THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Anti-Luck Epistemology and the Gettier Problem

Citation for published version:

Pritchard, D 2015, 'Anti-Luck Epistemology and the Gettier Problem', *Philosophical Studies*, vol. 172, no. 1, pp. 93-111. <https://doi.org/10.1007/s11098-014-0374-0>

Digital Object Identifier (DOI):

[10.1007/s11098-014-0374-0](https://doi.org/10.1007/s11098-014-0374-0)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Philosophical Studies

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

ANTI-LUCK EPISTEMOLOGY AND THE GETTIER PROBLEM

Duncan Pritchard

University of Edinburgh

ABSTRACT. A certain construal of the Gettier problem is offered, according to which this problem concerns the task of identifying the anti-luck condition on knowledge. A methodology for approaching this construal of the Gettier problem—*anti-luck epistemology*—is set out, and the utility of such a methodology is demonstrated. It is argued that a range of superficially distinct cases which are meant to pose problems for anti-luck epistemology are in fact related in significant ways. It is claimed that with these cases properly understood, anti-luck epistemology is able to offer a suitable diagnosis of them which doesn't threaten the necessity of the anti-luck condition for knowledge.

1. THE GETTIER PROBLEM

There are several issues that one could take to constitute the Gettier problem. One question which could fall under this description is whether knowledge is susceptible to analysis—that is, can one offer an informative (possibly non-circular) definition of knowledge?¹ Call this formulation of the Gettier problem the *analytical problem*. The analytical problem is raised by the fact that Edmund Gettier's (1963) famous paper, and the ensuing debate, highlights the point that knowledge does not seem to be susceptible to an informative analysis. Gettier himself demonstrated that knowledge cannot be defined as justified true belief, at least insofar as justification is a non-factive notion (i.e., such that one can be justified in having a false belief). But it also became apparent in the debate that

followed that there is no straightforward way of adding a further condition (or conditions) to either justified true belief or true belief which would suffice as an informative analysis. (Defining knowledge as justification plus whatever extra condition is required to eliminate Gettier-style cases is obviously not very informative).

My concern here is not the analytical problem. I've offered my own solution to this difficulty elsewhere.² Instead, my focus will be on a closely related formulation of the Gettier problem. Gettier's famous paper didn't just demonstrate that knowledge wasn't (non-factively) justified true belief. He also demonstrated that one could have a justified true belief which was nonetheless subject to knowledge-undermining epistemic luck. This raises a challenge regarding what condition or conditions must be imposed on knowledge in order to exclude such luck. Call this formulation of the Gettier problem, the *anti-luck problem*.

The kind of knowledge-undermining luck that is at issue is these days known as *veritic epistemic luck*, where this means that it is a matter of luck that one's belief is true.³ That knowledge is incompatible with veritic epistemic luck is widely accepted, to the extent that one can plausibly characterise it as a platitude about knowledge.⁴ Finding out that justified true belief is compatible with veritic epistemic luck was a surprise, however, in that one would antecedently suppose that the justification condition would suffice to exclude this type of epistemic luck. Typically, after all, justified true belief is not subject to veritic epistemic luck.

Another way of putting this point is that Gettier showed that two of our core epistemological commitments are in tension with one another—*viz.*, the anti-luck platitude, and the claim that justification, *qua* the epistemic constituent of knowledge (in addition to true belief), is non-factive, in the sense of being compatible with false belief in the proposition justifiably believed. That justification is compatible with false belief was, after all, meant to be an obvious claim too. Surely one can have an excellent epistemic basis for one's belief, sufficient to meet the justification condition on any plausible construal, and yet that basis be nonetheless compatible with the falsity of the target belief?⁵ But if justification doesn't guarantee truth, then that leaves an opening for veritic epistemic luck, and hence for Gettier-style cases, for one can then imagine cases in which the agent's justified belief is true where the truth of the belief is completely divorced from the agent's justification.

One reaction that one might draw from this point is that we should look for a justification condition on knowledge which is factive. But this would be a mistake. For even though a case can be made for the idea that the justification condition on knowledge is plausibly factive as regards

certain kinds of knowledge, such as paradigm cases of perceptual knowledge, if one is to use this move as a way of dealing with the Gettier problem then one needs to argue for the factivity of justification across the board.⁶ That is, one would need a notion of justified belief which was almost certainly equivalent to knowledge.⁷ That would be extremely strong. It would probably mean that we have no inductive knowledge, for example, as this is precisely knowledge where our justification for believing the target proposition does not guarantee the truth of this proposition.

There is an even worse problem facing this way of responding to the anti-luck problem. For notice that all that immediately follows from the tension noted above between the anti-luck platitude and the non-factivity of justification is that a factive notion of justification would be *sufficient* to evade this problem. It does not follow, however, that a factive notion of justification is *necessary* to evade the anti-luck problem. As we will see below, we can retain the idea that justification is non-factive (i.e., such that it is at least sometimes compatible with false belief) and yet still resolve the anti-luck problem. It follows that resorting to a factive notion of justification solely in order to deal with the anti-luck problem is a kind of theoretical overkill.

In any case, the anti-luck problem arises out of Gettier's discovery that justified true belief is compatible with veritic epistemic luck, coupled with the further realisation that there is no straightforward way of eliminating veritic epistemic luck. We thus have the general shape of the anti-luck formulation of the Gettier problem—*viz.*, we need to determine what condition excludes veritic epistemic luck from knowledge.⁸

Note that the anti-luck problem, although it takes Gettier-style cases as its starting point, has a scope which extends significantly beyond these cases. After all, it isn't just Gettier-style cases which appeal to veritic epistemic luck, as such appeals are widespread in epistemology—consider, for example, the lottery problem (which we will be looking at further below). A solution to the anti-luck problem is thus in an important sense going to be much more than an anti-Gettier condition on knowledge.

Finally, note that while the anti-luck problem and the analytical problem are closely related, they are not the same difficulty. For example, it could be that determining the right condition to eliminate veritic epistemic luck will be enough to offer a complete analysis of knowledge, but there is no necessity that this be so. Indeed, I'll be arguing below that once we determine the right condition in this regard, then it becomes apparent that we will need further conditions on knowledge in order to be in possession of a complete analysis.

2. ANTI-LUCK EPISTEMOLOGY

Elsewhere I have argued for an approach to epistemological issues, and to the theory of knowledge in particular, that I term *anti-luck epistemology*. As we will see, this approach is particularly relevant to resolving the anti-luck problem.⁹

The starting point of anti-luck epistemology is the idea that if we take the claim that knowledge is incompatible with veritic epistemic luck seriously then it is incumbent upon us to unpack this claim. In particular, we should (i) offer a theory of luck, and (ii) offer an account of how, specifically, knowledge is incompatible with luck, and then (iii) put the two together in order to determine the anti-luck condition on knowledge. By doing so, we gain an understanding of what sort of condition is required to deal with the anti-luck problem, an understanding that we won't get by simply engineering an anti-luck condition in response to the anti-luck problem.

