THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Anti-luck virtue epistemology

Citation for published version:

Pritchard, D 2012, 'Anti-luck virtue epistemology', *Journal of Philosophy*, vol. 109, no. 3, pp. 247-279.
<https://doi.org/10.5840/jphil201210939>

Digital Object Identifier (DOI):

[10.5840/jphil201210939](https://doi.org/10.5840/jphil201210939)

Link:

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

Document Version:

Peer reviewed version

Published In:

Journal of Philosophy

Publisher Rights Statement:

This is an author's Accepted Manuscript of the following article: © Pritchard, D. (2012) "Anti-Luck Virtue Epistemology", *Journal of Philosophy*, 109, 3, p. 247-279. The final publication is available at <http://www.journalofphilosophy.org/articles/index.htm>

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 VIRTUE EPISTEMOLOGY¹

DUNCAN PRITCHARD

University of Edinburgh

ABSTRACT. It is argued that there are two ‘master’ intuitions about knowledge—an *anti-luck* intuition and an *ability* intuition—and that these impose distinct epistemic demands. It is claimed that recognising this fact leads one towards a new proposal in the theory of knowledge—*anti-luck virtue epistemology*—which can avoid the problems that afflict other theories of knowledge. This proposal is motivated in contrast to two other ways of thinking about knowledge which are shown to be ultimately unsuccessful: *anti-luck epistemology* and *virtue epistemology*. Finally, a diagnosis is offered of why our concept of knowledge should have the kind of structure dictated by anti-luck virtue epistemology.

1. TWO MASTER INTUITIONS ABOUT KNOWLEDGE

Until relatively recently, a key task assigned to the epistemologist was to offer an adequate definition of knowledge, one that was informative, non-circular, and which could suitably accommodate our salient epistemological intuitions. Call this the *analytical project*. This project has fallen into disfavour recently, with many arguing that it is a hopeless task.² Given the lack of success that epistemologists have had on this score it is not surprising that a disconsolate mood should have set in amongst those working on this project. Nevertheless, such pessimism is premature. Indeed, I will be arguing that there is an adequate theory of knowledge available which fulfils the remit of the analytical project.

Central to my proposal is the idea that we need to reconsider two overarching intuitions which govern our thinking about knowledge; specifically, our thinking regarding what turns true belief into knowledge. The first will be very familiar indeed. This is the intuition that when one knows one’s cognitive success (i.e., one’s believing truly) is not a matter of luck. Call this the *anti-luck intuition*. Consult any introductory text in the theory of knowledge and you will find a statement of this intuition. If, for example, a commentator is asked to explain why mere true belief cannot suffice for knowledge, the standard response is to point out that in cases of mere true

belief, unlike knowledge, one's cognitive success can simply be a matter of luck.³ The role of this intuition in contemporary theory of knowledge is particularly apparent in the post-Gettier literature, where it is often stated that precisely the point of the Gettier cases is that they demonstrate that justified true belief is compatible with one's cognitive success being merely due to luck. The failure of the tripartite account to accommodate the anti-luck intuition is thus meant to be a decisive strike against it.⁴

The second intuition is not so universally expressed, but it is certainly discernible in much of our thinking about knowledge. This is the intuition that knowledge requires cognitive ability, in the sense that when one knows one's cognitive success should be the product of one's cognitive ability. Call this the *ability intuition*. Again, one finds a commitment to this intuition being presupposed in introductory discussions of why mere true belief does not amount to knowledge. The worry about mere true belief, we are standardly told, is that it needn't be formed in the "right way", where this means via a process that is appropriate to the acquisition of knowledge. But what is a belief-forming process which is appropriate to the acquisition of knowledge if it is not a cognitive ability?⁵ Thus, a commitment to this intuition is tacitly supposed.⁶

Interestingly, these intuitions are often run together, or at least a clear statement of one of them (typically the anti-luck intuition) tends to go hand-in-hand with remarks that entail a commitment to the other intuition. On reflection, there is good reason for this since there do seem to be very close connections between these two intuitions. What does it take to ensure that one's cognitive success is not due to luck? Well, intuitively anyway, that it is the product of one's cognitive ability. Conversely, insofar as one's cognitive success is the product of one's cognitive ability, then—again, intuitively—one would expect it to thereby be immune to knowledge-undermining luck. One might thus regard these two intuitions as two 'faces' of a single intuition. If that's right, then any epistemic condition on knowledge which is formulated in order to satisfy the anti-luck intuition (call this an *anti-luck condition*) will, if successful, thereby satisfy the ability intuition, and any epistemic condition which is formulated in order to satisfy the ability intuition (call this an *ability condition*) will, if successful, thereby satisfy the anti-luck intuition.

I will be arguing that this conception of how these two master intuitions about knowledge are related to each other is fundamentally flawed. In particular, I will be claiming that these two intuitions in fact impose independent epistemic demands on our theory of knowledge, and that it is only once one recognises this fact that one can offer a successful resolution of the analytical project. My strategy for demonstrating this will be to explore two popular contemporary approaches to the analytical project which each take one of these intuitions about knowledge as central to their approach. The first, what I call an *anti-luck epistemology*, takes the anti-luck intuition as core and then aims to formulate an anti-luck condition which can accommodate this intuition

and thereby offer us an adequate theory of knowledge, one that can also accommodate the ability intuition. The second, what I call a *virtue epistemology*, takes the ability intuition as core and then aims to formulate an ability condition which can accommodate this intuition and thereby offer us an adequate theory of knowledge, one that can also accommodate the anti-luck intuition.

I will argue that both proposals fail, and fail precisely because they are unable to fully accommodate the particular master intuition which they do not treat as core. What the failure of these two approaches demonstrates, I will claim, is that we need to conceive of the two master intuitions as imposing distinct epistemic demands, and hence as requiring independent epistemic conditions. The view that results, what I call *anti-luck virtue epistemology*, can deal with the full gamut of ‘test’ cases in the theory of knowledge. Moreover, I will argue that one can also offer a plausible account of why the epistemic component of knowledge might have this bipartite structure. Far from being a lost cause, the analytical project is shown to be back in business.

2. ANTI-LUCK EPISTEMOLOGY

As noted above, the chief moral that is generally extracted from the post-Gettier literature is that the justification condition does not suffice to exclude knowledge-undermining luck as had been widely supposed. In a nutshell, the justification condition in the tripartite account of knowledge does not suffice to accommodate the anti-luck intuition. The post-Gettier debate thus inevitably generates a discussion of what sort of epistemic condition or conditions could accommodate this intuition.

One proposal that came to the fore in the early literature was that a necessary condition on knowledge was that one’s true belief should be *sensitive*, in the following sense:

The Sensitivity Principle

If *S* knows that *p* then had *p* not been true *S* would not have believed that *p*.⁷

The foremost exponent of this principle was of course Robert Nozick, but one can find endorsements of very similar principles in the work of a number of important philosophers, and this principle is still defended today.⁸

The sensitivity principle has no trouble dealing with Gettier-style cases. Consider the following three Gettier-style cases that are often discussed in the literature, all of which involve a true belief which enjoys good epistemic support and yet which doesn’t amount to knowledge because of the presence of epistemic luck:

Edmund

Edmund forms a belief that Jones owns a Ford on excellent grounds. He then validly infers that either Jones owns a Ford or Smith is in Barcelona, and accordingly forms a belief in this entailed proposition solely on the basis of his grounds for believing the entailing proposition and the relevant deduction. As it happens, the entailing proposition is false; the entailed proposition, however, is true since it just so happens (and unbeknownst to Jones) that Smith is in Barcelona.⁹

Roddy

Using his reliable perceptual faculties, Roddy non-inferentially forms a true belief that there is a sheep in the field before him. His belief is also true. Unbeknownst to Roddy, however, the truth of his belief is completely unconnected to the manner in which he acquired this belief since the object he is looking at in the field is not a sheep at all, but rather a sheep-shaped object which is obscuring from view the real sheep hidden behind.¹⁰

Barney

Using his reliable perceptual faculties, Barney non-inferentially forms a true belief that the object in front of him is a barn. Barney is indeed looking at a barn. Unbeknownst to Barney, however, he is in an epistemically unfriendly environment when it comes to making observations of this sort, since most objects that look like barns in these parts are in fact barn façades.¹¹

In all three cases we have examples of cognitive success which is such that, had the relevant fact been otherwise (but everything else had remained the same, consistent with that change), then the agent would have continued to believe the target proposition regardless, and hence would have believed falsely. Had it been false that either Jones owns a Ford or Smith is in Barcelona—i.e., had Smith not in fact been in Barcelona but away visiting friends in Tarragona, say—then clearly Edmund would have continued to believe this proposition regardless, since his basis for this belief (his grounds for believing that Jones owns a Ford, and his knowledge of the relevant entailment) would be unchanged. Had it been false that there is a sheep in the field—i.e., if the sheep in question had wandered into a neighbouring field—Roddy would have continued to believe this proposition regardless, since his basis for this belief (the sheep-shaped object that he can see in the field) would be unchanged. And had it been false that the object that Barney is looking at is a barn—i.e., if it were a barn façade—Barney would have continued to believe this proposition regardless, since his basis for this belief (that he is presented with a plausible barn-shaped object) would be unchanged. All three cases thus involve an insensitive true belief, and hence the sensitivity principle has no trouble explaining why they do not amount to knowledge.

The sensitivity principle can also deal with other cases which trade on the anti-luck intuition too, such as the lottery case in which the agent lacks knowledge even while having a true belief which is supported by excellent grounds:

Lottie

Lottie has a ticket for a fair lottery with very long odds. The lottery has been drawn, although Lottie has not heard the result yet. Reflecting on the odds involved she concludes that her ticket is a loser. Lottie's belief that she owns a losing ticket is true.

