THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Radical Scepticism, Epistemic Luck, and Epistemic Value

Citation for published version:

Pritchard, D 2008, 'Radical Scepticism, Epistemic Luck, and Epistemic Value', *Aristotelian Society Supplementary Volume*, vol. 82, no. 1, pp. 19-41. <https://doi.org/10.1111/j.1467-8349.2008.00160.x>

Digital Object Identifier (DOI):

[10.1111/j.1467-8349.2008.00160.x](https://doi.org/10.1111/j.1467-8349.2008.00160.x)

Link:

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

Document Version:

Peer reviewed version

Published In:

Aristotelian Society Supplementary Volume

Publisher Rights Statement:

This is the pre-peer reviewed version of the following article: Pritchard, D. (2008). Radical Scepticism, Epistemic Luck, and Epistemic Value. *Aristotelian Society Supplementary Volume*, 82(1), 19-41, which has been published in final form at: <http://dx.doi.org/10.1111/j.1467-8349.2008.00160.x>

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.

RADICAL SCEPTICISM, EPISTEMIC LUCK AND EPISTEMIC VALUE

Duncan Pritchard

University of Edinburgh

ABSTRACT. It is argued that it is beneficial to view the debate regarding radical scepticism through the lens of epistemic value. In particular, it is claimed that we should regard radical scepticism as aiming to deprive us of an epistemic standing that is of special value to us, and that this methodological constraint on our dealings with radical scepticism potentially has important ramifications for how we assess the success of an anti-sceptical strategy.

0. Suppose we take it as a starting-point in our dealings with radical sceptical arguments that they are trying to demonstrate that the epistemic standing in question which, they claim, is unavailable to us—knowledge, typically—is an epistemic standing that is of distinctive value. This is a sensible enough working assumption, after all, since if the epistemic standing that is, it seems, snatched from us by the (successful) radical sceptical argument is not of special value, then it is hard to see why anyone should particularly care. Put another way, if the epistemic standing in question is not distinctively valuable, then why is our response to the sceptical argument not simply a resigned shrug?

On the face of it, such a working assumption may seem pretty benign, since we tend to take it for granted that the targets of radical sceptical arguments are indeed of special value to us. Suppose, however, that it could be demonstrated that the relevant epistemic standing is not of distinctive value, what then? As it happens, I am suspicious of the idea that knowledge *simpliciter* is distinctively valuable (at least in the way that we typically suppose), and I will sketch why I think this in a moment. Accordingly, this issue looms large for me. What I propose to explore here is how such a perspective on the value of knowledge impacts on our understanding of the radical sceptical argument. More generally, I want to argue that considerations about epistemic value can potentially have important ramifications for the debate regarding radical scepticism.

1. Following most contemporary epistemologists, I will take radical scepticism to be primarily concerned with knowledge. Such scepticism has been formulated in various ways, but the overarching formulation in recent times has been along the lines of the following template, where ‘E’ is some ‘everyday’ proposition (i.e., a proposition which we all tend to think that we know, such as that one has two hands), and ‘SH’ is a radical sceptical hypothesis which is logically inconsistent with the everyday proposition (e.g., that one is a brain in a vat (BIV)):

The Template Radical Sceptical Argument

- (S1) *S* does not know not: SH.
- (S2) If *S* does not know not: SH, then *S* does not know E.
- (SC) *S* does not know E.¹

The motivation for the first premise concerns the fact that there is, so the argument goes at any rate, nothing available to us to indicate that we are not the victims of a radical sceptical hypothesis. Since, for example, my experience of being a BIV who merely seems to have two hands would be subjectively indistinguishable from my experience of actually having two hands in normal circumstances, the thought goes that one cannot possibly know that one is not envatted. The same goes for other radical sceptical hypotheses.

The motivation for the second premise comes from the ‘closure’ principle for knowledge, which we can formulate as follows:

The Closure Principle for Knowledge

If *S* knows that *p*, and *S* competently deduces *q* from *p* (thereby coming to believe *q* while retaining her knowledge that *p*), then *S* knows that *q*.²

This principle is clearly highly intuitive. Moreover, given that we can legitimately suppose that any reasonably reflective agent is at least in a position to know the relevant entailment once prompted (i.e., that the target E-type proposition entails the target not-SH-type proposition), it follows from closure that our inability to know the denials of sceptical hypotheses is able to undermine our putative knowledge of E-type propositions. Indeed, given that the epistemic standing of our belief in just about any E-type proposition can be threatened in this way—one would just need to vary the sceptical hypothesis to suit—it follows that radical scepticism quickly ensues.

We thus have two highly intuitive premises as part of a valid argument which generates a sceptical conclusion. Given that the sceptical conclusion is so counterintuitive, what we have on our hands is thus a paradox.

2. There are various ways of responding to this argument. One approach, advanced by Fred Dretske (1970) and Robert Nozick (1981) amongst others, is to block the argument by denying the closure principle. Given the intuitive nature of this principle, however, this is more easily said than done, and this proposal has since fallen into disfavour. Accordingly, I won't explore it further here.³ With closure intact, however—at least in a form in which it can appropriately support (S2) of the sceptical argument—the options for evading radical scepticism of this sort become rather limited. Indeed, a few complications aside, I think most would now agree that the premise in the sceptical argument that needs to be denied is (S1).⁴ The anti-sceptical position that results is one that we might broadly call 'Moorean'.⁵

There are various ways of defending a Moorean anti-sceptical strategy, some more concessive to the sceptic than others. For example, at one end of the spectrum of concessiveness, one might argue that one has merely a standing entitlement to believe in the denials of sceptical hypotheses, even though one lacks any specific grounds to epistemically support such belief.⁶ Alternatively, at the other end of the spectrum, one might argue that one has factive epistemic support in favour of one's belief in this regard, such that the epistemic support one has for this belief actually *entails* the target proposition.⁷ One could also be more or less concessive to scepticism along a different axis in terms of the extent to which one wishes to accommodate our sceptical intuitions—in particular, our intuition, at least in contexts in which the problem of scepticism is under discussion, that there is *something* right about scepticism. One concessive option in this regard is to adopt some form of contextualism and so argue that in certain contexts what the sceptic claims is true.⁸ Still, the point remains that while there are a myriad of anti-sceptical approaches in play here, the step in the radical sceptical argument above that is ultimately objected to is (S1).

