The Value of Knowledge

By Duncan Pritchard

0. Introduction

IT IS WIDELY HELD THAT KNOWLEDGE IS OF DISTINCTIVE VALUE. PRESUMABLY, THIS IS THE REASON why knowledge—and not, say, justified true belief—has been the principal focus of generations of epistemological theorizing. Understanding just why knowledge is distinctively valuable, however, has proved elusive, and this has led some to question whether it is distinctively valuable at all. Call this the *value problem*.

In this paper I will be arguing that knowledge is not, in fact, of distinctive value, and hence I will be giving a negative response to the value problem. As I show, however, understanding just why knowledge is not distinctively valuable has important ramifications for how we think about the very nature of knowledge.

In §1, I unpack what is involved in the value problem. I argue that an adequate response to the value problem needs to be able to offer an explanation of why knowledge, unlike any epistemic standing which falls short of knowledge, is finally valuable. In §2, I show how a certain virtue-theoretic account of knowledge—one that understands knowledge as a cognitive achievement—can respond to the value problem by appeal to the putative final value of cognitive achievements. In §3 and §4, I argue that this thesis is on closer inspection untenable, in that an agent can exhibit a cognitive achievement while failing to have the corresponding knowledge, and can conversely have knowledge while failing to exhibit the corresponding cognitive achievement. I then diagnose, in §5, just why this conception of knowledge fails. In particular, I claim that the moral to be drawn from this discussion is that any adequate theory of knowledge needs to give equal weight to two master intuitions about knowledge—what I call the *ability* and *anti-luck* intuitions—and hence must incorporate both an anti-luck condition and an ability condition. In §6, I apply this moral to the value problem and argue that it is not knowledge which is distinctively valuable, but rather cognitive achievement, an epistemic category which is distinct from knowledge. Finally, in §7, I offer a diagnosis of the source of the intuition that knowledge is distinctively valuable.

Duncan Pritchard joined the University of Edinburgh's philosophy department in 2007 as the Chair in Epistemology. His work covers a range of questions in epistemology, including testimony, the history of skepticism, and virtue epistemology. These are addressed in his book Epistemic Luck *(Oxford University Press, 2005) as well as* The Nature and Value of Knowledge: Three Investigations *(Oxford University Press, 2010), co-written with Adrian Haddock and Alan Millar. He is the editor-in-chief of the soon-to-be launched* Oxford Bibliographies Online: Philosophy *(spring 2010). He is presently writing a book on epistemological disjunctivism, to be published by Oxford University Press in conjunction with the Aristotelian Society, as part of their* Lines of Thought *series.*

1. Unpacking the Value Problem

PART OF THE DIFFICULTY POSED BY THE VALUE PROBLEM INVOLVES GETTING CLEAR ABOUT just what it means to say that knowledge is distinctively valuable. One minimal reading of this claim is that knowledge is more valuable than mere true belief.[1] Call the challenge to explain why knowledge is more valuable than mere true belief the *primary value problem*.

On the face of it, there is a very straightforward answer to the primary value problem—namely, that knowledge is more valuable than mere true belief because it tends to be of greater practical value. Of course, there may be particular propositions which, for some special reason, one would prefer to merely truly believe rather than know (perhaps in knowing them one would incur a penalty which one would not incur if one merely true believed them), but in general you are more likely to achieve your goals with knowledge than with mere true belief.

Indeed, this was precisely the way that Socrates answers the primary value problem in the *Meno*. Why should you prefer knowledge of the correct way to Larissa rather than mere true belief, given that both will, on the face of it, ensure that you get to your destination? Socrates' answer is that knowledge has a 'stability' which mere true belief lacks. Mere true belief, argues Socrates, is like one of the untethered statues of Daedalus in that it is liable to be lost in the way that knowledge, like one of those statues tethered, will not. For while a mere true belief may well enable you to achieve your goals as well as knowledge, one will be far more insulated from failure by possessing knowledge. Suppose, for example, that the road to Larissa takes an unexpected course. Someone with mere true belief—where the belief is based on just a hunch, say—may well at this point lose all faith that she is on the right tracks and turn back. Someone who knows that this is the right way to go, however—perhaps because she consulted a reliable map before departure—will not be so shaken by this turn of events.[2,3]

Even supposing that we are able to respond to the primary value problem in this way, however, there would still be more to do to secure our intuition that knowledge is distinctively valuable. At the very least, we would need to answer the *secondary value problem* of explaining why knowledge is more valuable than any proper subset of its parts.

In order to see this, suppose that one answered the primary value problem by, for example, pointing to a necessary condition for knowledge which in general added practical value (the justification condition, say), but suppose further that the satisfaction of this condition, in conjunction with true belief, was not sufficient for knowledge. Perhaps, for example, when one knows that p, it is the fact that one's belief that p is thereby justified that ensures that knowledge has a greater practical value than mere true belief that p alone. One would thereby have answered the primary value problem while leaving the secondary value problem unanswered. Moreover, let us take it as given that there is no further feature of knowledge which is value-conferring, such that the secondary value problem is regarded not just as unanswered, but as unanswerable.

On the face of it, this lacuna might not seem that problematic, since just so long as one can show that knowledge is more valuable than mere true belief, then that would seem to satisfy our intuition that knowledge is of some special value to us (on this view it is, after all, the kind of thing that we should prefer to

mere true belief, all other things being equal). The problem, however, is that if the distinctive value of knowledge is due to some feature of knowledge which, with true belief, falls short of knowledge, then it seems that what we should seek is not knowledge as such, but rather that which falls short of knowledge (that is, true belief plus the value-conferring property X, in this case justification). But if that is right, then why do we regard knowledge as distinctively valuable at all?[4]

The primary value problem thus naturally leads to the secondary value problem, and it seems that both will need to be answered if we are to account for the distinctive value of knowledge. Even if we can offer a response to the secondary value problem, however, it is still not clear that we have accounted for the distinctive value of knowledge. This is because the secondary value problem leaves open the possibility that the difference of value at issue is merely one of degree rather than kind. To say that knowledge is of *distinctive* value, however, appears to suggest that the difference in value between knowledge and that which falls short of knowledge is not just a matter of degree, but of *kind*. After all, if one regards knowledge as being more valuable than that which falls short of knowledge merely as a matter of degree rather than kind, then this has the effect of putting knowledge on a kind of continuum of epistemic value, albeit further up the continuum than anything that falls short of knowledge. The problem with this 'continuum' account of the value of knowledge, however, it that it fails to explain why the long history of epistemological discussion has focused specifically on the stage in this continuum of epistemic value that knowledge marks rather than some other stage (such as a stage just before the one marked out by knowledge, or just after). Accordingly, it seems that accounting for the value of knowledge requires us to offer an explanation of why knowledge has not just a greater degree, but also a different kind of value than whatever falls short of knowledge. Call this the *tertiary value problem*.