The sensitivity principle can explain why Lottie lacks knowledge because her true belief, despite its

excellent epistemic pedigree, is insensitive. Had it been false that Lottie's ticket was a losing ticket—i.e., had Lottie won the lottery—then she would have continued to believe this proposition regardless, since her basis for this belief (gained by reflecting on the long odds involved in winning) would be unchanged.

Moreover, the sensitivity principle can also explain why Lottie's lack of knowledge in this case is compatible with the undeniable fact that Lottie could have gained knowledge of the target proposition in other ways where the supporting evidence would have been weaker from a probabilistic point of view. For example, Lottie can gain knowledge that she has a losing ticket by reading the result of the lottery in a reliable newspaper, and yet the chances of her forming a false belief on this basis are surely higher than the chances of her forming a false belief in the same proposition by reflecting on the long odds involved. No matter how reliable the newspaper is, it is surely the case that the odds that it contains a misprint in this regard are higher than the odds of winning your average national lottery. Sensitivity can account for what is going on here because had Lottie formed her belief in the target proposition by consulting a reliable newspaper then her belief would have been sensitive. For suppose that Lottie wins the lottery but everything else consistent with this stays the same. Given Lottie's basis for her belief, she would no longer believe that she had a winning ticket since in this counterfactual scenario she would be looking at the set of *winning* ticket printed numbers in her reliable newspaper.

The contrast between these two variations on the Lottie case demonstrates that a high probabilistic strength of one's evidence (at least if it falls short of 1) may not suffice to ensure that one's belief is sensitive. Sensitivity thus explains what is going on in these two cases by highlighting the surprising point that while what is required for knowledge is an epistemic basis which ensures the sensitivity of one's belief, the probabilistic strength of one's evidence, no matter how strong (bar a probabilistic strength of 1), may not suffice to supply such an epistemic basis.¹²

Sensitivity thus seems to be able to deal with the anti-luck intuition, and hence appears to be a strong contender to be the right way of thinking about the anti-luck condition on knowledge. This principle faces a range of problems, however, not least of which is that it is unable to deal with a certain kind of inductive knowledge.¹³ Consider the following case:

Ernie

Ernie deposits a rubbish bag into the rubbish chute in his high-rise flat. He has every reason to think that the chute is working correctly and so believes, a few minutes later, that the chute is in the basement. His belief is true.¹⁴

Intuitively, Ernie has knowledge in this case, since even though he hasn't seen the rubbish in the basement, he does have an excellent inductive basis for thinking that it is there. Clearly, though, Ernie's belief is not sensitive, since had the rubbish not made it to the basement, but everything

else had stayed the same (had a workman recently damaged the chute so that rubbish was getting stuck on the third floor for example), Ernie would have continued to believe what he does regardless, and so would have believed falsely.

It is problems like this that have led commentators to move away from the sensitivity principle and adopt a similar modal principle which seems better placed to capture our intuitions about knowledge:

The Safety Principle

If S knows that p then S 's true belief that p could not have easily been false.¹⁵

The safety principle is also able to deal with Gettier-style cases, since in all such cases the agent forms a true belief in such a way that she could have very easily been in error. Had Jones not been in Barcelona but in Tarragona instead then Edmund would have believed falsely; had the sheep wandered into a neighbouring field then Roddy would have believed falsely, and had the object before Barney been a barn façade rather than a barn then he would have believed falsely.

Moreover, the safety principle can also deal with the lottery case, since this too involves a belief that could very easily have been false, had Lottie happened to be in possession of the one winning ticket (we will return to this point). Indeed, the safety principle offers the same kind of explanation of why Lottie lacks knowledge as we saw being offered by the sensitivity principle, since what is important is not the probabilistic likelihood of error (which is of course very low), but rather the modal closeness of that error.

Where the safety principle has an advantage over the sensitivity principle is when it comes to the problem posed by the kind of inductive knowledge at issue in the Ernie case. For while Ernie's beliefs are not sensitive, they are safe. Given how he formed his belief, after all, it couldn't have easily been the case that his belief is false. Now, one might balk at this claim on the grounds that whether it really is the case that Ernie's true belief could very easily have been false is an open question given how the case is described. Perhaps, for example, there is a snag in the rubbish chute that Ernie's bag of rubbish could so very easily have snagged on? If so, then even despite the inductive basis for his belief, it is not safe since it could very easily have been false (after all, in such a case where it does get caught on the snag on the way down, Ernie would still believe on the same inductive basis that the rubbish is in the basement, and so believe falsely).

Interestingly, though, this issue in fact speaks in favour of safety rather than against it, since it highlights that examples like the Ernie case need to be understood in a certain way if we are to attribute knowledge to our protagonist. For note that while it is true that the presence of a snag in the rubbish chute that could so very easily have prevented the rubbish from getting to the basement suffices to make the target belief unsafe, if we understand the example in this way then

there is no temptation at all to think that Ernie has knowledge. In contrast, if we interpret the example in the way that we naturally tend to, so that there are no snags and such like to prevent the rubbish from getting to the basement, then the intuition that Ernie has knowledge returns, but so too does the safety of the target belief. If significant change in the actual circumstances is required to ensure that the rubbish fails to reach the basement (significant change, moreover, which is undetectable to Ernie), then his belief will be safe, since it couldn't easily have been false.

Considering the Ernie case in light of the Lottie case is also instructive in this regard. Very little about the actual circumstances is required to change to ensure that Lottie's true belief becomes undetectably false, and this explains why her belief is unsafe and so not a candidate for knowledge. In contrast, quite a lot has to change about the actual circumstances to ensure that Ernie's true belief becomes undetectably false, and this explains why his belief is safe and so, on this score at least, in the market for knowledge.

This illustrates an important point which fits quite snugly with our anti-luck intuition. In wanting our cognitive success to be immune to luck we are not thereby desiring that it be free from *any* possibility of error, no matter how remote. Accordingly, as the error becomes more remote—i.e., as more needs to change about actual circumstances for the agent to (counterfactually) form a false belief—so we become more tolerant of it, to the point where we no longer regard the counterfactual error as indicating that there was anything lucky about the target cognitive success. The anti-luck intuition thus manifests itself, in keeping with how we are reading the safety principle, with a complete intolerance of error in close counterfactual circumstances, a tolerance of error in remote counterfactual circumstances, and a sliding scale of tolerance between these two extremes. So in the Lottie case, for example, where the counterfactual error is very close, we have no hesitation in dismissing the possibility of knowledge on anti-luck grounds. In contrast, in a parallel Lottie case where the belief in the target proposition is formed in an appropriate way (e.g., by reading a reliable newspaper), and where as a result there isn't any close counterfactual error, we are happy to attribute knowledge (even though the probabilistic likelihood of error may be higher). As we have seen, in examples like the Ernie case one can fill out the details in such a way as to elicit different responses. If you make the counterfactual error close, as when one supposes that the chute has a snag in it, then one loses the intuition that the target belief is in the market for knowledge and it also ceases to be safe. If, on the other hand, one reads the example in the natural way such that the counterfactual error is remote, then one retains the intuition that this is a case of knowledge and the target belief is also safe.¹⁶

There is another *prima facie* problem that faces the safety principle but which turns out to be illusory, which concerns how the principle would deal with necessary propositions, or at least propositions which are true in all circumstances similar to the actual circumstances in which the

belief was formed. The worry is that any true belief in such a proposition is trivially guaranteed to be such that it could not easily be false, but not because of any epistemically relevant feature of the belief (such as its epistemic standing) but purely because of the nature of the proposition believed. Consider, for example, the following case:

Mathema

Mathema uses a calculator to find out the sum of 12×13 . As a result, he forms a true belief that $12 \times 13 = 156$. Unbeknownst to Mathema, however, his calculator is in fact broken and generating 'answers' randomly.

Clearly Mathema does not know the target proposition. And yet, given that this proposition is necessarily true, it appears that it can't be the case that his belief could easily have been false, and hence we seem committed to holding that this belief is safe. Rather than being a devastating counterexample to safety, however, cases like this highlight that we need to understand safety in a quite specific way.

When we talk of a safe belief that p being such that it could not have easily been false, it is tempting to suppose that this means that the agent's belief *that* p in similar circumstances would not be false. This reading of the safety principle is indeed susceptible to problem cases like Mathema, since where the proposition in question is such that it is not false in any circumstances (similar or otherwise) then clearly the agent cannot help but have a true belief in this propositions which couldn't easily be false, regardless of the epistemic standing of this belief. On reflection, however, it is clear that this is not the right way to read safety. For what we are interested in is rather how the agent forms her beliefs in similar circumstances and in response to the same stimuli. These beliefs may be beliefs *that* p , but equally they may be beliefs in distinct propositions too.

In order to see this point, consider the Mathema case again. While there is indeed no similar circumstance in which we can imagine Mathema forming a belief that $12 \times 13 = 156$ on the same basis and yet believing falsely, we can certainly imagine lots of similar circumstances in which Mathema forms her belief on the same basis and yet ends up with a false belief, such as the similar situation in which the calculator generates a different result. Mathema's belief is thus unsafe, and hence the safety principle is perfectly able to explain why Mathema lacks knowledge in this case, at least so long as we formulate that principle correctly.¹⁷

In the safety principle, then, we seem to have a way of thinking about the anti-luck condition on knowledge that is in keeping with our general intuitions in this regard and which isn't, on closer inspection, susceptible to some obvious problems.^{18,19} Moreover, notice that the safety principle isn't just responding to the problem posed by Gettier-style cases but also to other problem cases in epistemology too, such as the lottery case. With this in mind, one might become

tempted by the idea that all that is required of a theory of knowledge is a properly formulated anti-luck condition, such that knowledge is true belief that satisfies this anti-luck condition. Call this an *anti-luck epistemology*.²⁰ Is such a view tenable?

