

# THE VALUE OF KNOWLEDGE

DUNCAN PRITCHARD  
University of Edinburgh

ABSTRACT. It is commonly held that knowledge is of distinctive value to us. The goal of this paper is to unpack and examine the most promising defence of this claim found in the recent literature which argues that knowledge is distinctively valuable because, *qua* cognitive achievement, it has final value. While it is argued that this proposal is unsuccessful, it is also demonstrated that a correct understanding of why this proposal fails has important ramifications for our understanding not just of the value of knowledge, but also of knowledge itself.

## 1. THE VALUE PROBLEM

It is widely thought 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 theorising. Understanding just why knowledge is distinctively valuable to us, however, has proved elusive. Call this *the value problem*.

Notice that the value problem, as it stands, is ambiguous in an important respect. Are we simply seeking an explanation of why, *whether rightly or wrongly*, we regard knowledge as distinctively valuable? Or are we seeking an explanation of why knowledge *is* distinctively valuable which can in turn explain why we regard knowledge as distinctively valuable? The distinction is of course important, since one will respond to the value problem in very different ways depending on which reading one takes. For now, I want to read the value problem as seeking an explanation of the second sort. As we will see below, however, a failure to adequately respond to the value problem, so construed, could force a retreat into the other construal of this problem; indeed, a construal on which what we seek is an explanation of why we *wrongly* regard knowledge as distinctively valuable.

Part of the difficulty posed by the value problem involves getting clear about just

what it means to say that knowledge is distinctively valuable to us. One minimal reading of this claim is that knowledge is more valuable than mere true belief.<sup>1</sup> Call the challenge to explain why knowledge is more valuable than mere true belief the *primary value problem*. Now it is almost certain that we would need to do more to account for the distinctive value of knowledge than simply answer the primary value problem. Nevertheless, this problem is of interest to us because if we were unable to respond to even this challenge then the whole project of accounting for the distinctive value of knowledge would become hopeless.

Indeed, there is reason to think that we are unable to answer the primary value problem. Initially, we might appeal to the fact that knowledge appears to be of more instrumental value than true belief but, as Socrates famously notes in the *Meno*, this claim is far from obvious on closer inspection. After all, a true belief about the correct way to Larissa is surely of just as much practical use as knowledge of the way to Larissa—both will get us to our destination.

We need to be clear about the nature of this objection. One might be tempted to read it as saying that there exist cases in which knowledge that *p* is of no more value than mere true belief that *p*, and thus that the primary value problem is not met. This is only right, however, if what the primary value problem is read as demanding is an explanation of why knowledge that *p* is *always* more valuable than mere true belief that *p*. Call this the *austere reading*. It is far from clear why we should endorse the austere reading of the primary value problem. For one thing, it seems that the primary value problem, construed austere, is unanswerable, since there are surely bound to be cases in which it would be better to merely truly believe a proposition rather than know it (e.g., where one's life depends on it, for example). More importantly, however, the mere fact that there are such cases does not seem to decide the matter of whether knowledge is more valuable than mere true belief. If that's right, then there is something seriously amiss with construing the primary value problem in this severe fashion.

With this in mind, consider a *modest* reading of the primary value problem which merely demands an explanation of why we should, as a matter of general policy, *prefer* knowledge that *p* over mere true belief that *p*. For example, suppose one were able to show that knowledge is *typically* of more value than mere true belief. Intuitively, this would suffice to answer the primary value problem. After all, what would follow from this claim is that if one were faced with the choice between knowledge that *p* and mere true belief that

$p$ —where, crucially, one did not know in advance which proposition ‘ $p$ ’ referred to—there would be a solid rationale for preferring knowledge, even though it may well turn out that in this particular case knowledge that  $p$  is of no more value than a mere true belief that  $p$ .

Indeed, it may be possible to meet the modest reading of the primary value problem by demonstrating something weaker than the claim that knowledge is typically of more value than mere true belief. Suppose, for example, that knowledge is typically of no less value than mere true belief, and in a wide range of cases (though not typically) is of more value than mere true belief. Would this suffice to meet the primary value problem, modestly construed? Quite possibly. After all, this would almost certainly suffice to supply the needed rationale for preferring, for any non-specified proposition that one is presented with, knowledge of that proposition over mere true belief in that proposition. In what follows, however, we will read the primary value problem, modestly construed, as demanding a demonstration of the typical greater value of knowledge over mere true belief. Since this makes the primary value problem harder to answer rather than easier, this ought to be uncontentious.