Further support for the tertiary value problem comes from the fact that we often treat knowledge as being somehow *precious*, in the sense that its value is not merely a function of its practical import. Imagine that one were faced with the choice of either knowing a yet to be determined proposition or having an epistemic standing as regards that proposition that fell short of knowledge. Furthermore, suppose it was stipulated that there would be no practical costs or benefits whichever choice one opted for. Even so, is not the right thing to do to choose knowledge? But if that is correct, then the value of knowledge cannot simply be due to its greater practical value. Instead knowledge must be the kind of thing that, unlike that which falls short of knowledge, is valuable for its own sake: it must be non-instrumentally—that is, *finally*—valuable.[5]

Most of those who have explored the issue of the value of knowledge have tended to focus their attentions on the primary value problem, to the exclusion of the other two problems. There is a good rationale for a focus of this sort, since if one is unable to answer the primary value problem then, a fortiori, one will be unable to answer the secondary and tertiary problems as well. This rationale can be turned on its head, however, since it equally follows that if one could offer a response to the tertiary value problem, then one would thereby be able to deal with the primary and secondary value problems as well. This is precisely the possibility that I will be exploring here.[6]

2. Knowledge and Final Value

SO WHAT WE ARE LOOKING FOR IS AN EXPLANATION OF WHY KNOWLEDGE IS MORE VALUABLE, not only as a matter of degree but also as a matter of kind, than any epistemic standing that falls short of knowledge. This is certainly a tall order. As we just noted, offering such an explanation requires us to explain why knowledge, unlike that which falls short of knowledge, has final value.

As it happens, there is indeed an account of knowledge in the contemporary literature which seems able to explain why knowledge, unlike any epistemic standing which falls short of knowledge, is finally valuable (indeed, it is the *only* epistemological proposal in the literature which has the potential to offer such an explanation). This is the *robust* virtue-theoretic account of knowledge that is offered by, for example, Ernest Sosa (1988; 1991; 2007), Linda Zagzebski (1996; 1999) and John Greco (for example, 2002; 2007; 2009; *forthcoming*). What makes such a virtue-theoretic proposal robust is the fact that it attempts to exclusively analyze knowledge in terms of a true belief that is the product of epistemically virtuous belief-forming process.

Depending on what one builds into one's conception of an epistemic virtue, then one will be led to adopt a very different kind of virtue epistemology. As we will see in a moment, however, we do not need to worry about the specifics of different robust virtue-theoretic accounts since what is salient for our purposes is simply the *structure* of these proposals.

On the face of it, robust virtue epistemology does not look particularly promising because of the difficulty of specifying the virtue-theoretic condition on knowledge in such a way as to deal with the problem of knowledge-undermining epistemic luck—for example, of the sort found in Gettier-style cases. After all, no matter how reliable an epistemic virtue might be, it seems possible that it could generate a belief which is only true as a matter of luck.

Consider, for example, the case of Roddy.[7] Using his highly reliable cognitive faculties, Roddy the shepherd forms a true belief that there is a sheep in a field that he is looking at. Unbeknownst to Roddy, however, the item that he is looking at in the field is not a sheep at all, but rather a sheep-shaped object—a rock, say—albeit one which is obscuring from view a genuine sheep that is hidden behind (and which ensures that his belief is true). Here, then, we appear to have a true belief that is the product of the agent's epistemic virtue and yet which does not qualify as knowledge because of the presence of knowledge-undermining epistemic luck. With this problem in mind, one might naturally be tempted to opt for a *modest* virtue epistemology, one that does not try to completely analyze knowledge in terms of a virtue-theoretic condition, but which is willing to endorse in addition a further codicil that can deal with Gettier-style cases.[8]

In contrast, robust virtue epistemology attempts to get around this problem by, in effect, 'beefing-up' the virtue-theoretic demand on knowledge. Rather than allowing that knowledge is merely true belief that arises out of the agent's epistemic virtue—which, as we have just seen, is compatible with Gettier-style cases—the strengthened virtue-theoretic thesis is that knowledge only results when the agent's true belief is *because of* the operation of her cognitive abilities—that is, where it is *primarily* creditable to the agent that her belief is true.[9]

This strengthened proposal certainly deals with the case of Roddy, since

while his true belief is indeed produced by his cognitive abilities, it is not the case that his belief is true *because of* the operation of his cognitive abilities. Instead, his belief is true because of a helpful quirk of the environment—that there happened to be a sheep behind the sheep-shaped object that he was looking at. In contrast, had he actually been looking at a sheep (in normal circumstances), then his belief *would* have been true because of the operation of his cognitive abilities.

As Greco (2009) points out, a further advantage of understanding knowledge along these strengthened virtue-theoretic lines is that it seems to capture the idea of knowledge as being a kind of cognitive achievement: That is, we might broadly think of achievements as being successes that are attained because of one's ability, and virtue epistemology seems to be offering the epistemic analogue of this claim—on this view, knowledge is cognitive success that is because of one's cognitive ability.

In order to see the plausibility in this general account of achievement, consider the following case. Suppose that Archie the archer selects a target at random and uses his bow to fire an arrow at that target with the intention of hitting it. Suppose further that he does indeed hit the target. If, however, the success in question is purely a matter of luck—if, for example, Archie does not possess the relevant archery abilities—then we would say that this success is not an achievement on Archie's part. Similarly, even if Archie has the relevant archery abilities, and is in addition successful in hitting the target, we still would not count his success as an achievement if the success was not *because of* Archie's archery abilities. This is important because of the possibility that the success in question is 'Gettierized'. If, for example, a dog ran on to the range, grabbed the arrow (which was heading towards the target) in mid-flight, and proceeded to deposit it on the target, then we would not regard this successful outcome as Archie's achievement, even if the original firing of the arrow had been highly skilful. Instead, what is required for an achievement is that Archie's hitting of the target is *because of* the exercise of his relevant archery abilities.

That achievements are best understood this way, and that knowledge on the present proposal appears to be a kind of achievement, is important for our purposes because achievements are, plausibly, distinctively valuable. More specifically, it is plausible to hold that the kind of successes that count as achievements are valuable for their own sake because of how they are produced (that is, they are finally valuable because of their relational properties). If this is right, and we can show that knowledge (unlike that which falls short of knowledge) is a type of achievement, then we may be in a position to thereby show that knowledge has a kind of value which that which falls short of knowledge lacks, and hence show that it is distinctively valuable.[10]

In order to see why achievements are finally valuable, suppose now that Archie—in the manner of Robin Hood—is trying to escape from an adversary, and the target he is firing at is a mechanism which will drop the drawbridge in front him, thereby ensuring that he gets to safety. From a practical point of view, it may not matter whether the hitting of the target was because of Archie's archery abilities or through dumb luck (for example, by a lucky deflection). Either way, it still results in the dropping of the drawbridge, thereby enabling Archie to escape. Nevertheless, we would value Archie's success very differently if it were

the product of luck (even when the relevant ability is involved, but the success in question is 'Gettierized'), rather than being because of his ability such that it is an achievement. In particular, we would regard Archie's achievement of hitting the target through ability as, in this respect, a good thing in its own right, regardless of what other instrumental value it may accrue. And what goes here for Archie's achievement of hitting the target seems to be equally applicable to achievements more generally: achievements are finally valuable.