If it is tenable then it needs to be able to accommodate the ability intuition that we began with. On the face of it, this shouldn't be a problem since in all the cases we have considered so far, where the safety principle has been satisfied the agent concerned has exhibited the relevant cognitive ability. Indeed, one might think that reflecting on the anti-luck condition demonstrates that this is the more general intuition about knowledge. After all, in Gettier-style cases the agent's cognitive success is the product of ability and yet she lacks knowledge because she doesn't satisfy the anti-luck intuition. In contrast, in the cases we've looked at where the agent satisfies the anti-luck condition, and so has knowledge, the agent also satisfies the ability intuition. Thus it seems that there is a *prima facie* case for supposing that while a correct formulation of the ability condition will not be able to satisfy the anti-luck intuition (because of Gettier-style cases), a correct formulation of the anti-luck condition *will* be able to accommodate not just the anti-luck condition but also the ability condition.

Alas, anti-luck epistemology, despite its surface attractions, will not pass muster. Before we get to a decisive consideration against this view, I first want to consider a problem which I don't think is decisive at all (though which might look that way if one approaches anti-luck epistemology from a certain theoretical angle). The concern is that anti-luck epistemology is essentially wedded to epistemic externalism, and a very natural line to take is that no epistemically externalist theory of knowledge could ever fully accommodate the ability intuition.

There are various ways of formulating the epistemic externalism/internalism distinction, but most (if not all) epistemic internalists would agree that a necessary requirement for knowledge is that the knower has good reflectively accessible grounds in favour of her belief in the target proposition.²¹ Clearly, anti-luck epistemology does not demand this, since there is nothing about satisfying a modal condition like safety which would entail that one has such grounds. Anti-luck epistemology thus entails epistemic externalism. Does this fact suffice to demonstrate that anti-luck epistemology cannot satisfy the ability intuition?

I think not, or at least I don't think we should let an issue like this settle the matter. For while there is certainly a way of thinking about the ability intuition such that it demands that one's knowledge be due to cognitive ability in such a way that such ability is always accompanied by reflectively accessible grounds, such an interpretation is bound to be contentious because of its obvious negative implications for epistemic externalism. Indeed, epistemic externalists will surely respond by claiming that while it might well be desirable to possess reflective epistemic support for our beliefs, and while it might often be the case that the exercise of cognitive ability generates

such grounds (we are reflective creatures after all), it is wrong to build such an internalist requirement into the ability intuition in this way.

Interestingly, the kinds of examples that epistemic externalists and internalists argue over are precisely cases where the agent's cognitive success is held to be the product of her cognitive ability. In the notorious chicken-sexer case, for instance, it is built into the example that the agent has a highly reliable cognitive ability which enables her to form true beliefs about the sex of the chicks (it is just that our agent is held to be able to manifest such an ability while having no good reflectively accessible grounds in favour of the beliefs formed in this way). It is thus not meant to be in question that the agent satisfies the ability intuition (and, for that matter, the anti-luck intuition), though unsurprisingly the epistemic internalist will not be happy with an ascription of knowledge in this case. While it is of course open to the epistemic internalist to insist that a more demanding construal of the ability intuition should be adopted, one that is not satisfied in the chicken-sexer case, this would be a rather blunt dialectical move to make in this debate, one that would be of little concern to the epistemic externalist who will simply insist on their own more inclusive reading of the ability intuition.²²

In any case, I don't think we should let a controversial interpretation of the ability intuition decide the matter when it comes to evaluating anti-luck epistemology, for it is surely preferable to make such an evaluation on theory-neutral grounds if we can. As it happens, there is a type of case which poses a problem for anti-luck epistemology which both epistemic externalists and epistemic internalists should agree on. In particular, such a case shows that anti-luck epistemology cannot satisfy the ability intuition, regardless of whether we interpret that intuition along externalist or internalist lines.

Consider the following example:

Temp

Temp forms his beliefs about the temperature in the room by consulting a thermometer. His beliefs, so formed, are highly reliable, in that any belief he forms on this basis will always be correct. Moreover, he has no reason for thinking that there is anything amiss with this thermometer. But the thermometer is in fact broken, and is fluctuating randomly within a given range. Unbeknownst to Temp, there is an agent hidden in the room who is in control of the thermostat whose job it is to ensure that every time Temp consults the thermometer the 'reading' on the thermometer corresponds to the temperature in the room.

Intuitively, Temp cannot know the temperature in the room by consulting a broken thermometer in this way, even if his beliefs so formed are guaranteed to be true. In particular, what is wrong with Temp's beliefs is that they exhibit the wrong direction of fit with the facts, for while his beliefs formed on this basis are guaranteed to be true, their correctness has nothing to do with Temp's abilities and everything to do with some feature external to his cognitive agency. This means that what is underlying our intuition that Temp lacks knowledge in this case is the fact that

his beliefs fail to satisfy the ability intuition. Moreover, notice that it makes no difference whether we suppose that Temp has good reflectively accessible grounds in favour of his beliefs in this case, since either way they will not amount to knowledge. Thus, such an example does not trade on the fact that anti-luck epistemology is an externalist theory of knowledge.

Notice that whatever formulation of the anti-luck condition one opts for Temp will satisfy that condition. More generally, whatever one wishes to say about what is epistemically deficient in Temp's beliefs, it doesn't seem that his beliefs fail to satisfy the anti-luck intuition. After all, his beliefs are *guaranteed* to be true given how he is forming them, and hence it can hardly be the case that his cognitive success is merely a matter of luck. More specifically, while his cognitive success is not the product of his cognitive ability, that's not because it's simply a matter of luck.

We can bring this point out more clearly by considering how Temp's belief satisfies the safety principle.²³ This is ensured by the fact that the manner in which Temp is forming his beliefs, such that success is guaranteed, means that it can hardly be the case that he could easily have formed a false belief. Note too that the problem in play here isn't one that reveals a particular failing of our formulation of the safety principle, as if we could reformulate this principle in such a way as to ensure that a revitalised anti-luck condition could deal with this difficulty. For the underlying point demonstrated by this example is that no modal principle of the sort required to eliminate knowledge-undermining luck will be able to specify the kind of direction of fit that is required for a belief to satisfy the ability intuition. That is, in satisfying the relevant modal principle one ensures, across a suitable range of possible cases, that there is the right kind of correspondence between belief and fact; but what one doesn't ensure thereby is that a certain relationship between belief and fact obtains, one that cases like Temp indicate is essential to the manifestation of cognitive success which is the product of cognitive ability.

3. VIRTUE EPISTEMOLOGY

The upshot of the foregoing is that no plausible formulation of the anti-luck condition can fully accommodate the ability intuition, and hence anti-luck epistemology is under serious threat. One could, of course, react by denying or modifying the ability intuition, and thereby arguing that we should attribute knowledge in the Temp case. Before one takes such a desperate course of action, however, it is worthwhile to consider whether there are any less painful options available.

Given what we said earlier about how the agents in Gettier-style cases nonetheless satisfy the ability intuition, one might think that there is little to be gained by trying to define knowledge in terms of true belief and an ability condition. Interestingly, however, there is a potential way of

getting around the problem posed Gettier-style cases in this respect, and hence such an alternative proposal may be viable after all.

It is useful at this juncture to say a little more about what a cognitive ability involves. We noted above that a cognitive ability is a knowledge-conducive belief-forming process. We could of course think of cognitive abilities more generally than in terms of belief-forming processes, since they could be thought to have an output which is not doxastic but, say, emotional. But given that our primary concern is with the theory of knowledge it makes sense for us to focus specifically on cognitive abilities which have a doxastic output. One change that we should make to our earlier rough characterisation is that it is better to think of abilities in terms of dispositions rather than processes. After all, one retains one's cognitive abilities even when they are not exercised, but it is only when they are exercised that a belief-forming process is on display.

Given that we are engaged in the analytical project it is clearly not ultimately helpful to characterise cognitive abilities in terms of their conduciveness to knowledge. We thus need to ask what it is about the particular belief-forming dispositions that qualify as cognitive abilities that makes them knowledge-conducive. I take it that as a minimal requirement these belief-forming dispositions should be both reliable and suitably integrated with the agent's other belief-forming dispositions. The former requirement is needed if we are to think of these dispositions as genuinely akin to skills or abilities more generally, while the latter requirement is needed if we are to think of these dispositions as genuinely reflecting the agent's cognitive agency.

Note that as it stands any theory of knowledge which defines knowledge in terms of true belief that satisfies this conception of cognitive ability will be a form of epistemic externalism, just like anti-luck epistemology. This is because one can clearly manifest a cognitive ability in this sense while lacking any good reflectively accessible grounds in favour of one's belief in the target proposition. Indeed, the oft-cited example of the chicken-sexer that we described above fits this rubric, since this is an agent employing a reliable belief-forming disposition which is appropriately integrated with her other belief-forming dispositions. For while it is part of the example that the agent is lacking any good reflectively accessible grounds in favour of her beliefs so formed, we are clearly meant to suppose that her chicken-sexing ability works in concert rather than in tension with her other relevant belief-forming dispositions. Indeed, it is this feature of the example that makes it compelling, for if there were conflict between the various belief-forming dispositions in play—if, for example, her chicken-sexing ability generates a belief that the two chicks before her have different genders, and yet she otherwise perceives that there is no discernible difference between the chicks—then we would not find the example even remotely plausible as a case of knowledge.

So conceived, cognitive abilities are essentially the same as epistemic virtues, at least on a weak externalist construal of that notion.²⁴ In what follows we will talk interchangeably of cognitive abilities and epistemic virtues, though we will also consider below how a more restrictive internalist conception of epistemic virtue would relate to our discussion. Accordingly, we can call any view which defines knowledge in terms of true belief plus an ability condition, a *virtue epistemology*, at least provided we keep in mind that our usage of this title is partly stipulative. Given how we are characterising the ability condition, any version of virtue epistemology will be an externalist theory of knowledge.