Moreover, notice that the concession to the sceptic that all Moorean anti-sceptical strategies make is to allow that agents are unable to distinguish between everyday experiences and the sort of sceptical experiences that would be generated by being the victim of a (relevant) sceptical hypothesis. So even if one argues that we do have adequate evidence

in favour of our beliefs in the denials of sceptical hypotheses, one still faces the task of treading the delicate path of endorsing this claim while nonetheless conceding that agents lack the relevant discriminatory powers.

Of course, one might be tempted to take the heroic option here of maintaining that there is nothing more to telling the difference between two scenarios than knowing that the one scenario has obtained rather than the other. After all, with closure in play, it will follow on the Moorean view that one does know both the E-type and the not-SH-type proposition. For example, one will know that one has hands and that one is not envatted, and thus that one has hands rather than being handlessly envatted. The temptation, then, is to maintain that one can tell the difference between having hands and being a handless BIV after all, since one is able to know that the one possibility has obtained and the other hasn't. I take it that this sort of claim will not stand up to closer scrutiny, however, attractive though it may be given the potential theoretical pay-off involved.

In order to see this, forget about scepticism for a moment and consider instead an apparently normal situation in which one is at home with one's children. Suppose now that someone raises the error possibility that you are not at home with your children at all, but are instead sitting next to extremely life-like robots pretending to be your children who were substituted a few moments ago when you were temporarily distracted. Given everything that you know about the plausibility of this scenario, you are surely in a position to rationally dismiss it. Indeed, once you have rationally formed a view as to why you are entitled to dismiss this error-possibility we would describe your epistemic situation as being one in which (i) you know that the persons before you are your children, and (ii) you know that the persons before you are not extremely life-like robots. In short, you know that it is your children who are before you rather than extremely life-like robots.

Still, all this is entirely compatible with the idea that were there to be extremely life-like robots before you, you wouldn't be able to tell the difference. That is, your belief in this regard is *insensitive*—viz., you would continue to believe that it is your children before you in such a scenario even though what you are in fact looking at are life-like robots. After all, your grounds in support of your knowledge that it is your children before you rather than extremely life-like robots need have nothing to do with whether you can tell the difference between the two, but would in all likelihood instead concern other factors—as noted above, it would probably just concern your grounded judgement of the plausibility of this scenario

given the background information available to you. Moreover, given the implausibility of the error possibility in question, we surely wouldn't insist that you should have specific grounds to dismiss it in order to know that you are looking at your children. We wouldn't require, for example, that you should examine your children closely to ensure that they aren't life-like robots (run a Turing test on them, say).

There is thus more to being able to tell the difference between two competing scenarios than simply knowing that the one scenario has obtained rather than the other.⁹ With this in mind, one should resist the temptation of arguing that because on the anti-sceptical approach which denies (S1) one is able to know both the E-type proposition and the denial of the SH-type proposition that one can thereby perform the incredible feat of being able to tell the difference between normal experiences and their sceptical counterparts.

3. I noted above that I am suspicious of the idea that knowledge *simpliciter* is distinctively valuable, at least in the way that we typically suppose. I think it is pretty clear that we instinctively regard knowledge as distinctively valuable. If this were not the case, then it would be mysterious why so much of our epistemological theorising unhesitatingly focuses on knowledge rather than on other lesser epistemic standings. Moreover, various thought-experiments also lend support to this claim. For example, suppose you were faced with the choice between knowing a proposition and, say, merely truly believing it. Suppose further that it is stipulated in advance that there will be no practical implications resting on your decision. Still, wouldn't you prefer to have knowledge?

The latter observation in particular suggests that the way in which we instinctively value knowledge is as something which is good in its own right, regardless of any additional instrumental value it might have. That is, what makes knowledge distinctively valuable, on this view, is that it is non-instrumentally, or *finally*, valuable. Furthermore, notice that if it is right that knowledge—unlike lesser epistemic standings, like reliably-formed belief—is finally valuable in this way then this would surely explain why we unhesitatingly focus on knowledge in our epistemological theorising, rather than on other epistemic standings. After all, whereas these other epistemic standings may be of great instrumental value, it is only when we get to the epistemic standing picked out by knowledge that we have something

which is good in its own right. Put another way, the difference in value between knowledge and lesser epistemic standings is a difference in kind, and not merely a difference in degree.

In contrast, if knowledge were just instrumentally valuable then it is hard to see how this could explain the way in which we instinctively value knowledge. For example, one might have a picture on which knowledge lay on a continuum of instrumental epistemic value, with lesser epistemic standings generally enjoying lesser degrees of instrumental value and more enhanced epistemic standings generally enjoying greater degrees of epistemic value. On this ‘continuum’ picture of the value of knowledge, however, it is mysterious why we should fixate on the particular point on the continuum that knowledge marks out, rather than, say, a point just before or just after.

The value that we intuitively attach to knowledge is thus final value, a final value that lesser epistemic standings lack.¹⁰

4. Of course, that we treat something as finally valuable does not mean that we are right to do so (think, for example, of the miser). The question, then, is whether we are right to finally valuable knowledge in the way that we do. I am inclined to think not.