One worry that one might have about the idea that achievements are finally valuable is that some achievements seem to have very little value—perhaps even a negative value—because, for example, they are too trivial or just plain wicked. Are even achievements of this sort of final value? Note, however, that the claim is only that achievements have final value qua achievements. This is entirely consistent with the undeniable truth that some achievements may have no practical value, and may even accrue *dis*value, perhaps because of the opportunity cost incurred by seeking the trivial achievement over a more substantive achievement, or because of the wicked nature of the achievement in question. Indeed, there may well be situations in which the all-things-considered value of Archie's success of hitting the target when it is due to luck is much greater than the all-things-considered value of a corresponding success attained because of Archie's ability. It is important to recognize that the thesis that achievements qua achievements accrue final value is entirely consistent with this possibility.[11]

This point is also important when it comes to understanding the way in which this thesis that knowledge, qua cognitive achievement, accrues final value can contribute to an answer to the tertiary value problem. In particular, we need to note that the mere fact that knowledge (unlike that which falls short of knowledge) is, qua cognitive achievement, of final value will not necessarily be enough to resolve the tertiary value problem. This is because of the possibility that that which falls short of knowledge is generally of greater non-final value than knowledge. If this were so, then it could still be true that knowledge is generally of less all-things-considered value than that which falls short of knowledge, even granting the fact that knowledge, in contrast to that which falls short of knowledge, is finally valuable. Nevertheless, it is plausible to suppose that knowledge is not generally of *less* instrumental value than that which falls short of knowledge. And with this assumption in play, the final value of knowledge would ensure that the tertiary value problem is met and, with it, the primary and secondary value problems too. In what follows we will let this assumption stand.

On the face of it, then, we have a plausible account of knowledge according to which knowledge is a type of achievement and which, as such, accrues final value. If this is right, then—at least given the plausible assumption just noted—we seem to have a rather straightforward answer to the tertiary value problem and, with it, the value problem more generally. On this view it is no surprise that we regard knowledge as distinctively valuable because knowledge, unlike that which falls short of knowledge, is deserving of a special kind of value.

3. Achievement Without Knowledge

Despite the surface appeal of this proposal, it faces several fundamental problems. The difficulties lie, however, not with the claim that cognitive achievements accrue

final value—a thesis which I have argued is very plausible once it is understood correctly—but rather with the further thesis that knowledge is to be understood as a cognitive achievement. Without this further claim, however, the import of this virtue epistemological proposal to the value problem is moot.

Consider again the case of Archie, who selects a target at random from a target range and then successfully fires an arrow at that target. We noted above that if Archie lacks any kind of archery skill, such that his success was entirely lucky, then we would not count his success as being an achievement. Similarly, even if Archie has plenty of skill at archery but his success is 'Gettierized'—such that it is not *because of* his skill—then we would not count it as an achievement. So far so good. But now consider a third case in which Archie again selects a target at random, skillfully fires at this target, and successfully hits it because of his skill. On the account of achievement on the table, his hitting of the target is a genuine achievement. Suppose, however, that unbeknownst to Archie, there is a force field around each of the other targets such that, had he aimed at one of these targets, he would have missed. It is thus a matter of luck that he is successful, in the sense that he could very easily have not been successful. Notice, however, that luck of this sort does not seem to undermine the thesis that Archie's success is a genuine achievement. Indeed, we would still ascribe an achievement to Archie in this case even despite the luck involved. It is, after all, *because of* his skill that he is successful, even though he could very easily have not been successful in this case.

The problem that cases like this pose for the robust virtue epistemologist is that if we allow Archie's success to count as an achievement, then we seem compelled to treat cognitive successes which are relevantly analogous as also being achievements. On the virtue epistemological proposal under consideration, however—that is, the proposal under which knowledge is a type of achievement—this would mean that we would thereby be compelled to regard the cognitive achievement in question as knowledge, even despite the luck involved.

In order to see why this is a problem for those virtue epistemologists who defend the knowledge-as-achievement thesis, consider the case of Barney which is structurally analogous to that of 'Archie'. Barney forms a true belief that there is a barn in front of him by using his cognitive abilities. That is, unlike a Gettier-style case—such as the case of 'Roddy' described above—Barney does not make any cognitive error in forming his belief in the way that he does. Accordingly, we would naturally say that Barney's cognitive success is because of his cognitive ability and so we would, therefore, attribute a cognitive achievement to Barney. According to the knowledge-as-achievement thesis, then, we should also treat Barney as knowing that what he is looking at is a barn. The twist in the tale, however, is that, unbeknownst to Barney, he is in fact in 'barn façade county' where all the other apparent barns are fakes. Intuitively, he does not have knowledge in this case because it is simply a matter of luck that his belief is true.

Cases like that of 'Barney' illustrate that there is a type of knowledge-undermining epistemic luck—what we might call 'environmental' epistemic luck—which is distinct from the sort of epistemic luck in play in standard Gettier-style cases like that of 'Roddy'. In particular, the kind of epistemic luck in play in standard Gettier-style cases 'intervenes' between the agent and the fact, albeit in such a way that the agent's belief is true nonetheless (that is, Roddy is not

looking at a sheep at all, even though he reasonably believes that he is, but his belief that there is a sheep in the field is true nonetheless). In contrast, in cases of environmental epistemic luck like that involving Barney, luck of this intervening sort is absent—Barney really does get to see the barn and forms a true belief on this basis—although the epistemically inhospitable nature of the environment ensures that his belief is nevertheless only true as a matter of luck such that he lacks knowledge.

Robust virtue epistemology is thus only able to exclude Gettier-style epistemic luck, but not also environmental luck. The moral to be drawn is that there is sometimes *more* to knowledge than merely a cognitive achievement, contrary to what the robust virtue epistemologist (who defends the knowledge-as-achievement thesis) argues. Rather, there can be cases in which (environmental) knowledge-undermining luck is involved where the luck does not in the process undermine the achievement in question. Put another way, the conclusion that is warranted by cases like this is that sometimes knowledge is more than just a cognitive achievement because merely exhibiting a cognitive achievement will not suffice to exclude all types of knowledge-undermining epistemic luck.