Call any view which simply holds that knowledge is true belief that is the product of cognitive ability a *weak virtue epistemology*. Weak virtue epistemology can certainly deal with the Temp case, since as we noted above the cognitive success exhibited by the agent in this example has nothing to do with the exercise of his cognitive abilities, and everything to do with the assistance of his hidden helper. Weak virtue epistemology can also deal with other cases that an anti-luck epistemology would struggle with as well, such as cases of reliable cognitive malfunction. Consider, for example, the following case:

Alvin

Alvin has a brain lesion. An odd fact about the brain lesion that Alvin has, however, is that it causes the sufferer to form the (true) belief that he has a brain lesion. Accordingly, Alvin truly believes that he has a brain lesion.²⁵

Given how Alvin is forming his beliefs he is guaranteed to be right, and hence his beliefs will thereby satisfy any anti-luck condition such as safety. Clearly, though, Alvin does not have knowledge in this case, and the reason for this is that his beliefs are true *despite* his cognitive abilities and not because of them. That is, what explains why Alvin's beliefs do not amount to knowledge is the fact that they fail to satisfy the ability intuition, even though they do satisfy the anti-luck intuition.

Weak virtue epistemology, as an externalist theory of knowledge, will face the usual objections that are levelled against externalist theories by internalists, but I think we can legitimately set these concerns to one side for our purposes.²⁶ This is because, as noted above, such a view faces a more pressing problem that does not trade on the epistemic externalism/internalism dispute, which is its failure to deal with Gettier-style cases. These are examples, after all, in which the agent's cognitive success *is* the product of her (relevant) cognitive abilities. It seems to follow that since the agent in such cases does not satisfy the anti-luck intuition, and therefore lacks knowledge, so weak virtue epistemology cannot satisfy the anti-luck intuition either and thus cannot be a complete account of knowledge.

Recently, however, some commentators have argued that there is a way of re-thinking virtue

epistemology such that it might be able to deal with Gettier-style cases, and thus the anti-luck intuition, after all. The crux of the matter is the manner in which we think of the target cognitive success as being the ‘product’ of the relevant cognitive abilities. In weak virtue epistemology all that is demanded is that the target belief be the product of the relevant cognitive abilities and be, in addition, true, but this is not the only way of thinking about how a cognitive success can be the product of a cognitive ability.

In particular, it has been suggested that we should regard the cognitive success as being the product of the relevant cognitive abilities in the sense that the exercise of those cognitive abilities is the overarching *explanation* for the agent’s cognitive success, such that the cognitive success is primarily creditable to her cognitive agency.²⁷ In the Temp case, for example, while there is both cognitive ability on display and cognitive success, one wouldn’t regard the cognitive success as being in any sense explained by Temp’s cognitive ability, since what explains the cognitive success is rather something external to Temp’s cognitive agency, the hidden helper. We can gloss this point by saying that what is epistemically amiss about Temp’s cognitive success is that it is not *because of* his cognitive ability but rather because of something external to his cognitive agency, where the ‘because of’ here is given an explanatory reading.²⁸

Call the view that knowledge is cognitive success that is because of the exercise of the relevant cognitive abilities *strong virtue epistemology*.²⁹ The question in hand is whether strong virtue epistemology can deal with the Gettier-style cases, and thus accommodate the anti-luck intuition. On the face of it, it seems that it can.

Consider again the Edmund case. While Edmund is cognitively successful and exhibits the relevant cognitive abilities, his cognitive success is not explained by his cognitive ability, but rather by the good fortune that the other disjunct in the proposition that he infers is, unbeknownst to him, true. Or consider the Roddy case. While Roddy is cognitively successful and exhibits the relevant cognitive abilities, his cognitive success is not explained by his cognitive ability, but rather by the good fortune that there happens to be a sheep in the field hidden from his view.

Moreover, strong virtue epistemology also predicts the right result in a range of other cases too. For example, Mathema’s cognitive success is not because of her cognitive ability but rather because of the good fortune that the broken calculator she uses delivers the right result, and hence strong virtue epistemology correctly treats her as lacking knowledge. Strong virtue epistemology also generates the right result in the Ernie case, since his cognitive success is best explained by his cognitive ability, at least so long as we understand the detail of such a case in such a way as to secure the intuition that Ernie knows. It thus appears that strong virtue epistemology might be able to succeed where anti-luck epistemology failed and offer a fully adequate theory of knowledge, one that can accommodate both the anti-luck and the ability intuitions.³⁰

Unfortunately, matters are not so straightforward, since there is a group of cases that the view struggles to cope with. Worse, these problem cases seem to make conflicting demands on the view, such that there doesn't appear to be a principled way in which one could adapt the view to avoid it falling foul of these cases.

As with anti-luck epistemology, one class of problems facing the view will concern its commitment to epistemic externalism. I don't think we should worry too much about such cases, partly because they are bound to be contentious due to the controversial nature of the epistemic externalism/internalism distinction, and partly because, as we will see in a moment, there are more pressing concerns which do not trade on this distinction. It is worth noting, however, that although we have opted for an inclusive externalist construal of cognitive ability/epistemic virtue, and thus formulated an externalist virtue epistemology based around this notion, it is open to those persuaded by epistemic internalism to take the different route of defining cognitive abilities/epistemic virtues along more restrictive internalist lines and then formulating an internalist virtue epistemology on this basis. They might demand, for example, that only those reliable belief-forming dispositions which generate the required reflectively accessible epistemic support should be counted as cognitive abilities/epistemic virtues.³¹ Accordingly, in setting forth an externalist version of virtue epistemology here we are not taking sides in the epistemic externalism/internalism dispute, and hence we can safely bracket the qualms that epistemic internalists will have with strong virtue epistemology as it stands.

With this point in mind, let us turn to those problems that afflict strong virtue epistemology by both externalist and internalist lights. First off, notice how strong virtue epistemology struggles with the lottery case. After all, Lottie's cognitive success does seem to be explained by the exercise of her relevant cognitive abilities, doesn't it? Put another way, if it is not Lottie's cognitive ability that explains her cognitive success, then what does explain it? For note that the only plausible candidate here is the (epistemic) good fortune that her ticket is indeed a losing ticket. However, given the odds involved it is hard to see how this eventuality could be considered a matter of fortune at all.

Even if strong virtue epistemology is able to fend off this problem, a more serious difficulty lurks in the wings. We noted a moment ago that strong virtue epistemology seems to be able to deal with the Gettier-style cases, and we demonstrated this point by looking at the examples of Edmund and Roddy. Consider, however, how strong virtue epistemology fares when it comes to the Barney case. Unlike the Edmund and Roddy cases, the knowledge-undermining luck in this example is entirely environmental. Barney is, after all, really seeing a genuine barn, unlike, say, Roddy who merely thinks that he is seeing a genuine sheep. In a very real sense, then, Barney's cognitive abilities are putting him in touch with the relevant fact, unlike in standard Gettier-style

cases where there is a kind of fissure between ability and fact, albeit one that does not prevent the agent from having a true belief regardless. Nonetheless, Barney's environment is so epistemically unfriendly that he doesn't count as having knowledge, despite his cognitive ability and his genuine perception of a barn, because his belief is manifestly unsafe, and so offends against the anti-luck intuition. The problem, however, is that given that Barney does undertake, using his cognitive abilities, a genuine perception of the barn, it seems that his cognitive success *is* explained by his cognitive abilities, unlike in standard Gettier-style cases. More specifically, it seems that strong virtue epistemology cannot explain why the agents in Gettier-style cases involving environmental epistemic luck lack knowledge.³²

There are various lines of response that the proponent of strong virtue epistemology might make to this problem. One option, of course, is just to treat Barney as having knowledge, and at least one commentator has taken this route, though of course this is not without cost.³³ A superficially more attractive option is to try to exploit the fact that abilities are relative to environments in order to evade the problem. In this way, one might be able to argue that Barney does not exhibit the relevant cognitive abilities at all, since the relevant abilities would be those highly specialised barn-spotting abilities that are applicable to this very unusual environment containing barn façades and of course, *ex hypothesi*, Barney does not have *these* abilities at all. Unfortunately, while this line of response may seem initially appealing, it does not bear up to close scrutiny.

In the first place, notice that while we do ordinarily relativise abilities to environments, we don't tend to do this in a very fine-grained way unless there is a specific reason to do so. For example, the ability to play the piano is relativised to a broad class of normal environments, such that it wouldn't count as a fair test of your ability to play this instrument to be given the task of doing so whilst, say, placed underwater. Nonetheless, in, for example, playing the piano while outside on a sunny day one is surely exhibiting the very same ability that one exhibits while playing inside, even though there are some additional factors to take into account (more environmental noise, say). The sort of coarse-grained relativization of abilities to environments found in ordinary language is of no use to the defender of strong virtue epistemology, however, since it is vital to their proposed solution to the Barney problem that a fairly nuanced relativization be appropriate. While this does not of course rule-out a response cast along these lines, this does mean that strong virtue epistemology would be forced to saddle itself to a revisionistic view about the nature of abilities, and this is hardly desirable.³⁴

One might perhaps be willing to live with this sort of theoretical cost, if it were the only cost incurred. But there is a deeper problem lying in wait here. For not only does ordinary language not tend to relativise abilities to environments in a fine-grained way, neither does it relativise abilities

to environments in a way that would be relevant to the Barney case. Suppose, for example, that our hero is playing piano in a situation in which, unbeknownst to her, she could so very easily be underwater right now (but isn't). Perhaps, for instance, she is in a room that is surrounded by water where the walls could completely give way and let the water in at any moment. Intuitively, the ability that our hero is displaying in this case is just the very same ability she displays when she plays the piano in normal environments. But here is the rub: the Barney case is entirely analogous to this example. After all, although Barney could so very easily have been presented with a barn façade in our example, the fact of the matter is that he isn't. Similarly, just as our piano player could so very easily have been underwater, the fact of the matter is that she isn't. Accordingly, just as we should treat our piano player as exhibiting her ordinary ability to play piano in this case, so we should treat Barney as exhibiting his ordinary cognitive ability to spot barns.³⁵ So unless one is to offer a radically revisionary account of abilities, this way of defending strong virtue epistemology is a hopeless cause.

