinct. Imagine a fair lottery with *extremely* long odds (a billion-to-one, say). Now suppose that our agent—we'll call her 'Lottie'—is in possession of one of the tickets for this lottery, a ticket which, as it happens, is a losing ticket (though Lottie doesn't have any inkling of this yet). Lottie now reasons to herself that given that the odds against her winning are so high, her ticket must be a losing ticket. On this basis, she forms the (true) belief that she has lost the lottery and so tears up her ticket.

I think we would find Lottie's behaviour rather puzzling, and part of the reason for this is that intuitively Lottie cannot come to know that her ticket has lost simply by reflecting on the long odds involved even though her belief is in fact true. Indeed, we would probably say to Lottie that she shouldn't have ripped up her ticket because, for all she knew, she had won the lottery. This example is thus itself a Gettier-style case, in that it is an example of justified true belief which, intuitively, is not knowledge. This is not the only import of this example, though.

What is interesting about this case is that while Lottie is unable to come to know that she has lost the lottery simply by considering the odds involved, she can come to know that she has lost by reading the results in a reliable newspaper. What is mysterious about this, however, is that the probability that the newspaper has printed the wrong result is surely much higher than the probability of her winning. From the point of view of the probability that her belief is correct, then, there is a greater probability that her belief is true if she forms it by reflecting on the odds involved than by forming it by reading the result in a reliable newspaper. And yet Lottie can come to know this proposition by the second method but not by the first.

This is the so-called 'Lottery Puzzle'.⁵ What it demonstrates is the surprising fact that whether or not you know is not a function of the probabilistic strength of your supporting evidence—that is, one might have evidence in favour of one's belief which makes it very probable that one's belief is true and yet *lack* knowledge even though possessing evidence in favour of one's belief which does not make it as probable that it is true *can* suffice for knowledge. The way out of the puzzle is to recognise that whether or not your belief is only luckily true—and hence, in line with the anti-luck platitude, not a case of knowledge—is not directly related to the probabilistic strength of your supporting evidence.

In order to see this, notice that what is wrong with Lottie's true belief that she has lost when it is based on the consideration of the odds involved is that this belief very easily have been wrong.

Imagine, for example, that Lottie had been in possession of the winning ticket and formed her belief about whether her ticket had lost in this way. In such a case she would have ended up forming a false belief via this method. Matters are different, however, when it comes to forming one's belief by looking up the result in a reliable newspaper. This is because had Lottie won the lottery then we would expect Lottie to form a true belief via this method. After all, had Lottie won then different results would have been published in the reliable newspaper, results which corresponded to the numbers on her ticket.

So the probabilistic strength of the evidential support you have for your belief does not itself determine whether your true belief is only true by luck, since a very high probabilistic strength of evidential support is consistent with one's belief being only luckily true, while a relatively low probabilistic strength of evidential support might be enough to ensure that your belief is not luckily true. More generally, the lottery case further illustrates the point made above that the demands imposed by the ability and anti-luck platitudes are distinct. After all, whether Lottie forms her true belief by considering the odds involved or by reading the results in a newspaper, it is surely through ability that she forms her belief. Merely forming one's belief through ability is thus not enough to ensure knowledge, because even a belief so formed could still be only luckily true.

4. EXTERNALISM AND INTERNALISM ABOUT KNOWLEDGE

One way of responding to the challenge posed by the Gettier-style cases and the lottery case could be to say that what we need to do is simply define knowledge as non-lucky justified true belief. This would then explain why the agents in these cases lack knowledge, since in each case their justified true beliefs were only luckily true.

