THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Virtue Epistemology and Epistemic Twin Earth

Citation for published version:

Kallestrup, J & Pritchard, D 2014, 'Virtue Epistemology and Epistemic Twin Earth', *European Journal of Philosophy*, vol. 22, no. 3, pp. 335-357. <https://doi.org/10.1111/j.1468-0378.2011.00495.x>

Digital Object Identifier (DOI):

[10.1111/j.1468-0378.2011.00495.x](https://doi.org/10.1111/j.1468-0378.2011.00495.x)

Link:

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

Document Version:

Peer reviewed version

Published In:

European Journal of Philosophy

Publisher Rights Statement:

This is the post-peer reviewed version of the following article: Kallestrup, J. & Pritchard, D. (2011) "Virtue Epistemology and Epistemic Twin Earth", *European Journal of Philosophy*; which has been published in final form at <http://dx.doi.org/10.1111/j.1468-0378.2011.00495.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.

Forthcoming in European Journal of Philosophy.

VIRTUE EPISTEMOLOGY AND EPISTEMIC TWIN EARTH

Jesper Kallestrup & Duncan Pritchard

ABSTRACT. A popular form of virtue epistemology—defended by such figures as Ernest Sosa, Linda Zagzebski and John Greco—holds that knowledge can be exclusively understood in virtue-theoretic terms. In particular, it holds that there isn't any need for an additional epistemic condition to deal with the problem posed by knowledge-undermining epistemic luck. It is argued that the sustainability of such a proposal is called into question by the possibility of *epistemic twin earth* cases. In particular, it is argued that such cases demonstrate the need for virtue-theoretic accounts of knowledge to appeal to an independent epistemic condition which excludes knowledge-undermining epistemic luck.

1. VIRTUE EPISTEMOLOGY, MODEST AND ROBUST

One of the most influential proposals in epistemology in recent years has been *virtue epistemology*. Although there is a wide range of views which could plausibly be regarded as falling under this general heading, what they have in common is the idea that epistemology needs to make central appeal to the notion of cognitive abilities (i.e., epistemic virtues).¹ One of the central questions of traditional epistemology is, of course, the question of the nature of knowledge. While some virtue epistemologists have argued for a reorientation of traditional epistemology, such that this question is no longer central to the epistemological project, most virtue epistemologists do offer a virtue-theoretic response to this issue.²

Now one might think that any virtue-theoretic theory of knowledge must be inevitably *modest* in ambition on account of the problem of knowledge-undermining epistemic luck. After all, one can set-up Gettier-style cases in which the subject concerned is

stipulated as having appropriately exercised a reliable belief-forming process in forming the target belief, and yet is found to lack knowledge nonetheless on account of the knowledge-undermining epistemic luck present. In particular, we can construct Gettier-style cases in which even despite the fact that the subject's true belief is the product of the appropriate exercise of her relevant cognitive abilities, it is nonetheless the case that this cognitive success is lucky—and hence not an instance of knowledge—in that the subject could very easily have been mistaken.

Consider, for example, Roderick Chisholm's (1977: 105) famous 'sheep' example. In this case the subject concerned forms a belief that there is a sheep in the field before him by employing his reliable perceptual abilities, and thus his cognitive abilities. Moreover, the belief he forms is true. Crucially, however, this belief does not amount to knowledge because what the subject is in fact looking at is not a sheep at all, but rather a sheep-shaped object which is obscuring from view the genuine sheep behind. The reason why this deprives our subject of knowledge is that the belief so formed is only luckily true. That is, given how this belief was formed it could so very easily have been false (e.g., if the genuine sheep that was hidden from view had wandered off into another field).³

The moral of such cases appears to be that true belief that is the product of cognitive ability does not suffice for knowledge. Accordingly, it seems that even the most ambitious virtue-theoretic account of knowledge will need to incorporate a non-virtue-theoretic epistemic condition which can deal with the problem of knowledge-undermining epistemic luck, such as the condition that knowledge requires safe belief (i.e., a true belief that could not have easily been false), or that the basis on which the belief is presently formed be safe (i.e., that the proposition in question would be believed on the present basis only if that proposition is true).⁴ That, in any case, has been precisely the line taken by certain virtue epistemologists.⁵ Call such a proposal *modest virtue epistemology*.

Interestingly, however, some of the main virtue epistemologists who have tackled this issue have argued that a virtue-theoretic account of knowledge, at least when properly formulated, is perfectly able to deal with the problem of knowledge-undermining epistemic luck. Hence they have claimed that there is no need for a virtue-theoretic account of knowledge to appeal to an independent epistemic condition in order to deal with the problem of knowledge-undermining epistemic luck. One finds such a view defended by such key figures within contemporary virtue epistemology as Ernest Sosa (1988; 1991; 2007;

2009), Linda Zagzebski (1996; 1999) and John Greco (2003; 2007; 2008; 2009*a*; 2009*b*; 2009*c*). Call such a view *robust virtue epistemology*.⁶

Although there are differences between the proposals in play here (we will consider some of these differences presently), the basic move that is made by these virtue epistemologists is to argue that what is required for knowledge is not merely the conjunction of cognitive success and the exercise of relevant cognitive ability. Instead, what is required is cognitive success that is *because of* cognitive ability. More precisely, the claim here is not merely that knowledge comprises true belief and cognitive ability. That is consistent with the belief not being a product of exercising the ability. Nor is it the slightly stronger claim that knowledge is belief that is true and because of a cognitive ability. For the belief could result from exercising the ability without owing its truth to that exercise. The intended, even stronger, claim is that knowledge is belief that is true because of a cognitive ability. That is to say, the belief is true in virtue of resulting from the exercise of such ability. (Or, to use Sosa's terminology, knowledge is *apt* belief, where this is belief which is *accurate* (i.e., true) because *adroitly* (i.e., skilfully) formed).⁷

In order to see the attraction of such a proposal, consider again Chisholm's sheep case. While we have here not just the conjunction of cognitive ability and cognitive success, but also a cognitive state produced by exercising that ability, we don't intuitively have cognitive success that is because of the cognitive ability. Instead, the cognitive success seems to be attributable to the happenstance that there is a sheep hidden from the subject's view behind the sheep-shaped object that he is looking at. More generally, at least for standard Gettier-style cases anyway (the reason for the qualification will become apparent below), there seems every reason for supposing that adding the 'because of' relation to the virtue-theoretic account of knowledge obviates the need for introducing a separate epistemic condition in order to deal with the problem of knowledge-undermining luck. And that means that the path is cleared for robust virtue epistemology.

Of course, the devil is, as always, in the detail. In particular, we need to be told how to read the 'because of' relation in play here. As John Turri (2011) notes, some readings are unsuitable. For instance, Wayne Rooney scores amazing goals for Manchester United because of his exceptional football skills, but he is also paid £250,000 a week because of those skills. Only in the former case does Rooney exercise his football skills. In the latter case, he might be said instead to exercise his negotiation skills—e.g., by threatening to move

to Manchester City. We can say that Rooney scores goals *through* (or out of) his football skills but he is not paid that amount of money through those skills.

There are two main proposals in the literature as to how the 'because of' or 'through' relation is best understood in the context of robust virtue epistemology.⁸ The first, due to Greco (e.g., 2009a: ch. 5), understands it in causal explanatory terms. In particular, on this view when a subject has knowledge the overarching element in the causal explanation of her cognitive success will be her cognitive ability.⁹ To illustrate, Rooney's football skills feature centrally in a causal explanation of (most of) the goals he scores but only marginally or at least indirectly in a causal explanation of the astronomical salary he receives. Call this the *causal explanatory* version of robust virtue epistemology.