Here is an outline of how the anti-luck epistemology project pans-out. First, the account of luck. I defend what is known as the *modal account of luck*.¹⁰ Roughly, this holds that lucky events are those events which actually obtain but which don't obtain in a wide class of close possible worlds where the relevant initial conditions for that event are kept fixed.¹¹ So, for example, consider a paradigm case of a lucky event: a lottery win. On this view, this event is lucky because although it obtains in the actual world, it tends not to obtain in close possible worlds where the relevant initial conditions for the target event are kept fixed (i.e., the subject continues to buy a lottery ticket, the lottery continues to be fair, the odds of winning and the method of selecting the winner are not significantly different, and so on). In contrast, consider an event that is paradigmatically not lucky, such as the sun rising this morning.¹² This is an event which doesn't just obtain in the actual world, but which also obtains in all close possible worlds too.

There is a lot that can be said to motivate the modal account of luck, but I will confine myself to making a few key points. First, it is useful to note why it would be a mistake to treat lucky events as being, *contra* the modal account, merely low probability events which happen to obtain. The example of a lottery win might seem to suggest such a view, since this is of course an event that has astronomically long odds but which nonetheless obtains. But notice, however, that for most lotteries it's *both* the case that one's winning is a matter of luck *and* that one's losing is also a matter of luck (good luck in the first case, bad luck in the second).¹³ And yet in the latter case the event in question is a high probability event.

In this regard it is useful to compare lottery wins with normal bets that have astronomically long odds. Suppose, for example, that someone encouraged you to bet on whether I would win gold in the 100m sprint at the next Olympics, and let's stipulate that the odds of success are calculated such that they are roughly akin to the odds of a win in a national lottery (trust me, this would be about right!). If you won such a bet, this would certainly count as lucky. Losing such a bet wouldn't count as lucky though. Indeed, we would regard anyone who took such a bet as mad since it has no realistic chance of success. What is going on here?

The key point to note about these two cases is that although the odds in question are roughly identical, the modal profiles of the 'losing' event in each case are very different. Suppose one doesn't win the lottery. Given how lottery numbers are usually drawn, the possible world in which one wins the lottery is nonetheless very close—all that needs to change is that a few coloured balls need to fall in a slightly different configuration. In contrast, suppose one loses one's bet on me winning the next 100m Olympic sprint. The possible world in which I win this bet is very distant from the actual world, in that a great deal needs to change in order to get from the actual world to the winning world. (Indeed, it is hard to imagine what kind of strange run of events could transpire to enable me to win gold in the next Olympic sprint). This is why people 'bet' on lotteries but not on normal events with astronomically low odds, since it's a specific feature of lotteries that even despite the low odds in play the scenario in which one is a winner is nonetheless modally close.¹⁴

A second point to note about the modal account is that it can accommodate the fact that luck can come in degrees. Imagine that one is almost shot by a sniper, the bullet whizzing a few inches away from one's ear. Imagine now an identical scenario except that the bullet doesn't pass quite so close—it's a yard or so away rather than a few inches. Both events would be judged to be lucky, with the first event luckier than the first. These judgements would be confirmed by the modal account. In both cases, although one is not hit by the bullet in the actual world, there will be a wide class of close possible worlds in which one is shot. Moreover, the first event is luckier than the second because the possible worlds in which one is hit are closer to the actual world. In effect, our judgements about luck reflect a greater weighting on the closest possible worlds: an event which could so very easily not obtain will be luckier than one which could easily (but not very easily) not obtain.

This example also illustrates quite nicely how notions of luck and risk go hand-in-hand, with both of them captured by the modal account. In the scenario where the bullet passed by very closely, one was at greater risk of being shot than in the second scenario where it passed by further

away. Risky events are events that could very easily occur, with the degree of risk in play proportionate to just how easily they could happen. Our judgements about risk are thus tracking the degree of luck in play, just as we would expect.¹⁵

With the modal account of luck in hand, we can now turn to the issue of the specific sense in which knowledge excludes luck. We noted above that the problem of epistemic luck as highlighted by the Gettier-style cases is that knowledge is incompatible with veritic epistemic luck, which is luck that one's belief is true. With our account of luck in hand, we can be more specific about what we have in mind here. Treating the actual acquisition of the true belief—i.e., one's getting it right—as the target event, to get it right as a matter of luck is for it to be the case that in a wide class of close possible worlds where one acquires one's belief in the same way as in the actual world (i.e., the relevant initial conditions for the target event are kept fixed) one 'gets it wrong' (i.e., one forms a false belief). Moreover, the closer the possible worlds are where one forms a false belief on the same basis, the more veritically epistemically lucky one's belief in the target proposition is.¹⁶

It can be useful to think about this unpacking of the notion of veritic epistemic luck in terms of epistemic risk. On this proposal, a true belief is epistemically risky when there is a wide class of close possible worlds in which the very same actual basis for belief results in a false belief. Moreover, the degree of epistemic risk involved is determined by just how close the possible world is where the subject forms a false belief on the same basis.¹⁷

We are also now in a position to explain why certain kinds of epistemic luck are benign, in the sense of being compatible with knowledge. For example, take *evidential epistemic luck*. That it is a matter of luck that one has the evidence that one does is not incompatible with knowledge—one can perfectly well make lucky discoveries. Why is this form of epistemic luck compatible with knowledge? The answer lies in the fact that the kind of epistemic risk that we are trying to eliminate from knowledge with our anti-luck condition is a basis-relative notion. That is, we are interested in whether the basis we have for our true belief is epistemically risky in the sense that it would lead to false belief in close possible worlds. In a case where one gains knowledge from evidence acquired by luck, however, then although it is a matter of luck that one has the basis for belief that one does, it is not a matter of luck that forms a true belief on that basis—i.e., given one's basis, one will not tend to form false beliefs in close possible worlds. There thus need be no epistemic risk involved in cases of evidential epistemic luck.

We can get a sense of the utility of anti-luck epistemology by considering how it enables us to adjudicate between different candidate anti-luck conditions in the literature. In order to simplify matters, we will focus on two such conditions which have received a lot of critical attention: safety and sensitivity. Roughly, safety demands that one has a true belief that could not have easily been formed on the same basis and yet been false. In contrast, sensitivity roughly demands that one has a true belief such that, had what one believed been false, then one would not have believed it on the same basis.¹⁸

Note that both conditions have been advocated independently of the anti-luck epistemology as set out here. That is, they have each been put forward directly as plausible ways of eliminating knowledge-undermining epistemic luck of the kind found in Gettier-style cases. So construed, they both seem to fare equally well, at least initially. Consider a familiar Gettier-style case.¹⁹ A farmer, in good epistemic conditions (in good light, at close range, and so on) sees what he takes to be a sheep, and so forms the belief that there is a sheep in the field. While this belief is true, in that there is a sheep in the field, the farmer is not looking at a sheep but rather a sheep-shaped object (such as a hairy dog). The genuine sheep is hidden from view behind the sheep-shaped object. Intuitively, the farmer lacks knowledge even while having a justified true belief that there is a sheep in the field. The farmer's belief is manifestly unsafe. Given how he formed his belief, he could very easily have formed a false belief on the same basis—for example, if the sheep hidden from view had wandered out of the field. The farmer's belief is also insensitive. In the closest possible world where the target proposition is false, but where everything else stays fixed as much as possible—i.e., where there is no sheep hidden from view behind the sheep-shaped object that the farmer is looking at—the farmer will continue to believe this proposition on the same basis.