If this problem were not bad enough, there is another difficulty facing strong virtue epistemology, which concerns its ability to account for the epistemology of testimonial belief. Most accounts of the epistemology of testimonial belief allow that an agent can, in suitable circumstances, come to gain knowledge by, for the most part, trusting the word of a knowledgeable informant. To deny this, most claim, is to succumb to a kind of scepticism about testimonial knowledge, given how much of our testimonial knowledge is gained in this way. The problem, however, is that strong virtue epistemology seems unable to account for testimonial knowledge of this sort. Consider the following case:

Jenny

Jenny gets off the train in an unfamiliar city and asks the first person that she meets for directions. The person that she asks is indeed knowledgeable about the area, and helpfully gives her directions. Jenny believes what she is told and goes on her way to her intended destination.³⁶

On most views of the epistemology of testimony, Jenny gains knowledge in this case. Moreover, notice that insofar as we are willing to ascribe knowledge in this case then we will be understanding the details of the case such that the true belief so formed satisfies the anti-luck condition. For example, in crediting Jenny with knowledge in this case we are taking it as given that there is no conspiracy afoot among members of the public to deceive her in this regard, albeit one which is unsuccessful in this case. The problem, however, is that it is hard to see how crediting knowledge to Jenny can be squared with strong virtue epistemology. After all, given that she forms her belief by, for the most part, trusting the word of another, Jenny's cognitive success does not seem to be explained by her cognitive abilities specifically at all, but rather by her informant's. If that's right, then it seems that strong virtue epistemology should deny knowledge in

these cases, despite this being a counterintuitive result, and accordingly align themselves with the unpopular ‘sceptical’ view regarding the epistemology of testimony.³⁷

If proponents of strong virtue epistemology do not opt to take this line then one strategy they might try to exploit is to claim that Jenny is exhibiting a far greater degree of cognitive ability than at first meets the eye. Now it is certainly undeniable that Jenny is exhibiting some relevant cognitive ability in this case. Although she in fact asks the first person she meets for directions, we are clearly meant to be reading the example such that she does exercise some discretion in this respect. Had the first person she met been obviously mad, or a stereotypical tourist, for example, then we would expect her to move on to the next prospective informant down the street. Moreover, we expect Jenny to be exercising similar discretion when it comes to what she is told, and how she is told it. If the informant were to give her directions which were manifestly false, for instance, then we wouldn’t expect her to follow them regardless. Furthermore, if the manner in which the informant passed on the directions was clearly questionable—if the informant was vague, shifty, hostile and evasive, say—then we would expect our hero to exercise due caution. It is undeniable then that Jenny exhibits a fair degree of relevant cognitive ability, and hence one might try to argue on this basis that despite first appearances her cognitive success is because of *her* cognitive abilities after all.

The problem facing this line of counterargument is that it doesn’t really get to the nub of the problem. After all, no-one is denying that Jenny’s cognitive success is in part due to her cognitive abilities: this case is not meant to be a counterexample to the ability intuition. The point, however, is that the degree of trust involved in this case means that the cognitive success is not *because of* Jenny’s cognitive abilities—*viz.*, it is not primarily creditable to Jenny’s cognitive abilities that she is cognitively successful. To ensure that Jenny’s cognitive success would be because of her cognitive abilities it would be necessary to substantially reduce this degree of trust, but then one is back to the original counterintuitive response of arguing that Jenny lacks knowledge in this case after all.³⁸

Although the Jenny case causes a problem all by itself, the Jenny and Barney cases collectively pose a quite formidable difficulty. For notice that while the Barney case puts pressure on the proponent of strong virtue epistemology to *strengthen* her account of knowledge so that it excludes knowledge in this case, the Jenny case puts pressure on the proponent of strong virtue epistemology to *weaken* her account. Thus, the two types of case pull this account of knowledge in two opposing directions, with the potential of collectively pulling the view asunder. For if you strengthen the view in order to deal with the Barney case then you face an even tougher problem when it comes to the Jenny case; and if you weaken the view in order to deal with the Jenny case then you face an even tougher problem when it comes to the Barney case.^{39,40}

4. ANTI-LUCK VIRTUE EPISTEMOLOGY

Rather than review all the various ways in which one might go about defending either a virtue epistemology or an anti-luck epistemology I want to suggest that the difficulties facing these two views in fact demonstrate that we need an account of knowledge which synthesizes the best features of both. As we will see, such a view is not only able to evade the difficulties facing these alternative positions but can also be motivated on independent grounds. In short, there is no need to revise one's anti-luck or virtue epistemology in order to deal with the problems in hand, thereby incurring the various theoretical problems that such revisions will bring, since there is an alternative proposal available that does the job without need to resort to such manoeuvres.

In order to see what this proposal amounts to, let us return to the two intuitions that we began with. Whereas anti-luck epistemology takes the anti-luck intuition as primary and then tries to account for the ability intuition by appeal to an anti-luck condition, virtue epistemology treats the ability intuition as primary and then tries to account for the anti-luck intuition by appeal to an ability condition. Both proposals thus effectively treat these intuitions as essentially intertwined, such that properly responding to one of them thereby accommodates the other. This is their fundamental mistake, since the proper way to think about the relationship between these two intuitions is as treating them as posing two *distinct* demands on knowledge.

On reflection, this is not as surprising as it may at first seem. On the one hand, we have seen that dealing with the anti-luck intuition leads us to adopt an anti-luck condition which involves a modal principle. But if that's right, then it is to be expected that one could satisfy such a condition while not exhibiting any cognitive ability, since whatever modal requirement is imposed with imagination one could think of a way in which it can be satisfied in a manner that bears no relation to the agent's cognitive abilities. The relevant modal principle will simply demand a match between belief and fact in appropriate counterfactual cases, but a cognitive ability requires far more than this—*viz.*, it requires an appropriate direction of fit between belief and fact. On the other hand, on reflection there is no reason why forming a true belief through cognitive ability should suffice to meet the relevant anti-luck condition—such, after all, is the moral of Gettier-style cases. One can exhibit one's reliable cognitive abilities and yet form a true belief that could very easily have been false.

Furthermore, note that trying to 'beef-up' your anti-luck/ability condition in order to resolve these difficulties will only exacerbate the problem. As we saw in the last section, strong virtue epistemology is a good example of this point. Not only does the move to the more austere version of the view not deal with the anti-luck intuition, but it also now loses the ability to accommodate a common form of testimonial knowledge.

Recognising that the anti-luck and ability intuitions are epistemically distinct, and therefore require independent epistemic conditions, leads us away from anti-luck and virtue epistemology and towards a hybrid view, which I call *anti-luck virtue epistemology*:

Anti-Luck Virtue Epistemology

S knows that *p* if and only if *S*'s safe true belief that *p* is the product of her relevant cognitive abilities (such that her safe cognitive success is to a significant degree creditable to her cognitive agency).

This proposal incorporates both an anti-luck condition (the demand that the true belief be safe) and an ability condition. Note that the ability condition in play here is that proposed by a weak virtue epistemology rather than a strong virtue epistemology, except that in line with the hybrid nature of the account we are interpreting this condition as demanding that the agent's *safe* cognitive success should be to a significant degree creditable to her cognitive agency. As we will see, so long as we accompany this ability condition with an anti-luck condition then there is no need to opt for the stronger rendering of the ability condition as a means of dealing with cases that trade on the anti-luck intuition. The conception of safety in play in the anti-luck condition should be understood along the lines set out above in §2.

Anti-luck virtue epistemology can accommodate all the usual problem cases. The Gettier-style cases, including non-standard Gettier-style cases which involve environmental epistemic luck, like the Barney case, are dealt with by the anti-luck condition, since in all of these examples we have a true belief which is unsafe. The same goes for all the other cases which we noted above can be accommodated by the safety principle, such as the Lottie, Ernie and Mathema cases. In short, anti-luck virtue epistemology has all the advantages of an anti-luck epistemology.

Moreover, the kind of cases that anti-luck epistemology cannot deal with, such as the Temp and Alvin cases, are dealt with by the ability condition, since examples like this are precisely cases in which the agent has a safe belief which is not formed via the reliable cognitive traits that make up her cognitive character. In short, anti-luck virtue epistemology has all the advantages of a virtue epistemology. Crucially, however, since we have not 'beefed up' the ability condition in terms of adding a 'because of' clause in the manner of strong virtue epistemology, we do not need to worry about cases like the Jenny case. For while Jenny's cognitive success is not primarily creditable to her cognitive agency, her safe true belief is to a significant degree creditable to her cognitive agency and that is all that is required for knowledge on this view. Furthermore, the Barney case is dealt with by the separate anti-luck condition. Anti-luck virtue epistemology thus does not inherit the problems of strong virtue epistemology.

Finally, as with virtue epistemology, we do not need to concern ourselves here with how this view fares with regard to the epistemic externalism/internalism dispute. While the view as it is presently defined—with an inclusive externalist conception of cognitive ability in play—will be

controversial from an internalist perspective, there is nothing preventing the proponent of anti-luck virtue epistemology from ‘internalising’ the ability condition in order to satisfy her internalist sympathies.

We thus seem to have an entirely adequate theory of knowledge, in that it can account for our most fundamental intuitions about what it is to have knowledge. Are we, *qua* epistemologists, home and dry?⁴¹

5. DIAGNOSIS

One question that we might ask about anti-luck virtue epistemology is why knowledge has this two-part structure in the first place. Put another way, why do we have these two master intuitions about knowledge, and have them in such a fashion that they make distinct demands on our theory of knowledge? It may be helpful in this regard to undertake a thought-experiment that has been suggested by Edward Craig as a way of casting light on the nature of knowledge.⁴² Imagine an imaginary society which lacked the concept of knowledge. Why might they feel the need to introduce it? Well, notice first that it would be very practically useful to have some way of picking out good informants—i.e., informants who can help us to find out the truth on matters that interest us. We could thus imagine a concept very like knowledge—a kind of proto-knowledge concept—being employed for just this purpose. Call this proto-knowledge, ‘knowledge*’, and call anyone who possesses knowledge* a ‘knower*’. The idea is then that it would be useful to label good informants as knowers*, and accordingly to label the accurate information that they offer on subjects about which they are good informants, knowledge*.