There are moves that those virtue epistemologists who defend the knowledge-as-achievement thesis might make to respond to this sort of problem. For example, one might argue that there is something special about the cognitive achievement at issue in knowledge which ensures that it is resistant to even this type of luck, even though non-cognitive achievements are entirely compatible with this kind of luck. There may be a case that can be made for this, though it will obviously face the charge of being ad hoc. Alternatively, one might simply insist that achievements exclude luck, and thus that we should not, contrary to intuition, treat Archie's success as an achievement when his success is lucky in the relevant fashion. The problem facing this proposal, however, is to explain why our intuitions about achievements are so off the mark in this case.[12]

Perhaps the proponent of the knowledge-as-achievement could make one of these strategies (or some third strategy) stick. I do not think the result would be a happy one, but it is often the case that our theories force us to make awkward theoretical moves in order to save the theory, so that such a move is not that compelling need not be a decisive count against the view. The more fundamental problem is that there is a further difficulty on the horizon for a view of this sort. Once these two objections for the knowledge-as-achievement thesis are taken together, however, they suggest not a mere 'patching-up' of the original proposal, but a radical re-thinking.

4. Knowledge Without Achievement

Consider the following example, due to Jennifer Lackey (2007, §2). Our protagonist, who we will call Jenny, arrives at the train station in Chicago and, wishing to obtain directions to the Sears Tower, approaches the first adult passerby that she sees. Suppose further that the person that she asks has first-hand knowledge of the area and gives her the directions that she requires. Intuitively, any true belief that Jenny forms on this basis would ordinarily be counted as knowledge. Relatedly, notice that insofar as we are willing to ascribe knowledge in this case, we will be understanding the details of the case such that the true belief so formed is non-

lucky in all the relevant respects (that is, it is not subject to either Gettier-style or environmental epistemic luck). For example, we are taking it as given that there is no conspiracy afoot among members of the public to deceive Jenny in this regard, albeit one which is unsuccessful in this case.

The moral that Lackey draws from this example is that sometimes one can have knowledge without the success in question being of credit to the agent. I think this conclusion is ambiguous. In particular, we need to make a distinction between a true belief being *of credit* to an agent, in the sense that the agent is deserving of some sort of praise for holding this true belief, and the true belief being *primarily creditable* to the agent, in the sense that it is to some substantive degree down to her agency that she holds a true belief. Lackey's focus when employing this example is on the former claim,[13] and this is not surprising since a number of commentators—see, for example, Greco (2002) in particular—have expressed their view in such a way that it seems to straightforwardly support this thesis. That said, strictly speaking the robust virtue-theoretic proposal is the latter claim.

Now this may initially seem to be an idle distinction, in that one might naturally suppose that in every case in which the former description holds the latter description holds as well, and vice versa. The problem, however, is that closer inspection of these two formulations of the view reveals that they in fact make very different demands. Moreover, one kind of case in which they come apart is precisely scenarios like the 'Jenny' example where an agent gains knowledge, to a large degree, by trusting the word of another.

In order to see this, we just need to note that it is of *some* credit to Jenny that she has a true belief. It is, after all, a *person* that she asks for directions in this case, and not, say, a lamppost or a dog. Moreover, the person she asks is not a small child, or someone who one might reasonable expect to be unreliable on this score (for example, someone who is clearly a tourist). In addition, if the testimony which Jenny received was obviously false, then we would expect her to be sensitive to this fact. If, for example, the informant told her that she should get back on the train and go home to New York, we would expect her to treat these directions as entirely spurious. So the moral to be drawn from this case is not (as Lackey does) that sometimes knowledge can be possessed even though the cognitive success in question is of no credit to the agent concerned.[14] The real moral resides in the fact that it is not *primarily creditable* to Jenny that she has formed a true belief More specifically, that Jenny has a true belief in this case does not seem to be *because of* her cognitive abilities, but rather because of the cognitive abilities of the informant who knows this proposition on a non-testimonial basis. One can thus have a true belief that is deserving of credit and yet that true belief not be primary creditable to your cognitive agency.[15]

Given that the true belief needs to be primarily creditable to the agent in order for it to count as a cognitive achievement, it follows that while Jenny has knowledge in this case, she does not exhibit a cognitive achievement. Once more then, we have seen that there is a problem associated with the idea that knowledge is to be identified with cognitive achievement.

It is not obvious how the proponent of the knowledge-as–achievement thesis can respond to cases of this sort short of either denying that Jenny has

knowledge, or else maintaining that it is primarily creditable to her that her belief is true, and thus that she is exhibiting a bona fide achievement after all. Both strategies involve denying some pretty strong intuitions about this case and so anyone taking either line will face a tough up-hill struggle.[16] That the defender of the knowledge-as-achievement thesis must also simultaneously deal with the other problem outlined above—concerning the apparent possibility of cognitive achievements which are not cases of knowledge—indicates that the struggle in question will be particularly onerous.

5. Diagnosis

I THINK WE CAN OFFER A DIAGNOSIS OF WHAT HAS GONE WRONG HERE. THERE ARE TWO closely related 'master' intuitions regarding knowledge that inform much of the contemporary epistemological discussion. The first is that knowledge is the product of one's cognitive abilities, such that when one knows, one's cognitive success is, in substantial part at least, creditable to one. Call this the *ability intuition*. The second master intuition is that knowledge is incompatible with luck, in the sense that if one knows, then it ought not to be the case that one could have easily been wrong. Call this the *anti-luck intuition*.

On the face of it, one would think that these two intuitions are entirely distinct, in the sense that whatever epistemic condition one places on knowledge in order to accommodate the one intuition will not thereby accommodate the other intuition. After all, one would naturally suppose that a true belief that is formed as a result of cognitive ability is not thereby immune to epistemic luck because of the possibility of Gettier-style cases (that is, cases in which one exhibits the relevant cognitive ability and one has the relevant true belief, and yet one has nevertheless been 'Gettierized' such that one's true belief is epistemically lucky). Conversely, it seems antecedently plausible that there might be true beliefs which meet the relevant anti-luck condition—such that the agent could not have easily been wrong—which are not thereby formed as a result of a cognitive ability. After all, there could be all manner of reasons why one's belief tracks the truth in nearby possible worlds which have nothing to do with one's own cognitive ability (perhaps, for example, it is the facts that are changing in order to correspond with your belief, rather than vice versa).

Virtue epistemologists are clearly most impressed by the ability intuition. Given the foregoing, the natural way to develop the view as a theory of knowledge would be in such a way that one had both a virtue-theoretic condition which accommodated this intuition and, in addition, an anti-luck condition—such as, for example, the safety principle[17]—which accommodated the anti-luck intuition (and so dealt with the problem of knowledge-undermining epistemic luck). Indeed, we noted the appeal of such a proposal above. As we have seen, however, there is a prominent strand of virtue epistemic thought—robust virtue epistemology—which in effect tries to accommodate the motivation for the anti-luck intuition by responding to the ability intuition in a particularly robust manner. We saw just such a move above, where the virtue epistemologist demands that the cognitive success be *because of* the relevant cognitive ability as a means of eliminating the epistemic luck at issue in the Gettier-style cases. In making this claim, the virtue epistemologist moves away from the thesis that knowledge involves cognitive

ability in the sense that the cognitive success in question is *to some significant degree* creditable to the agent and towards the more robust claim that knowledge is an achievement such that the success in question is *primarily* creditable to the agent.

