

CHAPTER TWO

KNOWLEDGE AND FINAL VALUE

§2.1. INTRODUCTION

We saw in the last chapter that in order to answer the value problem we need to explain 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. We noted that offering such an explanation requires us to explain why knowledge, unlike that which falls short of knowledge, has *final* value. Moreover, we saw that if one could account for the final value of knowledge then one would thereby have a resolution to the swamping problem as well (in the form of either a monistic or pluralist response to this difficulty). The goal of this chapter is to examine the best—indeed, the *only*—response to the value problem in the contemporary literature which proceeds by arguing that knowledge has final value. As we will see, this response to the value problem is ultimately unsuccessful, though the failure of this account of epistemic value highlights some important epistemological morals, both as regards the problem of epistemic value and also regarding the very project of understanding knowledge.

§2.2. ROBUST VIRTUE EPISTEMOLOGY

The account of knowledge which has the best shot at accounting for the final value of knowledge is a type of virtue epistemology, what I will refer to as a *robust* virtue epistemology. This is the sort of virtue-theoretic account of knowledge that is offered by, for example, Ernest Sosa (1988; 1991; 2007), Linda Zagzebski (1996; 1999) and John Greco (e.g., 2002; 2007a; 2007b; *forthcoming*). What makes such a virtue-theoretic proposal robust is the fact that it attempts to exclusively analyse knowledge in terms of a true belief that is the product of epistemically virtuous belief-forming process.

The big attraction of virtue-theoretic accounts of knowledge is that they capture our strong intuition that knowledge is the product of one's reliable cognitive abilities.¹ A true belief that is gained in a way that is completely unconnected with one's cognitive abilities, even if that belief is reliably formed or would count as justified by the lights of certain conceptions of justification, just would not count as knowledge. We will explore the motivation for virtue epistemology in greater detail in the next chapter (and also examine some specific cases that lend support for the view). For now, however, this rather schematic presentation of the view will suffice even though it glosses over the many important differences between different types of virtue epistemic proposal (for example, depending on what one builds into one's conception of an epistemic virtue, one will be led to adopt a very different kind of virtue epistemology). As we will see in a moment, 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—e.g., 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'.² Using his highly reliable cognitive faculties, Roddy the shepherd forms a true belief that there is a sheep in the 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 cases like this in mind, one might naturally be tempted to opt for a *modest* virtue epistemology, one that does not try to completely analyse knowledge in terms of a virtue-theoretic condition but which is instead willing to endorse in addition a further codicil that can deal with Gettier-style cases.³

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 cognitive abilities—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.

How are we to read the 'because of' relation here? There is as yet no consensus amongst robust virtue epistemologists on this score, but the most developed view in the literature in this regard—due to Greco (2007a; 2007b)—takes the causal explanatory line that a true belief is because of an agent's cognitive abilities when it is *primarily* creditable to the agent that her belief is true.⁴ Although this way of reading the because of relation does generate some surprising results, we will set these potential problems to one side in order to give the account the best run for its money.⁵

So construed, 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 in the relevant sense, since we would not count his true belief as being primarily creditable to 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. Moreover, what goes for Roddy will intuitively also go for other Gettier-style cases, since they each share the same relevant properties (i.e., the cognitive success in question is not properly attributable to the agent's cognitive abilities, but rather to some other factor outwith our hero's cognitive agency).

It seems, then, that robust virtue epistemology may well have the resources to deal with the kind of knowledge-undermining epistemic luck in play in Gettier-style cases. If that's right, then there is no need to add a codicil to one's virtue-theoretic account of knowledge in order to make it Gettier-proof. From a theoretical point of view, this is very satisfying, since having such a codicil in one's account of knowledge looks *ad hoc*. Why should knowledge have this structure such that the virtue-theoretic component captures almost all the cases, but not quite? Robust virtue epistemology thus appears to have a lot going for it.

§2.3. KNOWLEDGE AND ACHIEVEMENT

As Greco (*forthcoming*) points out, a further advantage of understanding knowledge along robust 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 because of one's ability (i.e., primarily creditable to the exercise 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. As we will see, that knowledge turns out to be a type of achievement on this view is key to its defence of the final value of knowledge.

In order to see the plausibility of this general account of achievement, consider the following case. Suppose that our hero—let's call him 'Archie'—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 wouldn't count his success as an achievement if the success was not *because of* Archie's archery abilities (i.e., where his success is not primarily creditable to his archery abilities but rather to some further factor).