In any case, the point of the foregoing is not to suggest that Socrates was failing to engage with the primary value problem properly, since one can easily re-cast his concern in terms of the modest reading. After all, although the ‘Larissa’ example simply appeals to one case, the difficulty it adverts to is not peculiar to this case since we could substitute any number of parallel scenarios to make the same point. Accordingly, it poses a challenge even to the thesis that knowledge is typically of greater value than mere true belief. There is thus still a difficulty posed by the primary value problem even on the modest reading.<sup>2</sup>

Even supposing that we are able to respond to the primary value problem, modestly construed, 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 sub-set 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 value to a mere true belief (justification, say), but suppose further that the satisfaction of this condition, in conjunction with true belief, was not sufficient for knowledge. 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 typically 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 would generally prefer over mere true belief). 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 (i.e., true belief plus the value-conferring property X). But if that's right, then why do we regard knowledge as distinctively valuable at all?<sup>3</sup>

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. As before, the secondary value problem is subject to both an austere and a modest reading, but it is only the latter that we are interested in. That is, what we seek is some account of the greater value of knowledge over that which falls short of knowledge which could offer a rationale for preferring, for any non-specified proposition that one is presented with, knowledge of that proposition over an epistemic standing as regards that proposition which falls short of knowledge. If knowledge is only typically of greater value than that which falls short of knowledge, that will suffice to meet the secondary value problem on the modest reading.

Even if we can offer a response to the secondary value problem, 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, is that it fails to explain why the long history of epistemological discussion has focussed 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 but also a different kind of value than whatever falls short of knowledge. Call this the *tertiary value problem*.

As with the primary and secondary value problems, one can cast this problem in terms of either an austere or a modest reading, only the second of which is our concern here. That is, what we seek is only, at most, a defence of the thesis that knowledge is typically of more value, as a matter of kind rather than merely degree, than that which falls short of knowledge.

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. As noted above, 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.<sup>4</sup>

## 2. THE 'SWAMPING' PROBLEM

One natural response to the primary value problem is to argue that where knowledge and mere true belief crucially differ is that only the true belief at issue in the former is acquired via a reliable process, where a reliable process is one that is likely to lead to the truth. Insofar as we value true beliefs, it follows that we will value reliable belief-forming processes on account of how they enable us to acquire true beliefs. Hence, given the further plausible assumption that knowledge entails reliability, we should value knowledge more than mere true belief because only the true belief at issue when one has knowledge has the additional value that accrues in virtue of being reliably formed.

This line of argument faces a problem, however. On this account we only value the reliability of the way in which a belief is formed because reliability is conducive to attaining a certain good, true belief. If that good is already attained, however, as would be the case if one has a true belief, then that the true belief is in addition reliably formed seems neither here

nor there when it comes to assessing its value. Thus, the value conferred by being formed by a reliable process does not seem to enhance the value of a true belief after all, and so this response to the primary value problem appears to be in jeopardy.<sup>5</sup>

This problem is known as the “swamping problem” because of how the value of true belief seems to ‘swamp’ the value of the belief in question being reliably formed.<sup>6</sup> In order to illustrate this problem, consider the example—due to Linda Zagzebski (2003)—of a great cup of coffee which is produced by a reliable coffee making machine (i.e., one that usually produces great coffee). Clearly, such a cup of coffee is no more valuable than an identical cup of coffee produced by an unreliable coffee making machine. As regards both great coffee and true belief, argues Zagzebski, the reliability of the process generates value in virtue of its tendency to produce a certain valuable effect (great coffee/true belief), but this means that where the effect has already been produced—where one has a great cup of coffee or a true belief—then the value of the product is not improved upon by being produced in a reliable way.

One problem with this argument is that it appears to rest on a failure to recognise an important distinction that is common currency in value theory.<sup>7</sup> In particular, the value distinction that this argument ignores is that between *intrinsic value*—where an object’s value is determined by its intrinsic properties—and *final value*—where we value an object for its own sake (i.e., non-instrumentally).<sup>8</sup> The swamping problem in effect presupposes that the only way that the relational properties of a true belief—in particular, the way that the belief was formed—could contribute value would be instrumentally. But that it is to simply ignore the possibility that the relational properties of a true belief could contribute final value.

Consider, for example, a book printed on the first ever printing press and an exact replica produced by lasers. We would undoubtedly value the former more than the latter even though we can accept, for the sake of argument, that these two objects have the same relevant intrinsic properties. Moreover, we would value the former for its own sake, quite apart from what instrumental value it had (e.g., its monetary value). The difference in value, however, is clearly due to the relational properties of the objects concerned. The reason for this difference in value is that a book printed on the first ever printing press, unlike an exact replica, has final value—it is valuable for its own sake because of how it is produced (i.e., its relational properties).<sup>9</sup>

It therefore does not follow from the fact that a true belief produced by an unreliable process has the same relevant intrinsic properties as a true belief produced by a reliable process that they are equivalent in value, which is what the swamping argument assumes, since it could well be that a reliably formed true belief has final value that a mere true belief lacks.<sup>10</sup> That is, the swamping argument only shows, at most, that the additional value conferred by the reliability of a belief-forming process onto a true belief cannot be instrumental epistemic value; it does not show that that this is the only possible type of additional value that could be conferred.

