

CHAPTER THREE

ANTI-LUCK VIRTUE EPISTEMOLOGY

§3.1. INTRODUCTION

We saw in the last chapter that the most promising account available of why knowledge, unlike that which falls short of knowledge, is finally valuable—that offered by robust virtue epistemology—is ultimately unsuccessful. We thus concluded that knowledge is not of final value after all, and hence that the tertiary value problem cannot be positively answered. The emphasis now turns from offering a solution to the value problem in its strongest (tertiary) form, to explaining why a different account of the value of knowledge is desirable.

The goal of this chapter is to develop a theory of knowledge that can accommodate the intuitions that drive virtue-theoretic views but which does not succumb to the problems that they face. I call the view that I have in mind, which constitutes a new theory of knowledge, *anti-luck virtue epistemology*. As we will see, once we understand why knowledge has this structure, we will be able to diagnose where our intuition that knowledge is distinctively valuable comes from and also explain why this intuition should not be taken at face-value as motivating the tertiary value problem.

§3.2. *CONTRA* VIRTUE EPISTEMOLOGY

Let us begin by reminding ourselves where robust virtue epistemology went awry. Recall that we noted that what motivates virtue-theoretic proposals in epistemology is our strong intuition that knowledge is in some way the product of cognitive ability. Any cognitive success, even a cognitive success which is otherwise reliably produced or which would count (on some views at least) as justified, will not amount to knowledge if it is not appropriately related to the agent's cognitive abilities.

We didn't consider any examples to illustrate this point in the last chapter, but now is a good juncture to offer a suitable case. Imagine that our agent—let's call him 'Temp'—forms his beliefs about the temperature in his room by consulting a thermometer on the wall. Unbeknownst to Temp, however, the thermometer is broken and is fluctuating randomly within a given range. Nonetheless, Temp never forms a false belief about the temperature by consulting this thermometer since there is a person hidden in the room, next to the thermostat, whose job it is to ensure that whenever Temp consults the thermometer the temperature in the room corresponds to the reading on the thermometer.

Notice that it is entirely consistent with how we have set-up this example that Temp's method of forming beliefs about the temperature of the room (roughly, by consulting the thermometer) is perfectly reliable, in that whenever he forms a belief in this way his belief is true. Moreover, it is also consistent with how we have set-up the example that Temp's beliefs are justified. Perhaps he has no reason to think that the thermometer is broken, and has good reason for believing that it has been reliable in the past. On some views of justification, at least, this would suffice to enable his belief to be justified. Nonetheless, even though Temp's true belief formed in this way may have been reliably produced and justified, it is clearly not knowledge, and the reason for this is that his cognitive success is in no way a product of his cognitive abilities, but is rather due to a factor completely independent of his cognitive abilities (i.e., the person hidden in the room).

Adding a virtue-theoretic condition to one's theory of knowledge will deal with such cases, since one can now argue that they don't count as instances of knowledge because the virtue-theoretic condition is not met—the true belief in question, despite its other epistemic properties, is not the product of the agent's cognitive ability.

We also noted in the last chapter, however, that on the face of it adding a virtue-theoretic condition to one's theory of knowledge will not suffice by itself to deal with all cases. In particular, we pointed out that Gettier-style cases seem entirely compatible with the satisfaction of a virtue-theoretic condition on knowledge. The example we gave to illustrate this point was that of 'Roddy' who forms his true belief that there is a sheep in the field by using his highly reliable cognitive abilities, but who is subject to Gettier-style epistemic luck in that what he is looking at is not a sheep at all, but rather a sheep-shaped object which is obscuring from view a genuine sheep in the field. In this case Roddy lacks knowledge even though his reliable cognitive abilities are generating a true belief in the target proposition.

