THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Greco on knowledge: Virtues, contexts, achievements

Citation for published version:

Pritchard, D 2008, 'Greco on knowledge: Virtues, contexts, achievements', *The Philosophical Quarterly*, vol. 58, no. 232, pp. 437-447. <https://doi.org/10.1111/j.1467-9213.2008.550.x>

Digital Object Identifier (DOI):

[10.1111/j.1467-9213.2008.550.x](https://doi.org/10.1111/j.1467-9213.2008.550.x)

Link:

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

Document Version:

Peer reviewed version

Published In:

The Philosophical Quarterly

Publisher Rights Statement:

This is the pre-peer reviewed version of the following article: Pritchard, D. (2008). Greco on knowledge: Virtues, contexts, achievements. *The Philosophical Quarterly*, 58(232), 437-447, which has been published in final form at: <http://dx.doi.org/10.1111/j.1467-9213.2008.550.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.

GRECO ON KNOWLEDGE: VIRTUES, CONTEXTS, ACHIEVEMENTS

DUNCAN PRITCHARD
University of Edinburgh

ABSTRACT. This paper critically discusses John Greco's paper, 'What's Wrong With Contextualism?', in which he outlines of a theory of knowledge which is virtue-theoretic while also being allied to a form of attributer contextualism about 'knows'.

1. VIRTUE EPISTEMOLOGY AND CONTEXTUALISM

John Greco is one of the foremost exponents of a view in the theory of knowledge known as virtue epistemology. In general, virtue epistemologists hold that at least a necessary condition for knowledge is that one attains a true belief via the stable and reliable cognitive traits that make up one's cognitive character. What counts as a cognitive trait here—and, relatedly, what counts as a 'cognitive character'—will very much depend on the details of the view, though Greco is one of the more liberal interpreters of these notions.¹ In effect, one can read Greco's version of virtue epistemology as simply a refined type of reliabilism, what has been called an 'agent' or 'virtue' reliabilism.

There is an important difference between the sort of virtue-theoretic proposal that Greco endorsed in his earlier work and the kind of proposal that he now defends. This difference is significant for our purposes since it is what leads Greco to ally his virtue epistemological proposal to a form of attributer contextualism.² In order to get a sense of

¹ In contrast, for example, to Linda Zagzebski who offers a very restrictive account of the cognitive abilities relevant to knowledge possession. See, for example, L. Zagzebski, *Virtues of the Mind* (Cambridge: Cambridge UP, 1996).

² See, for example, J. Greco, 'Agent Reliabilism', *Philosophical Perspectives*, 13 (1999), pp. 273-96 and *Putting Skeptics in Their Place* (Cambridge: Cambridge UP, 2000). Note that Greco's attributer contextualism is very different from standard varieties, such as the sort of view endorsed by Keith DeRose, David Lewis and Stewart Cohen. See K. DeRose, 'Solving the Skeptical Problem', *Philosophical Review*, 104 (1994), pp. 1-52,

what this difference involves, it is worthwhile tracing the history behind Greco's current stance.

Consider a simple form of process reliabilism which holds that knowledge is essentially just reliably-formed true belief. One of the problems facing such a view is that it is far too permissive, in that there are clearly cases of reliably-formed true belief which are not knowledge. This is where a virtue reliabilism of the type that Greco defends can come to the rescue. Whereas a simple process reliabilism would in principle hold that just about any reliable process could count as knowledge-conducive, Greco's account restricts the range of knowledge-conducive reliable processes to those that make up the cognitive traits of an agent's cognitive character. There is a solid rationale for a restriction of this sort, in that our interest in reliability surely reflects a specific interest in cognitive ability. If it is due to some incidental feature of the environment that I am able to reliably form a true belief about a certain subject matter—and in no way connected to my cognitive abilities—then the mere reliability of my belief-forming processes will not suffice for knowledge (or, indeed, any robust epistemic standing that falls short of knowledge).