Of course, the mere fact that the swamping argument overlooks the distinction between intrinsic and final value does not mean that it fails to go through. Whether or not this is the case depends on the further issue of whether reliable true belief is the sort of thing which, unlike mere true belief, has final value. On the face of it, however, this does not look very plausible. Why should the mere fact that one's true belief was formed in a reliable fashion make that true belief something that is valuable for its own sake? The case of the reliably produced cup of coffee is instructive here. While we might value a book produced by the first ever printing press for its own sake because of how it is produced, we do not value a cup of coffee for its own sake because of the reliability of the coffee-making machine that produced it. The question we need to ask is: Why isn't reliable true belief in the relevant respects like reliably produced coffee rather than like a book produced on the first ever printing press?

### 3. KNOWLEDGE AND FINAL VALUE

Let us grant that the mere fact that a true belief is produced by a reliable belief-forming process will not suffice to confer final value. Nevertheless, that doesn't mean that the general strategy of arguing that knowledge is typically of final value because of the way in which it is acquired is hopeless. In particular, consider virtue-theoretic accounts of knowledge which hold that knowledge is to be understood as true belief that arises out of one's intellectual virtues. There are various ways of understanding the intellectual virtues, some more permissive than others, but the common thread is that such virtues are at least reliable cognitive traits which are stable parts of one's cognitive character.<sup>11</sup> Henceforth, we will refer

to these reliable cognitive traits that are stable parts of one's cognitive character as 'cognitive abilities'.

So construed, the reliabilist ancestry of virtue-theoretic accounts of knowledge is clear. There are, however, a number of good reasons for preferring a virtue-theoretic account of knowledge over a bare process reliabilist account, even though both views incorporate a reliability requirement. In particular, one of the key problems that faces a simple process reliabilist view is that it appears to be far too permissive. That is, such a view seems to predict knowledge in cases where, intuitively, we would not regard knowledge as being possessed.

To illustrate this point, consider the case of Temp who forms her beliefs about the temperature in a room by looking at a thermometer which she has every reason to believe is reliable. Suppose further that this is indeed a reliable way of forming beliefs about the temperature of the room, in the sense that every time Temp forms a belief about the temperature of the room in this way her belief is true. Nevertheless, the thermometer is, in fact, broken and is randomly fluctuating within a given range. The twist to the story is that there is someone hidden in the room next to the thermostat who, unbeknownst to Temp, is making sure that every time she forms a belief about the temperature of the room by looking at the broken thermometer the reading on the thermometer corresponds to the temperature of the room.

Clearly, one cannot gain knowledge of what the temperature of the room is by looking at a thermometer which is, unbeknownst to one, broken, even if this way of forming one's beliefs is reliable. What such cases illustrate is that what we have in mind when we take reliability to be important to knowledge is not a bare reliability of the sort that is exhibited in the 'Temp' case, but rather a more specific type of reliability which reveals a relevant cognitive sensitivity to the facts. What is going wrong in the Temp case is that the facts are, as it were, being sensitive to the agent's beliefs, rather than *vice versa*. In this way, the true beliefs in question are not arising out of Temp's cognitive abilities at all, but are instead due to the 'helpful' person hidden in the room. By restricting the range of reliable cognitive processes to those stable and reliable traits that make up one's cognitive character, one can thus deal with cases like this that plague a simple process reliabilism.

Merely moving from a simple process reliabilism to a virtue reliabilism will not deal with all cases, however. In particular, a virtue reliabilism seems to be susceptible to Gettier-

style cases in much the same way as reliabilism is. To illustrate this point, consider the case of Roddy who, upon looking into a field and seeing a sheep-shaped object, forms the true belief that there is a sheep in the field. Unfortunately for Roddy, however, what he is looking at is in fact not a sheep but a big hairy dog. Nevertheless, his belief is true since there is a sheep in the field, hidden from view behind the dog.<sup>12</sup> The problem is that in this case there is no obvious reason why we should deny that Roddy has indeed formed a true belief through the exercise of his cognitive abilities. After all, there is no suggestion in this case that he was being sloppy in forming his belief in this way—we can suppose, for example, that he had good reason for thinking that his eyesight was adequate to the task, and no reason for thinking that he should be wary of the possibility that there might be big hairy sheep-shaped dogs in the vicinity. Nevertheless, this is not a case of knowledge since it is just a matter of luck that Roddy's belief is true.