This is important because of the possibility that the success in question is 'Gettierized'. If, for example, a dog ran on to the range and 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, where this means that his success is primarily creditable to his abilities rather than to some factor independent of his abilities. Call this the *achievement* thesis.

There are some *prima facie* problems with the achievement thesis. In particular, there are grounds for thinking that as it stands it is too permissive. After all, we tend to think of achievements in such a way that they involve the overcoming of an obstacle of some sort, and

yet it seems consistent with the achievement thesis that the mere success-through-agency at issue in this thesis need involve nothing of the sort. Relatedly, it seems an essential part of achievements that they involve certain motivational states on the part of the subject with regard to the success in question—in particular, that the subject is actively seeking to bring this success about. But since the achievement thesis makes no mention of such motivational states it seems to allow that achievements could be entirely passive. More generally, the problem is that the achievement thesis seem to count as achievements successes which are just too easy to legitimately fall into this category.

We will come back to the problem of easy achievements in chapter four since it raises issues that are not directly relevant to our present concerns. For now, we will take it that the achievement thesis is on roughly the right lines in order to give the robust virtue epistemologist's argument for the final value of knowledge the best run for its money. What is important for our present purposes is that if this account of achievement is right then it follows that knowledge, by the lights of the robust virtue epistemologist at any rate, is just a specifically cognitive type of achievement. That is, achievements are successes that are because of ability and yet knowledge, according to the robust virtue epistemologist, is just cognitive success (i.e., true belief) that is because of cognitive ability (i.e., epistemic virtue, broadly conceived). The achievement thesis when combined with robust virtue epistemology thus entails the claim that knowledge is a type of achievement, what we will call the *knowledge-as-achievement* thesis, or $K=E$ for short.

The reason why the $K=E$ thesis is important for our purposes is 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 (i.e., 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—final value—which that which falls short of knowledge lacks, and hence show that it is distinctively valuable.⁶

In order to see why achievements might be thought to be finally valuable, consider again the case of 'Archie'. This time, though, suppose 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 is because of Archie's archery abilities or through dumb luck (e.g., 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 it 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.

Moreover, what goes here for Archie's achievement of hitting the target seems to be equally applicable to achievements more generally: achievements are finally valuable. Imagine, for example, that you are about to undertake a course of action designed to attain a certain outcome and that you are given the choice between merely being successful in what you set out to do, and being successful in such a way that you exhibit an achievement. Suppose further that it is stipulated in advance that there are no practical costs or benefits to choosing either way. Even so, wouldn't you prefer to exhibit an achievement? And wouldn't you be right to do so? If that's correct, then this is strong evidence for the final value of achievements.

Indeed, that achievements are valuable in this way is hardly surprising once one reflects that they constitute the exercise of one's agency on the world. A life lacking in such agential power, even if otherwise successful (e.g., one's goals are regularly attained), would clearly be severely impoverished as a result. A good life is thus, amongst other things, a life rich in achievement. Call the claim that achievements are finally valuable the *value of achievements* thesis.⁷

Now, if knowledge is simply a type of achievement, and achievements are finally valuable, then it immediately follows that knowledge has final value too. Robust virtue epistemology, when combined with a claim about the nature of achievements (the achievement thesis) and a claim about the final value of achievements (the value of achievements thesis), thus entails the thesis that knowledge has final value. More formally, we can express the reasoning in play here as follows:

From Robust Virtue Epistemology to the Final Value of Knowledge

- (P1) Achievements are successes that are because of ability. (Achievement thesis)
- (P2) Knowledge is a cognitive success that is because of cognitive ability. (Robust Virtue Epistemology)
- (C1) So, knowledge = cognitive achievement. (K=A thesis)
- (P3) Achievements are finally valuable. (Value of Achievements thesis)
- (C2) So, knowledge has final value.

Since the inferences in play here are clearly valid, if one wishes to object to this argument then one will need to deny one of the premises in play.

Let's start with the two premises concerning achievements more generally, (P1) and (P3). We have already noted that there is a *prima facie* worry regarding (P1) which concerns 'easy' achievements. In order to give this argument the best run for its money, however, we will let this premise stand for now (we will consider the status of (P1) in more detail in chapter four).

That brings us to (P3), the value of achievements thesis. One worry that one might have about this thesis is that some achievements seem to have very little value—or are even *disvaluable*—because, for example, they are pointless or just plain wicked. Are even achievements of this sort of final value? Note, however, that the value of achievements thesis, properly construed, 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 *disvalue*, perhaps because of the opportunity cost incurred by seeking the pointless 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 recognise that the value of achievements thesis when properly understood is entirely consistent with this possibility.⁸

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 help us answer 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.