Suppose, for example, that one is forming one's belief about the temperature of the room by looking at a thermometer. Suppose further that this is indeed a very reliable way of forming beliefs about this subject matter, but that its reliability is entirely unconnected with any cognitive ability of mine. For example, imagine that the thermometer is in fact broken and is fluctuating randomly within a certain range. Nevertheless, the reliability of the belief-forming process in question is preserved because there is someone hidden in the room, next to the room's thermostat, who is ensuring that every time one goes to check the reading on the thermometer, the temperature is adjusted to correspond with what it says on the thermometer.

Intuitively, even despite the reliability in play here, one does not know what one believes, and a natural explanation of this is that the reliability in question is of the wrong

D. Lewis, 'Elusive Knowledge', *Australasian Journal of Philosophy*, 74 (1995), pp. 549-67, and S. Cohen, 'Contextualism and Skepticism', *Philosophical Issues*, 10 (2000), pp. 94-107. This is because on Greco's view the context of someone other than the attributer's can be the relevant one. In this way Greco can accommodate some of the intuitions behind the subject-sensitive invariantist view defended by John Hawthorne and Jason Stanley. See J. Hawthorne, *Knowledge and Lotteries* (Oxford: Oxford UP, 2004), and J. Stanley, *Knowledge and Practical Interests* (Oxford: Oxford UP, 2005). Nevertheless, since Greco's view does allow the attributer's context to sometimes be the relevant context, I think it is acceptable to describe his view as a form of attributer contextualism, so long as one keeps this qualification in mind. Henceforth, I will describe his view in this way without qualification.

sort since it doesn't reflect any genuine cognitive ability on my part. Put another way, what we want when we are interested in reliability is to capture some sense in which, through ability, one's beliefs are sensitive to the facts whereas what we have in this case is a reliability that results when the facts are sensitive to what one believes. It should be clear that the move to a virtue reliabilist theory will deal with problems of this sort since the reliability in question in this case does not in any way reflect the cognitive character of the agent; it has nothing to do with the agent's cognitive abilities.

Interestingly, however, it was part-and-parcel of Greco's earlier view to concede that a virtue reliabilism would be unable to accommodate a range of difficult cases by itself. In particular, the thought was that one could form true beliefs through one's cognitive skill and yet the relationship between cognitive success and cognitive skill be impeded in such a way that it was a matter of luck that the belief in question was true. In such cases one would not know, even though one's true belief met the virtue-theoretic condition on knowledge laid down by Greco's virtue reliabilism.

For example, consider the case of Roddy who is looking into a field and sees a sheep-shaped object. Using his highly reliable cognitive abilities he forms a belief that there is a sheep in the field. Moreover, his belief is true. Nevertheless, he does not know that there is a sheep in the field because it is a matter of luck that his belief is true, even despite the cognitive abilities in play. This is because what Roddy is looking at is not a sheep at all, but rather a big hairy dog. However, there is, fortuitously, a sheep in the field, hidden from view behind the sheep-shaped object that Roddy is looking at.

Since the reliability in question here is clearly appropriately related to the relevant cognitive abilities of the agent, it seems that in order for virtue reliabilism to offer an adequate account of knowledge it must supplement the view with some further anti-luck condition, such as a 'safety' condition (roughly, that the belief in question could not have easily been wrong).³ In more recent work, however, Greco has seen a way around this problem that seems to be in keeping with the virtue-theoretic approach to knowledge.⁴

³ For more on safety, see E. Sosa, 'How to Defeat Opposition to Moore', *Philosophical Perspectives*, 13 (2000), pp. 141-54, and D. Pritchard, 'Resurrecting the Moorean Response to the Sceptic', *International Journal of Philosophical Studies*, 10 (2002), pp. 283-307, *Epistemic Luck* (Oxford: Oxford UP, 2005), at ch. 6, and 'Anti-Luck Epistemology', *Synthese*, 156 (2007).

⁴ See, for example, J. Greco, 'Knowledge as Credit for True Belief', in M. DePaul & L. Zagzebski (eds.), *Intellectual Virtue* (Oxford: Oxford UP, 2003), pp. 111-34, and 'The Value Problem', in A. Haddock, A. Millar & D. H. Pritchard (eds.), *The Value of Knowledge* (Oxford: Oxford UP, forthcoming).