In light of this point, one might be tempted to opt for a modest virtue epistemology which adds a codicil to the account of knowledge to exclude such cases and thereby make the view Gettier-proof. What was distinctive about robust virtue epistemology, however, was that it aimed to avoid the need for such a codicil by formulating the view in such a way that it could deal with Gettier-style cases after all. The way it did this was by insisting that it was not enough for knowledge that one exercised one's cognitive abilities and one was, in addition, cognitively successful. Rather, that cognitive success needed to be because of one's cognitive abilities, where this meant that it is primarily creditable to the operation of those abilities. This view had a number of advantages. For one thing, it offered an elegant account of knowledge, one that did not need to appeal to any non-virtue-theoretic component in the way that modest virtue epistemology does. Moreover, by understanding knowledge in this fashion it becomes a kind of achievement. Given the further claim that achievements are finally valuable—which we found in the last chapter to be quite plausible—this held out the promise of dealing with the value problem by demonstrating the final value of knowledge.

The problem, however, was that we found that knowledge and cognitive achievements are not the same thing. In particular, one can have knowledge while not exhibiting the corresponding cognitive achievement, and one can exhibit a cognitive achievement while failing to have knowledge. We gave the example of 'Jenny' to illustrate the former possibility. This was a case of testimonial knowledge where the agent gains this knowledge by, for the most part, trusting the word of another, and this meant that the true belief in question was not primarily creditable to her cognitive abilities (even though it was *bona fide* knowledge). And we gave the case of 'Barney' to illustrate the latter possibility. This is essentially the barn façade case, an example in which the agent fails to have knowledge because his belief is subject to knowledge-undermining epistemic luck but where, nonetheless (and unlike Gettier-style cases), the agent does exhibit a cognitive achievement (i.e., the agent's true belief is primarily creditable to his cognitive abilities).

Robust virtue epistemology is thus not a viable theory of knowledge, and this means that the response to the value problem that it advertises is not viable either. Still, nothing we have said so far undermines modest virtue epistemology, the view that knowledge is true belief that is the product of the agent's cognitive abilities and which, in addition, satisfies a suitably formulated 'anti-Gettier' condition. Even so, modest virtue epistemology will retain the key problem that it is an unmotivated account of knowledge. For why should knowledge

have this structure such that it consists of a virtue-theoretic condition plus an apparently arbitrary anti-Gettier condition?

Moreover, we still need to be told how modest virtue epistemology can deal with cases like that involving Jenny and Barney. In the case of the former, it will be important that the virtue-theoretic condition imposes a very weak demand, and yet if this condition is meant to capture the essence of knowledge—the anti-Gettier condition is just a codicil, remember—then that is *prima facie* implausible. Moreover, given the differences we have noted regarding the type of epistemic luck at issue in the Barney case when compared to standard Gettier-style cases—i.e., that it is ‘environmental’ epistemic luck rather than ‘intervening’ epistemic luck—we cannot simply take it as granted that the anti-Gettier condition imposed by modest virtue epistemology will deal with this case anyway. Indeed, insofar as we need to ‘beef-up’ the anti-Gettier condition in order to deal with environmental epistemic luck, then this raises the obvious question of why this condition is a mere codicil to the modest virtue epistemic proposal. Isn’t this condition now playing a substantive role in this theory of knowledge?

I think that once these problems are made explicit it becomes clear that we need to step back from this debate and reconsider the motivations for the virtue-theoretic account of knowledge. Indeed, as we will see, once we undertake this re-examination of what underlies the virtue-theoretic account of knowledge we find that the right way to think about knowledge is along different lines. Virtue epistemology—in both its guises, but especially in its modest guise—comes close to gaining a correct understanding of knowledge, but ultimately fails precisely because it misunderstands the fundamental intuitions that a theory of knowledge must answer to.

§3.3. TWO MASTER INTUITIONS ABOUT KNOWLEDGE

I noted earlier that virtue epistemology answers to a fundamental intuition about knowledge: that knowledge is the product of one’s cognitive abilities, such that when one knows one’s cognitive success is, in substantial part at least, creditable to one. Call this the *ability intuition* and call any epistemic condition that one imposes on one’s theory of knowledge in order to

account for this intuition an *ability condition*. Where virtue epistemology goes wrong, I will argue, is in failing to recognise that there is a second fundamental intuition about knowledge which imposes an independent constraint on one's theory of knowledge. This second master intuition is the intuition that knowledge is incompatible with luck, in the sense that if one knows then it ought not to be the case that one's true belief could easily have been false. Call this the *anti-luck intuition* and call any epistemic condition that one imposes on one's theory of knowledge in order to account for this intuition an *anti-luck condition*.