The second, due to Sosa (e.g., 2007: ch. 5), treats this relation on the model of disposition manifestation, such that knowledge involves the manifestation of a cognitive (i.e., truth-conducive) disposition on the part of the subject. To illustrate, the goals Rooney scores manifest his football abilities but his salary constitutes no manifestation of those abilities. Call this the *disposition manifestation* version of robust virtue epistemology.

To see how these two accounts can come apart, consider a glass that was broken as a result of someone dropping it on a wooden floor. Ordinarily, the most salient part of the causal explanation of why the glass broke will be that someone dropped it on the floor, and in this sense it will be true to say that the glass broke because it was dropped in this way. Note, however, that this is consistent with the claim that it was because of the glass's fragility that it broke, since here we are talking about the manifestation of a disposition and not offering a causal explanation. If fragility is the (second-order) property of having a property that causes breaking if dropped, then we cannot causally explain why the glass broke when dropped in terms of it being fragile. We can say the glass broke when dropped because the glass has a (first-order) molecular bonding property. What is causally responsible for the shattering is this micro-structural property together with the dropping. The dispositional property itself, thus understood, is causally inefficacious of the effects in terms of which it is defined.¹⁰

Both renderings of robust virtue epistemology face difficulties. In particular, one line of critique that has been applied against both versions of robust virtue epistemology is that they are unable to properly accommodate all cases of knowledge-undermining epistemic luck. In particular, it has been argued that once we move away from standard Gettier-style

cases of knowledge-undermining luck we find examples where the subject lacks knowledge on account of the knowledge-undermining luck involved, and yet it remains true that the subject's cognitive success is appropriately because of her cognitive ability. Hence it is argued that, at best, only a modest version of virtue epistemology is available, one that appeals to a further epistemic condition which deals with the problem of knowledge-undermining luck (such as a safety condition).¹¹

We will also be arguing for this conclusion, although via a completely new line of argument, one that makes appeal to what we call *epistemic twin earth* cases. The novelty of our argument is that such cases demonstrate the existence of epistemic differences between virtue-theoretic duplicates. As we shall show, on the plausible assumption that knowledge excludes a certain type of environmental luck, the agent on earth will have knowledge that her twin on twin earth lacks. Yet there is no scope for either stripe of robust virtue epistemology to consistently maintain that while the agent's belief is true because of ability, the twin's belief is not true because of ability. Both the agent and her twin possess the pertinent ability, both exercise that ability under suitable conditions, and their cognitive success are in both cases because of that very ability. Still, if knowledge is incompatible with the type of luck in play, then knowledge cannot possibly be ascribed to both of them.

2. VIRTUE EPISTEMOLOGY AND EPISTEMIC TWIN EARTH

Standard twin earth arguments run as follows.¹² Despite appearances there is no water on twin earth. Water is essentially H_2O , and all the watery stuff on twin earth has the different microstructure XYZ—earthlings call that 'twin-water'. When *S* on earth utters 'water is wet', she expresses the proposition that water is wet, but when *S*'s intrinsic physical duplicate on twin earth utters the same sentence, twin-*S* expresses the proposition that twin-water is wet. Since *S* and twin-*S* refer to different kinds of stuff when they token 'water' the truth-conditions of their respective utterances differ. Assuming the contents of their beliefs are fixed by the truth-conditional contents of the sentences that they use to express those beliefs, then these contents also fail to supervene on their intrinsic physical properties. Indeed if belief states are individuated in part by their contents, then what belief states *S* and twin-*S* are in fail to supervene on their intrinsic physical properties. Instead these states

depend partially for their individuation on which patterns of causal relations S and twin- S bear to their respective physical environments.

Consider now *epistemic twin earth* on which most watery stuff is H_2O . In between there is some scattered twin-water the exact location of which varies from case to case. Our contention is that an epistemic twin earth argument shows that whether a subject is in a perceptual knowledge state cannot merely be a question of getting things right through exercising her cognitive abilities in the way that robust virtue epistemology suggests.

Let's divide epistemic twin earth into three regions. The subject's *local environment* is where the subject is currently located. It contains the objects and properties that are the proximate causes of her current perceptual experiences. Take facts to be objects instantiating a property at a time. If the subject now perceives that p , then the fact that p (the ' p -fact') is one that concerns her local environment—i.e., it is a *local fact*. Other local features have to do with aspects of the perceptual process and various background conditions on perception—e.g., distorting noise, brightness, and so on.

The subject's *regional environment* is neither where the subject is currently located, nor where she typically forms any beliefs. Still, it contains the objects and properties with which she might easily have been causally connected. If the q -fact is such that if the subject had not now perceived that p then she would have perceived that q , then the q -fact is one that pertains to her regional environment—i.e., it is a *regional fact*. Regional facts, thus understood, are nearby perceptual possibilities, but they play no causal role in producing the subject's current perceptual experience on which she bases her belief that p .

Finally, the subject's *global environment* is where she is normally located although not at present. It contains the objects and properties with which she ordinarily causally interacts. The *global facts* thus comprise all the facts that extend in space-time beyond the regional facts. Assuming the subject now perceives the local fact that p , the fact that r is a global fact only if she would not have perceived that r had she not perceived that p . Given the subject's current location, global facts are not only distant perceptual possibilities, they are also causally inefficacious in producing her current perceptual experiences.

We can now mount an epistemic twin earth argument to the effect that robust virtue epistemology is an inadequate account of knowledge. The subject, S , is on earth where all watery stuff is H_2O . S 's perceptual apparatus is highly reliable in that a high frequency of S 's perceptual beliefs is both actually true and true across relevantly close worlds. Based on a

perceptual experience as of water, S forms the demonstrative belief that that's water. There is no question that S thereby comes to know just that.

On epistemic twin earth S has an intrinsic physical duplicate called 'twin- S '. S and twin- S are conceptually competent but chemically ignorant. On epistemic twin earth all watery stuff in twin- S 's global environment is H_2O . Not only is twin- S therefore able to entertain water-thoughts, a high frequency of twin- S 's water-beliefs as formed in her global environment is true both in actual fact and across relevantly close worlds. Twin- S 's perceptual apparatus is thus equally reliable. Moreover, all watery stuff in twin- S 's local environment is H_2O . When twin- S forms the demonstrative belief that that's water on the basis of a perceptual experience as of water, her belief is true. Yet, unbeknownst to twin- S , twin-water is abundant in her regional environment. The basis on which twin- S holds that belief is thus such that her belief is only luckily true, in that given the basis for her belief it could very easily have been the case that she would have formed a false belief (e.g., had she been interacting, unbeknownst to her, with twin-water). That is to say, very easily could twin- S have believed that that's water on the same basis—a perceptual experience as of water—without that being so.¹³ On the plausible assumption that knowledge excludes such environmental luck, it follows that twin- S lacks knowledge.

We can illustrate what is going on here with the following diagram:

The explanation robust virtue epistemology offers of why S has knowledge on earth is that her cognitive success is because of her cognitive ability. The challenge, however, is to explain why twin- S lacks knowledge on epistemic twin earth. The fact that S and twin- S are intrinsic

physical duplicates embedded in physically identical global environments means that one cannot possess a cognitive ability that the other lacks.

To use an analogy robust virtue epistemologists are fond of, suppose *S* is an expert archer. *S* possesses that ability in virtue of relevant bodily/psychological features and mostly occupying an environment that is conducive for her to frequently hit the innermost rings when dispatching arrows. Given that the latter are equally true of twin-*S*, she will be an expert archer too. And the fact that both *S* and twin-*S* currently occupy physically identical local environments means that their cognitive successes must arise in the very same way. To use the analogy, the ways in which *S* and twin-*S* propel their respective arrows into the yellow ring are identical. After all, fletching, bow strings, body positions, prevailing winds, distances to target, energy imparted to arrows, and so on, are identical in the two cases. Combining these two facts spells trouble for robust virtue epistemology, for it deprives proponents of this view of a principled basis on which they can treat the two cases differently. And yet there clearly is an epistemic difference between them, in that twin-*S*, unlike her counterpart *S*, lacks knowledge.