Moreover, as we will see, it is this change in his view that has led him to ally his virtue epistemology to a version of attributer contextualism.

The change in question is to insist that knowledge is not merely the conjunction of cognitive ability and cognitive success, but is rather cognitive success that is *because of* cognitive ability. Intuitively, the addition of this relation between cognitive success and cognitive ability will deal with the case just given since we would naturally say that while Roddy is exhibiting the relevant cognitive abilities and is cognitively successful, his cognitive success is not *because of* his cognitive abilities, but is rather due to some incidental feature in the environment (in this case that there happens to be a sheep hidden from view behind the sheep-shaped object that he is looking at). If this proposal can indeed deal with cases of this sort, then it seems that a virtue reliabilism may well constitute a more complete account of knowledge.

Interestingly, there may be an additional benefit to thinking of knowledge in these terms, as Greco makes clear elsewhere, in that knowledge is now being understood in such a way as to account for why we tend to regard knowledge as distinctively valuable.⁵ This is because knowledge is now being understood in terms of cognitive success that is because of cognitive skill. As Greco argues, however, a plausible definition of an achievement is success that is because of skill, and so knowledge, on his view, turns out to be a particular type of achievement (i.e., a cognitive achievement). This is important because achievements tend to be the sort of thing that accrue a distinctive form of value that is lacking in successes that do not constitute achievements (e.g., lucky successes). If this is right, then it is unsurprising that we tend to suppose that knowledge is distinctively valuable, since we regard achievements in general as distinctively valuable and knowledge is just a type of achievement which is specifically cognitive.⁶

According to Greco, the reason why the introduction of this ‘because of’ clause into the virtue reliabilist proposal naturally leads to the adoption of a version of attributer contextualism about ‘knows’ is that to say that cognitive success is because of cognitive ability is to say that the cognitive abilities in question were the salient factors in explaining

⁵ J. Greco, ‘The Value Problem’.

⁶ For further discussion of this virtue-theoretic proposal as regards the value of knowledge, see D. Pritchard, ‘Recent Work on Epistemic Value’, *American Philosophical Quarterly*, 44 (2007), pp. 85-110, at §4, ‘Virtue Epistemology and Epistemic Luck, Revisited’, *Metaphilosophy*, 38 (2007), and ‘Knowledge, Understanding and Epistemic Value’, in A. O’Hear (ed.), *Epistemology (Royal Institute of Philosophy Lectures)* (Cambridge: Cambridge UP, *forthcoming*).

the cognitive success. Put more simply, to say that cognitive success is because of cognitive ability is to say that the cognitive abilities in question *explain* the cognitive success. Moreover, Greco wants to argue that this isn't a point that is confined purely to epistemology, but reflects a wider fact applicable to achievements more generally. What determines that one's success is because of one's ability, and so is a *bona fide* achievement, is that one's ability explains one's success. Crucially, however, Greco also argues that "the semantics of causal explanation language requires a contextualist treatment."⁷ If this is right, then the connection between virtue reliabilism and attributer contextualism starts to become apparent.

Here is the example that Greco gives to illustrate this point, which is worth quoting in full:

Consider [...] the cause of a car crash at a major intersection. The police at the scene deem that the crash was caused by excessive speed, and accordingly they issue the driver a summons. Later in the year, city planners consider the crash along with several others that occurred at the same location. They determine that the crash was caused by difficult traffic patterns and recommend changes to the Board of Transportation. Who is right, the police or the planners?

One wants to say, of course, that they are both right. To say that excessive speed caused the crash is to pick out one necessary part of a much broader set of causal conditions. To say that difficult traffic patterns caused the crash is to do the same. Plausibly, the different interests and purposes governing traffic enforcement and traffic planning make different explanations appropriate. City police, in their capacity, appropriately blame the crash on excessive speed, considering the difficult traffic patterns to be part of the normal background. City planners, in their capacity, appropriately blame the crash on difficult traffic patterns, considering speeding drivers to be part of the normal background. What counts as "the explanation" or "the cause" of the crash is relative to these different contexts.⁸