I know of only one proposal in the literature that is able to offer a good argument for why knowledge is finally valuable. This is the virtue-theoretic account of the value of knowledge which gains its clearest expression in recent work by John Greco (e.g., 2002; 2007; *forthcominga*; *forthcomingb*). Essentially, Greco argues that once we understand knowledge along virtue-theoretic lines then we see that knowledge—unlike lesser epistemic standings—is a form of achievement and thus, like achievements more generally, is something that is valuable in its own right.¹¹

There are two claims in play here. The first, which we will call the *final value thesis*, is that achievements are finally valuable. The second, which we will call the *achievement thesis*, is that knowledge, unlike lesser epistemic standings, is a form of achievement. If both theses can be adequately defended, then that would suffice to support the claim that knowledge is distinctively valuable in the sense that the difference in value between knowledge and that which falls short of knowledge is a matter of kind and not merely degree.

Both theses are, in fact, highly plausible, at least once they are understood in the right way. Greco understands achievements to be successes that are because of ability, where the

‘because of’ relation is to be read in causal explanatory terms—i.e., roughly, that the ability is the best explanation of the success. This certainly seems right. Consider an archer—let’s call him ‘Archie’—taking aim at a target and successfully hitting the target. Suppose, however, that Archie lacks any relevant archery abilities, such that his success was simply down to luck. In such a case we would not regard Archie’s success as being because of his ability because he lacks the relevant ability. Equally, we would not regard his success as an achievement.

Interestingly, the same goes even if Archie has the relevant ability and is also successful but where the success is not appropriately related to—i.e., is not *because of*—his ability. Suppose, for example, that Archie’s success is ‘gettierized’, as would happen, for instance, if a freak gust of wind blew his arrow off-course but then a second freak gust of wind blew it back on course again. Even though in this case there is the conjunction of success and ability, the success clearly does not constitute an achievement, and the natural explanation for this is that the success is not because of the ability.

The idea that achievements should be understood as successes that are because of ability thus has a great deal to commend it. With achievements so construed, however, it is likewise plausible to suppose that achievements are finally valuable. In order to see this, suppose that the practical benefits of successfully hitting the target are exactly the same regardless of whether or not the success in question is because of one’s ability (and so a *bona fide* achievement). Nevertheless, wouldn’t one prefer to be successful because of one’s ability rather than because of luck? Given that the instrumental value at issue is the same either way, it follows that the greater value we attribute to achievements in such a case reflects a judgement on our part that achievements are finally valuable in a way that successes which fall short of achievements are not. Moreover, there seems no good reason why we shouldn’t take this commonsense judgment about the value of achievements at face-value.

The final value thesis thus looks compelling, at first pass at least. What about the achievement thesis? Here is where the virtue-theoretic element of the proposal comes to the fore, for virtue epistemologists like Greco argue that knowledge is to be understood as cognitive success (i.e., true belief) that is because of cognitive ability (i.e., intellectual virtue, broadly conceived). Most epistemologists will agree that there must be some ability component to knowledge, since cognitive success that is unrelated to one’s ability—and so

just down to luck—would clearly not constitute knowledge. What is distinctive about this proposal, however, is that it in effect exclusively defines knowledge in terms of ability, in that it argues that once the relationship between cognitive ability and cognitive success is understood correctly—in terms of the ‘because of’ relation—there is no need to introduce any further epistemic component into one’s analysis of knowledge.¹²

One can see the appeal of such a proposal by considering how it deals with standard Gettier-style cases. After all, such cases tend to involve agents who exhibit the relevant cognitive ability, and who are in addition cognitively successful, and yet who nonetheless lack knowledge because the cognitive success in question is not because of their cognitive ability. Imagine, for example, a farmer looking into a field who, by using her highly reliable perceptual faculties, comes to believe that there is a sheep in the field because she sees a sheep-shaped object there. Suppose further that she is cognitively successful in this regard, in that there really is a sheep in the field. Finally, however, suppose that the success is gettierized in the sense that what the farmer is in fact looking at is a big hairy dog which looks just like a sheep, and which is obscuring from view the real sheep in the field. We would not class such a success as a cognitive achievement, nor would we class it as knowledge.¹³ For Greco, the reason why is the same in both cases: the cognitive success in question is not because of the agent’s cognitive ability but is instead merely down to luck.

There is thus also a great deal of intuitive support for the achievement thesis as well. With both theses in play, the virtue-theoretic account of the distinctive value of knowledge seems to have much to commend it: knowledge, like achievements more generally, is rightly regarded as distinctively valuable because it is finally valuable.

5. Unfortunately, on closer inspection this proposal does not hold water. I think that many will probably find the final value thesis to be the weaker of the two claims in play here, on account of how there do seem to be some achievements—such as achievements which are very easy, or trivial, or just plain wicked—which are not valuable at all, much less finally valuable. This problem is, however, tractable, for notice that the claim is only that successes *qua* achievements are finally valuable. In particular, it is not part of the view that the overall (*ultima facie*) value of a success that constitutes an achievement should be particularly great.

To illustrate this point, consider the case of a trivial cognitive achievement. Suppose, for example, that one devotes one's time to pointlessly counting the pebbles on a beach, a task at which one is very successful. Now one might say that on this view there is nothing stopping such trivial successes counting as achievements, and thus as accruing final value. But notice that there is nothing problematic in itself about supposing that a trivial success, *qua* achievement, is finally valuable. The problem only emerges once one further supposes that if a success is finally valuable then it has a high overall value. This conditional, however, is false; indeed, cases like that of the trivial cognitive achievement represent clear counterexamples. After all, given the fact that the instrumental value of this success will inevitably be negligible, and given further that the opportunity cost of engaging in such an endeavour will almost certainly be very high—thereby generating quite a lot of disvalue—it follows that the overall value of this success will be very low (perhaps even negative), even despite the fact that it generates some final value.