Now one might reasonably respond to cases of this sort by arguing that the problem of eliminating the sort of knowledge-undermining luck in play in Gettier-style cases is a problem for nearly all theories of knowledge, and hence it is legitimate for virtue epistemologists to bracket this difficulty. According to this view, knowledge is to be understood as non-Gettierized true belief that arises out of the reliable cognitive traits that make up one's cognitive character.<sup>13</sup>

As some virtue theorists have noted, however—among them Sosa (1988; 1991; 2007), Zagzebski (1996; 1999) and John Greco (e.g., 2002; 2007a; 2007b; *forthcoming*)—there seems to be a virtue-theoretic response to the Gettier-style cases available. This response proceeds 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 seen, is compatible with Gettier-style cases—the strengthened virtue-theoretic thesis is that knowledge only results when the truth of an agent's belief is *because of* the operation of the agent's cognitive abilities—i.e., where it is *primarily* creditable to the agent that her belief is true.<sup>14</sup>

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 big hairy dog

that he was looking at. In contrast, had he actually been looking at a sheep, then his belief *would* have been true because of the operation of his cognitive abilities.

As Greco (*forthcominga*) 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 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. 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.

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 (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 which that which falls short of knowledge lacks, and hence show that it is distinctively valuable.<sup>15</sup>

In order to see why achievements are 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 was 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 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.

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, all things considered, the overall value of Archie’s success of hitting the target when it is due to luck is much greater than the overall value of a corresponding success attained because of Archie’s ability. It is important to recognise that the thesis that achievements *qua* achievements accrue final value is entirely consistent with this possibility.<sup>16</sup>

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 overall 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.

#### 4. ACHIEVEMENT WITHOUT KNOWLEDGE

Despite the surface appeal of this proposal, there are several fundamental problems facing it. The difficulties lie, however, not with the claim that cognitive achievements accrue final value—a thesis which I think is plausible—but 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 starts to become moot.

Consider again the case of Archie, who selects a target at random from a 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 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, this 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.

The problem that cases like this pose for the 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—i.e., 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 (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.

Interestingly, however, it seems that the virtue-theoretic proposal to understand knowledge as true belief that is because of cognitive ability 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 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 seems to be 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 in the case in which 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.<sup>17</sup>

Perhaps the proponent of the knowledge-as-achievement could make one of these strategies (or some third strategy) 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 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-think.

## 5. KNOWLEDGE WITHOUT ACHIEVEMENT

Consider the following case, due to Jennifer Lackey (2007, §2).<sup>18</sup> 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 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 case is that sometimes one can have knowledge without the success in question being in any way creditable to the agent. The example does not warrant this conclusion, however. After all, it is certainly to a certain degree creditable to Jenny that she has gained a true belief in this way, since she does appropriately exercise her relevant cognitive abilities. 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 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 of this case is not that sometimes knowledge can be possessed even though the cognitive success in question is not in any way creditable to the agent concerned.<sup>19</sup>

Nevertheless, it is far from obvious that it is *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. Accordingly, the *bona fide* knowledge that results in this case does not seem to be a cognitive achievement on Jenny's part. So if one takes this

example at face-value, it seems that one can have knowledge without that knowledge being 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 knowledge-as-achievement thesis can respond to cases of this sort short of simply maintaining, *contra* intuition, that it is primarily creditable to Jenny that her belief is true in this case, and thus that it is a *bona fide* achievement after all.<sup>20</sup> Perhaps a suitable rationale can be offered for such a thesis, though I think it is clear that one who wishes to defend such a line faces a tough up-hill struggle. 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 up-hill struggle in question is particularly onerous.

## 6. DIAGNOSIS

I think we can offer a diagnosis of what has gone wrong here. There are two closely related ‘master’ intuitions as regards knowledge that inform much of the contemporary epistemological discussion. The first is that knowledge is the product in some way of one’s cognitive abilities, such that when one knows one’s cognitive success is, in substantial part at least, creditable to one. It is this intuition that is behind the thought that if an agent forms a true belief in a reliable fashion which does not arise out of her cognitive ability—as in the ‘Temp’ case described above—we ought not to attribute knowledge. 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.

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 (i.e., 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 could be true beliefs which meet the relevant anti-luck condition—such that the agent couldn't have easily been wrong—which are not thereby formed as a result of a cognitive ability. (Indeed, the 'Temp' example described above—perhaps suitably modified—could be used to illustrate this point).

Virtue epistemologists are clearly most impressed by the first master 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 'ability' intuition and, in addition, an anti-luck condition which accommodated the second master intuition (and so dealt with Gettier-style cases). Indeed, we noted the appeal of such a proposal above. As we have seen, however, there is a prominent strand of virtue epistemic thought which in effect tries to accommodate the motivation for the 'anti-luck' master intuition by responding to the 'ability' master 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 have seen 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 master intuition as its lead. On this view, one tries to formulate an anti-luck condition on knowledge in such a robust way that one does not need an additional 'ability' condition to accommodate the other master intuition.