With this point in mind, it should be clear that insofar as knowledge is to be understood in such a fashion that it essentially appeals to the idea that the cognitive success in question be explained by the relevant cognitive abilities, then it will inherit the kind of context-sensitivity exhibited here. Or, as Greco puts the point, "we get the result that knowledge attributions are sensitive to context because they involve causal explanations, and causal explanations are sensitive to context."⁹

The example that Greco offers to illustrate the kind of context-sensitivity in knowledge ascriptions that he has in mind is that of the gambler. In the first case, the gambler says to his wife that he knew that a certain horse would win (a horse that he bet a great deal of money on) and is regarded with scepticism (i.e., as if he is speaking falsely). In the second

⁷ J. Greco, 'What's Wrong With Contextualism?', *Philosophical Quarterly* [this issue], at p. 5?

⁸ *Ibid.* [?]

⁹ *Ibid.* [?]

case, the gambler asserts the same sentence in the company of his friends and is treated as speaking truly. Greco's diagnosis of what is going on here is that in the first case the risk of losing a large sum of money is what is salient, thereby undermining the salience of the gambler's abilities in picking winning horses. In this context, his success is not *because of* his abilities. Accordingly, it does not constitute an achievement, and so isn't knowledge. Hence, his claim to know is regarded as false. In the second case, in contrast, only the gambler's abilities are treated as salient, and not the inherent riskiness of the bet, and so his success is because of his abilities. As a result, his success does constitute an achievement and thus his knowledge claim is rightly treated as true.

We thus get a form of attributer contextualism about 'knows', in the sense that whether or not a knowledge ascription expresses a truth can depend upon the context of the person making the ascription. If the gambler's wife is talking about her husband and says 'He didn't know he would win', she speaks truly, while if the gambler's friends are talking about the gambler and say 'He did know he would win', they would also speak truly. In this sense, then, 'knows' is a context-sensitive term.

Notice, however, that the contextualism in question is tempered by the additional story that Greco tells about the core function of our concept of knowledge, which is "to flag information and sources of information for use in practical reasoning."¹⁰ This tempers the form of contextualism in play because if knowledge is required for practical reasoning, then it is essential that knowledge be widely available, and so the standards for knowledge must not be too high. Equally, however, if informants are to be of any use to us, then the knowledge they convey must be suitably robust. Hence, the standards for knowledge can never be too low either. So although Greco allies his virtue epistemology to a form of attributer contextualism, the form of contextualism in play is rather modest.

So if Greco is right, then virtue reliabilism, properly understood, naturally leads to a

¹⁰ *Ibid.*, at p. 22. [?]

modest form of attributer contextualism. I will raise three worries regarding Greco's proposal. The first is ultimately just a quibble regarding the key example he uses to illustrate his view, that of the gambler. The second is effectively a request for further argumentation to explain why Greco's line of argument leads to attributer contextualism and not a form of epistemological contrastivism. Finally, the third worry I will raise is a more substantive concern regarding Greco's identification of knowledge with cognitive achievement, a thesis which is central to his argument.

2. THE 'GAMBLER' EXAMPLE

First off, let me just state that I don't like the gambler example that Greco offers one bit. For one thing, notice that the example is crucially underdescribed, since it leaves it open whether the gambler really does have the relevant ability to pick winners. We tend to think of gamblers—especially ones, as in this case, who bet large sums that they can ill-afford to lose—as rather self-deluded figures who credit themselves with gambling abilities they just do not have. If this is how we are to understand the case, however, then there is surely *no* context in which what is salient could ensure that the gambler's self-ascription of knowledge is properly treated as true. On this reading of the example, the gambler's friends are wrongly treating his self-ascription as true, rather than merely having a different conception of what is salient.

Alternatively, we can imagine a way of reading this case such that the gambler bucks the usual stereotype and genuinely has the relevant skill to pick winners. I think this is the only credible reading of the example. It is now plausible to suppose that the gambler's friends are right to treat his self-ascription of knowledge as expressing a truth given that they do not find the risk of him wrecking the household finances as salient. But what do we now say about his wife's unwillingness to treat his self-ascription as true? If this is due to her ignorance of her husband's *bona fide* ability to pick winners, then it is surely unfounded.