In any case, while we will reconsider the status of (P3) in more detail in chapter four when we look again at (P1), we have sufficient grounds for taking this premise to be well-founded. This leaves only (P2), which is the robust virtue epistemological account of knowledge itself. As we will see in a moment, this thesis faces some fairly severe problems. Primarily, we will be arguing against this claim by showing that the K=A thesis (i.e., (C1)) that it (along with (P1)) generates is false. Ultimately, however, the way that we will be arguing against (C1) should leave no-one in any doubt that whatever other difficulties the other premises in this argument might face, it is (P2) that is the key weak link in the robust virtue epistemologist's argument for the final value of knowledge. As a result, this argument fails to demonstrate its conclusion, (C2).

§2.4. INTERLUDE: IS ROBUST VIRTUE EPISTEMOLOGY A REDUCTIVE THEORY OF KNOWLEDGE?

We noted in the last chapter in our discussion of the swamping problem that there are two *prima facie* plausible responses to that problem which are consistent with the final value of knowledge. What is key to both responses is that they reject the epistemic value T-monism that is essential to the setting-up of the swamping problem and argue instead that knowledge is a fundamental epistemic good. The first proposal (the monistic response) advances a form of epistemic value monism—epistemic value K-monism—which treats knowledge as the only fundamental epistemic good. The second proposal (the pluralist response) is a form of epistemic value pluralism which treats knowledge as a fundamental epistemic good in addition to true belief.

We noted in the last chapter that if knowledge is finally valuable then it thereby follows that it is a fundamental epistemic good. As a result, if the robust virtue epistemic

defence of the final value of knowledge is successful, then it will on the face of it lend support to *both* of these responses to the swamping problem. (It will also constitute a decisive strike against the third response to the swamping problem that we considered—the practical response—since this was inconsistent with the final value of knowledge). This raises the question of whether the robust virtue epistemic account of knowledge is more naturally allied with one of these responses to the swamping problem over the other.

In order to answer this question, we need to decide whether robust virtue epistemology offers a reductive account of knowledge. That is, is the proposal meant to ‘decompose’ knowledge into constituent parts which can be understood independently of knowledge? The reason why this question is important in this regard is that we saw in the last chapter that epistemic value K-monism is naturally wedded to a ‘knowledge-first’ account of knowledge whereby knowledge is treated as a primitive relative to which other epistemic standings are defined (rather than *vice versa*, as is the case with reductive accounts of knowledge). This is not to suggest that it would be in principle incoherent to advance epistemic value K-monism while nevertheless endorsing a reductive account of knowledge. The point is rather that if one does not already accept knowledge-first epistemology then it is unclear where the motivation for this form of epistemic value monism comes from.

I take it that most, if not all, robust virtue epistemologists regard their account of knowledge as a reductive account, and this certainly seems the default reading to take of the view (i.e., the reading that we should take unless we have grounds for the contrary). If that is right, then this response to the value problem is naturally allied to the pluralist response to the swamping problem. Moreover, on the face of it it does seem right to conceive of robust virtue epistemology in this way. After all, *prima facie* at least, it does seem possible to define cognitive abilities independently of knowledge—e.g., as, roughly, the stable, reliable and cognitively integrated belief-forming traits of the agent.

Still, the devil is in the detail, and it certainly might well turn out on closer inspection that ultimately one is unable to define cognitive abilities without making appeal to knowledge. Indeed, in the extreme case, it may turn out that cognitive abilities need to be defined as those belief-forming traits which are knowledge-conducive, and if that’s the case then clearly robust virtue epistemology cannot be a reductive theory of knowledge. The robust virtue-theoretic response to the value problem would then be more naturally allied to

the monistic response to the swamping problem.

Notice, however, that we do not need to take a stand on this issue here, since either way so long as the robust virtue epistemological account can demonstrate the final value of knowledge, then it will be in a position to answer both the value problem and the swamping problem. The issue of whether the view constitutes a reductive account of knowledge merely influences what kind of conception of epistemic value is in play. That said, although this issue may not be important for our present purposes, one might still regard this issue as salient in light of broader epistemological concerns. For example, one might for various reasons be antecedently suspicious of reductive accounts of knowledge, and hence regard any view which was committed to such a reduction as being *prima facie* implausible. It is because of these broader epistemological concerns that I have flagged this issue here.