Indeed, one could read certain modal epistemological proposals in the recent literature as proposing just such a view.<sup>21</sup> That is, just as the virtue epistemologists who advance the knowledge-as-achievement thesis in effect try to accommodate the motivation for the anti-luck master 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—such as a suitably adapted version of the ‘Temp’ case—in which there are true beliefs which exhibit the relevant 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 master 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.

I will not be arguing for this claim here, however, since my interest is rather in whether the fact that cognitive achievements are distinctively valuable can help support the claim that knowledge is distinctively valuable. That is, supposing that one rejects the knowledge-as-achievement thesis, does the fact that cognitive achievements are finally valuable offer us any way of responding to the value problem?

## 7. BACK TO THE VALUE PROBLEM

Suppose that instead of opting for a theory of knowledge which aims to do away with either a separate anti-luck condition or a separate ability condition, one simply responds to the two master intuitions just outlined by opting for a theory which incorporates *both* an anti-luck and an ability condition. With the ability condition construed along virtue-theoretic lines, the proposal would be that it is essential for knowledge that one’s true belief 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. Rather than strengthening this condition in order to meet the Gettier-style cases, however, this ability condition is instead supplemented with an anti-luck condition. We do not need to get into the detail of how such a condition should be formulated here. What is important is only that it captures the idea that

the agent could not have easily been wrong.<sup>22</sup>

By weakening the ‘ability’ requirement, this proposal does not face the problem posed by cases like ‘Jenny’, in which the agent has a non-lucky true belief which is to some significant degree creditable to the agent—and which therefore qualifies as knowledge on this proposal—but where the agent does not exhibit a cognitive achievement. Given that Jenny is in fact in an epistemically friendly environment—for example, no-one is out to trick her on this score—the anti-luck condition is easily met.<sup>23</sup> Moreover, her true belief is also sufficiently creditable to her to ensure that she meets the ability condition. No wonder, then, that we regard her as having knowledge.

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, this view has a lot to recommend it. But can it enable us to make any progress towards answering the value problem? Recall a distinction that I drew in §1 between modest and austere responses to the value problem. I noted, for example, that it would be implausible to suppose that a response to the primary value problem should include a defence of the claim that every single instance of knowledge is of more value than the corresponding mere true belief. (Such a demand would, in any case, be unmeetable). Rather, what we are seeking in this case is some motivation for why we should prefer knowledge over mere true belief such that, for example, if faced with the choice between knowledge of  $p$  and a mere true belief that  $p$  (where we don’t know in advance which proposition ‘ $p$ ’ refers to), there is a solid rationale for opting for knowledge. I argued that to meet this challenge it ought to suffice to show that knowledge is *typically* of greater value than mere true belief. I also noted that what goes for the primary value problem on this score also applies to the secondary and tertiary value problem.

With this point in mind, then just so long as knowledge is *typically* a cognitive achievement, and thus of final value, then we may be in a position to answer the tertiary

problem after all and, with it, the value problem more generally (at least given the further plausible assumption that we outlined at the end of §3). Whether such a response to the value problem could be successful depends upon two further issues. Unfortunately, as we will now see, closer inspection of these issues reveals that this response to the value problem is in fact unlikely to be successful.

The first issue concerns just how extensive the cases of knowledge are which are not cognitive achievements. 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 every reason to hold that there may be non-testimonial cases that have the relevant features. After all, 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.<sup>24,25</sup>

Even if we can block this worry by arguing for a close relationship between knowledge and cognitive achievement, a second worry remains. I noted above that on a modest reading of the tertiary value problem one may well be able to answer this problem merely by demonstrating that knowledge is only *typically* a cognitive achievement, and so of final value. This strategy will not work, however, if there is something which falls short of knowledge which is *always* of final value, since then it would seem that what we should prefer is this lesser epistemic standing rather than knowledge. But that seems to be exactly the situation that we are left with on this view, in that, strictly speaking, it is cognitive achievements, rather than knowledge, which are distinctively valuable and yet there are cases of cognitive achievements which are not cases of knowledge. Accordingly, one might naturally argue that the moral to be drawn is that we should seek the former rather than the latter. But if that’s right, then how is one to answer the tertiary value problem by appeal to the fact that knowledge often involves a cognitive achievement?

One might respond to this line of argument by pointing out that cognitive achievements do not fall short of knowledge on the view under consideration, since it is true both that there are cognitive achievements which are not cases of knowledge *and* that there

are cases of knowledge which are not cognitive achievements. It could therefore still be the case that if knowledge is generally an achievement, and so of final value, then there is a solid rationale for preferring knowledge over that which falls short of knowledge.