Perhaps, though, the gambler's wife knows that her husband has these abilities, but she simply worries about the risks involved regardless and so treats the possibility that he might have wrecked their finances as salient (I think Greco may well have this reading in mind, judging by what he says in the text). Accordingly, on this reading although there is in

fact no serious risk that the husband may have wrecked their finances with his betting—indeed, given his track record, we would in fact reasonably expect him to substantially *improve* them with his betting—the fact that the gambler’s wife is an extremely cautious individual suffices to make the risk salient such that any assertion she made to the effect that her husband did not know that he would win would be true.

My concern about this reading of the example is that it seems to imply that there is no distinction being drawn between what the agent thinks is salient and what the agent properly or reasonably thinks is salient. If one followed through this reasoning, then it would seem to suggest that the *overly* cautious amongst us—i.e., the *unreasonably* cautious—could go around asserting that they and others do not know pretty much any proposition with impunity. But surely this is not a consequence that anyone would wish to sign up to.

Notice that I am not claiming here that Greco’s form of attributer contextualism is implausible, since my problem is specific to this particular case. My suspicion is that this is simply a bad example for Greco to use to illustrate his view. Indeed, given the further thesis that Greco offers regarding the main function of the concept of knowledge and how this function ensures that the attributer contextualism on offer will be mild in its consequences, it strikes me that even Greco ought to want to reject an example like this, since it implies a radicalism that is not part of his wider view. That is, given the core function of the concept of knowledge that Greco identifies, why would the standards for knowledge go so high as to allow the gambler’s wife’s assertion that he does not know to count as true? Alternatively, if the gambler really does lack the abilities in question, then why would the standards for knowledge drop so low as to allow the gambler’s friends’ assertions that he does know count as true? Greco would thus be wise to find another example to illustrate his view, one more congenial to his broader outlook.

3. KNOWLEDGE AND CONTRASTS

A second worry I have regarding Greco’s paper is a little more substantive. The concern is that it is unclear, even by Greco’s own lights, that the argument that he offers which ties virtue epistemology to attributer contextualism really goes through. The reason for this is that it is very plausible to suppose that causal explanations are best understood along contrastive

lines, and thus that if we are to evaluate knowledge ascriptions on the same model as causal explanations, then it appears that a *contrastivist* view, rather than an attributer contextualist one, is what is licensed.

In order to see this, think again about the example of the car crash that we saw Greco offering above. To begin with, it is worth noting that the example is, as it stands, far from uncontentious. For instance, it is hard to imagine why a traffic cop, in offering a causal explanation of a crash, would decide to ignore the vital fact that the junction can be dangerous within the range of speeds normally expected. Similarly, why would a traffic planner ignore the fact that the car involved was speeding in offering their causal explanation of the crash? (And note that if they weren't ignoring the relevant fact, but simply unaware of it, then that hardly helps matters. Why think that the causal explanation they offer is the right one if they are unapprised of all the relevant facts?).

Rather than try to rescue the example (which I'm sure can be done), let us focus on a less controversial case from the literature. Consider the following example, adapted from one described by Steven Rieber.¹¹ Only those people who have syphilis get paresis, but some people with syphilis do not get paresis. Now consider the question of what caused a particular person, Smith, to get paresis. Suppose one is in a context in which it is part of the background that one would not ordinarily expect someone to have syphilis. In such a context, it would seem natural to say that it is the fact that Smith has syphilis that explains why he has paresis. In contrast, however, suppose that one is in a context in which this is not part of the background—perhaps one is working in a special section of a hospital which deals with patients with syphilis and talking to another member of the medical staff. It would *not* now be natural to say that the fact that Smith has syphilis explains why he has paresis. We thus get a kind of context-sensitivity of the general sort that Greco is interested in, in that an assertion of the very same causal explanation sentence will be—rightly, it seems—judged to be true in one context and false in another.