Safety and sensitivity also fare equally well when it comes to other kinds of knowledge-undermining epistemic luck. Consider the lottery problem. This is the task of explaining why one can't know that one has a losing lottery ticket merely by reflecting on the long odds involved. This is puzzling, since one can gain knowledge in cases where the odds are not nearly so much in one's favour. Indeed, one can discover that one has a losing lottery ticket by reading the result in a reliable newspaper, where the chances of a misprint are higher than the chances of one winning the lottery. The upshot is that knowledge is not a straightforward function of the odds that one's belief is true, which is surprising and hence calls for an explanation.

If knowledge requires either safety or sensitivity, however, then an explanation—or the start of one anyway—is at hand. This is because forming one’s belief that one has lost the lottery merely by reflecting on the odds involved, as opposed to reading the result in a reliable newspaper, is both unsafe and insensitive. It is unsafe, because one could very easily have formed a false belief via this method (i.e., had one’s ticket happened to win). And it is insensitive because had one’s ticket won one would have formed a false belief on this basis.

Given that safety and sensitivity fare equally well in these regards, how is one to adjudicate between them? One strategy is to pursue a listing of relative ‘pros’ and ‘cons’. Safety fares better than sensitivity when it comes to certain kinds of cases, for example, such as regards inductive knowledge. Pick an event that an agent truly believes has obtained because she has an excellent (non-factive) rational basis for thinking that this is so, but where the event is such that the agent didn’t see for herself that it obtained. Furthermore, let’s stipulate that this excellent rational basis is all entirely correct. Accordingly, if there is inductive knowledge at all, then this ought to be such a case. But the agent’s belief, while safe, is not sensitive. After all, had the event not obtained, then the agent would, on the same basis, have formed a false belief. In contrast, the belief is safe, since the agent could not have easily formed a false belief on this basis, given the correctness of her rational grounds for believing that the event in question obtained.²⁰

But that doesn’t mean that safety is assured victory in this debate, since there are difficulties which afflict this principle too. In particular, it has been alleged that there is no consistent way of unpacking the principle such that it can accommodate cases of inductive knowledge while also delivering the right result on the lottery problem (in that the principle needs to be understood in a relatively liberal fashion in order to render the right verdict in the former case, and yet it needs to be understood in a quite restrictive fashion in order to deal with the latter difficulty).²¹ Moreover, safety faces the challenge of explaining why beliefs in necessary propositions—indeed, in any proposition, necessary or contingent, whose truth is modally stable, such that it couldn’t easily have been false—are not *de facto* safe.

Anti-luck epistemology offers us a way of navigating through these issues. In particular, anti-luck epistemology doesn’t just present us with a basis for preferring one of these principles over the other, but also affords us a way of construing the preferred principle such that it can evade the difficulties that have been levelled at it. Consider again the formulation of veritic epistemic luck that we offered above as a result of employing the anti-luck epistemology methodology. This held that a

true belief is subject to veritic epistemic luck when there is a wide class of close possible worlds in which the very same actual basis for belief results in a false belief. Moreover, we noted that the degree of veritic epistemic luck involved is determined by just how close the possible world is where the subject forms a false belief on the same basis. With this in mind, we can now see what kind of anti-veritic epistemic luck condition on knowledge will be imposed by anti-luck epistemology. This will require a true belief such that there is not a wide class of possible worlds where the agent forms a belief on the same basis and yet forms a false belief. In particular, given the weighting given to the closest possible worlds in this regard, the anti-veritic epistemic luck condition will display a complete intolerance to false belief (formed on the same basis) in the closest possible worlds, with that tolerance gradually increasing as one moves further out, modally speaking. Eventually, as one gets to non-close possible worlds, there will be a complete tolerance to false belief. In short, the anti-veritic epistemic luck condition captures our intolerance to high levels of epistemic risk, where this is proportionate to the closeness of the error in question (i.e., false belief formed on the same basis).

One can see straight away that anti-luck epistemology favours safety over sensitivity, as safety is closest in spirit to the anti-veritic epistemic luck condition we have formulated. In particular, notice that whereas sensitivity is concerned with the closest possible world where the believed proposition is false, no matter how modally far-off that world is, safety is only concerned with the modal neighbourhood (i.e., in what could easily have been the case). Anti-luck epistemology backs up safety on this score, since it highlights the point that epistemic risk is concerned with what is taking in place in one's modal neighbourhood, and does not concern the possibility of error in far-off possible worlds. This is why safety offers us the right result when it comes to inductive knowledge. Given that the subject has an excellent (albeit non-factive), and entirely correct, rational basis for her inductive belief, it follows that there is no close possible world in which she believes falsely on the same basis, and hence her belief is safe. But that is compatible with there being a non-close possible world where the target proposition is false and yet the subject continues to believe the target proposition on the same basis regardless, and this is why the belief is insensitive. According to anti-luck epistemology, however, it would be a mistake to not attribute knowledge in this case, as there is in fact no epistemic risk involved in the subject's true belief.

Furthermore, notice that anti-luck epistemology motivates a particular delineation of the safety principle, such that it can evade the two problems noted earlier. First, how strict should the principle be? The safety-based response to the lottery problem might lead us to suppose that safety should be construed robustly as being incompatible with any false belief in the target proposition

(and on the same basis) in a close possible world. But how is that to be squared with the putative compatibility of safety and inductive knowledge (which suggests a weaker construal, such that there can be some degree of false belief of this kind in close possible worlds, so long as it is not too extensive)?

The rendering of the anti-veritic epistemic luck condition provided by anti-luck epistemology presents us a way of dealing with this problem, on account of how it requires us to accord greater weight to error (i.e., false belief formed on the same basis) in the very closest possible worlds. We therefore do not need to choose between a restrictive and a liberal construal of safety. Instead, we just need to note that the closer the error, the greater the level of epistemic risk. In this regard the lottery case is very different to that involving inductive knowledge. The whole point of a lottery, after all, is that the possibility that one wins is very close indeed; as we saw above, that's why people play lotteries (but don't generally bet on events with astronomically long odds). The risk of error, given that the subject would believe falsely on the same basis in the winning world, is thus very close too. In contrast, even if we can envisage a possible world not that far removed from the actual world in which the subject in the inductive case ends up with a false belief formed on the same basis (this is in fact debatable, given how the case is described, but we will let this pass), such a world is clearly not going to be anything like as close as the corresponding world in the lottery case. When safety is thus properly understood by being motivated in terms of an anti-luck epistemology, it can thus deftly evade this problem.

Now consider the second problem for safety that we noted, which is that one's belief in modally stable true propositions will be *de facto* safe. In order for this concern to get a grip, we would need to suppose that safety demanded that one evaluate the safety of a subject's belief that p by considering whether there are close worlds in which the agent forms the false belief *that* p on the same basis. If that were the right way to construe safety, then the problem would immediately arise, since of course where p is modally stable then there can be no close possible world in which the agent falsely believes that p (on the same basis or otherwise).