Unfortunately, however, we saw in the last two sections that there is good reason to believe that such a 'beefed-up' account of the relationship between cognitive ability and cognitive success will lead to an unduly demanding account of knowledge. In particular, it cannot accommodate cases of knowledge (like the 'Jenny' case) where the true belief in question, while to some significant degree creditable to the agent, does not constitute a cognitive achievement. Moreover, we have also seen that the proposal is in any case ineffective, since there is good reason to think that one can exhibit a cognitive achievement without thereby possessing knowledge because of the presence of environmental epistemic luck, a variety of epistemic luck distinct from Gettier-style epistemic luck.

Now one might respond to the failure of this sort of account of knowledge by opting for the opposing radical thesis that takes the anti-luck intuition as its lead. On this view—what one might call a *robust anti-luck epistemology*—one tries to formulate an anti-luck condition on knowledge in such a demanding way that one does not need an additional ability condition in order to accommodate the other master intuition. Indeed, one could read certain modal epistemological proposals in the recent literature as proposing just such a view.[18] That is, just as the virtue epistemologists who advance the knowledge-as-achievement thesis in effect try to accommodate the motivatiou for the anti-luck intuition by offering a robust construal of their ability condition on knowledge, so proposals along these lines offer a robust construal of the anti-luck condition in order to accommodate the motivation behind the ability intuition.

Perhaps a view of this radical sort is possible, though I have my doubts. After all, as noted above, it seems inevitable that there will be cases in which there are true beliefs which exhibit the required modal properties to ensure that they are not lucky in the relevant sense and yet which are not formed as a result of the agent's cognitive abilities. Given the ability intuition, such cases will not count as knowledge, even though it will be the case that the agent in question could not have easily been wrong.

Consider the following example. Our agent, Temp, forms his beliefs about the temperature in his room by consulting a thermometer on the wall. Unbeknownst to Temp, however, the thermometer is broken and is fluctuating randomly within a given range. Nonetheless, Temp never forms a false belief about the temperature by consulting this thermometer since there is a person hidden in the room, next to the thermostat, whose job it is to ensure that whenever Temp consults the thermometer, the temperature in the room corresponds to the reading on the thermometer.

Clearly, Temp cannot gain knowledge of the temperature of the room by consulting a broken thermometer, and we can explain why by appeal to the ability intuition about knowledge—the truth of his belief in this case is in no substantive way the product of his cognitive ability. Even so, depending on the details of the case, it could nonetheless be true that Temp's belief is not subject to knowledge-undermining epistemic luck. After all, if the example is set up in the right way, then it could well be the case that his true belief could not have

very easily been false. Indeed, with the example suitably described, it may well follow that in *all* near-by possible worlds in which Temp continues to form his belief in this way his belief is true.

Cases like Temp show that an anti-luck condition all by itself cannot offer a complete account of knowledge—that is, they demonstrate that a robust anti-luck epistemology is unsustainable. But if both robust virtue epistemology and robust anti-luck epistemology do not work, then it seems that in order to accommodate the anti-luck and the ability intuitions, we will need to aim for an intermediate position between these two extremes. What we need, in short, is an *anti-luck virtue epistemology*: an account of knowledge which gives equal weight to both of the master intuitions and so incorporates both an anti-luck and an ability condition.

For example, let us describe any true belief that could not have easily been wrong as *safe*. The general structure of the account of knowledge offered by an anti-luck virtue epistemology can now be described as follows: knowledge is safe belief that arises out of the reliable cognitive traits that make up one's cognitive character, such that the cognitive success at issue is to some significant degree creditable to one.

Since anti-luck virtue epistemology imposes a less restrictive ability condition on knowledge, it does not face the problem posed by cases like 'Jenny'. After all, as we noted above, Jenny's true belief is *partly* creditable to her (this is why she is deserving of some credit for holding it), and the reason for this is that it does indeed arise out of reliable cognitive traits that make up her cognitive character. The crux is just that Jenny's cognitive success is not *primarily* creditable to her, but this need not be bar to her possessing knowledge on this view. In epistemically friendly environments of the sort that Jenny is in—environments in which the anti-luck condition is very easily met—one can gain knowledge even though one's true belief is not primarily creditable to one. But *that*, as the 'Jenny' case indicates, is entirely in accordance with intuition.

Furthermore, by adding the anti-luck condition, this proposal can deal with both the standard Gettier-style epistemic luck found in cases like that of 'Roddy' and also the more tricky environmental epistemic luck found in cases like that of 'Barney'. In all such cases, the agent could very easily have been wrong, and so does not count as knowing by the lights of this proposal. Finally, the ability condition enables the proposal to deal with cases like 'Temp' in which the agent has a non-lucky true belief and yet does not count as knowing because the non-lucky true belief does not arise out of the cognitive abilities of the agent.

Clearly, given the problems facing the more robust virtue-theoretic proposal, anti-luck virtue epistemology has a lot to recommend it.[19] But can it enable us to make any progress towards answering the value problem? In particular, once we have jettisoned the idea that knowledge is a kind of achievement, can the final value of achievements give us any purchase on the value problem?

6. Back to the Value Problem

On the face of it, one might think that there is a fairly straightforward way of resurrecting the knowledge-as-achievement account of the value of knowledge along these new lines. After all, while we have noted that there are cases of

knowledge where the agent does not exhibit a cognitive achievement, and cases of cognitive achievement where the agent does not possess knowledge, one can nonetheless consistently argue that knowledge is the kind of epistemic standing that *tends* to go hand in hand with cognitive achievement. And that might be enough to answer the tertiary value problem.

For example, suppose that it is true that the cases in which knowledge and cognitive achievement come apart are peripheral. Now let us ask our question that we posed at the end of §1. Recall that one has to decide between knowledge that p and an epistemic standing as regards p that falls short of knowledge, where it is stipulated in advance that there are no practical costs or benefits from choosing either way. We noted there that we have the intuition that in such cases one should choose knowledge, and that this is so reflects our view that knowledge is precious. But so long as knowledge tends to be the kind of epistemic standing that goes hand-in-hand with cognitive achievement then, given that cognitive achievements are finally valuable, are we not in a position to account for this intuition and thereby meet the tertiary value problem? That is, why can one not argue that since knowledge is the kind of thing that tends to be finally valuable (even though sometimes, when it is not a cognitive achievement, it is not), then it is no surprise that we regard knowledge as precious in this way?