Interestingly, however, many commentators on causal explanations do not conclude on this basis that we ought to adopt the relevant version of attributer contextualism—i.e., that the context of the agent uttering the causal explanation sentence determines what truth-value this sentence should have. Instead, they opt for a contrastivist treatment of causal

¹¹ S. Rieber, 'Skepticism and Contrastive Explanation', *Noûs*, 32 (1998), pp. 189-204, at p. 195.

explanations. As regards the example just given, for example, Rieber argues that what counts as the right causal explanation all depends on which contrast is salient. For example, if the salient contrast is with Jones, who does not have syphilis, then it would be natural to say that it is the fact that Smith has syphilis that explains why he has paresis. Alternatively, if the salient contrast is with Brown, who has syphilis but not paresis, then this explanation will be false.¹²

In effect, what a contrastivist account of causal explanation does is treat causal explanation sentences as elliptical, with the relevant contrast implicit. So, for example, it is true that the fact that Smith has syphilis explains why he *rather than someone who doesn't have syphilis* has paresis, but it is also true that the fact that Smith has syphilis does not explain why Smith *rather than someone who does have syphilis but not paresis* has paresis. What is crucial to such an account is that it leads to a very different kind of context-sensitivity to that at issue in attributer contextualism. *Whatever* context one is in, an assertion of the sentence 'Smith has paresis because he has syphilis' is true if the contrast is someone who doesn't have syphilis, and false if the contrast is someone who does have syphilis. Put another way, once we make the relevant contrast explicit, then there is no further context-sensitivity at issue.

The problem for Greco, however, is that if one extends this contrastivist treatment of causal explanation into the epistemological sphere then what one ends up with, it seems, is not attributer contextualism about 'knows' but rather an epistemological version of contrastivism. Given the problems noted above with the 'gambler' case that Greco offers, consider instead Fred Dretske's 'zebra' case to illustrate this point.¹³ That the agent in this example is able to tell zebras from, say, baboons or giraffes may well explain why she has a true belief that what she sees is a zebra *rather than* the salient contrasts of baboon or giraffe. In contrast, the fact that the agent has discriminatory abilities of this sort will not explain why she has a true belief that what she sees is a zebra *rather than* a cleverly disguised mule. As a number of epistemological contrastivists have noted, there is no need to opt for attributer contextualism once one has gone this far down the contrastivist road.¹⁴ Instead, one just needs to recognise that there is no such thing as knowing *p simpliciter*; rather, knowing is

¹² For another prominent example of a contrastivist treatment of causal explanation, see P. Lipton, 'Contrastive Explanation', in D. Knowles (ed.), *Explanation and its Limits* (Cambridge: Cambridge UP, 1990).

¹³ F. Dretske, 'Epistemic Operators', *Journal of Philosophy*, 67 (1970), pp. 1007-23.

¹⁴ See, for example, J. Schaffer, 'Contrastive Knowledge', *Oxford Studies in Philosophy*, 1 (2005), pp. 235-72.

always to be understood contrastively in terms of knowing that *p* rather than a set of contrasts *Q*. Relative to one set of contrasts one might be truly said to know *p*—meaning only that one knows *p* relative to that set of contrasts—whereas relative to another set of contrasts one might not be truly said to know *p*. Attributer contextualism doesn't come into it.

I am not meaning to endorse epistemic contrastivism here. My point is only that given the line of argument offered by Greco, it is odd that he thinks that the *natural* outcome of that argument is attributer contextualism and not contrastivism, especially when he never considers the contrastivist alternative.

4. KNOWLEDGE AND ACHIEVEMENT

I noted earlier that a key feature of the virtue-theoretical proposal that Greco offers which essentially appeals to causal explanations is that it in effect understands knowledge as a kind of achievement. That is, achievements are to be understood as successes that are *because of* ability, and knowledge is to be understood as a cognitive achievement—i.e., a cognitive success that is because of cognitive ability. This idea has a lot of *prima facie* appeal, but I do not think that it stands up to scrutiny, even given a contextualist semantics for the relevant causal explanations that are in play in knowledge ascriptions. The reason for this is that there seem to be cases of knowledge which clearly aren't (by anyone's lights) cognitive achievements, and also cases of cognitive achievements which clearly aren't (by anyone's lights) cases of knowledge.