But notice that this is not the version of safety which is generated by anti-luck epistemology. Our interest is instead more broadly on whether the same actual basis for belief leads to false belief in close possible worlds. Any false belief formed in close possible worlds on the same basis is thus relevant to the safety of the target belief, and not just false beliefs formed in the very same proposition as in the actual world. With safety so construed, this problem disappears. Imagine, for example, that a subject forms a true belief in a mathematical proposition by flipping a coin. Clearly

there is no close possible world in which she forms a false belief in this proposition on this basis, since the proposition in question is true in all possible worlds. But this belief is nonetheless intuitively epistemically risky, and this borne out by the fact that there are close possible worlds where this way of forming one's beliefs will lead to false beliefs, such as the close possible world where the coin-tossing leads the subject to form a belief in a false mathematical proposition.

So anti-luck epistemology gives us a principled way of responding to the anti-luck problem. Moreover, it offers us a better way of dealing with the traditional debates about the anti-luck condition by enabling us to think more deeply about why we would endorse a particular kind of anti-luck condition. We've seen this point in action by looking at the debate between safety and sensitivity. Anti-luck epistemology doesn't just positively favour the safety principle on this score, but it also leads to a particular rendering of this principle that avoids the standard problems that are levelled at this condition.

One final point is in order on this score, which is the contrast between a *modest* and a *robust* version of anti-luck epistemology. Robust anti-luck epistemology would be the bold thesis that by solving the anti-luck problem and thereby identifying the anti-luck condition on knowledge one has thereby also solved the analytical problem. That is, knowledge is to be identified with true belief which meets the anti-luck condition laid down by anti-luck epistemology. Modest anti-luck epistemology, in contrast, is merely the claim that a necessary condition of knowledge is that one's true belief satisfies this condition.²²

I think the prospects for robust anti-luck epistemology are not promising, and let me briefly offer one reason why. When one knows, there is a certain direction of fit between one's cognitive success (i.e., one's true belief) and one's cognitive agency, in that one's cognitive success is to some significant extent attributable to one's exercise of relevant cognitive agency. Elsewhere I have referred to this as the *ability intuition*.²³ The problem is that one can satisfy the anti-luck condition and yet have a cognitive success that has nothing at all to do with one's cognitive agency. The anti-luck condition, after all, merely ensures that one's true belief has a certain modal profile—*viz.*, that belief and fact match across a specific range of worlds. Crucially, however, no modal profile can ensure that one's belief has the direction of fit that is at issue when it comes to the ability intuition.

Imagine, for example, that a helpful demon is intent on ensuring that one's beliefs in a certain domain are all true, a domain where one would form false beliefs if left to one's own devices. One way that the demon could achieve this result is by ensuring that every time one forms a belief he changes the facts so that one's belief comes out as true. This would be a clear-cut case where

one's cognitive success has nothing at all to do with one's cognitive agency—*a fortiori*, it is a case where ability intuition is not satisfied. Notice, however, that the target belief is *not* subject to veritic epistemic luck. Given how the belief was formed, it could not help but be true, and hence there is absolutely no epistemic risk involved in one forming one's beliefs in this way.²⁴

The upshot is that one's beliefs can have the relevant modal profile to ensure that they are not subject to veritic epistemic luck without thereby being in the market for knowledge because they fail to satisfy the ability intuition. This means that while modest anti-luck epistemology is a promising research programme, the same cannot be said of robust anti-luck epistemology.

3. PURELY MODAL VERITIC EPISTEMIC LUCK

We are now in a position to evaluate the response to the anti-luck problem offered by anti-luck epistemology. What I want to focus on here is a family of cases that have been offered against safety as a necessary condition on knowledge that I think share some core characteristics.²⁵

Consider first the following counterexample to safety as a necessary condition on knowledge that has been offered by Ram Neta and Guy Rohrbaugh (2004):

WATER: "I am drinking a glass of water which I have just poured from the bottle. Standing next to me is a happy person who has just won the lottery. Had this person lost the lottery, she would have maliciously polluted my water with a tasteless, odorless, colorless toxin. But since she won the lottery, she does no such thing. Nonetheless, she almost lost the lottery. Now, I drink the pure, unadulterated water and judge, truly [...], that I am drinking pure, unadulterated water. But the toxin would not have flavored the water, and so had the toxin gone in, I would still have believed falsely that I was drinking pure, unadulterated water [...] Despite the falsity of my belief in the nearby possibility, it seems that, in the actual case, I know that I am drinking pure, unadulterated water."
(Neta & Rohrbaugh 2004, 399-400)

The claim is that the belief in question, while unsafe, is nonetheless a case of knowledge. I think that we should resist this verdict. In the circumstances as described, the subject's basis for belief—which will involve the sensory experience of a water-like substance—is such that she would form a false belief on this basis in close possible worlds. In particular, whereas in normal conditions such a basis for belief would not be subject to any epistemic risk, in these conditions this basis for belief is epistemically risky. So why would we attribute knowledge in such a case?

I think what is going on here is that cases like these are effectively purely modal versions of the famous barn façade case.²⁶ Recall that in the original barn façade case the subject's actual

environment was such that she could so very easily have been looking at a barn façade rather than a barn, and hence it was judged to be down to luck that she ended up with a true belief. Accordingly, since veritic epistemic luck is incompatible with knowledge, it was claimed that the subject lacked knowledge. The barn façade case is different from standard Gettier-style cases, however, in that the veritic epistemic luck in play is of an *environmental* rather than *intervening* variety.

Compare the barn façade case with the ‘sheep’ case discussed earlier. In the latter, the farmer forms her true belief that there is a sheep in the field not by looking at a genuine sheep but rather by looking at a sheep-shaped object. In contrast, in the barn façade case the agent is genuinely looking at a real barn—i.e., nothing intervenes between the belief and the target fact, as happens in standard Gettier-style cases—it is just that she is in an environment in which a belief based on the appearance of a barn will be nonetheless veritically epistemically lucky.²⁷

It should be granted that the intuition that one lacks knowledge is weaker when the variety of veritic epistemic luck in play is environmental rather than intervening. There is a good reason for this. Recall that we noted earlier in our discussion of the ability intuition that knowledge requires that one’s cognitive success be significantly attributable to one’s cognitive agency. In cases of environmental epistemic luck, however, the subject’s cognitive agency is playing more of a role in her cognitive success than in cases of intervening epistemic luck. It is thus unsurprising that cases of the former don’t elicit such a strong intuition that the agent lacks knowledge.

This phenomenon becomes even more marked once we focus on a purely modalized version of the barn façade case. Imagine now that we vary the case such that although there are in fact no barn façades in the subject’s environment, there could very easily have been, such that in a wide class of near-by possible worlds it is still the case that the agent forms a false belief on the same basis. This is a particular kind of environmental epistemic luck which specifically concerns the subject’s modal environment. The intuition that the subject in such a case lacks knowledge is even weaker, and the reason for this is that when the veritic epistemic luck is entirely due to factors in the subject’s modal environment, then the subject’s cognitive success can be as creditable to her cognitive agency as it is in normal cases of knowledge. Indeed, I’ve argued elsewhere that in such cases the subject exhibits a genuine *cognitive achievement*—i.e., a cognitive success that is sufficiently due to one’s exercise of cognitive agency that it counts as being *because of* one’s cognitive agency. There is thus, by everyone’s lights, a positive epistemic status that accrues to the agent’s belief which

is absent in other cases of veritic epistemic luck, particularly cases involving intervening epistemic luck.