If the epistemic twin earth argument goes through then robust virtue epistemology will be an inadequate account of knowledge. In particular, it will be necessary for virtue epistemologists to appeal to a separate epistemic condition which is devoted to excluding the problem of knowledge-undermining epistemic luck. One such condition, we suggested, is that knowledge requires safe belief or a safe basis for belief. In short, virtue epistemologists will have to retreat from robust virtue epistemology to modest virtue epistemology.

Note finally that the factivity of knowledge plays no role in the epistemic twin earth argument. Our claim is not that while *S* knows that *p* on earth, twin-*S*'s belief that *p* as produced by her cognitive ability fails to qualify for knowledge because that belief is false on epistemic twin earth. Everybody agrees that knowledge states are externally individuated in that trivial sense. Our claim is rather that no epistemic virtue-theoretic condition can *fully* individuate knowledge states even on the assumption that the corresponding belief states are true.¹⁴

3. RESPONSE I: NO ABILITY

We will explore the efficacy of our epistemic twin earth objection to robust virtue epistemology by considering how a proponent of this view might respond. In particular, we will look at four types of response that might be made on behalf of robust virtue epistemology, paying careful attention throughout to the two different formulations of the proposal that we noted above.

A first line of response that the proponent of robust virtue epistemology might mount could be to question whether the belief formed by twin-*S* really is the product of a relevant cognitive ability on the ground that twin-*S* simply fails to possess such ability.¹⁵

Consider first how this line of response might work on the disposition manifestation rendering of robust virtue epistemology. Here there seems very little room for manoeuvre. For notice that Sosa (2007: 29; 2009: 135) explicitly conceives of cognitive abilities in terms of cognitive dispositions which have a physical basis resident in whoever has those dispositions. If that's right, however, then it is difficult to see why twin-*S* should lack a cognitive ability that *S* possesses given that they are physically identical. For whatever physical basis is sufficient for *S* to possess her cognitive ability is a basis shared by twin-*S*. Of course, which cognitive abilities *S* possesses depend on the environment in which she is typically embedded in terms of its operative laws or law-like regularities and physical background conditions. But there are no nomological differences between earth and epistemic twin earth, and *S* and twin-*S* share physically identical global environments.

The analogy with other physical dispositions, such as solubility, is instructive here. After all, these dispositions are such that the instantiation by an object of the physical base property for the disposition physically necessitates the instantiation of the dispositional property. So, taking the case of solubility as an example, as long as the laws of physics are fixed, any intrinsic physical duplicate of a solute is also soluble. To find such a duplicate that is not soluble you must go to a world with deviant laws of physics. It is hard to see why cognitive dispositions should be so different as to come and go with hidden variations in particular physical facts in the regional environment.

On the face of it, the causal explanatory rendering of robust virtue epistemology might have more scope to mount a reply on this score, given that it is not wedded to the kind of metaphysical picture inherent to the disposition manifestation rendering of the view.

Indeed, Greco's particular version of robust virtue epistemology seems to have the resources to at least in principle maintain that we should deliver a different epistemic assessment with regard to the two counterpart subjects at issue in the epistemic twin earth case. The reason for this is that he argues for an account of abilities which relativizes them to *environments* and *conditions*, and then allows purely *pragmatic considerations* to fix which environments and conditions count as relevant, hence to have a bearing on whether a subject counts as possessing a certain ability.

Here is his account of abilities, where 'R' stands for the result which the ability is an ability to achieve and 'C' is a range of conditions relevant to the manifestation of that ability:

"*S* has an ability A(R/C) relative to environment E = Across the set of relevant close worlds W where *S* is in C and in E, *S* has a high rate of success in achieving R." (Greco 2009a: 77; cf. Greco 2007)

So, to employ one of Greco's own examples—see, e.g., Greco (2009a: ch. 5)—when saying that a baseball player has an ability to hit fastballs we are claiming that relative to a set of relevant conditions (e.g., when he doesn't have sand in his eyes), and in suitable environments (e.g., in normal baseball-playing environments), he can reliably manifest that ability (i.e., he has a high rate of success at hitting fastballs).

The relativization of abilities to environments may be thought to offer a principled basis for drawing the relevant epistemic distinction in the epistemic twin earth cases. Greco (2009c: 22) suggests that 'environment' picks out "sets of relatively stable circumstances" and 'conditions' picks out "shifting circumstances within an environment", but also notes that some states of affairs could fall under either heading. The thought might then be that while *S* has the ability to discriminate between water and non-water relative to her regional environment, twin-*S* lacks the ability to discriminate between water and non-water relative to her regional environment.¹⁶ Twin-*S* would have a high frequency of true water-beliefs if formed in her local environment, but abilities imply *reliability* which in turn is contingent upon a range of true beliefs in nearby worlds. And while twin-*S* correctly classifies the watery stuff in her local environment, her actual track-record is the result of good luck. Given the nature of twin-*S*'s regional environment, very easily could she misclassify the stuff she calls 'water'.

In response, one could point out that twin-*S* clearly has that discriminatory ability relative to her global environment, because she is genuinely reliable in that environment with

regard to the target result (in this case true beliefs regarding water). Greco (2009c: 21-22) would concede that point. But then twin-*S* should also possess the target ability relative to her regional environment, and this should suffice on this view to ensure that the subject concerned is genuinely manifesting that ability. Abilities are relative only to the stable environment in which whoever has them is typically located. That is to say, abilities are possessed as long as they are reliably manifested in global environments. These are the normal circumstances in which abilities are acquired through learning and sustained through practice. Temporary abnormal environments cannot rob a subject of an ability that she otherwise reliably manifests in the normal run of things. Due to deceiving regional facts exercising her ability in a local environment might be unreliable, but all that shows is that the subject would fail to meet some non-virtue-theoretic epistemic condition such as having a safe basis for belief which excludes knowledge-undermining epistemic luck.

Indeed, Greco's own baseball analogy confirms this point. For consider a baseball player who is in normal baseball-playing conditions (good light, not too much wind, no sand in his eyes, and so on), but who is in circumstances such that poor baseball-playing conditions are in the modal vicinity (e.g., there is a fierce storm raging near to where the baseball player is playing). In this case, does the baseball player either fail to exhibit any ability at all or exhibit a special ability, the ability (say) to hit fastballs when in a normal baseball-playing situation which could so very easily have been an adverse baseball-playing situation? Intuitively, the answer to both questions is 'no'—it is just the normal baseball playing ability on display here.¹⁷

In order to illustrate why relativizing abilities to regional environments is a non-starter let's fine-tune the target ability. If, on the one hand, to have the ability to discriminate between water and non-water is to have the ability to discriminate between water and twin-water, then neither *S* nor twin-*S* has that ability relative to *any* environment. Either could acquire that ability, but only by learning some chemistry. Certainly, the mere absence of twin-water on Earth does not bestow that ability on *S*. If, on the other hand, to have the ability to discriminate between water and non-water is to have the ability to discriminate between water and distinct liquids that are also superficially distinct, then both *S* and twin-*S* have that ability relative to *all* environments. In particular, the mere presence of twin-water on twin-earth does not prevent twin-*S* from being able to reliably tell water apart from petrol or beer.

There is a twist to Greco's view on this score, however, which might come to his rescue. For as Greco notes, there is a fair degree of vagueness in his account of abilities regarding how loosely or tightly one defines the relevant W, C, R and E in each particular case. Thus Greco (2009*c*: 22-24) takes his view to be faced with a version of the *generality problem* that afflicts any reliabilist theory. His response to this vagueness in his account is to let these parameters be set by purely pragmatic and contextual factors. In particular, he argues that it is the practical reasoning context of the subject to whom we are attributing the ability which sets these parameters.