One problem facing this proposal, however—which we alluded to above—is that it is far from obvious that we do think that in order to have knowledge you must be justified in what you believe, at least if by 'justification' here we mean that the agent concerned is able to cite good grounds in favour of what she believes. In order to see this, consider the following case. Imagine that our agent—we'll call him 'Chick'—possesses a highly reliable ability to tell the difference between male and female chicks. Chick believes that he is distinguishing between the chicks by using his sight and touch, but let us stipulate that he's mistaken in this regard and that he is actually doing this via his smell. Suppose

further that Chick doesn't have any good reason for thinking that he's reliable in this regard. For example, perhaps the reason why he believes that he's reliable is because someone he trusts told him this, but that this person was in fact trying to deceive him on this score but accidentally told him the truth nonetheless. It should be clear that Chick is unable to offer any good grounds in favour of what he believes. Nevertheless, it is far from obvious that Chick does not know that, for example, the two chicks he has in his hands are of a different gender. After all, he really does have a highly reliable ability to tell the two apart. Moreover, given that he has this ability, it is not a matter of luck that his belief is true—he couldn't have easily been mistaken—and thus he satisfies the constraints laid down by both the ability and anti-luck platitudes even though he lacks justification for his belief.

Intuitions about what to say about such cases differ widely, with some epistemologists arguing that Chick lacks knowledge and others arguing that he has knowledge. If you think that Chick does have knowledge then the conclusion to draw is that one can meet the constraints laid down by the ability and anti-luck platitudes without thereby meeting a justification condition. In particular, it seems that one can satisfy the ability platitude without thereby meeting the justification condition, such that it is only the satisfaction of the former that is essential for knowledge. On this view, then, the conclusion that one should draw is that knowledge should be defined as non-lucky true belief that is the product of ability. In contrast, if you think that Chick *lacks* knowledge then you are committed to holding that there is at least sometimes more to knowledge than a non-lucky true belief that is the product of ability.

Those epistemologists who ascribe knowledge in cases like that of Chick are called *externalists*, while those who deny knowledge to Chick are called *internalists*. Essentially, the debate between externalists and internalists boils-down to whether you think that knowledge requires justification, with internalists making this demand—and so denying knowledge to Chick—and externalists allowing that there are cases in which agents possess knowledge even while lacking justification for what they believe—thereby enabling them to attribute knowledge to Chick.⁶

Externalists hold that knowledge is often relatively easy to come by. Indeed, they often allow that very small children and intellectually sophisticated animals—neither of which are likely to have beliefs that satisfy a justification condition—can have knowledge. In contrast, internalists hold that knowledge is much harder to come by. Notice, though, that this isn't a strike against internalism in itself, since it isn't in any way absurd to suppose that perhaps we do know an awful lot less than we think we know.

The debate between externalism and internalism has seemed to many to be intractable; a straightforward clash of intuition that will not admit of a resolution. Nevertheless, one conciliatory proposal in this regard that has been relatively popular is to argue that externalists and internalists are in effect speaking past one another by focussing on different ‘grades’ of knowledge. That is, one might argue that we need to distinguish between a low-grade type of knowledge—what is sometimes called ‘brute’ or ‘animal’ knowledge—and a high-grade type of knowledge—what is sometimes called ‘reflective’ knowledge. The thought is that while it can suffice for low-grade knowledge to merely have a true belief which satisfies the constraints laid down by the anti-luck and ability platitudes, if one wishes to have high-grade knowledge then it is essential that one also possesses, in addition, a justification for one’s belief.⁷

The advantage of viewing the matter in this way is that one can do justice to both externalist and internalist intuitions. On the one hand, one accommodates the externalist intuition that Chick does count as possessing *bona fide* knowledge. On the other hand, one accommodates the internalist intuition that there is something epistemically deficient about the epistemic standing of Chick’s belief. We would, after all, prefer to possess high-grade knowledge rather than low-grade knowledge—*viz.*, it would be better to be like Chick but to have a justification for the target belief than to be like Chick and lack a justification. Perhaps, then, the choice between externalism and internalism in epistemology is not as stark as it first appears.