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. That is, that no formulation of the ability condition could fully accommodate the anti-luck intuition and thereby obviate the need for a separate anti-luck condition, and that no formulation of the anti-luck condition could fully accommodate the ability intuition and thereby obviate the need for a separate ability condition. Let's take these points in turn.

First, 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 target true belief, and yet one has nevertheless been 'Gettierized' such that one's true belief is epistemically lucky). More generally, that one's true belief is appropriately due to one's cognitive abilities does not seem to entail that it thereby has the required modal stability across the relevant near-by possible worlds to ensure that it is immune to knowledge-undermining epistemic luck. If that's right, however, then no formulation of the ability condition could obviate the need for an additional anti-luck condition.

Second, and conversely, it seems antecedently plausible that there could be true beliefs which meet the relevant anti-luck condition—such that the agent's true belief couldn't have easily been false—but which are not thereby formed as a result of a cognitive ability. After all, there could be all manner of reasons why one's belief tracks the truth in nearby possible worlds which have nothing to do with one's own cognitive ability (perhaps, for example, it is the facts that are changing in order to correspond with your belief rather than *vice versa*, as happens in the 'Temp' case). If that's right, then no formulation of the anti-luck condition could obviate the need for an additional ability condition.

In effect, however, virtue epistemologists fail to see that these two master intuitions are imposing distinct demands. In particular, they hold that the ability condition, properly formulated along virtue-theoretic lines, is, if not all that is required, then at least *pretty much* all that is required for a fully adequate account of knowledge. That is, they either hold, *qua* robust virtue epistemologists, that the virtue-theoretic ability condition—advanced, principally, in order to accommodate the ability intuition—can all by itself offer us a complete theory of knowledge. Or else they hold, *qua* modest virtue epistemologists, that a virtue-theoretic ability condition—again, advanced, principally, in order to accommodate the ability intuition—can *pretty much* offer a complete theory of knowledge by accommodating the ability intuition—all that is required is a mere codicil to handle those troublesome Gettier-style cases.

The failure of virtue-theoretic accounts of knowledge of both stripe thus reflects a failure to realise that the anti-luck intuition imposes a distinct constraints on knowledge, one that cannot be captured merely by the addition of a virtue-theoretic ability condition which is principally designed to accommodate the ability intuition. The adding of an anti-Gettier codicil to one's view simply fails to comprehend the important role that the anti-luck intuition plays in our thinking about knowledge. It is for this reason that an ability condition all by itself falls some way short of offering an adequate account of knowledge.

Now one might respond to the failure of this sort of account of knowledge by opting for the opposing radical thesis that takes the anti-luck intuition as its lead. On this view—what one might call a *robust anti-luck epistemology*—one tries to formulate an anti-luck condition on knowledge in such a demanding way that one does not need an additional ability condition in order to accommodate the other master intuition. Indeed, one could read certain modal epistemological proposals in the recent literature as proposing just such a view.¹ That is, just as robust virtue epistemologists try to accommodate the motivation for the anti-luck intuition by offering a more stringent construal of the virtue-theoretic ability condition on knowledge, so proposals along these lines offer a more stringent construal of the anti-luck condition in order to accommodate the motivation behind the ability intuition.

Such a proposal makes essentially the same mistake that robust virtue epistemology makes of failing to recognise that these two master intuitions about knowledge impose distinct constraints on one's theory of knowledge. For just as there is no formulation of the

ability condition that can obviate the need for an anti-luck condition, so there is no formulation of the anti-luck condition that can obviate the need for an ability condition. After all, as noted above, it is inevitable that there will be cases in which there are true beliefs which exhibit the required modal properties to ensure that they are not lucky in the relevant sense and yet which are not formed as a result of the agent's cognitive abilities. Given the ability intuition, such cases will not count as knowledge, even though it will be the case that the agent in question could not have easily been wrong.