We can illustrate Greco's proposal on this score by offering two versions of the 'barn façade' case.¹⁸ In the first, the subject concerned is a tax official charged with counting barns on a farm, where there is a tax advantage to a farmer having lots of barns on her property. In the second, the subject concerned is a new farmhand who has been charged to find a barn nearby to provide shelter for a cow. In both cases, although the subject is, unbeknownst to her, in barn façade county, she happens to be looking at the frontage of the one real barn in the vicinity (i.e., on this particular farm there is only one object that looks like a barn, and it's a genuine barn), and in both cases she lacks the ability to tell barns and barn-façades apart simply by observing their frontages.

According to Greco (2009*a*: 80), we are inclined to offer very different knowledge ascriptions in each case. In particular, the seriousness of the practical reasoning context of the subject in the first case is meant to prompt us to not ascribe knowledge, while the relative unimportance of the practical reasoning context of the subject in the second case is meant to prompt us to ascribe knowledge. Greco thus endorses a form of *pragmatic encroachment* in epistemology, the thesis that purely non-epistemic factors, such as the importance of a practical reasoning context, can have a bearing on the epistemic assessment we offer of a subject's belief.¹⁹

The import of this feature of Greco's view about the nature of abilities, and thus the nature of cognitive abilities (and hence knowledge, given that he is a robust virtue epistemologist), is that it is open to him to explain the epistemic difference between the two cases in terms of a relevant difference in the practical reasoning context of the two subjects concerned. Perhaps, that is, the difference between *S* and twin-*S* is that the latter is in a more demanding practical reasoning context, and that this explains why she lacks knowledge while her counterpart on earth possesses knowledge.

But even with Greco's proposal understood along these lines it is still hard to see how a response of this sort is going to be made to hold fast. After all, given the way we have set up the epistemic twin earth cases there seems no reason at all for thinking that the two counterpart subjects are in different practical reasoning contexts. Moreover, this is not merely a matter of the cases being underdescribed. For notice that no matter what further detail is added to the cases there seems no obvious reason why these additional details should *necessitate* a difference in practical reasoning context across the two cases. But if that's right then there is nothing preventing us from running the epistemic twin earth argument against even this formulation of robust virtue epistemology, since for this argument to be successful we only need to formulate one version of the argument where there is no difference in practical reasoning context across the two cases in question.

4. RESPONSE II: SUCCESS NOT BECAUSE OF ABILITY

A second line of response to the epistemic twin earth argument might be to question whether twin-*S*'s cognitive success is really *because of* her ability in the relevant sense. On either rendering of robust virtue epistemology, however, it is hard to see how such a line could be made to hold.

Consider first the disposition manifestation reading of robust virtue epistemology. First off, notice that the manifestation of purely physical dispositions is a local matter. A solute dissolves in a solvent even if there are plentiful superficially indistinguishable solvents around in which it could easily have been immersed and had it been immersed in any one of them it would not have dissolved. The same is true of skills that have both cognitive and non-cognitive components. When an expert archer hits the target she manifests her archery skill even if her competitor could easily have tampered with her bow, or even if a strong gust could easily have diverted the arrow off course. Whether the accuracy of the performance is through the exercise of such a skill is determined locally. But if that's right, then we should expect the manifestation of a cognitive disposition to be a local matter also. Accordingly, given that *S* and twin-*S* are physically identical and in identical local environments, we should expect them to be both manifesting the relevant cognitive disposition. On the disposition manifestation reading, then, there seems no basis on which one could argue that

the twin-*S*'s cognitive success is any less because of her cognitive ability than it is in the case of *S*'s cognitive success.

The causal explanatory reading doesn't fare any better on this score. For given that the local environments that *S* and twin-*S* inhabit are completely identical, it is hard to see how there should be a different causal explanation offered of their respective cognitive successes. Since the *explanandum* (i.e., true belief) is identical in the two cases, the explanations would differ only if a difference in *explanans* could be established. But what could that possibly be? Remember that, at this juncture, we take *S* and twin-*S* to be sharing all pertinent cognitive abilities. For while Greco (2009c: 21-22) is surely right that abilities are tied to conditions understood as local environments, it remains that only the outcomes of exercising abilities can shift with variations in local facts, and not whether abilities are possessed at all. For instance, the fact that a sudden, and entirely unexpected, gust diverts an arrow could mean that the expert archer misses the target altogether, but she is no less skilled for that. So, if one such ability features in the causal explanation of *S*'s true belief, it would also do so in the case of twin-*S*'s true belief. Likewise, no local physical facts could be cited in one causal explanation but not the other. True, causal explanations frequently involve extrinsic properties, but as *S* and twin-*S* also share global physical environments, Greco needs to show that the differences in their regional environments somehow is bound to bear on the causal explanations in question.

Let's explore that avenue. Although Greco can argue, as we saw above, that an extra-environmental factor such as the practical reasoning context of the subject can have a bearing on the causal explanation in play, this doesn't seem to offer any assistance when it comes to dealing with the specific problem posed by epistemic twin earth. Note that this is not to deny that it could well be true that *S* and twin-*S* are in relevantly different practical reasoning contexts. For example, perhaps twin-*S* is in a practical reasoning context which makes her regional environment relevant to the causal explanation of her cognitive success in a way that it is not relevant to *S*'s cognitive success. If that's right, then on Greco's view only *S*'s cognitive success will be because of her cognitive ability, and so only she will be in the market for knowledge on this score.

As we also noted above, however, Greco needs to do more than merely show that there is this possibility. In particular, what he needs to be able to show is that *S* and twin-*S* *must* be in a relevantly different practical reasoning context. For so long as he can't show this

then it will follow that there will be a formulation of the cases in play in the epistemic twin earth argument on which even Greco will be committed to holding that twin-*S*'s cognitive success is as much because of her cognitive ability as *S*'s cognitive success. But since one formulation of the cases is all we need to make the argument go through, it follows that this feature of Greco's view is unable to offer him a route out of this problem.

5. RESPONSE III: MODIFY THE PROPOSAL

A third line of response on the part of robust virtue epistemology to the epistemic twin earth argument might be to modify the view in the hope of thereby being able to handle such cases. One way in which this line of argument might go is suggested in certain passages in the work of Sosa. Consider, for example, this passage:

“[...] the requirement for [*knowledge*] is not just that one's belief be true, and derive from a competence. The requirement is rather that one believe correctly (with truth) through the exercise of a competence in its proper conditions.” (Sosa 2007: 33)

The important qualification consists in the requirement that the cognitive abilities be exercised in proper conditions or, as he glosses it elsewhere “appropriately normal conditions.” (Sosa 2007: 29) Thus the proponent of robust virtue epistemology might be tempted to adopt a view such that knowledge is cognitive success that is because of the exercise of a relevant cognitive ability in proper or normal conditions (where the ‘because of’ in question is given either a causal explanatory or disposition manifestation reading, depending on the variety of robust virtue epistemology in play). Call this rendering of the view, *normal conditions robust virtue epistemology*.

Sosa suggests at certain places that ‘normal/proper conditions’ picks out local facts, such as conditions of “distance, lighting, size of surface, etc.” (See especially Sosa 2007: 35-6) But on that reading of normal conditions robust virtue epistemology it faces the same difficulty as both versions of robust virtue epistemology in accounting for why twin-*S*'s cognitive success is not through cognitive ability when *S*'s cognitive success is. Since twin-*S*'s local environment is physically identical to *S*'s local environment, the conditions on perception that obtain in the former environment are normal/proper if the corresponding conditions that obtain in the latter environment are normal/proper. Consequently, insisting

that the conditions in which cognitive abilities are exercised be normal/proper, thus understood, cannot account for why our twins should differ epistemically.