5. ANTI-LUCK EPISTEMOLOGY

It was noted earlier that cases like Gettier-style cases and the lottery case demonstrate that having a true belief that is formed through ability does not suffice to ensure that one has a non-lucky true belief, and thus does not suffice for knowledge. The moral drawn from this was that the ability and anti-luck platitudes impose distinct demands on knowledge. One might wonder, however, whether having a non-lucky true belief might suffice for knowledge, in the sense that such a belief is of its nature acquired through ability and so satisfies the constraint laid down by the ability platitude. For if this were true, then it would seem that it is the anti-luck platitude which is the dominant epistemological platitude, with the ability platitude essentially just a product of the anti-luck platitude. That is, if this were true

then one could simply understand knowledge as non-lucky true belief, what we might call an *anti-luck epistemology*.⁸

Interestingly, however, simply meeting the constraint imposed by the anti-luck platitude will not suffice for knowledge, and it is important to understand why. Consider the following case. Imagine an agent—we'll call him 'Temp'—who is in a room and who regularly forms his belief about the temperature of the room by looking at the thermometer in the corner. Suppose further that this is a perfectly reliable way of forming beliefs about the temperature of the room, in the sense that every time he forms a belief in this way his belief is true. Here's the twist. Unbeknownst to Temp, the thermometer is broken and is simply fluctuating randomly within a certain range. That the thermometer is broken doesn't in any way undermine the reliability of the belief so formed, however, for the simple reason that there is someone hidden in the room next to the thermostat who ensures that every time that Temp goes over to the thermometer to find out the temperature, the reading on the thermometer matches up with the temperature in the room.

What is significant about this case is that Temp's true belief about the temperature in the room is clearly not lucky. After all, given the existence of the person hidden in the room, he is destined to form a true belief by forming his belief in this fashion; his true belief could not have easily been mistaken. Clearly, however, Temp does not have knowledge in this case, since one cannot come to know what the temperature of a room is simply by looking at a broken thermometer (any more than one can find out the time by simply looking at a broken clock). Moreover, the right diagnosis of why this is the case seems to be that Temp's true belief is not in any way a product of his abilities. Indeed, if anything, the truth of his belief is entirely the product of someone else's abilities—i.e., that of the person hidden in the room helpfully adjusting the thermostat. It thus follows that we cannot simply regard the constraint laid down by the ability platitude as a consequence of the constraint laid down by the anti-luck platitude. More generally, given the conclusion that we drew earlier that we could not treat the constraint laid down by the anti-luck platitude as flowing from the constraint laid down by the ability platitude, we can conclude that the constraints laid down by these platitudes are independent of one another in *both* directions. We are thus back to the thesis that knowledge is non-lucky true belief that is the product of ability.

As it happens, cases like the Temp case also illustrate why a certain kind of radical externalism about knowledge is unsound. One such view is reliabilism, and it holds that there is nothing more to knowledge than reliably formed true belief (i.e., true belief that is formed in a way that is more likely

to lead to the truth than to falsehood).⁹ On this view, while having a justification for one's belief might be epistemically advantageous, in that justified beliefs are more likely to be reliably formed beliefs, they are not essential, since what is important is just that the belief is reliably formed. For example, such a view, in keeping with other externalist proposals about knowledge, can allow that Chick has knowledge since his belief is, after all, being formed in a reliable fashion.

Notice, though, that Temp's belief is also formed in a highly reliable fashion, and yet he does not count as having knowledge. The same diagnosis for why Temp lacks knowledge also explains where the reliabilist account goes awry. It is not reliability *per se* that we are interested in when it comes to knowledge, but rather the specific kind of reliability that is directly connected to the agent's cognitive abilities. That's why Chick can count as having knowledge—at least by externalist lights—while Temp cannot: Chick's true belief, but not Temp's, is the product of his reliable cognitive abilities, and not simply a true belief that is reliable.