§2.5. ACHIEVEMENT WITHOUT KNOWLEDGE

In any case, despite the surface appeal of the robust virtue epistemologist's argument for the final value of knowledge, it faces some critical problems. In particular, the key concern lies with the intermediate conclusion (C1), the K=A thesis. As we will see, this claim is highly problematic on closer inspection. In particular, there are instances of knowledge which don't involve the corresponding cognitive achievement, and there are cognitive achievements which aren't also instances of knowledge. Moreover, as we will see, the case against the K=A thesis in no way depends on the specific account of achievement in play, and thus although there are two premises involved in the argument for the K=A thesis—the achievement thesis, (P1), and the robust virtue epistemological account of knowledge, (P2)—it is the second premise that is the problem. Crucially, however, without the K=A thesis, the robust virtue epistemologist's argument for the final value of knowledge will fail to go through. In this section we will be focussing on the left-to-right entailment that makes up the K=A thesis—*viz.*, the idea that it is sufficient for knowledge that one exhibit the corresponding cognitive achievement.

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 is 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 wouldn’t count it as an achievement. So far so good.

But now consider a third case in which Archie again selects a target at random, skilfully 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 forcefield 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. That is, his success in this case is still primarily creditable to his archery abilities, even despite the luck involved in that success.

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. Given the K=A thesis, however, 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 K=A thesis, consider the case of ‘Barney’ which is structurally analogous to the ‘Archie’ case. 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. That is, his cognitive success in this case is primarily creditable to his cognitive abilities. According to the K=A 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’.⁹ 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 (i.e., 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.

In short, then, robust virtue epistemology is only able to exclude Gettier-style epistemic luck and not also environmental luck. The moral to be drawn is thus that there is sometimes *more* to knowledge than merely a cognitive achievement, contrary to what the robust virtue epistemologist (who defends the K=A thesis) argues. That is, 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. Merely exhibiting a cognitive achievement will not suffice to exclude all types of knowledge-undermining epistemic luck. Call this the problem of *environmental luck*.

How might the defender of the K=A thesis respond to this problem? One response might be to try to evade it by reformulating the achievement thesis, and thus (P1). On the face of it, this might seem like a viable way of dealing with this issue since there are, after all, two premises—the achievement thesis, (P1), and the robust virtue epistemic account of knowledge, (P2)—being used to generate the K=A thesis, (C1). In order to see why such a strategy would be hopeless, however, we only need to note that the ‘Barney’ case is a counterexample *both* to the K=A thesis and to the more specific robust virtue epistemological claim that knowledge is cognitive success that is because of cognitive ability (i.e., (P2)). After all, it is both true that (i) Barney exhibits a cognitive achievement but does not possess the corresponding knowledge, and that (ii) Barney’s cognitive success is because of his

cognitive ability and yet he lacks the corresponding knowledge. There may well be good reasons to reformulate the achievement thesis—indeed, we will consider some reasons on this score in chapter four—but this issue is by-the-by here given that however the achievement thesis is formulated a key premise in the robust virtue epistemologist’s argument for the final value of knowledge is blocked.

A second response to the problem that is superficially appealing is to argue that abilities need to be understood relative to suitable environments in a far more fine-grained than we standardly suppose. One could thus argue that neither Archie nor Barney exhibit an achievement in cases where there is environmental luck in play since contrary to intuition they are not exercising the relevant abilities. Accordingly, the ‘Barney’ case can pose no problem for the K=A thesis.¹⁰

Now it is undeniable that abilities should be understood relative to suitable environments. In crediting you with the ability to play the piano, for example, we are not thereby supposing that you can play the piano underwater. Even so, I take it that we tend to understand what constitutes a suitable environment in a very coarse-grained fashion. Intuitively, for example, Archie is employing the very same archery ability in the case in which there is environmental luck present as he does in corresponding cases where such luck is absent, and the same point goes for Barney’s exercise of his barn-spotting ability in barn façade and corresponding non-barn façade environments.

Still, the proponent of the K=A thesis might well extract from the problem of environmental luck the moral that a more fine-grained conception of the relativization of abilities to suitable environments is required in order to deal with this problem. I think this would be a very theory-driven way of responding to the problem of environmental luck, and I also think that it would ultimately generate a very counterintuitive conception of abilities, one that is in the final analysis extremely fine-grained indeed.¹¹ Still, if the proposal worked then this might be a price worth paying in order to have a response to the problem. The key difficulty facing this response to the problem of environmental luck, however, is that it completely fails to understand what the source of this difficulty is.