So, for example, imagine that John lives on a hill and so has a particularly good view of what is happening in the valley below (and that he is generally truthful and helpful, etc.). He would thus be a good informant when it comes to a range of propositions concerning what is happening in the valley. It would clearly be practically useful for us to flag the fact that John is a good informant in this regard, and we can do this by calling him a knower* as regards these propositions, in that his true beliefs in these propositions amount to knowledge*.

Note that knowledge* is not yet like our concept of knowledge. For one thing, the concept only applies to other people’s true beliefs, while we also use the concept of knowledge to classify our own beliefs. In addition, in deciding whether an agent is a knower* we are only assessing how good an informant she is relative to the actual circumstances that she finds herself in—i.e., the ‘live’ error-possibilities that are in play in her environment. In the case of John just described, for

example, all that is at issue is whether he has a good view of the valley and the ability (and inclination) to make use of this advantage. The salient error-possibilities when it comes to the question of whether he has knowledge* are thus things like whether there is something in his environment which is obscuring his view (a heavy fog, say). Note, however, that our concept of knowledge treats a far greater range of error-possibilities as salient. In particular, it is also responsive to potential error-possibilities, even if they are not actual. (Recall the anti-luck intuition, one consequence of which is that the fact that you could very easily have been deceived is a ground to deny you knowledge, even if in fact you were not deceived.)

Interestingly, however, we would expect the proto-concept which is designed to pick out good informants to evolve over time so that it begins to resemble our concept of knowledge. For example, we could imagine knowledge* ultimately being used to classify oneself and not just others, and the application of the concept being ‘stretched’ so that it is responsive to non-actual but potential error-possibilities, and not just the actual ones (a process that Craig calls ‘objectification’). Over time, then, knowledge* would evolve into knowledge.

Many philosophers find this genealogical account of the source of the concept of knowledge to be very persuasive.⁴³ In particular, virtue epistemologists often cite this story as offering support for their view.⁴⁴ After all, this account of the source of the concept of knowledge puts good informants at the heart of the story, and one natural way of thinking about what constitutes a good informant is that she is an informant who is exercising a reliable cognitive ability (think, for example, of the case of John just described). This would thus appear to suggest that what is most central to the concept of knowledge is the ability intuition which is the primary concern of virtue epistemology. Interestingly, however, once we start to think about this account of the source of the concept of knowledge in more detail, it becomes apparent that it actually lends greater support to anti-luck virtue epistemology than it does to virtue epistemology.

In particular, while there is clearly something right about the suggestion that a good informant is a reliable informant, this does not play into the hands of the virtue epistemologist in the manner that it may at first appear. The reason for this is that there is an important ambiguity in the very notion of a reliable (and hence good) informant. In one sense, it means an informant who possesses a reliable cognitive ability with regard to the target subject matter (and who is willing to sincerely communicate what she believes, something that we will take for granted in what follows). In another sense, it means an informant whom one can rely on (i.e., whose information will not lead you astray).

Now one might naturally think that this is a distinction without a difference, in that informants who possess reliable cognitive abilities in the sense just specified are thereby informants one can rely on, and *vice versa*. Closer inspection, however, reveals that first appearances

are deceptive on this score. In order to see this, we just need to notice that it can be appropriate to rely on an informant who is forming her true belief via an unreliable cognitive ability, and also that it can be inappropriate to rely on an informant who nevertheless is forming a true belief via a reliable cognitive ability.

First, consider a potential informant who possesses a reliable cognitive ability as regards a certain subject matter but who is in an environment in which there exists a misleading defeater, one which you know about, but the prospective informant does not, and one which moreover you are unable to defeat. An example might be an informant who is a reliable barn detector, but where you have been given a misleading ground (e.g., false testimony from a good source) for supposing that the informant is in barn façade county. Given that this is a misleading defeater, the informant is in fact a *reliable* informant about the relevant subject matter. But given also that you know about the misleading defeater, and are aware that you are unable to defeat that defeater, would you be able to *rely on* this informant? Surely not.

The converse point also holds. In particular, we can imagine a case where there are compensating factors in play, known only to us, which mean that we can rely on the information presented to us by an informant even though this information is not the product of the informant's reliable cognitive abilities. Imagine, for example, an informant who thinks that they have clairvoyant powers, but in fact is mistaken on this score (and we know this). Suppose further that we also know that this informant's wife is a very powerful person who wants her husband to continue to believe that he has this power and hence does what she can, where possible, to make sure that events turn out in the way that her husband predicts. Finally, suppose that we know that the informant's wife can fix the result of any horse race. With this knowledge in hand, the testimony of the informant regarding who will win tomorrow's horse race would certainly be information that one could rely on, even though the informant's true belief in this regard is in no way the product of a reliable cognitive ability.

In general, what is key to both of these kinds of cases is the role that luck is playing. In cases in which the informant's relevant cognitive abilities are reliable but where we are nonetheless unable to rely on the information she provides, the problem is that a dose of bad epistemic luck is cancelling out the good epistemic luck that our informant possesses the relevant reliable cognitive abilities (and thus is in this sense a good informant). In the case described above, for example, this bad epistemic luck is the presence of the undefeated misleading defeater regarding the barn façades. In contrast, in cases in which the informant lacks the relevant reliable cognitive abilities but is nonetheless providing us with information that we can rely on, a dose of good epistemic luck on our part is cancelling out the poor epistemic luck that our informant lacks the relevant reliable cognitive abilities (and thus is in this sense a bad informant). In the 'clairvoyant' case

described above, for example, this good epistemic luck is our knowledge of the compensating factors in play.

With this point in mind, it ought to be clear why this ambiguity in the idea of a reliable informant explains why the concept of knowledge that evolves from the proto-concept will generate *both* the anti-luck and the ability intuition. For as the range of cases which the concept of knowledge is meant to apply to widens, so the distinction will open up between good informants who are reliable and good informants that we can rely on, and we would expect the concept of knowledge that results to respect both sides of this distinction. In particular, examples where an agent possesses the relevant reliable cognitive abilities but where the presence of epistemic luck means that we would not be able to rely on this agent *qua* informant would not be counted as cases of knowledge. Similarly, those cases in which an agent forms a true belief in an epistemically friendly environment—such that any true belief so formed would not be subject to epistemic luck—would not be counted as cases of knowledge so long as the agent concerned failed to exhibit the relevant reliable cognitive abilities (even though we could rely on this agent *qua* informant). In short, the concept of knowledge that results will both (i) disallow cases of true belief as knowledge where the belief isn't appropriately due to the relevant cognitive abilities on the part of the agent, and (ii) disallow cases of true belief as knowledge where the truth of the belief is substantively due to luck and hence unsafe.

A very plausible and popular story about the genealogy of the concept of knowledge thus lends support to anti-luck virtue epistemology after all, despite first appearances. In fact, if I am right that the goal of picking out reliable informants is ambiguous in the way just described, then contrary to the prevailing wisdom on this score, this 'just so' account of the concept of knowledge actually *favours* anti-luck virtue epistemology over rival proposals, such as virtue epistemologies.

6. CONCLUDING REMARKS

We have thus seen that there is an account of knowledge available—anti-luck virtue epistemology—which is able to deal with a wide range of problem cases in epistemology, and which, more importantly, is able to deal with those cases in a more satisfactory way than competing proposals. Moreover, we have seen that there is a plausible diagnostic story available as to why knowledge should have the structure dictated by this proposal. There is thus good reason for supposing that anti-luck virtue epistemology is the right theory of knowledge to opt for. Far from being a lost cause, the analytical project is in fact alive and kicking.

NOTES

¹ Earlier versions of this paper have been presented at the *Midwest Epistemology Workshop* in St. Louis in September 2009; the *Iberian-American Philosophy Congress* in Lisbon in October 2009; the *Knowledge, Luck, Value and Virtue: Themes from the Philosophy of Duncan Pritchard* conference at the Institut Jean-Nicod in Paris in April 2010; at a Royal Institute of Philosophy lecture at the University of Hull in May 2010; and at the Open University Summer School at the University of Bath in August 2010. I am grateful to the audiences on these occasions for helpful discussions. In addition, thanks go to Evan Butts, J. Adam Carter, Julien Dutant, Georgi Gardiner, Mikkel Gerken, Emma Gordon, John Greco, Adam Green, Adrian Haddock, Jesper Kallestrup, Christoph Kelp, Joseph Kuntz, Alan Millar, Wayne Riggs and Ernie Sosa. This paper was written while I was in receipt of a Phillip Leverhulme Prize.

² The most prominent defender of this line of thought is of course Timothy Williamson. See, especially, chapter 1 of his *Knowledge and its Limits* (Oxford: Oxford UP, 2000).

³ Consider this passage from one of the main overviews of contemporary theory of knowledge:

“Why not say that knowledge is true belief? The standard answer is that to identify knowledge with true belief would be implausible because a belief that is true just because of luck does not qualify as knowledge.” (Matthias Steup, “The Analysis of Knowledge”, *Stanford Encyclopaedia of Philosophy*, Edward Zalta, ed., <http://plato.stanford.edu/entries/knowledge-analysis/index.html>, 2006, §1.2)

⁴ Consider, for example, this passage from an influential contemporary introduction to epistemology:

“This was just the point of the Gettier counter-examples; nothing in the tripartite definition excluded knowledge by luck.” (Jonathan Dancy, *Introduction to Contemporary Epistemology* (Oxford: Blackwell, 1985), p. 134)

⁵ I discuss the notion of a cognitive ability further below.