6. VIRTUE EPISTEMOLOGY

Whereas reliabilism is not attractive as an account of knowledge, there is a closely related view which retains much of the spirit of reliabilism but which is not susceptible to some of the same problems. This view is known as *virtue epistemology*. The most basic form of virtue epistemology holds, in essence, that knowledge is non-lucky true belief that is formed via the reliable cognitive abilities of the agent.¹⁰ So construed, the view responds very directly to the two platitudes that we have been discussing here. According to this proposal, reliability in one's belief-forming processes is important, but the mere fact that a process is reliable will not suffice to ensure that an agent has knowledge, even if one adds the further proviso that the agent's belief is not luckily true. Rather, what is required is that the reliability be directly related to the cognitive abilities of the agent. In effect, what this form of virtue epistemology does is make explicit what is already implicit in the ability platitude—*viz.*, that when we think of an agent's abilities as being knowledge-conducive, we are already thinking of them as reliable (that is, an unreliable cognitive ability is not a *bona fide* cognitive ability at all).

We noted earlier that Gettier-style cases and the lottery case demonstrate that merely having a true belief formed through (reliable) ability will not suffice for knowledge since it will not suffice to exclude the possibility that the belief is only true as a matter of luck. Some virtue epistemologists,

however, have argued that there is a way of dealing with this problem which ensures that we can treat the constraint imposed by the anti-luck platitude as simply flowing from the constraint imposed by the ability platitude after all.

According to this version of virtue epistemology, knowledge is to be understood, roughly, as true belief that is because of cognitive ability. Notice that all mention of the true belief being non-lucky has been dropped. The thought is that so long as the agent's true belief is properly attributable to her cognitive ability—i.e., *because of* her cognitive ability—then this will suffice by itself to eliminate any knowledge-undermining epistemic luck.¹¹

On the face of it, this proposal might seem quite appealing. Take Edmund's belief regarding the time. Although he possesses the relevant reliable cognitive abilities—he knows how to tell the time, for example—it is not because of these abilities that his belief is true, but rather down to the good fortune that he happened to look at the clock at the only time in the day in which it was displaying the right time. Or consider Lottie's belief that she has lost the lottery. Again, although this belief is the product of her reliable cognitive abilities—she has worked out the ramifications of the odds of her winning flawlessly, for example—it is not because of these abilities that her belief is true since had she been holding a winning ticket right now she would have continued to believe that she had lost.

It thus seems that we don't need to think of the anti-luck platitude as imposing a separate constraint on knowledge after all, just so long as we understand the relationship between true belief and cognitive ability correctly.

Unfortunately, this more robust form of virtue epistemology which dispenses with a separate anti-luck constraint on knowledge—while certainly offering a very elegant account of knowledge—does not pass muster. The reason for this is that there are cases of knowledge where the agent's true belief is not because of her cognitive ability, and cases in which the agent's true belief is because of her cognitive ability which are not thereby cases of knowledge.

The best way of illustrating the first claim is by considering cases of testimonial knowledge. What is significant about testimonial knowledge is that because of its social dimension it is knowledge that one can acquire by riding 'piggy-back' on the cognitive abilities of others. Suppose, for example, that our agent—we will call her 'Jenny'—gets off the train in an unfamiliar town and asks the first person she meets for directions. Suppose further that this person has first-hand knowledge of the area and communicates this to Jenny, thereby enabling her to form a true belief about where she needs to go. Intuitively, we would say that Jenny *knows* the way to go. Indeed, if one cannot gain testimonial

knowledge in this way then it seems that we know an awful lot less than we think we do. Interestingly, however, it isn't at all right to say that Jenny's true belief is because of *her* cognitive abilities, as opposed, for example, to it being due to her *informant's* cognitive abilities (or at least their combined cognitive abilities). It thus appears to be a case in which an agent has knowledge even while having a true belief that is not because of her cognitive ability.¹²

It is important to be clear about the claim that is being made here. The thesis is not, for example, that Jenny is not exercising her cognitive abilities in a relevant fashion *at all*—i.e., this example is not a counterexample to the ability platitude.¹³ After all, in order to maintain the intuition that Jenny has knowledge in this case, we need to suppose that she is indeed exercising a great deal of judgement. For one thing, we would expect her to be discriminating about who she asks for directions—i.e., if the first person she met had turned out to be a small child or someone who was clearly a tourist, then we would have expected her to find another potential informant. For another, we would expect Jenny to exercise discrimination when it comes to evaluating the truth of the testimony provided to her by the informant. If this testimony were clearly false, for example, then we would expect her to recognise this and disregard it accordingly.