Accordingly, I don't think the problem with the virtue-theoretic account of the value of knowledge lies with the final value thesis, at least once that thesis is understood correctly.¹⁴ Instead, I think the real weak point for this proposal lies with the achievement thesis. In particular, on closer analysis it turns out that knowledge is not a cognitive achievement at all. This is because one can possess knowledge without exhibiting a cognitive achievement, and exhibit a cognitive achievement while lacking knowledge.

Cases of the former are easy to come by, but the most straightforward cases concern testimonial knowledge. Consider an example put forward by Jennifer Lackey (2007), albeit to illustrate a slightly different point. Imagine someone getting off the train in an unfamiliar city and asking the first person she meets for directions to a local landmark. Suppose that the person our hero speaks to has first-hand knowledge of the area and communicates this information to our agent who promptly forms a true belief about the right way to go. Unless we know an awful lot less than we think we do, then we often gain testimonial knowledge in this fashion. Notice, however, that we wouldn't say that it is because of our *agent's* cognitive abilities that she is cognitive successful—i.e., that the explanation for her cognitive success is specifically her cognitive abilities. (We might say that it is because of her informant's cognitive abilities that she is successful, but that's a different matter). Testimonial knowledge of this sort, then, seems to be available even in the absence of a cognitive achievement.

It's worth spending a little time explaining exactly what the point being made here is. In particular, it is not being denied that our hero is exercising relevant cognitive abilities in this case, such that her cognitive abilities are playing *some* role in that success. We would expect her to be appropriately sensitive to potential defeaters, for example, such as an informant displaying body language which would indicate that she was not trustworthy. The point is only that such cases illustrate that one can gain knowledge by exhibiting a cognitive success which is not best explained in terms of one's own cognitive ability (and which is thus not because of that cognitive ability). There is thus sometimes less to knowledge than cognitive achievement.

Significantly, there is also sometimes more to knowledge than a cognitive achievement too. In order to see this, consider again the case of Archie described above. This time, though, imagine that Archie selects his target from a range of potential targets entirely at random before skilfully firing his arrow and hitting the target. Let us stipulate that the success here is not subject to Gettier-style luck, in that nothing intervenes between the ability and the success, such as the two freak gusts of wind which fortuitously cancel each other out that were described above. Nevertheless, the success is lucky because, unbeknownst to Archie, all of the other targets that he could have fired at contained forcefields which would usually repel arrows. Fortunately, Archie just happened to fire at the one target that lacked such a forcefield. His success is thus lucky in the sense that he could very easily have been unsuccessful.

Interestingly, however, I take it that we have a strong intuition in this case that even though the success in question is lucky, this does not prevent Archie from exhibiting a genuine achievement. Indeed, we would naturally say that his success—his hitting of the target—is because of his skills at archery. What this highlights is that there are two distinct ways in which a success could be lucky, only one of which is incompatible with achievements. On the one hand, there is the kind of luck found in Gettier-style cases in which something intervenes between the ability and the success to ensure that the success is not because of the ability. On the other hand, there is the kind of 'environmental' luck which does not intervene between ability and success in this way but rather simply reflects some unfortunate feature of the environment which would usually prevent success but which does not do so in this case. Environmental luck, it turns out, is compatible with genuine achievements, unlike Gettier-style luck.

The importance of this distinction between two types of luck and their relationship to achievements becomes apparent once one considers the epistemological analogue of the above case. Indeed, the famous ‘barn façade’ example is structurally analogous to the case involving Archie just described. In particular, unlike a standard Gettier-style case, such as that involving the farmer described above, in the barn façade case nothing intervenes between the agent’s cognitive ability and her cognitive success. She really does see a barn after all (and not, say, something that merely looks like a barn but which is obscuring from view a genuine barn, or something like that). Nevertheless, the agent’s belief is still lucky in the sense that she could very easily have been mistaken. The luck in question is thus environmental luck. Had the agent been looking at any other barn-shaped object in the area then she would have formed a false belief; it is just fortunate that despite the epistemically unfriendly nature of the environment, she happens nonetheless to be cognitively successful. Moreover, notice that just as in the Archie case involving environmental luck where we are inclined to say that a genuine achievement is being exhibited, similarly there seems no reason why we should deny that the agent in this case is not exhibiting a genuine cognitive achievement. Wouldn’t we say that her cognitive success is because of her cognitive abilities? Even so, she lacks knowledge.

The moral is that there is sometimes more to knowledge than exhibiting a cognitive achievement, and I suggest that this point has been overlooked by proponents of the virtue-theoretic proposal because they have failed to notice the distinction between Gettier-style and environmental epistemic luck (indeed, between Gettier-style and environmental luck more generally).

6. Now I don’t want to deny that there are moves available to someone like Greco when it comes to responding to these problems. He could insist, for example, that the testimonial case described above is an example in which the agent’s knowledge does reflect a cognitive achievement, and thus deny our intuitions in this case. Equally, he could also argue that there are relevant disanalogies between the Archer case and the barn façade case which can enable him to resist the conclusion that a cognitive achievement is exhibited in the latter situation even though an achievement is exhibited in the former. And, indeed, there are other potential moves available.

Space prevents me from exploring these responses here.¹⁵ In any case, the line of argument put forward in the last section is not meant to constitute a knockdown argument against the view anyway. Rather the point of this argument is to show that intuition in fact points towards a very different conception of knowledge, one on which knowledge and cognitive achievements come apart (in both directions).¹⁶

If the above line of argument is accepted then it should become clear why I am suspicious about the kind of distinctive value that we tend to ascribe to knowledge. Assuming that no further argument can be supplied in defence of the thesis that knowledge is finally valuable, then it seems that all we are entitled to suppose is that cognitive achievement, an epistemic standing which can fall short of knowledge, is finally valuable. But if that's right, then there is nothing particularly special about knowledge; in particular, knowledge is not more valuable, as a matter of kind and not merely degree, than that which falls short of knowledge.