This line of reasoning, while superficially appealing, is, however, ultimately unhelpful. After all, while it is true that cognitive achievements do not fall short of knowledge, it is nevertheless the case that the more specific category of ‘cognitive achievements which are not cases of knowledge’ *does* fall short of knowledge. Moreover, on the present account there are no grounds for holding that we should think that knowledge is more valuable than this lesser standing as a matter of kind and not merely degree.

The present proposal thus lacks the resources for answering the tertiary problem as it stands and will thus require supplementation. Alas, it is far from clear what such a supplementation would involve.

## 8. CONCLUDING REMARKS

At the start of this paper I noted that we need to make a distinction between the question of why, whether rightly or wrongly, we regard knowledge as distinctively valuable, and the question of why we *rightly* regard knowledge as distinctively valuable. Thus far we have focussed on the first of these questions. Unfortunately, we have found that the most promising response to this problem currently on the market—which identifies knowledge with cognitive achievements and argues that it is finally valuable on this basis—is unsuccessful. Moreover, we have also just seen that an adapted proposal to resolve the tertiary value problem which is in the same spirit—one which retains the thesis that cognitive achievements are finally valuable, but argues instead for the weaker claim that *typically* cases of knowledge are also cases of achievement—is also highly suspect on closer analysis. The prospects of providing an answer to the first of these questions therefore do not look good.

There is, however, a key difference in the problems that each of the two responses to the value problem that we have looked at face. For notice that the problem with the first view was not that it was unable to answer the value problem, but rather that it was an unsatisfactory account of knowledge which could not deal with some key counterexamples. In contrast, the second view does appear to be a plausible account of knowledge; it is just that

it is unable, it seems, of answering the value problem.

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.<sup>26</sup> This is surely too strong, however. Instead, what is presumably required is *either* that the theory can answer the value problem *or* that the 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 a theory of knowledge can answer the second of the two questions just identified, then that seems to suffice. Interestingly, the account of knowledge offered here may well be able to offer a good answer to this second question.

Perhaps, for example, 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 is also a cognitive achievement.<sup>27</sup> This certainly seems plausible, since 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 we would naturally suppose that knowledge is distinctively valuable. Under closer inspection, however, we have found that the relationship between knowledge and cognitive achievement is not as straightforward as it first appears, and thus that there is good reason to think that knowledge is not distinctively valuable at all, even though it seems to be.

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 *isn't* 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 theorising, on account of the light that it casts on the epistemological project.<sup>28</sup>

## REFERENCES

- Axtell, G. (1997). 'Recent Work on Virtue Epistemology', *American Philosophical Quarterly* 34, 1-26.  
 Baehr, J. (Forthcoming). 'Is There a Value Problem?', *Epistemic Value*, (eds.) A. Haddock,

- A. Millar, & D. H. Pritchard, Oxford University Press, Oxford.
- Brogaard, B. (2007). 'Can Virtue Reliabilism Explain the Value of Knowledge?', *Canadian Journal of Philosophy* 36, 335-54.
- Chisholm, R. (1977). *Theory of Knowledge* (2<sup>nd</sup> Ed.), Prentice-Hall, Englewood Cliffs, New Jersey.
- DePaul, M. (*Forthcoming*). 'Does an Ugly Analysis Entail that the Target of the Analysis Lacks Value?', *Epistemic Value*, (eds.) A. Haddock, A. Millar, & D. H. Pritchard, Oxford University Press, Oxford.
- Goldman, A., & Olsson, E. (*Forthcoming*). 'Reliabilism and the Value of Knowledge', *Epistemic Value*, (eds.) A. Haddock, A. Millar, & D. H. Pritchard, Oxford University Press, Oxford.
- Greco, J. (1999). 'Agent Reliabilism', *Philosophical Perspectives* 13, 273-96.
- (2000). *Putting Sceptics 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.
- (2007a). *Epistemic Evaluation: A Virtue-Theoretic Approach*, typescript.
- (2007b). 'The Nature of Ability and the Purpose of Knowledge', typescript.
- (*Forthcominga*). 'The Value Problem', *Epistemic Value*, (eds.) A. Haddock, A. Millar, & D. H. Pritchard, Oxford University Press, Oxford.
- (*Forthcomingb*). 'What's Wrong With Contextualism?', *The Philosophical Quarterly*.
- Hawley, P. (2006). 'Scepticism and the Value of Knowledge', typescript.
- Kaplan, M. (1985). 'It's 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, UK.
- (*Forthcoming*). 'Responses', *Epistemic Value*, (eds.) A. Haddock, A. Millar, & D. H. Pritchard, Oxford University Press, Oxford.
- Lackey, J. (2007). 'Why We Don't Deserve Credit for Everything We Know', *Synthese* 156.
- Percival, P. (2003). 'The Pursuit of Epistemic Good', *Metaphilosophy* 34, 29-47; and reprinted in *Moral and Epistemic Virtues*, (eds.) M. S. Brady & D. H. Pritchard, 29-46, (Blackwell, 2003).
- 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.
- (2007a). 'Anti-Luck Epistemology', *Synthese* 156.
- (2007b). 'Recent Work on Epistemic Value', *American Philosophical Quarterly* 44, 85-110
- (2007c). 'Virtue Epistemology and Epistemic Luck, Revisited', *Metaphilosophy* 38.
- (*Forthcominga*). 'A Defence of Quasi-Reductionism in the Epistemology of Testimony', *Philosophica* (special issue on *The Epistemology of Testimony*).
- (*Forthcomingb*). 'Epistemic Value', *Stanford Encyclopædia of Philosophy*, (ed.) Ed