If instead the different regional facts on earth and epistemic twin earth means that the conditions in which *S* exercises her cognitive ability are proper/normal while the conditions in which twin-*S* exercises her identical cognitive ability are improper/abnormal, then, obviously, the proponent of normal conditions robust virtue epistemology can explain why *S* and twin-*S* differ epistemically when both they and their local and global environments are physically identical. Understanding ‘proper’ or ‘normal’ conditions in this way is in effect to impose on knowledge the requirement that *S*’s belief be true as a result of *S* exercising her cognitive ability in conditions such that they would not lead to knowledge-undermining epistemic luck. That is, it is tantamount to insisting that *S* should not have attained her cognitive success in conditions such that this cognitive success could very easily have been cognitive failure.

By way of illustration of this point, let’s dwell on Turri’s (2011) refinement of Sosa’s (2007) theory of performance-assessment. As noted above, Sosa takes knowledge to be *apt* belief, where an apt belief is a belief that is *accurate* (i.e., true) because *adroit* (i.e., skilfully formed). But given the problems we have raised for a view of this sort, one might be tempted to modify the proposal by arguing that knowledge should be identified with something more stringent than mere apt belief. Turri suggests two such strengthenings.

The first is that the belief in question be *adept*, in the sense that its accuracy is not merely because of its adroitness, but where its accuracy *manifests* adroitness. Recall the ‘Rooney’ case from §1. Both the brilliant goals that Rooney scores and his extravagant salary are in a sense ‘because of’ his prodigious footballing ability. But only the former are *manifestations* of this ability. In the same way, when it comes to cognitive dispositions Turri argues that we should focus on cases where cognitive success manifests, and is not merely because of, cognitive ability. Adeptness is therefore a more demanding form of aptness, and hence entails aptness.

But that’s not all. In response to ‘barn façade’ cases, Turri proposes that knowledge also requires that the target belief be *ample*. A belief is ample just in case it is safe (i.e., it is true and could not have easily been false) and where its safety manifests adroitness. The problem with the target belief in the barn façade case, argues Turri, is that while it is adept, it

is not ample (since it is not safe). Note that, so defined, amplitude entails adeptness and therefore also aptness.

Turri's promising proposal is then that knowledge is ample belief: *S* knows that *p* just in case the safety of *S*'s belief manifests her cognitive dispositions. This is promising because *S* and twin-*S* will differ epistemically on this proposal. While their respective beliefs are both adept, only *S*'s belief is also ample and hence only *S* counts as having knowledge. The reason twin-*S*'s belief is not ample, and hence the reason she lacks knowledge, is that her belief is simply unsafe. Given the nature of twin-*S*'s regional environment, very easily would she have been saddled with a false belief.

Turri admits that identifying knowledge with ample belief involves a radical change to the original virtue-theoretic idea, particularly as that idea appears in Sosa's work (on which Turri's proposal is based). We agree, but we also contend that the proposal deprives the robust virtue epistemologist of the original analogies with which the view was explicated, if not motivated. Thus Sosa explicitly adopts the disposition manifestation version of robust virtue epistemology on the model of various skills that are not purely cognitive. But the modal profile of the kind of skills that Sosa often appeals to resembles ordinary physical dispositions in that neither requires safety.

This was in effect the upshot of the examples involving solubility and archery from §4. A solute dissolves in a solvent even if there are plentiful superficially indistinguishable solvents around in which it could easily have been immersed and had it been immersed in any one of them it would not have dissolved. That the solute might easily not have dissolved had it been immersed in a different solvent has no bearing on whether the resulting solution manifests the solubility of the solvent. Similarly, when an expert archer hits the target she manifests her archery skill even if a strong gust could easily have diverted the arrow off course. That the arrow could easily have landed in a neighbouring field had this sudden gust intervened has no bearing on whether its hitting the target manifests the skill of archery. But if that's right then no appeal to these analogies can be made in the course of motivating or even explicating why the disposition manifestation version of robust virtue epistemology should adopt Turri's proposal.

What is called for is thus an identification of alternative analogies which share the modal profile with ample belief.²⁰ We believe the prospects look dim. Not only have we already seen cases of unsafe adeptness—and hence which lack amplitude—there are also

cases of safe adeptness without amplitude. Suppose an expert archer dispatches an arrow which then hits the target. A sudden, strong gust might easily have diverted the arrow off course, but had that happened a guardian angel would then have intervened to ensure that the arrow was brought back on course again. The success is thus both safe and adept, but it is clearly not safe because it manifests adroitness (since its safety is down to the guardian angel), and hence it is not ample.²¹ But if safety and adeptness can come apart in such ways that are not down to the manifestation of dispositions, then it is by no means obvious how to construe safety as an integral part of disposition manifestation.

Instead we propose that the thesis that knowledge is ample belief—(or normal conditions robust virtue epistemology, where this is understood in this way)—is essentially nothing more than a notational variant of modest virtue epistemology. These ‘robust’ virtue epistemologists are in effect simply adding an additional non-virtue-theoretic condition to their account of knowledge which deals with the problem posed by knowledge-undermining epistemic luck. At the very least, the onus is on the proponent of this line of response to the epistemic twin argument to explain why their view does not now collapse into modest virtue epistemology.

6. REPONSE IV: BITE THE BULLET

A fourth line of response to the epistemic twin earth argument might be to try to pinpoint a palatable way of biting the bullet. Here is one such tack, which might seem to be available to a proponent of the disposition manifestation version of robust virtue epistemology. On Sosa’s view, knowledge *simpliciter* (i.e., cognitive success that is because of the relevant cognitive ability) constitutes *animal knowledge* (see, e.g., Sosa 1997: 24). Reflective knowledge, in contrast, is more demanding in that it requires second-order knowledge (i.e., it is essentially animal knowledge that one has animal knowledge).²² Elsewhere in his work Sosa exploits this distinction between animal and reflective knowledge in order to explain why on his view subjects have knowledge in cases where epistemologists often deny that the subject concerned is a knowing subject. That is, in such cases he contends that what the subject is lacking is not knowledge *simpliciter* (i.e., animal knowledge), but specifically reflective knowledge.²³ Perhaps, then, Sosa could extend this strategy to the epistemic twin earth

argument and hence contend that while S on earth has both animal and reflective knowledge, twin- S on epistemic twin earth qualifies only for the former. In this way, he could argue that while it is true that if S has knowledge of the target proposition then so does twin- S , there is nonetheless an epistemic difference between the two subjects just as intuition dictates, albeit an epistemic difference which is at the level of reflective knowledge rather than knowledge *simpliciter*.

One immediate problem with this line of argument is that it does not seem at all essential to the epistemic twin earth argument that either of these subjects should even have the relevant second-order belief. But insofar as they lack this belief then they are thereby *both* excluded from having reflective knowledge, and hence one cannot explain away the intuition that there is an epistemic difference between these two subjects by appealing to the thought that one of them has reflective knowledge that the other lacks. Furthermore, note that this problem does not trade on the particular rendering of robust virtue epistemology that is in play.

Even if we ignore this problem and allow that the subjects concerned have the relevant second-order beliefs, however, this line of response to the epistemic twin earth argument still seems to flounder. The reason for this is that we appear to be perfectly able to run a corresponding epistemic twin earth argument which is specifically focussed on reflective knowledge, regardless of the rendering of robust virtue epistemology that is in play.