To appreciate this point, let us first reflect on what is involved in non-cognitive achievements. Greco is surely right that we should think of achievements as, essentially at least, successes that are because of ability. Suppose that I select a target, fire at that target with my arrow with the intention of hitting it, and hit the target. If I lack the relevant abilities, then this success is just dumb luck, and so not an achievement. Equally, even if I have the relevant abilities, but the success is 'Gettierized'—perhaps because a dog jumps up and grabs the arrow in flight (which would have hit the target), and then proceeds to place the arrow in the target—then this doesn't count as an achievement either since the success clearly isn't *because of* the relevant abilities.

But now consider the following case. Suppose that as before I select a target, fire at

that target with my arrow with the intention of hitting it, and hit the target. Moreover, let us stipulate that I have all the relevant abilities, and that my success is not ‘Gettierized’, at least not in the usual fashion—that is, nothing interferes with the arrow during its journey to the target. Would this count as an achievement? Surely we would usually say that it did. Now suppose further that there is a twist to this case. The target I (randomly) selected is, unbeknownst to me, the only target on that range that is not fitted with a hidden forcefield which would repel any arrow aimed at it. My success is thus clearly lucky, in the sense that I could very easily have not been successful in this case. But is it any less of an achievement on this score? I suggest that it isn’t.

The moral of such cases is that achievements are entirely consistent with a certain kind of lucky success. Interestingly, however, knowledge—which Greco regards as a cognitive achievement—is not consistent with lucky success of this sort. In order to see this, one needs only to note that the barn façade case has almost exactly the same structure as the ‘archer’ case just described. Accordingly, by parity of reasoning, one should, it seems, regard this as a case of cognitive achievement—it is certainly cognitive success through cognitive ability that is not ‘Gettierized’ in the normal way. Indeed, wouldn’t we naturally say that the cognitive success in this case is *because of* the agent’s cognitive abilities? The problem is, however, that while this cognitive success does seem to count as a *bona fide* cognitive achievement, it doesn’t seem to be a case of knowledge because of the epistemic luck involved—the agent could very easily have been wrong. More specifically, it is hard to imagine a context in which one would ascribe knowledge to the agent in the barn façade case and speak truly. Hence it seems to be that knowledge is intolerant to luck in a way that achievements are not, and this creates a *prima facie* problem for any proponent of the view that knowledge is a kind of achievement.¹⁵

Moreover, there are also problems for Greco’s view from the opposite direction, since there seem to be cases of knowledge which are not cases of cognitive achievements. Consider, for example, the following case offered by Jennifer Lackey.¹⁶ In this example our hero, let’s call her ‘Jenny’, steps off the train in an unfamiliar city seeking directions to a

¹⁵ I explore the problem posed by barn façade-style cases for Greco’s view further elsewhere. See, for example, ‘Virtue Epistemology and Epistemic Luck, Revisited’ and ‘Knowledge, Understanding and Epistemic Value’. See also J. Kvanvig, ‘Responses to Critics’, in A. Haddock, A. Millar, & D. H. Pritchard (eds.), *The Value of Knowledge*, (Oxford: Oxford UP, *forthcoming*).

¹⁶ J. Lackey, ‘Why We Don’t Deserve Credit for Everything We Know’, *Synthese* 156 (2007), at §2.

local landmark. Walking up to the first adult passer-by that she sees, she asks for directions. The informant is indeed knowledgeable in this regard and conveys this knowledge to Jenny, who then locates the landmark in question. In cases like this we surely want to say that Jenny *knows* the way to the landmark. Clearly it is of some credit to Jenny that she is cognitively successful in this case—she wouldn't ask *anyone* after all, and she wouldn't believe *anything* that she is told—but would we really say that Jenny's cognitive success is a cognitive achievement of *hers*, as opposed, for example, to being a cognitive achievement on the part of the informant? Moreover, is there any context in which we would rightly regard Jenny's cognitive success as being *because of* her cognitive abilities, rather than, at least in substantial part, the cognitive abilities of the informant in question?