It should be clear how WATER is akin to a modalized version of the barn façade case, in that it essentially involves modal environmental epistemic luck.²⁸ After all, the lottery winner doesn't actually poison the water, so there's nothing taking place in the actual environment which undermines the safety of the agent's belief so formed. But the modal environment is nonetheless such that our agent could very easily on this basis have formed a false belief. That's why her belief is subject to veritic epistemic luck. Given this analogy between the two cases, it follows that we should treat the agent's true belief in WATER as enjoying a kind of positive epistemic status which is lacking in normal Gettier-style cases, even though it is nonetheless subject to veritic epistemic luck.

Indeed, notice that cases of modal environmental epistemic luck are effectively *epistemic Frankfurt-style cases*. Recall that in the original cases offered by Harry Frankfurt (1969), an agent opts for a course of action but, had they opted for a different course of action, an external factor would have intervened to ensure that they would act as they actually acted. As a result, in such cases the agent cannot act other than she does. Notably, however, in the actual case, since nothing in fact intervenes to affect her actions, it does seem entirely right to say that the agent freely chose the course of action that she adopts. It follows that one's actions can be appropriately due to one's agency, and hence free, even when one couldn't have acted other than one did.²⁹ More generally, what Frankfurt-style cases demonstrate is that what is going on in one's purely modal environment has no bearing on attributions of agency.³⁰

The WATER case is not completely analogous (albeit cast along epistemic lines) to a Frankfurt-style case, of course, in that it is not true that had the subject been disposed to believe other than she did then the relevant intervention would have occurred to prevent this from happening. But it is analogous to the extent that in the actual case there is in fact no intervention taking place, with this intervention merely being something which occurs in close possible worlds. The same goes for the modalized version of the barn façade case, which in this regard shares its characteristics with WATER. What all these cases share is the fact that they involve veritic epistemic luck the source of which is purely modal.

We can easily imagine an epistemic case which is completely analogous to a Frankfurt-style case. Indeed, in a recent paper Christoph Kelp (2009) presents just such a scenario, which runs as follows:

DEMON: A demon wants Chris to form the belief that the time is 8.22am when he comes down the stairs first thing in the morning (the demon doesn't care whether the belief is true). Since he is a demon, with lots of special powers, he is able to ensure that Chris believes this proposition (e.g., by manipulating the clock). Now suppose that Chris happens to come downstairs that morning at exactly 8.22am, and so forms the belief that the time is 8.22am by looking at the accurate clock at the bottom of the stairs. Accordingly, the demon achieves what he wants without having to do anything.

In DEMON we have all the features of a Frankfurt-style case, albeit adapted to apply to the epistemic realm. Chris could not do otherwise than believe that it's 8.22am, given the way he is forming his belief in this regard (i.e., by consulting the clock). That is, in all close possible worlds he will believe on this basis that it is 8.22am, and so he has a belief which is subject to veritic epistemic luck. Note, however, that although the demon will intervene across a wide range of close possible worlds to ensure that Chris believes this proposition (and on the same basis), in the actual world he does not intervene.

The moral to be drawn from DEMON is essentially the same as we should draw from other cases of purely modal veritic epistemic luck (such as WATER etc.), which is that Chris's cognitive success is as creditable to his cognitive agency as it would be in a corresponding case where there is nothing amiss taking place in the modal environment. In particular, Chris's cognitive success is sufficiently creditable to his cognitive agency to count as a cognitive achievement. Nonetheless, we should also register that, just as with the other cases just listed, Chris's cognitive success is also subject to veritic epistemic luck. Indeed, Chris is effectively forming his belief about the time by consulting a clock which is for practical purposes a stopped clock, which is obviously a very epistemically risky way of finding out the time.

Once we grant that the subjects in these cases are exhibiting a cognitive achievement, however, then the interesting question becomes whether we should go further and insist that knowledge is present in such cases, even despite the presence of veritic epistemic luck. I think that this is a key 'decision point' in the contemporary debate about the nature of knowledge. In particular, if one is impressed by virtue-theoretic treatments of knowledge, to the extent that one is attracted to the idea of having a theory of knowledge which is cast exclusively along virtue-theoretic lines (and which hence does not feature an independent anti-luck condition), then one might well be inclined to attribute knowledge in such cases. According to this proposal—which, in line with our contrast between modest and robust forms of anti-luck epistemology, is a form of *robust virtue epistemology*—knowledge is to be understood essentially a kind of cognitive achievement (i.e., as a cognitive success which is *because of* cognitive agency).³¹ Since even proponents of anti-luck

epistemology should (I claim) grant that the agent's belief in such cases amounts to a cognitive achievement, the key move here is the equating of a cognitive achievement with knowledge, something which the anti-luck epistemologist will dispute (due to the veritic epistemic luck in play), but which the robust virtue epistemologist will insist upon.

It is important to note how theory-driven the robust virtue epistemologist's move here is, however. For notice that in attributing knowledge in these cases we are thereby allowing that genuine knowledge can nonetheless involve high levels of epistemic risk—i.e., cognitive success that could very easily have been cognitive failure. As noted above, in the DEMON case one is effectively forming one's belief about the time by consulting what is for all practical purposes a stopped clock. And things are no better in the other cases. In the WATER case, for example, one is in an environment where what tastes just like ordinary water could very easily be something very different.

Furthermore, notice that anti-luck epistemology is in a position to explain away our conflicting intuitions in this regard, in that it grants that there is something positive to be said about the epistemic standing of the subject's belief, unlike in standard cases involving veritic epistemic luck (i.e., that it involves a genuine cognitive achievement). It is no wonder then that our initial judgements about whether the agent has knowledge tend to be so conflicted, since while the presence of a cognitive achievement pulls us toward attributing knowledge, the presence of the veritic epistemic luck inclines us not to attribute knowledge.

Moreover, notice that it makes a big difference to how we respond to such a decision point how we understand the theoretical trade-off in question. For example, if we cast it as a conflict between a particularly robust form of virtue epistemology (i.e., one on which knowledge is to be understood exclusively in virtue-theoretic terms as a cognitive achievement) and a specific modal condition on knowledge (i.e., safety), then it can seem like much-of-a-muchness which way one decides to jump. In contrast, once we understand the safety condition as motivated by anti-luck epistemology, such that it is capturing the idea that knowledge excludes significant levels of epistemic risk, then the costs of choosing the robust virtue-theoretic option become much more apparent.³²

I think we can put further pressure on the robust virtue-theoretic way of responding to modal environmental epistemic luck and epistemic Frankfurt-style cases by considering a scenario which shares the same essential features as these cases but where there is no veritic epistemic luck involved. Consider the following variation on the DEMON case:

DEMON*: A demon wants Chris to form a true belief about what the time is when he comes down the stairs first thing in the morning. The demon knows, however (but Chris does not), that the clock has been linked to the national lottery such that it will only be working that morning if Chris's lottery ticket is a winner; otherwise it will be malfunctioning and showing a time which, while plausible, is wrong. The demon is poised so that if the clock is malfunctioning, then he will intervene to ensure that it is showing the correct time when Chris consults it. As it happens, however, Chris does win the lottery, and hence the demon doesn't need to intervene. Chris consults the clock and forms a true belief as a result.