- Zalta.
- (Forthcoming). ‘Greco on Knowledge: Virtues, Contexts, Achievements’, *The Philosophical Quarterly* (special issue on *Moral Contextualism*).
- (Forthcoming), ‘Knowing the Answer, Understanding and Epistemic Value’, *Grazer Philosophische Studien* (special issue on *Knowledge and Questions*).
- (Forthcoming). ‘Knowledge, Understanding and Epistemic Value’, *Think* 11; to be reprinted in *Epistemology (Royal Institute of Philosophy Lectures)*, (ed.) A. O’Hear, Cambridge University Press, Cambridge, UK.
- 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, UK.
- (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, UK.
- (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).

## NOTES

<sup>1</sup> Note that where I draw these contrasts between, as in this case, knowledge and mere true belief, I have in mind, unless specified otherwise, 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.

<sup>2</sup> 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 (*forthcoming*, §1).

<sup>3</sup> Kaplan (1985) famously argues, for example, that the conclusion to be drawn from the Gettier counterexamples to the tripartite account of knowledge (i.e., 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.

<sup>4</sup> For a recent survey of research on the value problem, see Pritchard (2007*b*). See also Baehr (*forthcoming*) and Pritchard (*forthcomingb*).

<sup>5</sup> Note that one might also object to the reasoning of the preceding paragraph in another way by disputing the ‘compositional’ account of value that it seems to presuppose. That is, it does not seem to be in general a good way of finding out whether something is valuable to consider the value that some of its parts have when they obtain independently of the whole, since the value of the parts independently of the whole could be very different from the contribution they make when part of the whole.

<sup>6</sup> See especially Kvanvig (2003), who attributes the problem to Swinburne (1999; 2000). Note that while the swamping problem is most clearly a difficulty for reliabilist responses to the primary value problem, it will likewise affect any response to the primary value problem which has the same salient features (in particular, which argues that knowledge, unlike mere true belief, has a property which is instrumentally valuable relative to the epistemic good of true belief).

<sup>7</sup> This is not the only problem facing this argument either. For example, one could claim that the additional instrumental value conferred on a true belief in virtue of being formed by a reliable process is not simply epistemic value, but also a practical non-epistemic value. Additional value of *this* kind would obviously not be swamped by the value of true belief. For a compelling defence of a reliabilist response to the primary value problem in the light of the swamping argument, see Goldman & Olsson (*forthcoming*).

<sup>8</sup> For a contemporary discussion and defence of this distinction, see Rabinowicz & Roennow-Rasmussen (1999; 2003).

<sup>9</sup> Brogaard (2007) makes this point about final value, and an embryonic version of this point also seems to be present in Percival (2003), who puts forward the example of a book printed on the first ever printing press as a counterexample to Zagzebski’s (2003) version of the swamping argument. Notice that it doesn’t follow from the fact that we value something for its own sake that we are right to do so (to take a famous example, a miser values money for its own sake, but this just reflects an error on her part). In what follows, when I talk of something being of final value I should be read as saying that we *rightly* value it for its own sake.

<sup>10</sup> Or, at least, a *greater degree* of final value than a mere true belief. Whether this is a theoretical option depends, of course, on the further issue of whether mere true belief has final value. In what follows I will follow orthodoxy and take it as given that it doesn’t. For a defence of the claim that true belief is finally valuable (although he mischaracterises the value in question as intrinsic value), see Kvanvig (2003, ch. 2).

<sup>11</sup> At one end of the spectrum, for example, there is the agent reliabilism advanced by Greco (e.g. 1999; 2000; 2007*a*) which effectively treats *any* reliable belief-forming process that is a stable part of the agent’s cognitive character as an intellectual virtue. At the other end of the spectrum there is the neo-Aristotelian view advanced by Zagzebski (1996) which understands intellectual virtues in a far more exclusive fashion. For an excellent survey of recent work on virtue epistemology, see Axtell (1997).