In what follows, I will take it as given that no further argument in defence of the final value of knowledge is available, and thus conclude that the failure of the virtue-theoretic defence of the final value of knowledge entails that knowledge is not finally valuable. Moreover, I will also take it as given that cognitive achievements are finally valuable. I will then explore what consequences these two claims have for the problem of radical scepticism.

7. As explained there, the standard way to resolve the radical sceptical argument is by arguing—in a broadly ‘Moorean’ fashion—that we are able to know the denials of sceptical hypotheses after all, even though we are unable to distinguish between ordinary circumstances and the corresponding sceptical scenario. Given that we grant that knowledge is not distinctively valuable in the manner that we pre-theoretically suppose, the interesting question we now need to ask is whether the success of such a Moorean anti-sceptical strategy would ensure that we have thereby rescued from the sceptic's grasp a distinctively valuable epistemic standing. In particular, assuming the success of the Moorean strategy, is the kind of knowledge that is preserved potentially of the distinctively valuable sort such that it involves genuine cognitive achievement?

We can get a better handle on what is at issue here by reformulating the template sceptical argument outlined above specifically in terms of a type of knowledge which

essentially involves cognitive achievement—let's call it 'knowledge+'—and considering each of the reformulated premises in turn. Here, then, is the new version of the sceptical argument:

The Template Radical Sceptical Argument Reformulated

(S1+) *S* does not know+ not: SH.

(S2+) If *S* does not know+ not: SH, then *S* does not know+ E.

(SC+) *S* does not know+ E.

The crux of the matter is that such an argument still has the potential to pose a worrisome sceptical challenge, since while its conclusion is consistent with our having widespread knowledge of E-type propositions, it nevertheless has the capacity to generate a form of scepticism which deprives us of an epistemic standing that is of special value to us.

In order to give the Moorean approach the best run for its money, let us grant for the sake of argument the most robust form of Mooreanism available—i.e., the view that our knowledge, even of the denials of sceptical hypotheses, is typically supported by factive grounds. With this in mind, let us examine first the plausibility of (S1+).

Given the form of Mooreanism in play, there are certainly good grounds for supposing that the agent in question is able to know that she is not the victim of a sceptical hypothesis. In normal circumstances, there can be no question that the truth of the target belief—i.e., that one is not the victim of a sceptical hypothesis—is not a matter of luck. After all, in such circumstances a belief of this sort could not very easily be false, since an awful lot would need to change about the world in order to ensure the falsity of this belief. Accordingly, it is plausible to suppose that whatever anti-luck condition one imposed on one's theory of knowledge would be met.¹⁷ Moreover, if the belief is supported by factive grounds, then we ought to be able to allow that this belief would satisfy other constraints we might want to lay down on knowledge possession. The belief would surely be adequately supported by evidence, for example. There is thus a strong *prima facie* case for agreeing with the robust Moorean that one can know the denials of sceptical hypotheses.

Nevertheless, since we have already seen that one can have knowledge without thereby exhibiting a cognitive achievement—and that it is cognitive achievements which are distinctively valuable—that the conditions for knowledge are satisfied does not ensure that the kind of knowledge in question is distinctively valuable. In short, it does not ensure that one has knowledge+, which is the kind of knowledge specifically at issue in (S1+).

Moreover, there is every reason for supposing that the kind of knowledge in play here is not knowledge+. For recall, in order to exhibit a cognitive achievement the truth of one's belief needs to be because of the exercise of one's relevant cognitive abilities. It is, however, wholly unintuitive to suppose that the truth of an agent's belief that she is not, say, a BIV is *because of* her cognitive ability (even if we are willing to grant that her cognitive ability plays some substantive role in this cognitive success). Indeed, if anything, it seems that the cognitive success in question here is more creditable to the good fortune of being in an epistemically friendly environment than it is to the agent's cognitive abilities.

That we lack the relevant discriminatory abilities is key to understanding why this is so. If one is unable to discriminate between ordinary scenarios and sceptical scenarios, and one is in addition unable to rationally adduce independent epistemic support in favour of the target belief, then it can hardly be the case that one's cognitive success when one believes truly that one is not the victim of a sceptical hypothesis is because of one's own cognitive ability.

This claim requires some development, and part of the reason for this is that one might think that whether or not this is so depends on the kind of Moorean proposal in question. That is, one might argue that provided one's Mooreanism is of the robust type that we are interested in, such that it allows that one's beliefs in the denials of sceptical hypotheses could be supported by factive grounds, then there is no reason to deny that one's cognitive success in such a case is because of one's cognitive ability.

In order to explore this issue it will be helpful to return to the example given in §2 which was meant to illustrate that one can have better evidence for believing that one scenario has obtained rather than an alternative scenario, and thereby for knowing that the one scenario has obtained rather than the alternative scenario, even though one cannot discriminate between the two scenarios.¹⁸ The example I offered concerned being able to come to know, in normal circumstances, that one is presently looking at one's children rather than extremely life-like robots. The thought was that given the excellent background evidence one possesses which indicates that the error-possibility in question is so remote as to be discountable, it follows that one is in a position to rationally dismiss this error-possibility and thereby have adequate grounds to support one's knowledge that one is presently looking at one's children rather than extremely life-like robots. Crucially, however, it was pointed out that knowledge of this sort can be possessed—and, indeed, will usually be

possessed—even though one is unable to discriminate between one's children and extremely life-like robots.