In order to see this, consider again the ability to play the piano. We noted above that we would not evaluate the possession of such an ability relative to an environment in which the agent is underwater. Whatever ability our agent is exhibiting, or trying to exhibit, in such

a case, it is not her ordinary ability to play piano. Imagine, however, that, unbeknownst to our agent, she is in an environment in which she could very easily have been underwater right now but in fact is not. It is thus a matter of luck that she is not presently underwater. (Perhaps, say, she is standing in an empty chamber which in most near-by possible worlds is full of water right now). While standing there, she sits down at her piano and begins to play. What ability is she exhibiting? Intuitively, the ability on display here is the very ability to play the piano that she exhibits in normal circumstances. After all, although she could very easily be underwater right now, in fact she is not.¹²

The point of this case is that no matter how fine-grained we might want to make the relativization of abilities to suitable environments, we surely do not want to hold that our piano player is not manifesting her ordinary piano-playing abilities in a case like this. What this demonstrates is that while it is undoubtedly true that abilities should be understood relative to suitable environments, however that point is to be understood it must be compatible with the fact that it can be a matter of luck that one is in a suitable environment to exercise one's ability in the first place. Critically, however, that is just to allow that the presence of environmental luck is compatible with one exercising one's normal abilities (i.e., the abilities one exercises in corresponding cases which don't involve environmental luck). After all, what is key to cases of environmental luck is that while circumstances were indeed as it happens propitious for the exercise of the relevant ability, they could so very easily have not been. Properly understood, then, the issue in hand is not the degree to which abilities should be understood relative to suitable environments, but only whether it is possible for it to be a matter of luck that one is in suitable conditions to exercise one's ability. As we have seen, it can be, and hence however one relativizes abilities to suitable conditions one must allow that environmental luck—and thus environmental *epistemic* luck—is compatible with the exercise of the target ability. Responding to the problem of environmental luck by appeal to the relativization of abilities to suitable environments is thus a theoretical dead-end.¹³

With these options rejected, the prospects for the defender of the K=A thesis look somewhat dim. One possibility might be to 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.¹⁴

Perhaps the robust virtue epistemologist who wishes to retain the K=A thesis could make one of these strategies—or some third strategy, such as denying the intuition that knowledge is incompatible with environmental epistemic luck¹⁵—stick. I don't 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, however, is that there is a further difficulty on the horizon for a view of this sort. Once these two objections for the K=A thesis are taken together, however, they suggest not a mere 'patching-up' of the original proposal, but a radical re-think.

§2.6. KNOWLEDGE WITHOUT ACHIEVEMENT

Consider the following example, due to Jennifer Lackey (2007, §2). Our protagonist, who we'll call 'Jenny', arrives at the train station in Chicago and, wishing to obtain directions to the Sears Tower, approaches the first adult passer-by 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 then we will be understanding the details of the case such that the true belief so formed is non-lucky in all the relevant respects (i.e., 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,¹⁶ 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, and *vice versa*—*viz.*, that when your belief is primarily creditable to your cognitive agency then it is of credit to you, and where it is of credit to you then it is primarily creditable to your cognitive agency. The problem, however, is that closer inspection of these two formulations 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 by to a large degree 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 in this case. It is, after all, a *person* that she asks for directions, and not, say, a lamppost or a dog. Moreover, the person she asks is not a small child, or someone who one might reasonably expect to be unreliable on this score (e.g., someone who is clearly a tourist). In addition, if the testimony which Jenny received were 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, then we would expect her to treat these directions as entirely spurious. So the moral to be drawn from this case is not that sometimes knowledge can be possessed even though the cognitive success in question is of no credit to the agent concerned.¹⁷

Nevertheless, what is true is that it is not *primarily creditable* to Jenny that she has formed a true belief in this case, and this is where the true moral of these cases resides. 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 primarily creditable to one's cognitive agency.¹⁸

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. Again, 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 $K=A$ thesis can respond to cases of this sort. As before, notice that there is no mileage in trying to pin the blame here on the achievement thesis, (P1). This is because, as with the ‘Barney’ case, the ‘Jenny’ case is a counterexample *both* to the $K=A$ thesis and to the more specific robust virtue epistemological claim that knowledge is cognitive success that is because of cognitive ability (i.e., (P2)). After all, it is both true that (i) Jenny has knowledge while failing to exhibit the corresponding cognitive ability, and that (ii) Jenny has knowledge even though her cognitive success is not because of her cognitive ability. As we noted above, there may well be good reasons to reformulate the achievement thesis but this issue is entirely by-the-by here given that however the achievement thesis is formulated a key premise in the robust virtue epistemologist’s argument for the final value of knowledge is blocked.