In order to answer this question, we first need to form a view about just how extensive the cases of knowledge are which are not cognitive achievements. After all, although the testimonial case we have examined might initially seem quite peripheral, on reflection one might plausibly contend that quite a lot of our testimonial knowledge is gained in this fashion. Moreover, there is also good reason to hold that there may be non-testimonial cases that have the relevant features. For example, one might claim that just as there is a substantive degree of ungrounded trust of others involved in the 'Jenny' case offered above, so there is a substantive degree of ungrounded *self*-trust involved in much of our other knowledge, such as an ungrounded trust in the reliability of our faculties. If this is right, then it may turn out that very little of our knowledge, if any, involves a cognitive achievement. The prospects for meeting the value problem with a proposal of this sort would then be dim indeed.[20]

Even if we can block this worry by arguing for a close relationship between knowledge and cognitive achievement, a second, and more substantive, worry remains. For while the most modest claim about the value of knowledge might appear to meet the 'question' test just posed, this appearance vanishes on closer inspection. In order to see this, one just needs to note that on the view under consideration there *is* an epistemic standing that falls short of knowledge which is no less valuable as a result—the category of cognitive achievements that are not also cases of knowledge. If that is right, then on this view one has no basis on which to prefer knowing that p to having a belief that p which constitutes a cognitive achievement (albeit one that falls short of knowledge). Indeed, since on this view cognitive achievements are *always* finally valuable, one should if anything prefer the epistemic standing that falls short of knowledge over knowledge.

So once one rejects the idea that knowledge is a kind of achievement, the final value of achievements is no longer able to offer us a way of responding to the tertiary value problem.

7. Concluding Remarks

AT THE START OF THIS PAPER, I ARGUED THAT THE ONLY WAY TO CAPTURE OUR INTUITION that knowledge is distinctively valuable is by demonstrating that it is finally valuable. Unfortunately, we have found that the only prima facie plausible account of why knowledge might be finally valuable—that offered by robust virtue epistemology—does not stand up to closer scrutiny. Along the way, however, we have diagnosed the problems facing robust virtue epistemology as representing a failure to give due weight to the ability and anti-luck master intuitions about knowledge. Out of the ashes of robust virtue epistemology, therefore—and robust anti-luck epistemology as well for that matter—has risen a new epistemological proposal: anti-luck virtue epistemology. Accordingly, while the overarching message of this paper is negative—in that we are wrong to suppose that knowledge is distinctively valuable—an important positive epistemological conclusion has also been generated by this discussion.

Now one might argue that if one's theory of knowledge is unable to answer the value problem, then that is itself a pretty definitive strike against it, and thus that the fact that anti-luck virtue epistemology is unable to account for the distinctive value of knowledge entails that the view is a lost cause.[21] But this way of thinking about the value problem and its role in the theory of knowledge is surely too strong. Instead, what is presumably required is *either* that one's theory of knowledge can answer the value problem *or* that the theory is able to provide some plausible account of why knowledge is not really distinctively valuable after all, even though it appears to be. That is, provided one's theory of knowledge can answer the second of the two questions just identified, then that should suffice. Interestingly, the account of knowledge offered here may well be able to offer a good answer to this second question.

After all, it seems undeniable that the *paradigm* cases of knowledge that spring to mind when we reflect on the nature of knowledge tend to be cases in which the knowledge at issue also constitutes a cognitive achievement.[22] When we think of paradigm cases of knowledge we surely do not have barn façade-style cases in mind. If this is so, however, then it is hardly surprising that we instinctively regard knowledge as distinctively valuable, since it is not in dispute that in these cases the knowledge in question does go hand-in-hand with a cognitive achievement, and we have granted that the latter is finally valuable. Under closer inspection, however, we have found that the relationship between knowledge and cognitive achievement is not as straightforward as it first appears. In particular, once we move away from the paradigm cases of knowledge that most immediately spring to mind when we think about the nature of knowledge, and instead cast our net much wider, we discover that there is a range of cases where the two epistemic categories come apart.

We began this paper with the noble intention of explaining why we rightly regard knowledge as distinctively valuable, and have ended it by offering an account of why knowledge *is not* distinctively valuable, even though it appears to be. While this conclusion is of course pessimistic, notice that our examination of the value problem has generated the important side-effect of indicating what kind of direction a theory of knowledge should take. It may be that offering a positive resolution to the value problem is not a pre-condition of an adequate

theory of knowledge, but there is good reason to think that an examination of the value problem is at least an essential part of epistemological theorizing, on account of the light that it casts on the epistemological project.[23,24] φ

Notes

[1] Note that where I draw these contrasts between, as in this case, knowledge and mere true belief, I have in mind a mere true belief that appropriately corresponds to the true belief at issue in the instance of knowledge. For example, one is comparing a situation in which an agent knows a proposition with an exactly analogous situation in which that agent merely truly believes this proposition.

[2] If you think, like Sartwell (1992), that knowledge just is true belief, then it follows that knowledge cannot be more valuable than mere true belief. For other sources of scepticism about the claim that knowledge is more valuable than mere true belief, see Stich (1990, 122-3), Hawley (2006), and Goldman & Olsson (2009, §1).

[3] Many discussions of what I am calling the primary value problem take the so-called 'swamping problem' to be a version of this problem. I think this is a mistake, which is why I do not discuss this problem here. As I argue in Pritchard (*forthcoming*), the swamping problem, at least when properly understood, poses a more specific difficulty such that an adequate response to the primary value problem need not thereby be an adequate response to the swamping problem. For the key discussions of the swamping problem, see Swinburne (1999; 2000), Kvanvig (2003) and Zagzebski (2003).

[4] Kaplan (1985) famously argues, for example, that the conclusion to be drawn from the Gettier counterexamples to the tripartite account of knowledge (that is, knowledge as justified true belief) is that knowledge is not a distinctively valuable epistemic standing. Instead, what we should seek is justified true belief, something which the Gettier counterexamples demonstrate falls short of knowledge. Given the further claim that justified true belief is more valuable than mere true belief, Kaplan thus answers the primary value problem while leaving the secondary value problem unanswered. In a similar vein, Kvanvig (2003) has recently argued that knowledge is not more valuable than any proper sub-set of its parts because there is no Gettier counterexample-excluding theory of knowledge available which could account for the greater value of knowledge over any non-Gettier counterexample-excluding proper sub-set of its parts.

[5] Note that final value is not the same thing as intrinsic value. See endnote 10.

[6] For a recent survey of research on the value problem, see Pritchard (2007*c*; cf. Pritchard 2007*b*). See also Baehr (2009).

[7] This example is adapted from one offered by Chisholm (1977, 105).

[8] In early work, Greco (1999; 2000) took just this line. As we will see in a moment, these days he advocates a robust form of virtue epistemology.

[9] Those familiar with the literature in this respect will recognise that often virtue epistemologists like Greco (for example, 2002) make a stronger claim in this regard. That is, they do not simply argue that the true belief in question is primarily creditable to the knowing agent, but also that it is *of credit* to the knowing agent that she believes truly (that is, that she is deserving of some sort of praise, at least when assessed from a purely epistemic point of view). I think this is a mistake, and I explain why in §4.