Such a case is akin to cases of modal environmental epistemic luck and epistemic Frankfurt-style cases in that the intervention in question exclusively concerns the subject's modal environment, and not the actual conditions under which he forms his belief. (It is also analogous to these cases in that in those close possible worlds where the demon intervenes the subject would nonetheless form his belief on the same basis as in the actual world). Since we have already noted that a subject's purely modal environment has no bearing on attributions of cognitive agency (this is a point on which anti-luck epistemology and robust virtue epistemology agree), it follows that Chris's cognitive success amounts to a cognitive achievement on his part, just as in the other cases we have looked at, like WATER. The key difference between this case and the others, however, is that Chris's belief is not subject to veritic epistemic luck, in that it couldn't easily have been false, given how it was formed.

Does Chris have knowledge? For consistency, the robust virtue epistemologist would have to say so, since it involves a cognitive achievement. We have already noted above the awkwardness of this strategy, in that it commits the robust virtue epistemologist to allowing that knowledge is compatible with high levels of epistemic risk. The robust virtue epistemic line ought to be more secure when it comes to a case like DEMON*, however, since although there is epistemic 'funny business' going on in the modal environment, this is such that it ensures that the target belief is *not* subject to epistemic risk. Indeed, as the DEMON* case is described there is no epistemic risk involved in Chris forming his beliefs in the way that he does, in that he is guaranteed to end up with a true belief. If we judge that Chris lacks knowledge, then, it can't be because he fails to satisfy the anti-luck condition laid down by anti-luck epistemology.

Even so, I think we have a strong intuition that Chris does not acquire knowledge in this case, even despite exhibiting a cognitive achievement and even despite not being subject to veritic epistemic luck. I think the diagnosis for why this is so is that although Chris ends up with a safe true belief, that he has a safe true belief has nothing to do with the exercise of his cognitive agency. That is, while the fact that he has a true belief is sufficiently attributable to his cognitive agency for it to count as a cognitive achievement, that he has a safe true belief is not significantly attributable to his

cognitive agency at all. Instead, it largely due to the fact that the demon is poised to intervene across close possible worlds.

If this diagnosis is right, then I think it tells us something very important about knowledge, and indeed about how we should understand the ability intuition. When we attribute knowledge we do not merely want the agent's cognitive agency to play a significant role in her cognitive success, but we further want the agent's cognitive agency to play a significant role in her safe (i.e., non-lucky) cognitive success. This further reinforces the point, central to anti-luck epistemology, that knowledge requires cognitive success which is not subject to veritic epistemic luck. In doing so, it offers an additional reason for being sceptical of the alternative epistemological picture that is set out by robust virtue epistemology.^{33,34}

REFERENCES

- Becker, K. (2007). *Epistemology Modalized*, London: Routledge.
- Black, T., & Murphy, P. (2007). 'In Defense of Sensitivity', *Synthese* 154, 53-71.
- Black, T. (2008). 'Defending a Sensitive Neo-Moorean Invariantism', *New Waves in Epistemology*, (eds.) V. F. Hendricks & D. H. Pritchard, 8-27, Basingstoke: Palgrave Macmillan.
- Chisholm, R. (1977). *Theory of Knowledge*, (2nd ed.), Englewood Cliffs, NJ: Prentice-Hall.
- Comesaña, J. (2005). 'Unsafe Knowledge', *Synthese* 146, 393-402.
- (Forthcoming). 'Sosa on Safety and Epistemic Frankfurt Cases', *Virtuous Thoughts: The Philosophy of Ernest Sosa*, (ed.) J. Turri, Dordrecht, Holland: Springer.
- DePaul, M. (2009). 'Ugly Analyses and Value', *Epistemic Value* (eds.) A. Haddock, A. Millar & D. H. Pritchard, 112-38, Oxford: Oxford University Press.
- Dretske, F. (1970). 'Epistemic Operators', *Journal of Philosophy* 67, 1007-23.
- (1971). 'Conclusive Reasons', *Australasian Journal of Philosophy* 49, 1-22.
- Dutant, J. (2010). 'Knowledge, Methods and the Impossibility of Error', University of Geneva PhD Thesis (available at: <http://archive-ouverte.unige.ch/unige:23574>).
- Engel, M. (1992). 'Is Epistemic Luck Compatible with Knowledge?', *Southern Journal of Philosophy* 30, 59-75.
- Frankfurt, H. (1969). 'Alternate Possibilities and Moral Responsibility', *Journal of Philosophy* 46, 829-39.
- 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. (2007). 'Worries About Pritchard's Safety', *Synthese* 158, 299-302.
- Hetherington, S. (2013). 'There Can be Lucky Knowledge', *Contemporary Debates in Epistemology* (2nd Ed.), (eds.) M. Steup & J. Turri, §7, Oxford: Blackwell.
- Kallestrup, J., & Pritchard, D. H. (2011). 'Virtue Epistemology and Epistemic Twin Earth', *European Journal of Philosophy* (Online First, DOI: DOI: 10.1111/j.1468-0378.2011.00495.x).
- (2012). 'Robust Virtue Epistemology and Epistemic Anti-Individualism', *Pacific Philosophical Quarterly* 93, 84-103.
- (2013). 'Robust Virtue Epistemology and Epistemic Dependence', *Knowledge, Virtue and Action: Putting Epistemic Virtues to Work*, (eds.) T. Henning & D. Schweikard, ch. 11, London: Routledge.
- Kelp, C. (2009). 'Knowledge and Safety', *Journal of Philosophical Research* 34, 21-31.
- Lewis, D. (1973). *Counterfactuals*, Oxford: Blackwell.
- Luper, S. (1984). 'The Epistemic Predicament', *Australasian Journal of Philosophy* 62, 26-50.
- (2003). 'Indiscernability Skepticism', *The Sceptics: Contemporary Essays*, (ed.) S. Luper, 183-202, Aldershot: Ashgate.
- McDowell, J. (1995). 'Knowledge and the Internal', *Philosophy and Phenomenological Research* 55, 877-93.
- Nozick, R. (1981). *Philosophical Explanations*, Oxford: Oxford University Press.
- Neta, R., & Rohrbaugh, G. (2004). 'Luminosity and the Safety of Knowledge', *Pacific Philosophical Quarterly* 85, 396-406.
- Pritchard, D. H. (2002). 'Resurrecting the Moorean Response to the Sceptic', *International Journal of Philosophical Studies* 10, 283-307.
- (2004). 'Epistemic Luck', *Journal of Philosophical Research* 29, 193-222.
- (2005a). *Epistemic Luck*, Oxford: Oxford University Press.
- (2007a). 'Anti-Luck Epistemology', *Synthese* 158, 277-97.
- (2007b). 'Knowledge, Luck, and Lotteries', *New Waves in Epistemology*, (eds.) V. F. Hendricks