But if this response to the ‘Jenny’ case doesn’t work, then what other options are there? None that are particularly palatable, that’s for sure. On the one hand, one might bite the bullet and concede that Jenny lacks knowledge after all. On the other hand, one might try to resist this counterintuitive commitment by 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 (and hence has knowledge too). Both strategies involve denying some pretty strong intuitions about this case and so anyone taking either line will face a tough up-hill struggle.

Indeed, notice that 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 this view does have some adherents—most notably Elizabeth Fricker (e.g., 1995)—it is very unpopular, and most in the literature on the epistemology of testimonial belief regard it as a position to be avoided at all costs.¹⁹ It is perhaps for this reason that Greco (2007*b*) 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.

Given that we grant the account of achievement in question—such that achievements are nothing more than successes that are primarily creditable to one’s agency—then I think that Greco’s claim that this easy goal constitutes an achievement is correct (though remember that we will be questioning the adequacy of this account of achievement in the next chapter). The problem, however, is that this case is not relevantly analogous to the case of ‘Jenny’. After all, what is crucial to that example is not that someone appropriately skilful helps Jenny, but rather more specifically that Jenny gains her true belief by (for the most part at least) *trusting* this other person. This is why, for example, other cases in which we depend on the skills of others—as when one takes an inner city road sign a face-value—do not generate the same epistemological moral. In such cases my knowledge depends on—i.e., is made easy by—the skills of others, but it is not that I am merely trusting what the sign tells me: I have all kinds of independent grounds for believing what inner city road signs tell me.

An example that would be relevantly analogous to the ‘Jenny’ case is someone who lacks archery abilities who is being assisted by a skilled archer in firing an arrow and is thereby successful. (For example, the skilled archer helps the novice to take aim, steadies his arm, corrects her posture, and so on). 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, creditable to the combined efforts of the unskilled archer and the skilled archer). We certainly wouldn’t regard the success in question as being primarily creditable to the novice archer. On this basis, then, we would maintain that the unskilled archer’s success does not constitute a *bona fide* achievement, and hence cases like this should give us no cause to reconsider our original assessment of the ‘Jenny’ example as one in which the agent likewise does not exhibit an achievement.

In any case, whatever the defender of the K=A thesis says in response to the ‘Jenny’ example, remember that she must also simultaneously deal with the other problem outlined above—concerning the apparent possibility of cognitive achievements which are not cases of knowledge. Indeed, notice that it is significant that these two problems pull robust virtue epistemologists who endorse the K=A thesis in two different directions. Whereas the ‘Jenny’ case puts pressure on them to *weaken* their robust virtue epistemology and thus allow cases of knowledge which this view would ordinarily exclude, the ‘Barney’ case, in contrast, puts

pressure on them to *strengthen* their account in order to explain why merely exhibiting a cognitive achievement does not suffice for knowledge. This is why when these two problems are expressed in tandem they pose such a tricky difficulty for the robust virtue epistemologist.

It seems, then, that the K=A thesis is unsustainable. Moreover, the source of the problem with the K=A thesis that we have explored here is clearly the robust virtue epistemological account of knowledge, (P2). Without the K=A thesis, however, the robust virtue epistemologist loses the ingenious basis on which she argued for the final value of knowledge.²⁰

§2.7. BACK TO THE VALUE PROBLEM

Does that mean that the response to the value problem offered by robust virtue epistemology must be completely abandoned? Perhaps not. On the face of it, one might think that there is a fairly straightforward way of resurrecting the K=A 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. Since we have granted the *prima facie* plausibility of the thesis that achievements, and thus cognitive achievements, are finally valuable, the fact that knowledge at least tends to go hand-in-hand with cognitive achievement would suffice to show that knowledge at least tends to be finally valuable, even if it is sometimes not of final value. Would that be enough to answer the tertiary value problem?

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.²¹

Even if we can block this worry by arguing for a close relationship between knowledge and cognitive achievement, however, a second, and more substantive, worry remains. Recall that to say that knowledge is distinctively valuable is to say that it is more valuable, not just as a matter of degree but of kind, than that which falls short of knowledge. On this view, however, there is an epistemic standing which falls short of knowledge and which is no less valuable (indeed, which is in its nature finally valuable): cognitive achievements that are not also cases of knowledge. If that's right, then even if knowledge is the kind of thing that tends to be finally valuable, it still won't follow that knowledge is in the relevant sense distinctively valuable.