The relevance of this example for our present discussion is that an agent who knows, on this evidential basis, that she is not presently looking at extremely life-like robots which are masquerading as her children surely *does* exhibit a cognitive achievement. While it is undeniable that her cognitive success is not because of the operation of her relevant discriminative abilities, that doesn't undermine the fact that the truth of her belief (that she is not presently looking at extremely life-like robots) does seem to be because of her cognitive ability more generally. After all, she has weighed-up the evidence available to her in order to form an informed judgement on the matter. If this is right, however, then one might naturally wonder why a Moorean response to the sceptical problem which allowed our knowledge of the denials of sceptical hypotheses to be supported by adequate evidence could not similarly argue that such knowledge also involves a cognitive achievement.

There is, however, a crucial disanalogy between the two cases, and this is that the agent in the 'robot' case, while lacking the relevant discriminatory capacities, is nevertheless able to rationally adduce independent grounds in favour of her belief. Such independent grounds are not, however, rationally available to the agent in the sceptical case, even assuming a robust form of Mooreanism.

In the robot case, even though the agent cannot discriminate between the two alternative scenarios she is nonetheless able to rationally enhance the epistemic support for her belief that the target error possibility does not obtain by reflecting on relevant background information that is available to her and which is not itself called into question by that error possibility (and which is thus in this sense *independent* of that error possibility). Indeed, were our agent to be unable to rationally adduce independent grounds in favour of her belief, then we would not regard her belief as being true because of her cognitive abilities at all, and hence we would not regard it as a cognitive achievement. After all, such an agent would be in a position of being aware of the target error possibility and believing that it does not obtain, and yet regarding herself as lacking any independent rational basis for holding this belief. In effect, she would simply be groundlessly trusting that she is in the kind of epistemically friendly environment that she takes herself to be in, but if that's right then the truth of her belief cannot be because of her cognitive abilities.

Crucially, in the sceptical case the agent lacks both the discriminatory capacities *and* is unable to undertake the kind of rational process that is possible in the robot case. We have already noted that the discriminatory capacities are lacking, so the crux of the matter is the unavailability of the relevant rational process. The reason why this is unavailable is that sceptical error possibilities by their nature call into question the evidential support one possesses for one's beliefs *en masse*. Accordingly, there is no rational route by which the agent can adduce independent grounds in favour of her belief that she is not the victim of a sceptical hypothesis, since whatever background evidence she might appeal to will be itself problematized by the sceptical hypothesis.

From a rational perspective, then, all one can do is simply reflect that so long as one is indeed in the normal circumstances that one takes oneself to be in, then one has excellent evidence in support of one's belief that one is not a victim of a sceptical hypothesis. But that is just to grant that whether or not one's belief in this regard is true is to a substantive degree down to the good fortune of being in an epistemically friendly environment rather than being due to one's own cognitive ability. In effect, one's epistemic position is akin to that of the agent in the robot case who is aware of a certain error possibility but who is unable to rationally form a judgment as to why she is entitled to dismiss it. Just as this agent's true belief that she is not a victim of this error possibility does not constitute a cognitive achievement, neither does one's true belief that one is not the victim of a sceptical hypothesis.

So in order to exhibit a cognitive achievement, one must either appropriately exercise one's discriminative capacities or else be in a position to rationally adduce independent grounds in favour of what one's believes in order to compensate for the lack of such discriminatory capacities. Since one is unable to meet either condition when it comes to one's belief that one is not the victim of a sceptical hypothesis, it follows that such a belief, when true, does not constitute a cognitive achievement. So even if the Moorean is right that one can know such propositions, it does not follow that one can have knowledge+ of them, and hence (S1+) is on strong ground.

8. Still, one might argue that even though one's knowledge that one is not the victim of a sceptical hypothesis does not involve a cognitive achievement, it doesn't follow that one's

knowledge of E-type propositions also doesn't involve a cognitive achievement. Couldn't it be the case that the cognitive success involved in the latter case does constitute a cognitive achievement, even though the cognitive success in the former case doesn't? In effect, the point in play here is that it is (S2+) that is the contentious premise in the reformulated sceptical argument.

One could regard (S2+) as being motivated by a sister principle to the closure principle that motivated (S2) which explicitly focuses on knowledge+:

The Closure Principle for Knowledge+

If S knows+ that p , and S competently deduces q from p (thereby coming to believe q while retaining her knowledge+ that p), then S knows+ that q .

With this principle in play, it would not be hard to make a strong case for (S2+). Moreover, there does seem to be a great deal of intuitive support for such a principle. After all, if the truth of one's belief in a (known) proposition is because of one's cognitive ability, and so constitutes a cognitive achievement, then it is hard to see why any further knowledge which one competently deduces from that knowledge should involve a belief the truth of which is related to one's cognitive ability to any lesser degree. Indeed, if anything, shouldn't the presence of the competent deduction actually *enhance* the extent to which the truth of one's belief is creditable to one's own cognitive abilities?

One might think that one could adapt the robot case to generate a counterexample to this principle. Suppose, for example, that at time t_1 our hero knows+ that the persons before her are her children. Suppose now, however, that, at t_2 , she becomes aware of the error possibility that she might instead be looking at extremely life-like robots and that she lacks the background information that would enable her to rationally adduce independent grounds in favour of her belief in the denial of this error possibility. Nevertheless, she does recognise that this error possibility is inconsistent with something that she believes (indeed knows+), and so she competently deduces on this basis that this error possibility is false. Given what we have said earlier, while she might come to know the denial of the error possibility on this basis, she cannot come to know+ this proposition because her true belief would not constitute a cognitive achievement. Given that by hypothesis she does know+ the antecedent proposition, however, then it seems we have a counterexample to the closure principle for knowledge+.