Jenny's belief is thus a product of her cognitive abilities. The key point, however, is that the social nature of testimonial knowledge of this sort entails that it is not because of her cognitive abilities that her belief is true, and thus cases like this do count against the kind of virtue epistemology at issue.

The best way of illustrating the second kind of problem for this type of virtue epistemology—i.e., that there are cases in which the agent has a true belief that is because of cognitive ability but which is not thereby a case of knowledge—is via the following sort of scenario. Suppose that our agent—we'll call him 'Barney'—gets a good clear view of a barn in good cognitive conditions (e.g., good lighting etc.) and on this basis forms the belief that there is a barn in front of him. Suppose further that Barney has lots of relevant cognitive abilities which are working to enable him to form this belief and that his belief is also true—he is indeed looking at a barn. Here's the twist. Imagine that, unbeknownst to Barney, he is in 'barn façade county', a county where all the barn-shaped objects bar this one are not in fact barns at all but fakes (perhaps, for example, there is some elaborate tax-dodge taking place here). If Barney had been looking at one of these fakes, then he would have formed the false belief that what he is looking at is a barn rather than the true belief that he actually forms. Does Barney know that what he is looking at is a barn? Surely not. His true belief, after all, is just too lucky to count as knowledge since he could very easily have been mistaken in this regard. Notice, however,

that the truth of Barney's belief, while due to luck, *does* seem to be because of his cognitive abilities, in that it is these cognitive abilities that have led him to form a true belief.¹⁴

We can further emphasise this point by noting that the sort of epistemic luck in play in this case is very different from that in play in standard Gettier-style cases. In Gettier-style cases, such as the case involving Roddy described above, it is plausible to suppose that the truth of the agent's belief is not because of his cognitive abilities, and the reason for this is that something intervenes between the agent's belief and her cognitive abilities, albeit in such a way that does not prevent the agent from having a true belief. In the case of Roddy, for example, his cognitive abilities do not hook-up with the target of his belief—the sheep in the field—at all, but instead are led astray by the big hairy dog that is standing in front of the sheep. All Gettier-style luck is of this 'intervening' sort.

Notice, however, that the kind of epistemic luck in play in the example involving Barney is not of this intervening sort. After all, Barney really does see a barn in the sense that his cognitive abilities do indeed put him in touch with the target of his belief, the barn. Instead, the epistemic luck in play here is of an 'environmental' variety, in that it simply concerns the fact that this is not an epistemically friendly environment—i.e., it is not an environment in which one's cognitive abilities can easily enable one to have a true belief. However, it is because the epistemic luck in play is not of the intervening sort that it seems entirely right to say that Barney's true belief is because of his cognitive ability, unlike in Gettier-style cases, like that involving Roddy.

It thus follows that one can have a true belief that is because of one's cognitive ability and yet lack knowledge. Again, then, we find that we need to respect both the anti-luck and the ability platitudes.

7. CONCLUDING REMARKS

The less robust version of virtue epistemology was thus right all along: knowledge is non-lucky true belief that is the product of the agent's reliable cognitive abilities. We have thus answered one of the central questions of epistemology—*viz.*, 'What is knowledge?'. Notice, however, that we have left many more central issues of epistemology unanswered. For example, why does knowledge have this structure? It is, after all, far from obvious why knowledge should have these properties. A related question in this regard concerns why we regard knowledge as such an important philosophical notion,

a question that we might expect our analysis of knowledge to throw some light on.¹⁵ And perhaps an even more pressing epistemological issue that has not been engaged with here is the question of whether we have any knowledge.¹⁶

Nevertheless, while there are many questions that have not been answered in this essay, the hope is that we have learned enough about epistemology to gain a grip on what this key area of philosophy involves and thus provide a foundation for further explorations in this direction.¹⁷

FURTHER READING

I have offered further readings that are specific to particular points raised in this essay in the endnotes as we have gone along, so the purpose of this section is to guide you towards some general readings in epistemology that you may find useful. The literature in this regard is vast, so rather than aim for comprehensiveness I will instead simply focus on directing you towards a few key texts within a number of relevant categories. (For a comprehensive survey of further readings in epistemology, see Pritchard (2006, 163-6)).