⁶ See, for example, this passage in a recent survey of epistemology:

“Knowledge, then, requires factual belief. However, this does not suffice to capture the nature of knowledge. Just as knowledge requires successfully achieving the objective of true belief, it also requires success with regard to the formation of that belief. In other words, not all true beliefs constitute knowledge; only true beliefs arrived at in the right way constitute knowledge.

What, then, is the right way of arriving at beliefs? [...] We might begin by noting that sound reasoning and solid evidence seem to be the way to acquire knowledge. By contrast, a lucky guess cannot constitute knowledge. Similarly, misinformation and faulty reasoning do not seem like a recipe for knowledge, even if they happen to lead to a true belief.” (David Truncellito, “Epistemology”, *Internet Encyclopaedia of Philosophy*, Bradley Dowden and James Fieser, eds., www.iep.utm.edu/e/epistemo.htm, 2007, §1.c)

⁷ Strictly speaking, we should be relativizing this principle—and the safety principle offered below—to the actual basis for the agent’s belief in *p*, since otherwise it will be susceptible to ‘grandmother’-style counterexamples (see Robert Nozick, *Philosophical Explanations* (Oxford: Oxford UP), 179). In what follows for ease of expression I will take such relativization to be assumed.

⁸ See Nozick (*ibid.*), ch. 3. See also Fred Dretske, “Epistemic Operators”, *Journal of Philosophy* 67 (1970): 1007-23, and “Conclusive Reasons”, *Australasian Journal of Philosophy* 49 (1971): 1-22; and Alvin Goldman, “Discrimination and Perceptual Knowledge”, *Journal of Philosophy* 73 (1976): 771-91. For some recent texts which sympathetically explore the sensitivity principle, see Sherrilyn Roush, *Tracking Truth: Knowledge, Evidence and Science* (Oxford: Oxford UP, 2005); Kelly Becker, *Epistemology Modalized* (London: Routledge, 2007); Tim Black and Peter Murphy, “In Defense of Sensitivity”, *Synthese* 154 (2007): 53-71; and Tim Black, “Defending a Sensitive Neo-Moorean Invariantism”, *New Waves in Epistemology*, Vincent Hendricks and Duncan Pritchard, eds., 8-27, (Basingstoke: Palgrave Macmillan, 2008).

⁹ This is, of course, one of the two cases offered by Edmund Gettier, “Is Justified True Belief Knowledge?”, *Analysis* 23 (1963): 121-3.

¹⁰ This case is adapted from one offered by Roderick Chisholm, *Theory of Knowledge* (2nd ed.), (Englewood Cliffs, NJ: Prentice-Hall), p. 105.

¹¹ The barn-façade case was first put forward in print by Goldman (*op. cit.*), who credits the example to Carl Ginet.

¹² That what we care about when it comes to ascriptions of knowledge is the modal closeness of the error rather than the probability of the error accords with some of the recent empirical literature on subjects’ judgements about risk and luck. For it seems that subjects assign far more risk to an event that is modally close than to one that is modally distant even when they are fully aware that the two events have an equal probability of occurring. For a survey of the relevant empirical literature on risk and luck and its relevance to epistemology, see Pritchard and Matthew Smith, “The Psychology and Philosophy of Luck”, *New Ideas in Psychology* 22 (2004): 1-28.

¹³ For discussion of some of the other problems facing sensitivity, see Williamson (*op. cit.*, ch. 7) and Pritchard, “Sensitivity, Safety, and Anti-Luck Epistemology”, *The Oxford Handbook of Skepticism*, John Greco, ed., 437-55, (Oxford: Oxford UP, 2008).

¹⁴ This case is due to Ernest Sosa, “How to Defeat Opposition to Moore”, *Philosophical Perspectives* 13 (1999): 141-54.

¹⁵ For a key defence of the safety principle, see Sosa (*op. cit.*). Similar principles have been advanced by Steven Luper, “The Epistemic Predicament: Knowledge, Nozickian Tracking, and Skepticism”, *Australasian Journal of Philosophy* 62 (1984): 26-50, and “Indiscernability Skepticism”, *The Skeptics: Contemporary Essays*, Luper, ed., 183-202, (Aldershot:

Ashgate, 2003); Mark Sainsbury, “Easy Possibilities”, *Philosophy and Phenomenological Research* 57 (1997): 907-19; Williamson (*op. cit.*, ch. 5); and Pritchard, “Resurrecting the Moorean Response to the Sceptic”, *International Journal of Philosophical Studies* 10 (2002): 283-307, *Epistemic Luck* (Oxford: Oxford UP, 2005), and “Anti-Luck Epistemology”, *Synthese* 158 (2007): 277-97.

¹⁶ Earlier versions of the safety principle did not make this feature of the principle explicit, and this gave rise to the objection that safety cannot simultaneously deal with both the Lottie and the Ernie cases. For a clear statement of this objection, see Greco, “Worries About Pritchard’s Safety”, *Synthese* 158 (2007): 299-302. For a response in line with the account of the safety principle put forward here, see Pritchard, “Anti-Luck Epistemology” (*op. cit.*).

¹⁷ For further discussion of this point, see Pritchard, “Safety-Based Epistemology: Whither Now?”, *Journal of Philosophical Research* 34 (2009): 33-45, and “There Cannot Be Lucky Knowledge”, *Contemporary Debates in Epistemology* (2nd ed.), Steup and John Turri, eds., (Oxford: Blackwell, forthcoming).

¹⁸ For further discussion of some of the problems facing the safety principle, see Ram Neta and Guy Rohrbaugh, “Luminosity and the Safety of Knowledge”, *Pacific Philosophical Quarterly* 85 (2004): 396-406; Juan Comesaña, “Unsafe Knowledge”, *Synthese* 146 (2005): 393-402; Greco (*ibid.*); and Avram Hiller and Neta, “Safety and Epistemic Luck”, *Synthese* 158 (2007): 303-14. I offer a fuller defence of the safety principle in Pritchard (*ibid.*), to which Stephen Hetherington responds in “There Can Be Lucky Knowledge” (Steup and Turri, *op. cit.*).

¹⁹ Although I have opted for the safety principle over the sensitivity principle here as the best way of thinking about the anti-luck condition, I am sympathetic to the idea that there is a way of thinking about the sensitivity principle such that it is equivalent to the safety principle, although it should be noted that the sensitivity principle, so conceived, would be a very different beast to that put forward by folk such as Nozick. For more on this point, see Pritchard, *Epistemic Luck* (*op. cit.*), ch. 6. For a more general discussion of the relative merits of the safety and sensitivity principles, see Pritchard, “Sensitivity, Safety, and Anti-Luck Epistemology” (*op. cit.*).

²⁰ For a sustained defence of such a view, see Pritchard, *Epistemic Luck* (*op. cit.*) and “Anti-Luck Epistemology” (*op. cit.*). (This paper thus represents an apostasy on my part).

²¹ While most (if not all) epistemic internalists will accept this claim, there is a significant group of internalists who maintain that we should not define epistemic internalism in this ‘accessibilist’ fashion. See especially Earl Conee and Richard Feldman, *Evidentialism* (Oxford: Oxford UP, 2004). For discussion of this issue, see Pritchard, “Evidentialism, Internalism, Disjunctivism”, *Evidentialism and its Discontents*, Trent Dougherty, ed., 362-92, (Oxford: Oxford UP, 2011). For more general discussion of the epistemic externalism/internalism distinction, see the papers collected in Hilary Kornblith, ed., *Epistemology: Internalism and Externalism* (Oxford: Blackwell, 2001).

²² For more discussion of the chicken-sexer example, see Linda Zagzebski, *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge* (Cambridge: Cambridge UP, 1996), §2.1 and §4.1; Robert Brandom, “Insights and Blindspots of Reliabilism”, *Monist* 81 (1998): 371-92; and Pritchard, *Epistemic Luck* (*op. cit.*, *passim*).

²³ And, for that matter, the sensitivity principle as well. Had what Temp believed not been the case then the helper would have altered the thermostat accordingly to ensure that Temp formed a true belief regardless. Hence it is true that had what he believed been false, he wouldn’t have believed it, and so sensitivity is satisfied.

²⁴ On Greco’s agent reliabilist view, for example, epistemic virtues are essentially the same as cognitive abilities as we are understanding this notion here. See, for example, Greco, “Agent Reliabilism”, *Philosophical Perspectives* 13 (1999): 273-96, and *Putting Skeptics in Their Place: The Nature of Skeptical Arguments and Their Role in Philosophical Inquiry* (Cambridge: Cambridge UP, 2000). (Indeed, the account of cognitive abilities that I offer is largely inspired by Greco’s work in this regard). In contrast, some commentators have explicitly distinguished epistemic virtues and cognitive abilities, most notably Zagzebski (*ibid.*, *passim*).

²⁵ This case is due to Alvin Plantinga, *Warrant: The Current Debate* (Oxford: Oxford UP, 1993), p. 199.

²⁶ Even if, as is arguably the case, epistemic externalism fares better in this respect when it is cast along virtue-theoretic lines. For a recent detailed defence of an externalist rendering of virtue epistemology in light of internalist objections, see Greco, *Achieving Knowledge* (Cambridge: Cambridge UP, 2009), pt. 1.

²⁷ Note that I say ‘creditable’ rather than ‘of credit’. These two notions clearly are not the same—for example, one’s cognitive success could be creditable to one’s cognitive agency without being at all of credit to one (perhaps the cognitive success is the result of an inquiry that one ought not to be pursuing because, say, there are epistemically more desirable inquiries that one should be focussing on instead). Unfortunately, one often finds these two notions equated in the literature—see, for example, Greco, “Knowledge as Credit for True Belief”, *Intellectual Virtue: Perspectives from Ethics and Epistemology*, Michael DePaul and Zagzebski, eds., 111-34, (Oxford: Oxford UP, 2003)—and this has generated a fair amount of confusion. For further discussion of this point, see Pritchard, Alan Millar and Adrian Haddock, *The Nature and Value of Knowledge: Three Investigations* (Oxford: Oxford UP, 2010), §2.4.