On closer inspection, however, such a putative counterexample fails to convince. After all, if we really are to suppose that the agent is unable to rationally adduce independent grounds in support of her belief in the denial of the error possibility, then why would we continue grant that she has knowledge⁺ of the antecedent proposition? Indeed, wouldn't the fact that someone was struck dumb by the presentation of an error possibility of this sort indicate that their knowledge of the antecedent proposition (insofar as we continue to grant that they do have knowledge of this proposition) is far more 'brute' than it first appeared? (It's not as if such independent grounds are hard to come by, after all).

With this point in mind, there is a lot of *prima facie* support for the knowledge⁺ closure principle. The problem, of course, is that with a closure principle of this sort in play, then (S2⁺) immediately follows. Moreover, since we have already noted that there is a strong case in favour of (S1⁺), then the sceptical conclusion, (SC⁺), also follows.

9. No doubt there will be moves available to the anti-sceptic in this respect; the point of this paper has not been to supply a definitive basis for a new form of scepticism but rather to demonstrate that viewing the problem of scepticism through the lens of epistemic value highlights a very different sceptical challenge that needs responding to—*viz.*, that while we might well have much of the knowledge that we typically suppose ourselves to have, the type of knowledge that we possess may not be of the specially valuable variety. Perhaps this is a conclusion that we can ultimately grant to the sceptic; granting it certainly won't be as intellectually disastrous as allowing that widespread knowledge is impossible. Nevertheless, it is clearly a conclusion that is, to say the least, uncomfortable. By definition, if the sceptic deprives us of something that is of special value to us, then that is something that should concern us.¹⁹

REFERENCES

- Chisholm, R. (1977). *Theory of Knowledge* (2nd Ed.), Prentice-Hall, Englewood Cliffs, New Jersey.
- Cohen, S. (2000). 'Contextualism and Skepticism', *Philosophical Issues* 10, 94-107.
- DeRose, K. (1995). 'Solving the Skeptical Problem', *Philosophical Review* 104, 1-52.
- Dretske, F. (1970). 'Epistemic Operators', *The Journal of Philosophy* 67, 1007-23.

- (2005a). 'The Case Against Closure', *Contemporary Debates in Epistemology*, (eds.) E. Sosa & M. Steup, 13-26, Blackwell, Oxford.
- (2005b). 'Reply to Hawthorne', *Contemporary Debates in Epistemology*, (eds.) E. Sosa & M. Steup, 43-6, Blackwell, Oxford.
- Greco, J. (1999). 'Agent Reliabilism', *Philosophical Perspectives* 13, 273-96.
- (2000). *Putting Skeptics in Their Place: The Nature of Skeptical Arguments and Their Role in Philosophical Inquiry*, Cambridge University Press, Cambridge.
- (2002). 'Knowledge as Credit for True Belief', *Intellectual Virtue: Perspectives from Ethics and Epistemology*, (eds.) M. DePaul & L. Zagzebski, Oxford University Press, Oxford.
- (2007). 'The Nature of Ability and the Purpose of Knowledge', *typescript*.
- (Forthcominga). 'The Value Problem', *The Value of Knowledge*, (eds.) A. Haddock, A. Millar, & D. H. Pritchard, Oxford University Press, Oxford.
- (Forthcomingb). 'What's Wrong With Contextualism?', *Philosophical Quarterly*.
- Hawthorne, J. (2004). *Knowledge and Lotteries*, Clarendon Press, Oxford.
- (2005). 'The Case for Closure', *Contemporary Debates in Epistemology*, (eds.) M. Steup & E. Sosa, 26-43, Blackwell, Oxford.
- Kvanvig, J. (2003). *The Value of Knowledge and the Pursuit of Understanding*, Cambridge University Press, Cambridge.
- Lackey, J. (2007). 'Why We Don't Deserve Credit for Everything We Know', *Synthese* 156.
- Lewis, D. (1996). 'Elusive Knowledge', *Australasian Journal of Philosophy* 74, 549-67.
- McDowell, J. (1995). 'Knowledge and the Internal', *Philosophy and Phenomenological Research*, 55, 877-93.
- Nozick, R. (1981). *Philosophical Explanations*, Oxford University Press, Oxford.
- Pritchard, D. H. (2002a). 'Recent Work on Radical Skepticism', *American Philosophical Quarterly* 39, 215-57.
- (2002b). 'Resurrecting the Moorean Response to the Sceptic', *International Journal of Philosophical Studies* 10, 283-307.
- (2005a). *Epistemic Luck*, Oxford University Press, Oxford.
- (2005b). 'The Structure of Sceptical Arguments', *Philosophical Quarterly* 55, 37-52.
- (2007a). 'Anti-Luck Epistemology', *Synthese* 158, 277-97.
- (2007b). 'Contrastivism, Evidence and Scepticism', *Social Epistemology* 21.
- (2007c). 'Knowledge, the Value of', *Stanford Encyclopaedia of Philosophy*, (ed.) E. Zalta, <http://plato.stanford.edu/entries/knowledge-value/>.
- (2007d). 'Recent Work on Epistemic Value', *American Philosophical Quarterly*, 44, 85-110.
- (2007e). 'Virtue Epistemology and Epistemic Luck, Revisited', *Metaphilosophy* 38.
- (Forthcoming). 'Greco on Knowledge: Virtues, Contexts, Achievements', *Philosophical Quarterly*.
- (2008a). 'Anti-Luck Virtue Epistemology', *typescript*.
- (2008b). 'Relevant Alternatives, Perceptual Knowledge and Discrimination', *typescript*.
- (2008c). 'McDowellian Neo-Mooreanism', *Disjunctivism: Perception, Action, Knowledge*, (eds.) A. Haddock & F. Macpherson, Oxford University Press, Oxford.
- (2008d). 'The Value of Knowledge', *typescript*.
- Schaffer, J. (2005). 'Contrastive Knowledge', *Oxford Studies in Epistemology*, (eds.) T.