To begin with, there are some freely available articles on the internet that provide a useful introduction to the theory of knowledge. The best of these is almost certainly Steup (2005). If you are a member of an academic institution, then you may also have free access to Klein (2005) as well (if not, then viewing this article will require a subscription). For a general introduction to epistemology which does not presuppose any prior familiarity with the field, see Morton (1997) or Pritchard (2006). For a more advanced textbook on epistemology that offers a non-partisan overview of the central issues, see Audi (1998). Note, however, that there are some first-rate, albeit more idiosyncratic, advanced textbooks available. See, for example, Craig (1990), Williams (2000) and Welbourne (2002). When it comes to reference works, the only dictionary of epistemology that I am aware of is Blauw & Pritchard (2005). There are, however, a number of encyclopaedias, the best of which is currently Dancy & Sosa (1993). There are also some excellent volumes available which, while not being encyclopaedias as such, serve a similar purpose. See, for example, Greco & Sosa (1999), Moser (2002) and Steup & Sosa (2005). Finally, there are a number of superb anthologies available in epistemology. The most accessible is probably Bernecker (2006), but see also Bernecker & Dretske (2000) and Sosa & Kim (2000).

REFERENCES

- Audi, R. (1998). *Epistemology: A Contemporary Introduction to the Theory of Knowledge*, Routledge, London.
- Bernecker, S. (ed.). (2006). *Reading Epistemology*, Blackwell, Oxford.
- Bernecker, S. & Dretske, F. (eds.). (2000). *Knowledge: Readings In Contemporary Epistemology*, Oxford University Press, Oxford.
- Blaauw, M. & Pritchard, D. H. (2005). *Epistemology A-Z*, Edinburgh University Press, Edinburgh.
- Chisholm, R. (1977). *Theory of Knowledge* (2nd Ed.), Prentice-Hall, Englewood Cliffs, New Jersey.
- Craig, E. (1990). *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*, Clarendon Press, Oxford.
- Dancy, J. & Sosa, E. (eds.). (1993). *A Companion to Epistemology*, Blackwell, Oxford.
- Gettier, E. (1963). 'Is Justified True Belief Knowledge?', *Analysis* 23, 121-3.
- Goldman, A. (1976). 'Discrimination and Perceptual Knowledge', *The Journal of Philosophy* 73, 771-91.
- (1986). *Epistemology and Cognition*, Harvard University Press, Cambridge, Massachusetts.
- Greco, J. (1999). 'Agent Reliabilism', *Philosophical Perspectives* 13, 273-96.
- (2000). *Putting Skeptics in Their Place: The Nature of Skeptical Arguments and Their Role in Philosophical Inquiry*, Cambridge University Press, Cambridge, UK.
- (2002). 'Knowledge as Credit for True Belief', *Intellectual Virtue: Perspectives from Ethics and Epistemology*, (eds.) M. DePaul & L. Zagzebski, Oxford University Press, Oxford.
- Greco, J. & Sosa, E. (eds.). (1999). *The Blackwell Guide to Epistemology*, Blackwell, Oxford.
- Klein, P. (2005). 'Epistemology', *Routledge Encyclopedia of Philosophy*, (ed.) E. Craig, Routledge, London, <http://www.rep.routledge.com/article/P059>.
- Hawthorne, J. (2004). *Knowledge and Lotteries*, Oxford University Press, Oxford.
- Kornblith, H. (ed.) (2001). *Epistemology: Internalism and Externalism*, Blackwell, Oxford.
- Kvanvig, J. (2003). *The Value of Knowledge and the Pursuit of Understanding*, Cambridge University Press, Cambridge, UK.
- Lackey, J. (2007). 'Why We Don't Deserve Credit for Everything We Know', *Synthese* 156.
- Morton, A. (1997). *A Guide Through the Theory of Knowledge*, Blackwell, Oxford.
- Moser, P. K. (ed.). (2002). *The Oxford Handbook of Epistemology*, Oxford University Press, Oxford.
- Pritchard, D. H. (2002). 'Recent Work on Radical Skepticism', *American Philosophical Quarterly* 39, 215-57.
- (2005). *Epistemic Luck*, Oxford University Press, Oxford.
- (2006). *What is this Thing Called Knowledge?*, Routledge, London.
- (2007). 'Anti-Luck Epistemology', *Synthese* 156.
- Russell, B. (1948). *Human Knowledge: Its Scope and its Limits*, George Allen & Unwin, London.
- Shope, R. K. (1983). *The Analysis of Knowing: A Decade of Research*, Princeton University Press, Princeton, New Jersey.
- Sosa, E. (1991). *Knowledge in Perspective: Selected Essays in Epistemology*, Cambridge University Press, Cambridge, UK.
- (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, Oxford University Press, Oxford.
- Sosa, E. & Kim, J. (eds.). (2000). *Epistemology: An Anthology*, Blackwell, Oxford.
- Steup, M. (2005). 'The Analysis of Knowledge', *Stanford Encyclopaedia of Philosophy*, (ed.) E. Zalta, <http://plato.stanford.edu/entries/knowledge-analysis/>.