[10] While epistemic virtue theorists are aware that they may be able to account for the distinctive value of knowledge by appeal to the value of an achievement, they unfortunately mischaracterise the kind of value in question, since they hold that it is *intrinsic value* rather than final value that is at issue. See, for example, Greco (2009, §4). Crucially, however, intrinsic value is not the same as final value. This is because intrinsic value concerns only the value generated by the intrinsic properties of the target item, and yet something can be finally—that is, non-instrumentally—valuable because of its *relational* (and hence non-intrinsic) properties. Think, for example, of the first book produced on the first ever printing press. Moreover, it is important to our discussion that we focus on final value rather than intrinsic value because, on the account

of the value of knowledge under consideration, it is clearly because of the *relational* properties of the true belief in question—that is, it is true belief that is skilfully attained—that it constitutes a cognitive achievement, and hence on this view, accrues a distinctive kind of value. Thus, the additional value that is generated is final value, not intrinsic value. For further discussion of the intrinsic value/final value distinction, see Rabinowicz & Roennow-Rasmussen (1999; 2003).

[11] An alternative way of dealing with this problem would be to argue that it can be in the nature of something to be finally valuable even though sometimes it is not—for example, one might argue that pleasure is in its nature finally valuable, even though some pleasures (that is, the 'bad' ones) lack final value. According to this proposal, then, it would be in the nature of achievements to be finally valuable even though some of them (that is, the wicked or trivial ones) lack final value. I am grateful to Mike Ridge for this suggestion.

[12] Greco is the only virtue epistemologist to really confront this problem. In Greco (2002, §3; cf. Greco 2009, §5), for example, he takes the line that achievements are by their nature luck-excluding (with the consequence, presumably, that Archie's success is not an achievement in the case in which the other targets have the arrow-excluding force fields around them). Elsewhere, in Greco (2007, §5), he argues that there is something peculiar about knowledge which ensures that it is luck-excluding in a more exacting fashion than non-cognitive achievements, while also offering a more fine-grained account of what constitutes a cognitive ability which would exclude someone like Barney from possessing the relevant cognitive ability in the first place. I discuss Greco's treatment of these issues further in Pritchard (2008*c*; cf. Pritchard 2008*a*). See also Kvanvig (*forthcoming*).

[13] The title of the paper in which this example appears is 'Why We Do not Deserve Credit for Everything We Know'.

[14] A second type of case that Lackey (2007) offers—that of innate knowledge—might fare better in this regard. After all, if there is such a thing as innate knowledge, then it would presumably be such that it involves a true belief which is *neither* of credit to the agent nor primarily creditable to the agent. For my own part, I do not hold that innate knowledge is even possible, but this is an issue that cannot be usefully engaged with here.

[15] More generally, the distinction between credit-worthy true belief and true belief that is primarily creditable to one also comes apart in the other direction. There could, after all, be true beliefs that are primarily creditable to one's cognitive agency and yet for which you are deserving of no credit at all (for example, where the cognitive achievement in question is very easy).

[16] Indeed, taking the latter line will almost certainly commit one to a very restrictive account of testimonial knowledge, a view that is usually known as global reductionism. Although unpopular, this view does have some adherents, most notably Fricker (for example, 1995). It is perhaps for this reason that Greco (2007) opts for the former line, although he does not make a very strong case for it. By analogy, he argues that one might score a very easy goal as a result of that goal being set-up by a display of tremendous skill. He maintains that the skill involved in setting up this easy goal does not undermine the achievement of the agent who scores the goal. Greco is surely right about this, but the example is not relevantly analogous to Lackey's testimonial case. An example that would be relevantly analogous is someone who lacks archery abilities being assisted by a skilled archer in firing an arrow and being thereby successful. While the unskilled archer's abilities might have played *some* role in the successful outcome—such that it is *to some degree* creditable to him that he is successful—we would surely say that this success is primarily creditable to the skilled archer (or, at least, the combined efforts of the unskilled archer and the skilled archer). On this basis, we would surely maintain that the unskilled archer's success does not constitute a bona fide achievement.

[17] This anti-luck condition is usually characterised in terms of a 'safety' principle. For some key defences of safety, see Sosa (1999) and Pritchard (2002; 2005; 2007*a*).

[18] To take the two most prominent examples of this tendency, see the sensitivity-based theories of knowledge offered by Dretske (for example, 1970) and Nozick (1981). For a more general defence of the idea of a genuinely anti-luck epistemology, see Pritchard (2005; 2007*a*).

[19] I explore this proposal at greater length in Pritchard (2008*b*).

[20] This would constitute one way of recasting the sceptical problem in value-theoretic terms. That is, the primary target of the sceptical argument would not be knowledge simpliciter, but rather a distinctively valuable epistemic standing. The advantage of reading the sceptic in this

way is that it would clearly be irrelevant to respond to the sceptic by offering an account of knowledge on which knowledge was not distinctively valuable (indeed, this would constitute a kind of capitulation). I explore this possibility at length in Pritchard (2008*e*).

[21] One finds an assumption of roughly this sort expressed in a number of works. See, for example, Zagzebski (1999), Williamson (2000, ch. 1) and Kvanvig (2003, ch. 1). For a critical discussion of this assumption, see DePaul (2009).

[22] I am grateful to Chris Hookway for urging this line on me.

[23] For a fuller discussion of the question of epistemic value, including what ramifications the conclusions reached here have for epistemology more generally, see Haddock, Millar & Pritchard (2010, chs. 1-4).

[24] Earlier versions of this paper were presented at the *Royal Institute of Philosophy* in London as part of the *Epistemology* lecture series in October 2006, at the *Basic Knowledge* conference at the University of St. Andrews in November 2006, at the *Dialectica* annual conference in Geneva in December 2006, at a conference on *Epistemology* at the University of Leuven in February 2007, at an Pacific APA invited symposium on *Epistemic Value* in San Francisco in April 2007, and at departmental talks at the Universities of Sussex, Edinburgh, Leeds Sheffield and St. Andrews during 2006, 2007 and 2008. I am grateful to the audiences on these occasions, and to other people who have commented on this work, including Sharar Ali, Kelly Becker, Brit Brogaard, Ross Cameron, Matthew Chrisman, Andy Clark, John Divers, Igor Douven, Pascal Engel, Paul Faulkner, Gordon Finlayson, Lizzie Fricker, John Greco, Patrick Greenough, Dominic Gregory, Stephen Grimm, Lars Gundersen, Bob Hale, John Hawthorne, Allan Hazlett, Chris Hookway, Robert Hopkins, Jennifer Hornsby, Carrie Jenkins, Jesper Kallestrup, Matthew Kieran, Jonathan Kvanvig, Bob Lockie, Anne Meylan, Andy McGonigal, Michael Morris, Antony O'Hear, Erik Olsson, Diana Raffman, Murali Ramachandran, Mike Ridge, Wayne Riggs, Bruce Russell, Jenny Saul, Sarah Sawyer, Barry Smith, Jason Stanley, Bob Stern, Helen Steward, Kathleen Stock, Robbie Williams, Tim Williamson, René van Woudenberg, and Crispin Wright. Special thanks to Adam Carter, Adrian Haddock, Chris Kelp, Alan Millar, and Ernie Sosa. This paper has benefited from an AHRC Research Leave award.