- D. H. Pritchard, 28-51, London: Palgrave Macmillan.
- (2008). 'Sensitivity, Safety, and Anti-Luck Epistemology', *The Oxford Handbook of Scepticism*, (ed.) J. Greco, 437-55, Oxford: Oxford University Press.
- (2009a). 'Apt Performance and Epistemic Value', *Philosophical Studies* 143, 407-16.
- (2009b). *Knowledge*, London: Palgrave Macmillan.
- (2009b). 'Knowledge, Understanding and Epistemic Value', *Epistemology (Royal Institute of Philosophy Lectures)*, (ed.) A. O'Hear, 19-43, Cambridge: Cambridge University Press.
- (2009d). 'Safety-Based Epistemology: Whither Now?', *Journal of Philosophical Research* 34, 33-45.
- (2012a). 'Anti-Luck Virtue Epistemology', *Journal of Philosophy* 109, 247-79.
- (2012b). *Epistemological Disjunctivism*, Oxford: Oxford University Press.
- (2012c). 'In Defence of Modest Anti-Luck Epistemology', *The Sensitivity Principle in Epistemology*, (eds.) T. Black & K. Becker, 173-92, Cambridge: Cambridge University Press.
- (2012d). 'The Methodology of 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.
- (2014a). 'The Modal Account of Luck', *Metaphilosophy* 45, ??-??.
- (2014b). 'Risk', *unpublished manuscript*.
- (Forthcominga). 'Legal Risk and The Arithmetic of Criminal Justice', *Jurisprudence*.
- (Forthcomingb). 'Knowledge, Luck and Virtue', *The Gettier Problem*, (eds.) C. Almeida, P. Klein & R. Borges, Oxford: Oxford University Press.
- Pritchard, D. H., Millar, A., & Haddock, A. (2010). *The Nature and Value of Knowledge: Three Investigations*, Oxford: Oxford University Press.
- Pritchard, D. H., & Smith, M. (2004). 'The Psychology and Philosophy of Luck', *New Ideas in Psychology* 22, 1-28.
- Pritchard, D. H. & Whittington, L. (eds.) (2015). *The Philosophy of Luck*, Oxford: Wiley-Blackwell.
- Roush, S. (2005). *Tracking Truth: Knowledge, Evidence and Science*, Oxford: Oxford University Press.
- Sainsbury, R. M. (1997). 'Easy Possibilities', *Philosophy and Phenomenological Research* 57, 907-19.
- Sosa, E. (1999). 'How to Defeat Opposition to Moore', *Philosophical Perspectives* 13, 141-54.
- (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, Oxford: Oxford University Press.
- Stalnaker, R. (1968). 'A Theory of Conditionals', *Studies in Logical Theory*, (ed.) N. Rescher, 98-112, Oxford: Blackwell.
- Unger, P. (1968). 'An Analysis of Factual Knowledge', *Journal of Philosophy* 65, 157-170.
- Widerker, D., & McKenna, M. (2006). *Moral Responsibility and Alternative Possibilities: Essays on the Importance of Alternative Possibilities*, London: Ashgate.
- Williamson, T. (2000). *Knowledge and its Limits*, Oxford: Oxford University Press.
- Zagzebski, L. (1994). 'The Inescapability of Gettier Problems', *Philosophical Quarterly* 44, 65-73.
- (1999). 'What is Knowledge?', *Blackwell Guide to Epistemology*, (eds.) J. Greco & E. Sosa, 92-116, Oxford: Blackwell.
- (2001). 'Must Knowers be Agents?', *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*, (eds.) A. Fairweather & L. Zagzebski, Oxford: Oxford University Press.

NOTES

¹ I take it that what we are after when we seek an analysis of knowledge is one which is informative. Since circular analyses can nonetheless be informative, it follows that non-reductive analyses can potentially be adequate analyses. In principle, then, even a proponent of knowledge-first epistemology, like Williamson (2000), could offer a positive reply to this formulation of the Gettier problem. For more on the issue of what is required of an analysis of knowledge, see Zagzebski (1999), DePaul (2009), and Pritchard (2012*d*).

² I call this proposal *anti-luck virtue epistemology*. See, in particular, Pritchard (2009*b*; 2012*a*) and Pritchard, Millar & Haddock (2010, ch. 3). I further develop this proposal in Pritchard (forthcoming*b*).

³ See Unger (1968) for what is probably the first account in the literature of veritic luck, though note that this nomenclature is, I believe, due to Engel's (1992) discussion of Unger's paper. See also Pritchard (2004, 2005*a*).

⁴ For discussion of this claim, see Pritchard (2004; 2005*a*; 2007*a*; 2012*a*; 2012*c*) and the recent exchange between Hetherington (2013) and Pritchard (2013).

⁵ See Dutant (2010) for critical discussion of just how widespread a non-factive conception of the justification condition was at the time of Gettier's article.

⁶ Following McDowell (e.g., 1995)—who may well be inclined to defend the general claim that justification (*qua* condition on knowledge) is factive—in Pritchard (2012*b*) I defend the thesis that paradigm cases of perceptual knowledge are supported by reflectively accessible factive rational support.

⁷ I say 'almost' because of some potential complications here, such as the possibility that certain types of defeaters might make even factive justification insufficient for knowledge, or the possibility that while factively justified belief suffices for knowledge, there is nonetheless a kind of knowledge which doesn't require justification at all.

⁸ Of course, it might not be a single condition but rather several conditions. For ease of expression, however, I will write as if it is a single condition that we are after (and, as we will see, in the end it is a single condition which solves this problem).

⁹ See especially Pritchard (2005*a*). See also Pritchard (2007*a*; 2012*c*).

¹⁰ I defend the modal account of luck in a number of places. For my most recent defence of this view, see Pritchard (2014*a*). For my earlier defences of this proposal, see Pritchard (2004; 2005*a*) and Pritchard & Smith (2004). For a helpful set of papers on the philosophy of luck, including some pieces which are critical of the modal account, see Pritchard & Whittington (2015).

¹¹ Notice that we are understanding the ordering of possible worlds in the standard way in terms of their similarity to the actual world. See especially Stalnaker (1968) and Lewis (1973).

¹² Notice that 'not lucky' is not the same thing as 'unlucky'. The latter is itself a kind of luck, albeit bad luck. See also endnote 13.

¹³ Note that I bracket here the further issue of when a lucky event counts as good or bad luck. Thus, when I say that an event is lucky this is not to characterize it either way on this score. My own view is that it is not part of the role of a theory of luck to pick out when lucky events constitute good or bad luck, though it would take me too far afield to argue for this claim here. See Pritchard (2014*a*) for more on this point. See also endnote 12.