²⁸ Note that there is an alternative reading of the ‘because of’ in play in this version of virtue epistemology that has been put forward by Sosa. See his *A Virtue Epistemology: Apt Belief and Reflective Knowledge* (Oxford: Clarendon Press, 2007), and *Reflective Knowledge: Apt Belief and Reflective Knowledge* (Oxford: Clarendon Press, 2007). I comment on this alternative reading below—see endnotes 33, 35 and, especially, 39.

²⁹ The most prominent defender of strong virtue epistemology in the recent literature has been Greco. See, especially, his *Achieving Knowledge* (*op. cit.*), but also “Knowledge as Credit for True Belief” (*op. cit.*) and “The Nature of Ability and

the Purpose of Knowledge”, *Philosophical Issues* 17 (2007): 57-69. Whether Greco intends to offer an analysis of knowledge in these terms is not so clear, however. In the introduction to *Achieving Knowledge*, for example, he claims that he is merely offering necessary conditions for knowledge, although elsewhere—in chapter five of that same book for instance—he does express his view as offering necessary and sufficient conditions for knowledge. In any case, many have read him as a proponent of strong virtue epistemology even if in the final analysis this is not his settled view.

³⁰ Strong virtue epistemology also seems to possess another key attraction, which is its ability to account for the distinctive value of knowledge. See, for example, Sosa, *A Virtue Epistemology* (*op. cit.*), ch. 4, and Greco, “The Value Problem”, *Epistemic Value*, Haddock, Millar and Pritchard, eds., 313-21, (Oxford: Oxford UP, 2009). For critical discussion of this claim, see Pritchard, “Knowledge, Understanding and Epistemic Value”, *Epistemology (Royal Institute of Philosophy Lectures)*, Anthony O’Hear, ed., 19-43, (Cambridge: Cambridge UP, 2009); “The Value of Knowledge”, *Harvard Review of Philosophy* 16 (2009): 2-19; and Pritchard, Millar and Haddock (*op. cit.*), ch. 2. For an overview of the issues in play here, see Pritchard, “Recent Work on Epistemic Value”, *American Philosophical Quarterly* 44 (2007): 85-110; and “The Value of Knowledge”, *Stanford Encyclopedia of Philosophy*, Zalta, ed., <http://plato.stanford.edu/entries/knowledge-value/> (2007).

³¹ For example, Zagzebski’s (*ibid.*) neo-Aristotelian version of virtue epistemology incorporates an internalist conception of epistemic virtue along broadly these lines. See also Sosa, who argues for a version of virtue epistemology which maintains that all human knowledge requires a reflective perspective. See Sosa, *Knowledge in Perspective: Selected Essays in Epistemology* (Cambridge: Cambridge UP, 1991), *A Virtue Epistemology* (*op. cit.*), and *Reflective Knowledge* (*op. cit.*).

³² For more on this point, and on environmental luck more generally, see Pritchard, Millar and Haddock (*op. cit.*), ch. 2.

³³ This is the line taken by Sosa. See, for example, his *A Virtue Epistemology* (*op. cit.*), ch. 5. See also endnotes 35 and 39.

³⁴ For more discussion of this point, see Pritchard, “Greco on Knowledge: Virtues, Contexts, Achievements”, *Philosophical Quarterly* 58 (2008): 437-47, and Jonathan Kvanvig, “Responses to Critics”, *Epistemic Value* (Haddock, Millar and Pritchard, *op. cit.*), 339-51. See also Pritchard, “Virtue Epistemology and Epistemic Luck, Revisited”, *Metaphilosophy* 39 (2008): 66-99.

³⁵ For further discussion of this point, see Pritchard, Millar and Haddock (*op. cit.*), §2.5. One proponent of a strong virtue epistemology who recognises this difficulty is Sosa, and this is part of the reason why he allows knowledge in the Barney case. See, for example, his *A Virtue Epistemology* (*op. cit.*), ch. 5. See also endnotes 33 and 39.

³⁶ This case is adapted from one offered by Jennifer Lackey, albeit to demonstrate a slightly different point. See Lackey, “Why We Don’t Deserve Credit for Everything We Know”, *Synthese* 158 (2007): 345-61.

³⁷ Despite being counterintuitive, this view of the epistemology of testimony—known as *reductionism*—is not without its adherents. See, for example, Elizabeth Fricker, “Telling and Trusting: Reductionism and Anti-Reductionism in the Epistemology of Testimony”, *Mind* 104 (1995): 393-411.

³⁸ In “The Nature of Ability and the Purpose of Knowledge” (*op. cit.*), Greco offers a defence of this line of argument, though he does not make a very strong case for it. By analogy, he argues that one might score a very easy goal as a result of that goal being set-up by a display of tremendous skill. He maintains that the skill involved in setting up this easy goal does not undermine the achievement of the agent who scores the goal. The problem, however, is that this case is not relevantly analogous to the Jenny case. After all, what is crucial to that example is not that someone appropriately skilful helps Jenny, but rather more specifically that Jenny gains her true belief by (for the most part at least) *trusting* this other person. This is why, for example, other cases in which we depend on the skills of others—as when one takes an inner city road sign a face-value, for example—do not generate the same epistemological moral. In such cases my knowledge depends on—i.e., is made easy by—the skills of others, but it is not that I am merely trusting what the sign tells me: I have all kinds of independent grounds for believing what inner city road signs tell me. I discuss this point further in Pritchard, Millar and Haddock (*op. cit.*), §2.6.

³⁹ Note too that opting for a different construal of the ‘because of’ relation is of little help here. Sosa opts for a reading of this relation that understands it on the model of the manifestation of a power rather than in explanatory terms. See his *A Virtue Epistemology* (*op. cit.*), and *Reflective Knowledge* (*op. cit.*). So, for example, while it might be correct in the explanatory sense to say that the glass broke because I threw it on the floor, it is also correct to say that it broke because it was fragile, where this involves the alternative ‘power manifestation’ reading. As noted above (see endnotes 33 and 35), Sosa is committed to ascribing knowledge in the Barney case, contrary to intuition, and this is a direct result of the kind of strong virtue epistemology that he puts forward (it is, after all, undoubtedly the case that Barney manifests the relevant cognitive powers in acquiring his true belief). Sosa also struggles with the Jenny case too, since it is far from clear why her cognitive success should be credited to her cognitive powers, specifically. As a result, he argues that such cases display a kind of ‘distributed’ knowledge, where the cognitive success is creditable to more than one agent (in this case, Jenny and her informant). I discuss Sosa’s proposal in detail in Pritchard, “Apt Performance and Epistemic Value”, *Philosophical Studies* 143 (2009): 407-16.

⁴⁰ One potential problem for strong virtue epistemology which I haven’t engaged with here is the way in which the causal explanatory reading of the ‘because of’ relation might require the view to endorse a fairly radical form of

contextualism about knowledge ascriptions. For more discussion of this point, see Greco, “What’s Wrong With Contextualism?”, *Philosophical Quarterly* 58 (2008): 416-36, and Pritchard, “Greco On Knowledge” (*op. cit.*).

⁴¹ One perennial epistemological issue that I haven’t explored here is how anti-luck virtue epistemology—or, for that matter, anti-luck epistemology or virtue epistemology—fares when it comes to the problem of radical scepticism. Since this is a problem for all theories of knowledge, it ought to be uncontentious to put such an issue to one side for our purposes. That said, I do think that anti-luck virtue epistemology is potentially better placed to deal with this problem than rival views. On the one hand, anti-luck epistemology seems to get a victory over the sceptic that is too easy, since our beliefs that we are not, say, brains in vats will very easily satisfy the relevant modal condition if the sceptical possibility in question is indeed remote. On the other hand, strong virtue epistemologies seem to intensify the sceptical problem, since it is hard to see how our truly believing that we are not the victims of sceptical hypotheses (assuming these beliefs are true) could be because of our cognitive abilities. (For more discussion of this point, see Pritchard, “Radical Scepticism, Epistemic Luck and Epistemic Value”, *Proceedings and Addresses of the Aristotelian Society* (suppl. vol.) 82 (2008): 19-41). By treading an intermediate path between these two views, anti-luck virtue epistemology might thus be able to avoid the pitfalls that they each face. The development of this anti-sceptical line will need to be left for another occasion however.

⁴² See Craig, *Knowledge and the State of Nature: An Essay in Conceptual Synthesis* (Oxford: Clarendon Press, 1990).

⁴³ Interestingly, although Craig’s proposal has been highly influential within contemporary epistemology, there are in fact relatively few published discussions of it. For some notable exceptions, see Melissa Lane, “States of Nature, Epistemic and Political”, *Proceedings of the Aristotelian Society* 99 (1999): 211-24; Bernard Williams, *Truth and Truthfulness: An Essay in Genealogy* (Princeton, NJ: Princeton UP, 2002); Neta, “Epistemology Factualized: New Contractarian Foundations for Epistemology”, *Synthese* 150 (2006): 247-80; Greco, “The Nature of Ability and the Purpose of Knowledge” (*op. cit.*), and “What’s Wrong With Contextualism?” (*op. cit.*); Miranda Fricker, *Epistemic Injustice: Power and the Ethics of Knowing* (Oxford: Oxford UP, 2007), and “Scepticism and the Genealogy of Knowledge: Situating Epistemology in Time”, *Social Epistemology*, Haddock, Millar and Pritchard, eds., 51-68, (Oxford: Oxford UP, 2010); Martin Kusch, “Testimony and the Value of Knowledge”, *Epistemic Value* (Haddock, Millar and Pritchard, *op. cit.*), 60-94; and Klemens Kappel, “On Saying that Someone Knows: Themes from Craig”, *Social Epistemology* (Haddock, Millar and Pritchard, *op. cit.*), 69-88.

⁴⁴ See, for example, Greco, “The Nature of Ability and the Purpose of Knowledge” (*op. cit.*), and “What’s Wrong With Contextualism?” (*op. cit.*).