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. But given that this approach to the value problem constituted the best—indeed, the *only*—response to the value problem that seemed able in principle to support the key claim that knowledge is finally valuable, this means that the prospects for answering the value problem now appear dim indeed.

That leaves the swamping problem. Here the more modest strategy of arguing that knowledge tends to be finally valuable may gain more purchase. For recall that we noted in chapter one that so long as knowledge is at least *sometimes* finally valuable then it will constitute a fundamental epistemic good. If that's right, however, then that would suffice to block the swamping argument, since that argument essentially trades on a commitment to epistemic value T-monism. Of course, as we noted earlier in this chapter, that one blocks the swamping argument by appeal to the fact that knowledge is a fundamental epistemic good leaves it open whether one is committed to a monistic or a pluralist response to that problem. But which way one jumps on this issue rests on further theoretical claims which are of secondary importance to us here.

So even though the robust virtue epistemic defence of the final value of knowledge is ultimately unsuccessful, so long as we accept the thesis that cognitive achievements are

finally valuable (a claim that we will explore in more detail in chapter four) there may be scope to use the fact that knowledge at least tends to go hand-in-hand with cognitive achievements as a means of answering the swamping problem. Whether or not this is the best way to deal with the swamping problem is an issue that we will return to.

§2.8. CONCLUDING REMARKS

Given what is at stake in answering the value problem—and thus the tertiary value problem—one might argue that we simply cannot leave matters at that, and that instead we must continue to seek a resolution to this problem. Indeed, it has been suggested by some that it is an adequacy condition on any theory of knowledge that it is able to account for the distinctive value of knowledge, in the sense that if one's theory is unable to do this then this is a definitive strike against one's view.²²

This way of thinking about the value problem and its role in the theory of knowledge is, however, surely too strong. Instead, what is presumably required is *either* that one's theory of knowledge can answer the value problem *or* that one's theory is able to provide some plausible account of why knowledge isn't 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 desiderata just identified, then that should suffice.

In the next chapter we will be exploring a new theory of knowledge—what I call *anti-luck virtue epistemology*—which has many of the advantages of robust virtue epistemology and none of its failings. Although this account of knowledge is unable to account for the final value of knowledge, it does enable us to gain an understanding of why knowledge may initially appear to be finally valuable. This diagnostic story regarding our intuition about the distinctive value of knowledge is further reinforced in chapter four where we will explore an epistemic standing—understanding, or at least a core kind of understanding anyway—which is finally valuable, and which has a close relationship to knowledge. As we will see, this discussion will lead us to rethink some of the claims about the nature and value of achievements that we have looked at in this chapter. So although the main conclusion drawn in this chapter is a negative one, from the ashes of the failure of robust virtue epistemology to

adequately respond to the value problem we will be extracting some important positive epistemological conclusions.

NOTES

¹ Note that henceforth I will take it as given that a genuine cognitive ability is reliable.

² This example is adapted from one offered by Chisholm (1977, 105).

³ For example, 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.

⁴ Those familiar with the literature in this respect will recognise that often virtue epistemologists like Greco (e.g., 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 (i.e., 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 §2.4.

⁵ One surprising consequence—noted in Greco (2007*b*; 2008)—is that robust virtue epistemology becomes committed to a kind of attributer contextualism due to the context-sensitivity of causal explanations. As it happens, Greco welcomes this result, but as I have argued elsewhere—e.g., Pritchard (2008*d*)—he is unwise to do so. But if one does not analyse the because of relation in causal explanatory terms, then what are the alternatives? This is unclear. Zagzebski (e.g., 1999) ultimately treats this relation as a primitive, though she notes that a good approximation is the sensitivity principle (i.e., if the proposition believed had not been true then the agent would not have believed it). Perhaps the best alternative account on offer is that put forward by Sosa (2007) who argues that we should think of the relation in terms of the exercise of a power. For further discussion of Sosa's proposal, see footnote 20.

⁶ 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 (*forthcoming*, §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—i.e., 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—i.e., that 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 more on this point see Pritchard (2008*g*, §2). Brogaard (2007*a*) is one commentator who has recognised this point. See also Percival (2003). For two recent, and influential, discussions of the intrinsic value/final value distinction, see Rabinowicz & Roennow-Rasmussen (1999; 2003).

⁷ I discuss the relevance of achievements to the problem of the meaning of life in Pritchard (2008*a*).

⁸ 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 isn't. For example, one might argue that pleasure is in its nature finally valuable even though some pleasures (i.e., 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 (i.e., the wicked or trivial ones) lack final value. I am grateful to Mike Ridge for this suggestion.