¹⁴ This is not to say that betting on a lottery is rational of course, only that it is not as irrational as betting on a normal event with astronomically long odds. Note that the slogan for the British National Lottery for a long time was 'It could be you!'. This is clearly not the 'could' of probability, since in a probabilistic sense it *couldn't* be you. It is rather the 'could' of modal nearness—i.e., if you play the lottery, then not much needs to change about the actual world in order for you to be a winner. This point was reinforced by the advertising campaign that used to go with this slogan, which featured various lottery ticket holders going about their daily lives with a golden finger continually hovering in the sky above them preparing to zap the lucky winner. The moral is that playing lotteries has the effect of making the possible world in which one is very rich a lot closer than it would ordinarily be. See Pritchard (2007*b*) for an exploration of one epistemological upshot of this point.

¹⁵ I discuss the notion of risk in its own right, including how it relates to the notion of luck, in Pritchard (2014*b*). See also Pritchard (forthcoming*a*).

¹⁶ Note that for ease of expression I've put this point in terms of the *acquisition* of a true belief, but of course we can also talk of a belief being only luckily true in terms of what *sustains* that belief (i.e., where what leads to the acquisition of the true belief is very different to what at a later time sustains that true belief). This makes no material difference to the account on offer, and so I set this complication to one side in what follows. Note also that I am taking it for granted that when it comes to veritic epistemic luck we should understand a subject's 'not getting it right' as equivalent to 'getting it wrong'. While this is clearly not above dispute, it would take me too far afield to explore this issue here (but see endnote 16 for an indication of why I think this is the right way to proceed on this score).

¹⁷ Indeed, I now think that the notion of epistemic risk is more fundamental to our thinking about knowledge than veritic epistemic luck, in that it is because we wish to avoid the former in our beliefs that we are concerned to eliminate the latter. See Pritchard (2014*b*) for more on this point. (I think this also explains why when we evaluate whether a true belief is subject to veritic epistemic luck we naturally consider whether there are close possible worlds where the same basis for that true belief leads to false belief, as opposed to merely non-belief. On this point, see also endnote 16).

¹⁸ Safety-style principles have been offered by a number of authors, including Luper-Foy (1984; cf. Luper 2003), Sainsbury (1997), Sosa (1999), Williamson (2000), and Pritchard (2002, 2005*a*, 2007*a*). Sensitivity-style principles have been offered by Dretske (1970, 1971), Nozick (1981), Roush (2005), Becker (2007), Black & Murphy (2007), and Black (2008). For an overview of the relative merits and demerits of these principles, see Pritchard (2008).

¹⁹ Originally due to Chisholm (1977, 105).

²⁰ The claim that sensitivity principles struggle with inductive knowledge is originally due to Sosa (1999).

²¹ See Greco (2007) for a clear statement of this dilemma.

²² I first introduced the distinction between modest and robust anti-luck epistemology—or ‘weak’ and ‘strong’ anti-luck epistemology, as I have sometimes expressed this distinction—in Pritchard (2009*b*). See also Pritchard, Millar & Haddock (2010, ch. 3) and Pritchard (2012*a*).

²³ See Pritchard (2009*b*, 2012*a*) and Pritchard, Millar & Haddock (2010, ch. 3).

²⁴ For more on this point, see Pritchard (2009*b*, 2012*a*) and Pritchard, Millar & Haddock (2010, ch. 3).

²⁵ I’ve dealt with a range of other problems facing anti-luck epistemology in this regard—particularly in terms of putative counterexamples to the safety condition on knowledge—elsewhere. See especially Pritchard (2007*a*; 2009*d*; 2012*c*).

²⁶ See Goldman (1976), who credits the example to Carl Ginet.

²⁷ I first drew this distinction between environmental and intervening epistemic luck in Pritchard (2009*b*, chs. 3-4). For further discussion of this distinction, see Pritchard (2009*c*; 2012*a*), Pritchard, Millar & Haddock (2010, chs. 2-4), Kallestrup & Pritchard (2011; 2012; 2013).

²⁸ I think that the same diagnosis applies to Comesaña’s (2005, 397) ‘Halloween party’ case, in that it is also essentially a case of modalized environmental epistemic luck. My response to the WATER case is thus applicable, *mutatis mutandis*, to this case as well.

²⁹ See Widerker & McKenna (2006) for a recent helpful collection of papers on Frankfurt-style cases and their philosophical implications.

³⁰ For an argument in support of this general claim as it applies to cognitive agency in particular, see Kallestrup & Pritchard (2011). This offers an *epistemic twin-earth case* that considers two counterpart agents, one on earth and one on twin-earth, who are microphysical duplicates, and who have identical causal histories. These counterparts occupy not just identical causal environments but also environments which are identical in terms of what is normally the case. Both agents form a true belief that *p* on the same basis. The only difference between earth and twin earth is that on twin earth the modal environment is such that the subject could easily form a false belief on the same basis as in the actual world (this is not so on earth). We argue that there is no difference between the two subjects in terms of the extent to which their true beliefs are creditable to their cognitive agency. The moral, as in the Frankfurt-style cases, is that one’s purely modal environment has no bearing on attributions of (cognitive) agency.

³¹ This is indeed just what Sosa (2007, ch. 5) does. See in particular his discussion of the ‘jokester’ case. I critically discuss Sosa’s reasons for ascribing knowledge in this case in Pritchard (2009*a*). See also Zagzebski (2001), who describes an epistemic Frankfurt-style case and goes on to attribute knowledge to the agent concerned. (Though note, as Comesaña (*forthcoming*) correctly points out, Zagzebski’s formulation of the example is unfortunate in that the subject’s basis for belief would be different in the case where external factors intervene. It is thus not a counterexample to anti-luck epistemology, since this involves a basis-relative formulation of safety). I introduce and discuss the distinction between robust and modest virtue epistemology—or, as I sometimes term the distinction, between ‘weak’ and ‘strong’ virtue epistemology—in Pritchard (2009*b*, chs. 3-4; 2009*c*; 2012*a*) and Pritchard, Millar & Haddock (2010, chs. 2-4).

³² As I explain in Pritchard (2012*c*), there is also a further diagnosis available to explain why we might have conflicting intuitions about cases like WATER, which is that read a certain way such cases involve not veritic epistemic luck but rather evidential epistemic luck, which as we noted above is compatible with knowledge.

³³ I describe the particular kind of anti-luck epistemology which holds that, roughly, knowledge is to be understood as safe cognitive success which is appropriately attributable to cognitive agency, as *anti-luck virtue epistemology*. See, for example, Pritchard (2009*b*; 2012*a*) and Pritchard, Millar & Haddock (2010, ch. 3). I think that more can be said to motivate anti-luck epistemology over robust virtue epistemology. In particular, I’ve argued elsewhere that robust virtue epistemology (but not anti-luck epistemology) struggles to accommodate what I call the *epistemic dependence* of knowledge,

which is the extent to which knowledge can be dependent upon factors outwith a subject's cognitive agency. See especially Kallestrup & Pritchard (2013; cf. Kallestrup & Pritchard 2011; 2012).

³⁴ Earlier versions of this paper were presented at a conference on safety at Oxford University and at the 'Gettier at 50' conference at the University of Edinburgh, both in 2013. Thanks to Julien Dutant, John Hawthorne, Clayton Littlejohn, Lisa Miracchi, Ernie Sosa, and Tim Williamson. Special thanks to Allan Hazlett for detailed comments on an earlier version of this paper.