- Steup, M. & Sosa, E. (eds.). (2005). *Contemporary Debates in Epistemology*, Blackwell, Oxford.
- Welbourne, M. (2002). *Knowledge*, Acumen, Aldershot.
- Williams, M. (2001). *Problems of Knowledge: A Critical Introduction to Epistemology*, Oxford University Press, Oxford.
- Williamson, T. (2000). *Knowledge and its Limits*, Oxford University Press, Oxford.
- Zagzebski, L. (1996). *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*, Cambridge University Press, Cambridge, UK.

NOTES

- ¹ Some—most notably Williamson (2000)—have argued that this task cannot be completed, and thus that we should regard knowledge as unanalyzable. For my own part, I am far more optimistic on this score, as we will soon see.
- ² The example of the stopped clock comes from Russell (1948, 170-1), though he did not himself recognise that it was an example of justified true belief which is not a case of knowledge.
- ³ The *locus classicus* for discussions of Gettier-style cases is Shope (1983).
- ⁴ This example is adapted from one offered by Chisholm (1977, 105).
- ⁵ For more on the lottery puzzle, see Hawthorne (2004).
- ⁶ For more on the externalism/internalism distinction, see Kornblith (2001).
- ⁷ The *locus classicus* for discussions of ‘animal’ and ‘reflective’ knowledge is Sosa (1991).
- ⁸ For further discussion of anti-luck epistemology, see Pritchard (2005; 2007).
- ⁹ The *locus classicus* for discussions of reliabilism is Goldman (1986).
- ¹⁰ See, for example, Greco (1999; 2000), who describes this type of virtue epistemology as “agent reliabilism”.
- ¹¹ Versions of virtue epistemology of roughly this sort can be found in Sosa (e.g., 1991; 2007), Zagzebski (1996) and later work by Greco (e.g., 2002).
- ¹² This example is adapted from one offered by Lackey (2007), albeit to support a slightly different thesis.
- ¹³ This is the moral that Lackey (2007) draws from this example.
- ¹⁴ This example is due to Ginet, but it first appeared in print in Goldman (1976).
- ¹⁵ The *locus classicus* for discussions of the value of knowledge is Kvanvig (2003).
- ¹⁶ According to the radical sceptic, we have very little knowledge, if any. For an overview of the contemporary literature on radical scepticism, see Pritchard (2002).
- ¹⁷ I am grateful to John Shand for some very helpful comments on an earlier draft.