Now I don't doubt that there are moves that Greco can make to cases like this. As regards the Lackey case, for example, he could dig his heels in and insist that insofar as this is *bona fide* knowledge, then it must be because of Jenny's cognitive abilities that her belief is true. Similar moves could be made with the barn façade case, perhaps on the grounds that to say that there is luck involved in the cognitive success is *thereby* to say that the cognitive success is not because of cognitive ability. Alternatively, Greco could contend that the agent doesn't really have the cognitive ability in this case on account of how abilities are to be understood relative to environments, and this just is the wrong environment on this score. Or, finally, Greco could contend that there is something special about the cognitive achievement at issue when we consider knowledge—perhaps because of the core function of the concept of knowledge that we have discussed—which ensures that it is more resistant to luck than other types of achievement.¹⁷

I have my doubts about all these moves, but this is not the place to explore them.¹⁸ Rather I want to suggest that these problems should persuade Greco to re-evaluate the merits of the earlier view that he abandoned, the view on which knowledge has both a virtue-

¹⁷ Greco in fact makes *all* of these moves at some point or other (which is not to suggest that these moves are not complementary, rather than competing). At a number of junctures in 'What's Wrong With Contextualism?', for example, he argues that achievement simply excludes luck, the implication presumably being that one can't make sense of a 'lucky' achievement in this case (presumably, then, Greco also wants to hold that the archer's success in the corresponding case does not constitute a *bona fide* achievement either). Moreover, in 'The Value Problem' he argues that the agent in the barn façade case lacks the relevant cognitive ability. Finally, in an unpublished manuscript entitled 'The Nature of Ability and the Purpose of Knowledge', he explicitly offers the argument that there is something special about the cognitive achievement involved when one knows, such that it is more resistant to luck than other cognitive achievements.

¹⁸ I discuss them further in a number of places. See, for example, 'Virtue Epistemology and Epistemic Luck, Revisited' and 'Knowledge, Understanding and Epistemic Value'.

theoretic component to accommodate the intuition that the reliability in question must reflect the agent's cognitive abilities, and also an anti-luck component that is designed to deal with the problem posed by knowledge-undermining epistemic luck (and, thereby, the Gettier-style cases).

Recall that the key motivation for moving to the later view was that it could deal with cases of knowledge-undermining epistemic luck without having to appeal to a separate anti-luck condition, since cognitive achievements were (or so the thought went) of their nature luck-excluding in the relevant sense. The trouble is, however, that the moral of the cases just described seems to be that *both* conditions are required for knowledge. In barn façade-style cases, forming one's belief via one's cognitive abilities, even to the extent that it qualifies as a cognitive achievement, does not suffice for knowledge because of the malignant epistemic luck in play (so the virtue-theoretic condition on the later view does not do all the anti-luck work that is required of it after all). Conversely, in Lackey-style cases, knowledge is possessed even in the absence of a cognitive achievement since it involves a belief that is formed through ability in a non-lucky fashion. Thus, the earlier view has, it seems, the resources to meet the problems facing the later view that Greco defends.

Of course, if Greco returns to the earlier view, and so abandons the knowledge-as-cognitive-achievement thesis, then the appeal to knowledge being true belief that is because of cognitive ability will be lost too. And if that goes, then so too does Greco's argument, at least as it stands, that connects virtue epistemology with attributer contextualism, since causal explanations will no longer be playing the required role. A fundamental issue facing Greco's argument for contextualism is thus whether he can adequately motivate the adoption of the later view over the former view, and to deal with this problem he will need to do more to convince us of the knowledge-as-cognitive-achievement thesis. Perhaps Greco will be able to make the compelling case that is required in this regard, but until he does, I think the jury will be out on whether virtue epistemology should be allied with attributer contextualism, at least in the way (and on the grounds) that Greco suggests.¹⁹

¹⁹ An earlier version of this paper was presented at the 'Moral Contextualism' conference at the University of Aberdeen in July 2006, and I am grateful to the audience at this event. Thanks also go to Peter Baumann, J. Adam Carter, John Greco, Kent Hurtig and an anonymous referee for *The Philosophical Quarterly*. This paper was written whilst in receipt of an Arts and Humanities Research Council research leave award.