References

Baehr, J. (2009). 'Is There a Value Problem?', *Epistemic Value*, (eds.) A. Haddock, A. Millar, & D. H. Pritchard, 42-59, Oxford University Press, Oxford.

Chisholm, R. (1977). *Theory of Knowledge* (2nd Ed.), Prentice-Hall, Englewood Cliffs, New Jersey.

DePaul, M. (2009). 'Ugly Analyses and Value?', *Epistemic Value*, (eds.) A. Haddock, A. Millar, & D. H. Pritchard, 112-38, Oxford University Press, Oxford.

Fricker, E. (1995). 'Telling and Trusting: Reductionism and Anti-Reductionism in the Epistemology of Testimony', *Mind* 104, 393-411.

Goldman, A., & Olsson, E. (2009). 'Reliabilism and the Value of Knowledge', *Epistemic Value*, (eds.) A. Haddock, A. Millar, & D. H. Pritchard, 19-41, Oxford University Press, Oxford.

Greco, J. (1999). 'Agent Reliabilism', *Philosophical Perspectives* 13, 273-96.

——(2000). *Putting Skeptics in Their Place: The Nature of Skeptical Arguments and Their Role in Philosophical Inquiry*, Cambridge University Press, Cambridge.

——(2002). 'Knowledge as Credit for True Belief', *Intellectual Virtue: Perspectives from Ethics and Epistemology*, (eds.) M. DePaul & L. Zagzebski, Oxford University Press, Oxford.

——(2007). 'The Nature of Ability and the Purpose of Knowledge', *Philosophical Issues* 17, 57-69.

——(2008). 'What's Wrong With Contextualism?', *Philosophical Quarterly* 58, 416-36.

——(2009). 'The Value Problem', *Epistemic Value*, (eds.) A. Haddock, A. Millar, & D. H. Pritchard, 313-21, Oxford University Press, Oxford.

——(*Forthcoming*). *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*, Cambridge University Press, Cambridge.

Haddock, A., Millar, A., & Pritchard, D. H. (2010). *The Nature and Value of Knowledge: Three Investigations*, Oxford University Press, Oxford.

Hawley, P. (2006). 'Scepticism and the Value of Knowledge', *typescript*.

Jones, W. (1997). 'Why Do We Value Knowledge?', *American Philosophical Quarterly* 34, 423-40.

Kaplan, M. (1985). 'It is Not What You Know That Counts', *Journal of Philosophy* 82, 350-63.

Kvanvig, J. (2003). *The Value of Knowledge and the Pursuit of Understanding*, Cambridge University Press, Cambridge.

——(2009). 'Responses', *Epistemic Value*, (eds.) A. Haddock, A. Millar, & D. H. Pritchard, 339-52, Oxford University Press, Oxford.

Lackey, J. (2007). 'Why We Do not Deserve Credit for Everything We Know', *Synthese* 158, 345-61.

Pritchard, D. H. (2002). 'Resurrecting the Moorean Response to the Sceptic', *International Journal of Philosophical Studies* 10, 283-307.

——(2003). 'Virtue Epistemology and Epistemic Luck', *Metaphilosophy* 34, 106-30; and reprinted in *Moral and Epistemic Virtues*, (eds.) M. S. Brady & D. H. Pritchard, 210-34, (Blackwell, 2003).

——(2005). *Epistemic Luck*, Oxford University Press, Oxford.

——(2007*a*). 'Anti-Luck Epistemology', *Synthese* 156, 277-97.

——(2007*b*). 'Recent Work on Epistemic Value', *American Philosophical Quarterly* 44, 85-110.

——(2007*c*). 'The Value of Knowledge', *Stanford Encyclopedia of Philosophy*, (ed.) Ed Zalta, http://plato.stanford.edu/entries/knowledge-value/.

——(2008*a*). 'A Defence of Quasi-Reductionism in the Epistemology of Testimony', *Philosophica* 78, 13-28.

——(2008*b*). 'Anti-Luck Virtue Epistemology', *manuscript*.

——(2008*c*). 'Greco on Knowledge: Virtues, Contexts, Achievements', *Philosophical Quarterly* 58, 437-47.

——(2008*d*). 'Knowing the Answer, Understanding and Epistemic Value', *Grazer Philosophische Studien*, 77, 325-39

——(2008*e*). 'Radical Scepticism, Epistemic Luck and Epistemic Value', *Proceedings and Addresses of the Aristotelian Society* (suppl. vol.) 82, 19-41.

——(2008*f*). 'Virtue Epistemology and Epistemic Luck, Revisited', *Metaphilosophy* 39, 66-99.

——(2009). 'Knowledge, Understanding and Epistemic Value', *Epistemology (Royal Institute of Philosophy Lectures)*, (ed.) A. O'Hear, 19-43, Cambridge University Press, Cambridge.

——(*Forthcoming*). 'What is the Swamping Problem?', *Reasons for Belief*, (eds.) A. Reisner & A. Steglich-Petersen, Springer, Dordrecht, Holland.

Rabinowicz, W., & Roennow-Rasmussen, T. (1999). 'A Distinction in Value: Intrinsic and For its Own Sake', *Proceedings of the Aristotelian Society* 100, 33-49.

——(2003). 'Tropic of Value', *Philosophy and Phenomenological Research* 66, 389-403.

Sartwell, C. (1992). 'Why Knowledge is Merely True Belief', *Journal of Philosophy*, 89, 167-80.

Sosa, E. (1988). 'Beyond Skepticism, to the Best of our Knowledge', *Mind* 97, 153-89.

——(1991). *Knowledge in Perspective: Selected Essays in Epistemology*, Cambridge University Press, Cambridge.

——(1999). 'How to Defeat Opposition to Moore', *Philosophical Perspectives* 13, 141-54.

——(2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, Oxford University Press, Oxford.

Stich, S. (1990). *The Fragmentation of Reason*, MIT Press, Cambridge, MA.

Swinburne, R. (1999). *Providence and the Problem of Evil*, Oxford University Press, Oxford.

——(2000). *Epistemic Justification*, Oxford University Press, Oxford.

Williamson, T. (2000). *Knowledge and Its Limits*, Oxford University Press, Oxford.

Zagzebski, L. (1996). *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*, Cambridge University Press, Cambridge.

——(1999). 'What is Knowledge?', *The Blackwell Guide to Epistemology*, (eds.) J. Greco & E. Sosa, 92-116, Blackwell, Oxford.

——(2003). 'The Search for the Source of the Epistemic Good', *Metaphilosophy* 34, 12-28; and reprinted in *Moral and Epistemic Virtues*, (eds.) M. S. Brady & D. H. Pritchard, 13-28, (Blackwell, 2003).