<sup>12</sup> This example is due to Chisholm (1977, 105).

<sup>13</sup> See Greco (1999; cf. Greco 2000, 251-2) for an example of a virtue epistemologist who was content to take this line. Note, however, that in more recent work, as we are now about to see, he offers a virtue-theoretic response to Gettier-style cases.

<sup>14</sup> 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). The second claim is stronger than the first because it does not follow from the fact that the true belief in question is primarily creditable to the agent that she is *thereby* deserving of credit (though it is plausible that the converse holds). Since the stronger claim is not essential to the virtue epistemic proposal under consideration, in what follows I will focus on the weaker reading. I am grateful to an anonymous reviewer from *Mind* for pressing me on this issue.

<sup>15</sup> 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 (*forthcominga*, §4). This mistake is unsurprising, since it is quite common to confuse these two types of value. Nevertheless, the point remains that it is because of the *relational* properties of the knowledge 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.

<sup>16</sup> 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—e.g., 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.

<sup>17</sup> Greco is the only virtue epistemologist to really confront this problem, and he offers three (potentially complementary) strategies for dealing with it. In Greco (*forthcomingb*), 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 (2007b), 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. Finally, in Greco (*forthcominga*, §5), he takes a third approach of offering a more fine-grained account of what constitutes a cognitive ability in order to argue that someone like Barney does not have the relevant cognitive ability in the first place. I discuss Greco's treatment of these issues further in Pritchard (*forthcomingc*). See also Kvanvig (*forthcoming*).

<sup>18</sup> Note, however, that Lackey uses this example to motivate a slightly different claim to the one that I have in mind—*viz.*, that not all knowledge is credit-worthy in the relevant sense. I comment on this difference below.

<sup>19</sup> A second type of case that Lackey (2007) offers in this regard—that of innate knowledge—might be better placed to demonstrate this possibility, though this example raises tricky issues regarding the nature of such knowledge that would not be usefully engaged with here.

<sup>20</sup> Greco (2007b) opts for this response, but 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.

<sup>21</sup> To take the two most prominent examples of this tendency, see the sensitivity-based theories of knowledge offered by Dretske (e.g., 1970) and Nozick (1981). For a more general defence of the idea of a genuinely anti-luck epistemology, see Pritchard (2005; 2007a).

<sup>22</sup> 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; 2007a).

<sup>23</sup> Indeed, if the environment had not been epistemically friendly in this regard—if there had been a real possibility that her informant might try to deceive her, for example—then it would have been much harder for Jenny to satisfy the anti-luck condition. In particular, she would need to exhibit a greater degree of relevant cognitive ability in order to count as a knower in this case. That the degree of cognitive ability required in order to know is dependent upon the 'friendliness' of one's environment in this way is, I suggest, entirely in accordance with intuition. I discuss the relevance of this point to the debate regarding the epistemology of testimony in Pritchard (*forthcominga*; cf. Pritchard *forthcomingd*; *forthcominge*).

<sup>24</sup> 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).

<sup>25</sup> A second issue that would be important in this regard is the dependence, if any, of knowledge that is not a cognitive achievement on knowledge which is. The testimonial case involving 'Jenny', for example, suggests a picture on which Jenny is only able to gain knowledge in this way because there is someone further up the testimonial chain—in this case the informant—whose knowledge does constitute a cognitive achievement.

<sup>26</sup> 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 (*forthcoming*).

<sup>27</sup> I am grateful to Chris Hookway for urging this line on me.

<sup>28</sup> Earlier versions of this paper were presented at the *Royal Institute of Philosophy* 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, Sheffield and St. Andrews during 2006 and 2007. I am grateful to the audiences on these occasions, especially Sharar Ali, Kelly Becker, Brit Brogaard, Matthew Chrisman, Andy Clark, Igor Douven, Pascal Engel, Paul Faulkner, Gordon Finlayson, Lizzie Fricker, John Greco, Patrick Greenough, Dominic Gregory, Lars Gundersen, Bob Hale, John Hawthorne, Allan Hazlett, Chris Hookway, Robert Hopkins, Jennifer Hornsby, Carrie Jenkins, Jesper Kallestrup, Jonathan Kvanvig, Bob Lockie, Anne Meylan, 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, Kathleen Stock, Tim Williamson, René van Woudenberg and Crispin Wright. Thanks also to Adam Carter, Adrian Haddock, Chris Kelp, Alan Millar, Ernie Sosa and an anonymous referee for *Mind*. This paper has benefited from an AHRC Research Leave award.