Consider first the disposition manifestation rendering of robust virtue epistemology which Sosa favours. Animal knowledge and reflective knowledge differ on this view in terms of the different cognitive dispositions that are being manifested in each case. Whereas animal knowledge involves the manifestation of a first-order cognitive disposition, reflective knowledge involves the manifestation of a second-order cognitive disposition which is tracking the cognitive success of the relevant first-order cognitive disposition. Presumably, however, all dispositions, whether first or second-order, have physical bases which are resident in whoever has the disposition. Certainly, if, as Sosa acknowledges, first-order cognitive dispositions have physical bases, then the claim that second-order cognitive dispositions lack physical bases would seem *ad hoc*. But then how is the proponent of this hypothetical line of response to the epistemic twin earth argument to explain how S on earth and twin- S on epistemic twin earth could enjoy different second-order epistemic statuses

when they and their local and global environments are physically identical? This proponent could pursue either of two lines: deny that twin-*S* possesses the second-order cognitive disposition, or else deny that twin-*S*'s meta-cognitive success involves manifesting such a disposition. But neither looks promising, because everything said above (in §§3 and 4) in response to the objection that twin-*S* either lacks the first-order cognitive disposition or else her cognitive success is not through the manifestation of such a disposition carries over, *mutatis mutandis*, to second-order dispositions.

The foregoing may be thought slightly unfair to Sosa. In response to his kaleidoscope case, Sosa (2007: 33; cf. Sosa 2009: 238-9) proposes that *S*'s meta-belief is true because of a competence only if “it derives from the exercise of that competence in appropriate conditions for its exercise, and that exercise in those conditions would not then too easily have issued a false belief.” The latter supplementation is certainly fit for purpose. Since the identical, local environments in which *S* and twin-*S* are embedded are identical to their global environments, both form true beliefs in conditions appropriate for exercise of their respective cognitive abilities. But while *S*'s exercise of her cognitive ability results in true beliefs across nearby worlds, twin-*S*'s exercise would indeed all too easily have issued a false belief. For that reason *S* has animal knowledge that she has animal knowledge, but twin-*S* lacks such reflective knowledge. Assuming with Sosa (2007: 31-40) that first-order animal knowledge does not also require that the exercise of the pertinent cognitive ability in normal conditions would not very easily have issued a false belief, twin-*S* can retain such knowledge.

It should be pretty obvious by now why this proposal offers no comfort to the robust virtue epistemologist. The explicit recommendation is that reflective knowledge requires the exercise of a cognitive ability in conditions that would not lead to such knowledge being undermined by epistemic luck. For what prevents twin-*S* from having reflective knowledge is that the conditions in which she exercises her meta-cognitive ability are such that her cognitive success could very easily have been cognitive failure. A non-virtue-theoretic condition designed to deal with knowledge-undermining epistemic luck is thus explicitly built into reflective knowledge. Consequently, the proposal cannot explain away the intuition that *S* and twin-*S* differ epistemically in a way that is consistent with robust virtue epistemology.

The causal explanatory rendering of robust virtue epistemology doesn't fare any better on this score. For insofar as we grant the argument offered above to the effect that there is no difference in the causal explanation of *S*'s cognitive success and twin-*S*'s corresponding cognitive success, then it is hard to see what principled basis there would be for holding that there is a difference in the causal explanation of *S*'s meta-cognitive success and twin-*S*'s meta-cognitive success. If (as we argued in §§3 and 4), the mere fact that *S* and twin-*S* inhabit physically distinct regional environments cannot furnish causal-explanatory resources for accounting for *S*'s cognitive success that are unavailable in an account of twin-*S*'s cognitive success, then going second-order is to no avail. Bear in mind that no other physical difference between the two cases exists. Certainly, the onus is on the proponent of the causal explanatory brand of robust virtue epistemology to demonstrate which features of the agents' meta-cognitive abilities or respective environments make for relevantly different causal explanations in this regard. The prospects, we submit, look dim.

7. CONCLUDING REMARKS

Given the problems facing the four lines of response that we have considered to the epistemic twin earth argument, we conclude that robust virtue epistemology is in serious trouble. At the very least, it is incumbent upon the proponent of robust virtue epistemology to either show how one of these four lines of response is viable after all, or else demonstrate that there is an additional line of response available that we have not considered. Given that, as we noted earlier, robust virtue epistemology faces objections from other quarters which also purport to demonstrate the need for a separate epistemic condition which handles knowledge-undermining epistemic luck—such as a safety condition—we conclude that the current state-of-play on this score favours the view that any virtue-theoretic account of knowledge must be at best of the modest, rather than robust, variety.²⁴

Jesper Kallestrup & Duncan Pritchard
School of Philosophy, Psychology & Language Sciences
University of Edinburgh
Dugald Stewart Building

3 Charles Street

Edinburgh

EH8 9AD

Scotland, UK

jesper.kallestrup@ed.ac.uk & duncan.pritchard@ed.ac.uk

REFERENCES

- Axtell, G. (1997). 'Recent Work in Virtue Epistemology', *American Philosophical Quarterly* 34: 410-30.
- Baehr, J. (2004). 'Virtue Epistemology', *Internet Encyclopaedia of Philosophy*, (eds.) B. Dowden & J. Fieser, www.iep.utm.edu/virtueep/.
- (2008). 'Four Varieties of Character-Based Virtue Epistemology', *Southern Journal of Philosophy* 46: 469-502.
- Chalmers, D. (1996). *The Conscious Mind: In Search of a Fundamental Theory*, Oxford: Oxford University Press.
- Chisholm, R. (1977). *Theory of Knowledge*, (2nd ed.), Englewood Cliffs, NJ: Prentice-Hall.
- Code, L. (1987). *Epistemic Responsibility*, Hanover, NH: University Press of New England.
- Craig, E. (1990). *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*, Oxford: Clarendon Press.
- Fantl, J., & McGrath, M. (2010). 'Pragmatic Encroachment', *Routledge Companion to Epistemology*, (eds.) S. Bernecker & D. H. Pritchard, 558-68, London: Routledge.
- Frankfurt, H. (1969). 'Alternate Possibilities and Moral Responsibility', *Journal of Philosophy* 46: 829-39.
- Goldberg, S. (2010). *Relying on Others: An Essay in Epistemology*, Oxford: Oxford University Press.
- Goldman, A. (1976). 'Discrimination and Perceptual Knowledge', *Journal of Philosophy* 73: 771-91.
- 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: Cambridge University Press.
- (2002). 'Virtues in Epistemology', *Oxford Handbook of Epistemology*, (ed.) P. Moser, 287-315, Oxford: Oxford University Press.
- (2003). 'Knowledge as Credit for True Belief', *Intellectual Virtue: Perspectives from Ethics and Epistemology*, (eds.) M. DePaul & L. Zagzebski, 111-34, Oxford: Oxford University Press.
- (2007). 'The Nature of Ability and the Purpose of Knowledge', *Philosophical Issues* 17: 57-69.
- (2008). 'What's Wrong With Contextualism?', *Philosophical Quarterly* 58: 416-36.
- (2009a). *Achieving Knowledge*, Cambridge: Cambridge University Press.
- (2009b). 'The Value Problem', *Epistemic Value*, (eds.) A. Haddock, A. Millar & D. H. Pritchard, 313-21, Oxford: Oxford University Press.
- (2009c). 'Knowledge and Success From Ability', *Philosophical Studies* 142: 17-26.
- Greco, J., & Turri, J. (2009). 'Virtue Epistemology', *Stanford Encyclopaedia of Philosophy*, (ed.) E. Zalta, <http://plato.stanford.edu/entries/epistemology-virtue/>.
- Hetherington, S. (Forthcoming). 'There Can Be Lucky Knowledge', *Contemporary Debates in Epistemology* (2nd Ed.), (eds.) E. Sosa, M. Steup & J. Turri, Oxford: Blackwell.
- Hookway, C. (2003). 'How to Be a Virtue Epistemologist', *Intellectual Virtue: Perspectives from Ethics and Epistemology*, (eds.) M. DePaul & L. Zagzebski, 183-202, Oxford: Oxford University Press.
- Horgan, T., & Timmons, M. (1991). 'New Wave Moral Realism Meets Moral Twin Earth', *Journal of Philosophical Research* 16: 447-65.