- Gendler & J. Hawthorne, Oxford: Oxford University Press.
- Sosa, E. (1988). 'Beyond Skepticism, to the Best of our Knowledge', *Mind* 97, 153-89.
- (1991). *Knowledge in Perspective: Selected Essays in Epistemology*, Cambridge University Press, Cambridge.
- (1999). 'How to Defeat Opposition to Moore', *Philosophical Perspectives* 13, 141-54.
- (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, Oxford University Press, Oxford.
- Stanley, J. (2005). *Knowledge and Practical Interests*, Clarendon Press, Oxford.
- Williamson, T. (2000). *Knowledge and its Limits*, Oxford University Press, Oxford.
- Wright, C. J. G. (2004). 'Warrant for Nothing (and Foundations for Free)?', *Proceedings of the Aristotelian Society* 78 (supp. vol.), 167-212.
- Zagzebski, L. (1996). *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*, Cambridge University Press, Cambridge.
- (1999). 'What is Knowledge?', *The Blackwell Guide to Epistemology*, (eds.) J. Greco & E. Sosa, 92-116, Blackwell, Oxford.

NOTES

¹ For further discussion of this template radical sceptical argument, and its role in contemporary epistemology, see Pritchard (2002a). For a critical discussion of whether the sceptical argument is best expressed in this fashion, see Pritchard (2005b).

² This formulation of the closure principle is essentially that offered by Williamson (2000, 117) and Hawthorne (2005, 29).

³ For a recent discussion of the status of the closure principle, see the exchange between Dretske (2005a; 2005b) and Hawthorne (2005).

⁴ One complication in this regard is the contrastivist approach as defended, for example, by Schaffer (2005). According to this proposal, there is no such thing as knowing a proposition *simpliciter*; rather, all knowledge is to be understood as relative to a specific set of contrasts (i.e., not-*p* propositions), and this has important ramifications for how one should understand the sceptical argument. I have not the space to explore the contrastivist response to radical scepticism here. For further discussion, see Pritchard (2007b).

⁵ Even though, depending on the details of the specific rendering of the view in question, the position may not in practice resemble G. E. Moore's own response to the sceptic very much at all.

⁶ See, for example, Wright (e.g., 2004) for an anti-sceptical proposal which is cast along these lines.

⁷ This would be one way of reading the brand of anti-scepticism propounded by McDowell (e.g., 1995), for example, although it is debatable whether McDowell himself would ever put the point in quite these terms. See Pritchard (2008c) for further discussion of this reading. See also Williamson (2000, ch. 8).

⁸ The kind of contextualism that I have in mind here is attributer contextualism, of the sort defended by, for example, DeRose (1995), Lewis (1996) and Cohen (2000). Other types of view could also be understood as aiming to accommodate sceptical intuitions along broadly contextualist lines, however, such as a subject contextualism of the sort defended (under a very different name) by Hawthorne (2004) and Stanley (2005).

⁹ I expand on this point, and its epistemological ramifications, in Pritchard (2008b).

¹⁰ I develop this line of argument in more detail in Pritchard (2008d, §1). For more on the general issues surrounding the value of knowledge, see Kvanvig (2003) and Pritchard (2007c; 2007d).

¹¹ Greco, like other virtue epistemologists who have pursued this general line, actually argues that knowledge is *intrinsically* valuable rather than finally valuable (e.g., Greco *forthcominga*, §4). It is clear from what he says about the value of knowledge, however, that it is specifically final value that he has in mind, and so this is how I have expressed his view here.

¹² See also Sosa (1988; 1991; 2007) and Zagzebski (e.g., 1996; 1999) for virtue-theoretic accounts of knowledge that are cast along roughly these lines.

¹³ This example is due to Chisholm (1977, 105).

¹⁴ A second way of defending the final value thesis against this line of objection is to opt for a kind of holism about value. Just as one might hold that pain is the sort of thing which is generally bad but sometimes good, so one might argue that knowledge is the kind of thing which is generally finally valuable but sometimes—e.g., where the cognitive success in question is trivial etc.,—not finally valuable. The reason why such an approach might work is that all the virtue-theoretic account of the value of knowledge is aiming to account for is why knowledge is distinctively valuable in the way that we think it is, and demonstrating that it is *generally* finally valuable would probably suffice in this regard (at least so long as there is no lesser epistemic standing which is also generally finally valuable). I am grateful to Mike Ridge for this suggestion. For further discussion of what is required of an adequate account of the distinctive value of knowledge, see Pritchard (2007c; 2007d; 2008d).

¹⁵ I describe some of the moves that Greco makes in the light of this line of objection and evaluate their merits in Pritchard (2007e; 2008d; *forthcoming*).

¹⁶ In particular, I argue elsewhere that considerations like this point towards what I term an ‘anti-luck virtue epistemology’. I expand on what such a view involves in Pritchard (2008a; 2008d). Interestingly, although this view is motivated in a very different way, such a proposal is very close in spirit to the sort of view that Greco defended in earlier work, and abandoned in favour of the more robust virtue-theoretic proposal discussed here. See, for example, Greco (1999; 2000).

¹⁷ In particular, the belief would be ‘safe’. For more on safety and the extent to which it captures the ‘anti-luck’ requirement on knowledge, see Sosa (1999) and Pritchard (2002b; 2005a; 2007a).

¹⁸ Recall that for the purposes of this example we were setting sceptical worries to one side.

¹⁹ I am grateful to Adam Carter and John Turri for detailed comments on an earlier version of this paper. Thanks also to Jennifer Lackey and Mike Ridge.