⁹ In Pritchard (2005, ch. 5) I delineate the core kind of knowledge-undermining epistemic luck and label it 'veritic luck'. Both environmental epistemic luck and Gettier-style 'intervening' epistemic luck fall under the more general category of veritic luck.

¹⁰ This strategy forms a key part of the response to the problem of environmental luck offered by Greco (2007*b*, §5).

¹¹ For more discussion of this point, see Pritchard (2008*d*; cf. Pritchard 2008*h*) and Kvanvig (*forthcoming*).

¹² Moreover, notice that since our agent is unaware of how modally close she is to disaster, it is not as if her awareness of this danger could have a bearing on whether this environment is suited for the exercise of this ability.

¹³ One can strengthen this point by noting that it is in fact incidental to both the ‘Archie’ and ‘Barney’ cases that the relevant deception is *actually* occurring in the subject’s environment, albeit in such a way that it does not affect the exercise of the target ability. That is, one could redescribe both cases as involving no deception in the actual world but only in most near-by possible worlds and the cases would still demonstrate the same point. Interestingly, Sosa (2007) is one philosopher who has recognised that what is at issue here is the ‘fragility’ of the exercise of one’s abilities (i.e., that it can be a matter of luck that one is in a position to exercise them). His response has been to argue that an agent like Barney *does* have knowledge—i.e., he claims that knowledge is entirely compatible with environmental epistemic luck of this sort. Given the strong intuitions which support the barn façade case, such an approach will inevitably be highly contentious. For further discussion of Sosa’s view in this regard, see Pritchard (*forthcominga*). See also footnote 20.

¹⁴ Greco (2002, §3; cf. Greco *forthcoming*, §5) 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 (2007b, §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. For further discussion of Greco’s response(s) to the problem of environmental luck, see Pritchard (2008d; cf. Pritchard 2008h) and Kvanvig (*forthcoming*).

¹⁵ As noted in endnote 13, since is the line taken by Sosa (2007).

¹⁶ The title of the paper in which this example appears is ‘Why We Don’t Deserve Credit for Everything We Know’.

¹⁷ 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.

¹⁸ Moreover, 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 (e.g., where the cognitive achievement in question is very easy).

¹⁹ This is not the place to explore this issue in more detail. For more on the epistemology of testimony, see Adler (2006).

²⁰ One might think that the alternative robust virtue-theoretic account offered by Sosa (2007) could potentially offer a way out of this problem. Rather than understanding the ‘because of’ relation that is key to robust virtue epistemology in causal explanatory terms, Sosa instead understands it in terms of the manifestation of a power. To see how these two accounts can come apart, consider a glass that was broken as a result of someone deliberately smashing it against a wall. Ordinarily, the most salient part of the causal explanation of why the glass broke will be that someone smashed it against the wall, and in this sense it will be true to say that the glass broke because it was smashed against the wall. Note, however, that this is consistent with the claim that it was because of the glass’s fragility that it broke, since here we are talking about the manifestation of a power and not offering a causal explanation. Sosa’s idea is that when the robust virtue epistemology claims that knowledge is cognitive success that is because of cognitive ability, it is the ‘manifestation of a power’ reading that we should adopt, and not the causal explanatory reading. Although Sosa is offering a genuine alternative to Greco’s reading of the robust virtue-theoretic account of knowledge, I think it should be clear that his view is no less susceptible to the problem posed here. In the Barney case, for example, it is surely even clearer that Barney’s cognitive success constitutes the manifestation of his cognitive powers than that it is primarily creditable to his cognitive abilities. Indeed, Sosa recognises this, which is why—as noted in footnote 13—he argues that, contrary to intuition, environmental epistemic luck is compatible with knowledge possession. Moreover, Sosa’s view will also struggle with the Jenny case too, since again it is surely even clearer that Jenny’s cognitive success is not the manifestation of *her* cognitive powers than that her cognitive success is not primarily creditable to her cognitive ability. This is not to say that Sosa’s proposal is a complete non-starter as a type of robust virtue epistemology, since for one thing it at least avoids some of the counterintuitive consequences of Greco’s view that we noted in footnote 5. The point is rather that adopting such a reading of robust virtue epistemology doesn’t offer any easy resolution of the problem in hand.

²¹ 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). The relevance of the problem of epistemic value to radical scepticism is discussed more fully in chapter four. See also Pritchard (2008*e*).

²² One finds a claim 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 (*forthcoming*). It is important to note, however, that these authors almost certainly have a different view of what it would take for knowledge to be distinctively valuable to that which has been argued for here.