- (1992). 'Troubles on Moral Twin Earth: Moral Queerness Revived', *Synthese* 92: 221-60.
- Jackson, F. (1996). 'Mental Causation', *Mind* 105: 377-413.
- Kallestrup, J. (2011). *Semantic Externalism*, London: Routledge.
- Kvanvig, J. (1992). *The Intellectual Virtues and the Life of the Mind*, Savage, MD: Rowman & Littlefield.
- (2010). 'Virtue Epistemology', *Routledge Companion to Epistemology*, (eds.) S. Bernecker & D. H. Pritchard, 199-207, London: Routledge.
- Luper, S. (1984). 'The Epistemic Predicament', *Australasian Journal of Philosophy* 62: 26-50.
- (2003). 'Indiscernability Skepticism', *The Sceptics: Contemporary Essays*, (ed.) S. Luper, 183-202, Aldershot: Ashgate.
- Montmarquet, J. (1993). *Epistemic Virtue and Doxastic Responsibility*, Lanham, MD: Rowman & Littlefield.
- Pritchard, D. H. (2002). 'Resurrecting the Moorean Response to the Sceptic', *International Journal of Philosophical Studies* 10: 283-307.
- (2005). *Epistemic Luck*, Oxford: Oxford University Press.
- (2007a). 'Anti-Luck Epistemology', *Synthese* 158: 277-97.
- (2007b). 'Recent Work on Epistemic Value', *American Philosophical Quarterly* 44: 85-110.
- (2007c). 'The Value of Knowledge', *Stanford Encyclopædia of Philosophy*, (ed.) E. Zalta, <http://plato.stanford.edu/entries/knowledge-value/>.
- (2008a). 'Greco on Knowledge: Virtues, Contexts, Achievements', *Philosophical Quarterly* 58: 437-47.
- (2008b). 'Sensitivity, Safety, and Anti-Luck Epistemology', *The Oxford Handbook of Scepticism*, (ed.) J. Greco, 437-55, Oxford: Oxford University Press.
- (2009a). 'Apt Performance and Epistemic Value', *Philosophical Studies* 143: 407-16.
- (2009b). 'The Value of Knowledge', *The Harvard Review of Philosophy* 16: 2-19.
- (2011). 'Anti-Luck Virtue Epistemology', *Journal of Philosophy* 108: ??-??.
- (Forthcoming). 'There Cannot be Lucky Knowledge', *Contemporary Debates in Epistemology* (2nd Ed.), (eds.) M. Steup & J. Turri, Oxford: Blackwell.
- Pritchard, D. H., Millar, A., & Haddock, A. (2010). *The Nature and Value of Knowledge: Three Investigations*, Oxford: Oxford University Press.
- Sainsbury, R. M. (1997). 'Easy Possibilities', *Philosophy and Phenomenological Research* 57: 907-19.
- Sosa, E. (1988). 'Beyond Skepticism, to the Best of Our Knowledge', *Mind* 97: 153-89.
- (1991). *Knowledge in Perspective: Selected Essays in Epistemology*, Cambridge: Cambridge University Press.
- (1999). 'How to Defeat Opposition to Moore', *Philosophical Perspectives* 13: 141-54.
- (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, Oxford: Clarendon Press.
- (2009). *Reflective Knowledge: Apt Belief and Reflective Knowledge*, Oxford: Clarendon Press.
- Timmons, M. (1999). *Morality Without Foundations: A Defence of Ethical Contextualism*, Oxford: Oxford University Press.
- Turri, J. (2011). 'Manifest Failure: The Gettier Problem Solved', *Philosophers' Imprint* 11: 1-19.
- Turri, J., & Sosa, E. (2010). 'Virtue Epistemology', *Oxford Bibliographies Online: Philosophy*, (ed.) D. H. Pritchard, <http://oxfordbibliographiesonline.com/display/id/obo-9780195396577-0123>.

- Williamson, T. (2000). *Knowledge and its Limits*, Oxford: Oxford University Press.
- Zagzebski, L. (1996). *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*, Cambridge: Cambridge University Press.
- (1999). ‘What is Knowledge?’, *The Blackwell Guide to Epistemology*, (eds.) J. Greco & E. Sosa, Oxford: Blackwell.

NOTES

¹ For some helpful surveys of the literature on virtue epistemology, see Axtell 1997, Greco 2002, Baehr 2004, Greco & Turri 2009, and Kvanvig 2010. See also Baehr 2008. Note that the characterisation of epistemic virtues in terms of cognitive abilities is potentially contentious, since on some versions of virtue epistemology the epistemic virtues are explicitly understood *in contrast* to cognitive abilities. For an example of a virtue-theoretic proposal on which cognitive abilities can count as knowledge-producing, see the agent reliabilist position defended by Greco 1999; 2000. For a very different neo-Aristotelian virtue-theoretic proposal, one which explicitly denies that mere cognitive abilities can be knowledge-producing, see Zagzebski 1996.

² For further discussion of the idea that virtue epistemology might be best thought of as reorienting the concerns of traditional epistemology rather than simply responding to those concerns, see Code 1987, Kvanvig 1992, Montmarquet 1993, Hookway 2003, and Roberts & Wood 2007.

³ For more on the problem of knowledge-undermining epistemic luck, see Pritchard 2005; 2007*a*; 2008*b*.

⁴ Versions of the safety condition for knowledge have been offered by a number of authors, including Luper 1984; cf. Luper 2006, Sainsbury 1997, Sosa 1999, Williamson 2000, and Pritchard 2002; 2005; 2007*a*. For a comprehensive and up-to-date discussion of the merits of the safety condition for knowledge, see the exchange between Pritchard (*forthcoming*) and Hetherington (*forthcoming*).

⁵ One clear statement of such a view can be found in early work by Greco (e.g., 1999; 2000), where he develops his own variant of virtue epistemology known as agent reliabilism. For example, after setting out the problem posed by Gettier cases and noting that his position cannot deal with this problem, he writes:

“Accordingly, the [*virtue-theoretic*] conditions that agent reliabilism sets down as necessary for knowledge are not also sufficient for knowledge; something else must be added.” (Greco 2000: 251-2)

⁶ Note that a further motivation for robust virtue epistemology, which we will not be exploring here, is the idea that it is uniquely placed to deal with the problem of accounting for the special value of knowledge. See, for example, Greco 2009*b*. For a critical discussion of this claim, see Pritchard 2009*b* and Pritchard, Millar & Haddock 2010: chs. 1-2. For a recent survey of the problem of epistemic value, see Pritchard 2007*c*; cf. Pritchard 2007*b*.

⁷ See Sosa 2007: *passim*.

⁸ A third proposal is offered by Zagzebski (e.g., 1999) who treats the ‘because of’ relation as an undefined (perhaps indefinable) primitive. She notes that the best analysis available of this relation is in terms of counterfactuals. So, for example, to say that a particular brilliant goal was because of Rooney’s prodigious footballing skill is to say something like, had Rooney not exhibited this prodigious skill, then that particular brilliant goal would not have occurred. But she also argues that such an analysis is ultimately unsatisfactory because there is a range of cases that it cannot handle. For example, she writes that “there are no counterfactual circumstances in which a bachelor is not unmarried, but it would not be true to say that he is a bachelor *because* he is unmarried.” (Zagzebski 1999: 111) While it is of course a theoretical option to regard the ‘because of’ relation as primitive, we here confine our attention to explications of this relations. See Turri 2011 for a recent discussion of Zagzebski’s proposal.

⁹ For example, Greco 2009*a*: 12 writes that “the term “because’ [...] marks a causal explanation.” Later on (*ibid.*) he makes clear that the agent’s abilities must be the overarching element in the causal explanation in question when he states that “in cases of knowledge, *S*’s believing the truth is explained by *S*’s abilities, as opposed to dumb luck, or blind chance, or something else.”

¹⁰ Even so, dispositional properties frequently enter into causal explanations in science. For instance, Richard Feynman explained why the space shuttle Challenger blew up in terms of the failure of an O-ring in one of the solid rocket boosters to expand at lift-off. Jackson 1996: 397 suggests that causal explanations by dispositional properties provide two kinds of information: (i) the effect was caused by the categorical basis of the disposition, and (ii) the effect is one of the outputs in terms of which the disposition is defined. (ii) is required because some base properties can ground more than one disposition—e.g., electrical and thermal conductivity in metals share the same categorical basis. What Feynman discovered was that the categorical basis of the rigidity caused the disaster, and the disaster resulted from the kind of output distinctive of rigidity: the O-ring failed to expand after compression and its failure lead to the disaster.

¹¹ For a thorough discussion of the problems facing robust virtue epistemology (of either variety) on this score, see Pritchard, Millar & Haddock 2010: chs. 3 & 4 and Pritchard 2011.

¹² See, for example, Putnam 1975. For a recent discussion of standard twin-earth arguments, see Kallestrup 2011: ch. 3.

¹³ If twin-S uttered ‘that’s water’ while demonstrating twin-water she would express the false proposition that that’s water. We assume that the concept of water as deployed on both earth and epistemic twin earth is a natural kind concept that applies to all and only H₂O. One might envisage a loophole here for the robust virtue epistemologist if twin-S’s utterance has the purely descriptive truth-condition: ‘that’s water’ is true *iff* that’s watery stuff, or the disjunctive truth-condition: ‘that’s water’ is true *iff* that’s water or that’s twin-water. We find both views implausible. The presence of small amounts of twin-water on epistemic twin earth implies neither that water is a functional kind in the way that, say, vitamin is, nor that water is a disjunctive kind in the way that, say, jade is. Even those with descriptivist or semantic internalist leanings insist that, to a first approximation, ‘water’ in someone’s mouth picks out the dominant watery stuff of their acquaintance. XYZ is neither dominant nor stuff with which adequate causal connections are sustained. In fact, Chalmers 1996: 58 is explicit that if the watery stuff in our world turned out to be a mixture of 95% H₂O and 5% XYZ, the primary intension of ‘water’ would pick out only H₂O. For more details, see Kallestrup 2011: chs. 3 & 4.

¹⁴ Note that this epistemic twin earth argument is significantly different to the moral twin earth argument that has been proposed by Horgan & Timmons 1991; 1992 and Timmons 1999: ch. 2, and which calls into question a certain form of moral naturalism. That said, one could presumably run an epistemic twin earth argument which is structurally analogous to the moral twin earth argument, and which thereby creates problems for a corresponding version of epistemic naturalism.

¹⁵ To use Sosa’s own terminology, this would be to question whether the *accurate* (i.e., true) belief in question was *adroitly* (i.e., skilfully) formed on the basis of lacking *adroitness*, which for him is a necessary requirement for *apt* belief, and hence knowledge (since on his view knowledge is apt belief, where this is a belief which is accurate because adroit). See, for example, Sosa 2007: *passim*. Of course one could also in theory challenge whether the belief was adroitly formed on some other basis, but we shall not discuss that possibility any further here.

¹⁶ Greco 2009c: 21 offers a similar reply to the ‘barn façade’ case described in Goldman 1976 (and credited by Goldman to Carl Ginet).

¹⁷ For more on this point about the nature of abilities, see Pritchard, Millar & Haddock 2010: ch. 2 and Pritchard 2011: §3.

¹⁸ Greco 2009a: 80 also illustrates his view by offering two versions of the barn façade case. Unfortunately, he illicitly varies important features of the example across the two cases, features which by anyone’s lights will be epistemically relevant (such as the number of fake barns in the vicinity). As a result, the examples that Greco uses to illustrate his proposal are not that helpful, in that one can consistently grant, with Greco, that one should deliver a different judgement about whether the subject concerned has knowledge across the two cases without thereby assenting to Greco’s proposal regarding the relevance of the subject’s practical reasoning context for the purposes of knowledge ascriptions. We have stuck to Greco’s discussion in this regard as closely as possible while keeping all relevant features of the two cases fixed.

A further awkwardness in Greco’s discussion of these examples which is worthy of note is that he focuses entirely on the issue of whether the practical reasoning context of the subject has a bearing on our *knowledge* ascriptions to this subject. But notice that one could grant this claim without thereby granting that our ascription of the target ability to the subject is sensitive to which practical reasoning context she is in. For perhaps there are independent reasons why knowledge ascriptions should be sensitive to the practical reasoning context of the subject. (Indeed, the story that Greco tells in support of why knowledge ascriptions are sensitive to practical reasoning contexts in this way explicitly adverts to the genealogical account of the concept of knowledge offered by Craig 1990, so this is far from being a purely academic point). In any case, we take it that it is Greco’s view that *both* ability ascriptions and knowledge ascriptions are sensitive to the practical reasoning context of the subject, even if it is unclear whether, for example, Greco means to suggest that there are *additional* reasons for supposing that knowledge ascriptions are sensitive to the practical reasoning context of the subject in this way.

¹⁹ For an excellent overview of the issues with regard to pragmatic encroachment in epistemology, see Fantl & McGrath 2010. See Pritchard 2008a for a specific discussion of Greco’s own commitment to pragmatic encroachment, as presented in Greco 2008.

²⁰ Turri 2011 proposes that *overwhelming* an opponent in competition shares a modal profile with ample belief. But it is not obvious that overwhelming requires amplitude. Suppose that a novice and an expert archer are up against each other in an archery competition. They each dispatch a dozen arrows. While the novice is lucky to hit the target once, the expert propels the arrow into the target every single time. The expert secures an overwhelming victory. But unbeknownst to both, a sudden, strong gust could easily have diverted the expert’s arrows off course, and had it done so the novice would very surprisingly have won the competition instead of

suffering an overwhelming defeat. To overwhelm an opponent by the manifestation of some such skill thus requires repeated success but not necessarily a margin of safety manifesting that skill.

²¹ Note that this case is effectively an epistemic analogue of the kind of case offered (with regard to the debate about free will) in Frankfurt 1969.

²² As noted above, Sosa has a particular kind of terminology that he employs to characterise his view. First-order knowledge—or *animal knowledge*—is on his view *apt* belief, where an apt belief is a belief which is *accurate* (i.e., true) because *adroitly* (i.e., skilfully) formed. On this terminology, what he calls *reflective knowledge*—i.e., second-order animal knowledge—is “defensibly apt” belief, or apt belief that is “aptly noted” where this means that it is apt belief that one has an apt belief (e.g., Sosa 2007: *passim*).

²³ For example, Sosa (e.g., 2007: ch. 5) grants that on his view the subject in the barn façade case can have knowledge, even though most epistemologists deny this. His explanation of what is going on here is that the subject concerned merely has animal knowledge of the target proposition, and not also reflective knowledge. For a specific discussion of Sosa’s view on this score, see Pritchard 2009a.

²⁴ This paper was written while DHP was in receipt of a Philip Leverhulme Prize. An earlier version of this paper was presented (by JK) at a workshop on epistemology at the University of Copenhagen in November 2010, and the authors are grateful to the participants at this event for helpful feedback. Thanks also to Evan Butts and Orestis Palermos.