Indeed, we can adapt the example involving Temp that we offered above to make this point vivid. Clearly, Temp cannot gain knowledge of the temperature of the room by consulting a broken thermometer, and as noted above we can explain why by appeal to the ability intuition since the truth of his belief in this case is in no substantive way the product of his cognitive ability. Even so, depending on the details of the case it could nonetheless be true that Temp's belief is not subject to knowledge-undermining epistemic luck. After all, if the example is set up in the right way then it could well be the case that his true belief could not have very easily been false. Indeed, with the example suitably described it may well follow that in *all* near-by possible worlds in which Temp continues to form his belief in this way his belief is true.

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

§3.4. ANTI-LUCK VIRTUE EPISTEMOLOGY

What is essential to anti-luck virtue epistemology is thus that it incorporates two conditions on knowledge, an anti-luck condition and an ability condition, and that it accords each condition equal weight in the sense that they are each answering to a fundamental intuition about knowledge.

Let us describe any true belief that couldn't have easily been false as *safe*.² The general structure of the account of knowledge offered by an anti-luck virtue epistemology can now be described as follows: knowledge is safe belief that arises out of the reliable cognitive traits that make up one's cognitive character, such that one's cognitive success is to a significant degree creditable to one's cognitive character. The safety element of the view is the anti-luck condition, while the virtue-theoretic clause is the ability condition.

To begin with, notice how the ability condition is formulated. In particular, notice that it does not demand that the cognitive success in question must be *because of* the agent's cognitive ability, only that it should be to a significant degree creditable to one's cognitive character, which is a weaker claim. Accordingly, there is no suggestion on this view that knowledge should be thought of on the model of achievements, and hence this proposal does not attempt to answer the tertiary value problem in the manner of robust virtue-theoretic proposals (we will come back to the issue of what anti-luck virtue epistemology does have to say about the value problem later on).

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

Indeed, that we are willing to allow the exercise of relatively little cognitive ability result in the acquisition of knowledge in epistemically friendly environments lends support to the thinking behind anti-luck virtue epistemology. For consider what happens when we re-run the 'Jenny' case with our protagonist in an epistemically *unfriendly* environment, one where testimonial deception is common. In such a case we would not ascribe knowledge to our

hero, even when the target belief formed is true and gained from a knowledgeable and sincere source. Moreover, notice that in order for our subject to gain knowledge in this case it would be necessary for her to exhibit far more cognitive ability than is required by the corresponding hero in the standard ‘Jenny’ case. We would expect her, for example, to seek independent support for relying on this informant, such that she is no longer trusting this informant in the relevant sense to a substantive degree. Only then would her belief be safe and therefore satisfy the anti-luck condition. The degree of cognitive ability that is required in order to know thus varies in line with how epistemically friendly one’s environment is. But this result, of course, is just what an anti-luck virtue epistemology would predict.

Anti-luck virtue epistemology also fares very well across the range of problem cases beloved by epistemologists. For example, by incorporating 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’s true belief could very easily have been false, and so will not count as knowledge by the lights of this proposal since the anti-luck condition is not met.³

Moreover, the anti-luck condition also enables anti-luck virtue epistemology to deal with so-called ‘lottery’ cases. The problem posed by such cases is to explain why one cannot come to know that one has a losing lottery ticket for a lottery with long odds simply by reflecting on the odds in question. The reason why this is puzzling is that one can come to know that one has lost the lottery by reading the result in a reliable newspaper, and yet it may be far more likely that the newspaper contains a misprint on this score than that one wins the lottery even despite the long odds involved.

An anti-luck condition can deal with such cases because even though the true belief in question is the result of the exercise of the agent’s cognitive abilities, it nevertheless remains the case that this is a belief that could very easily have been false. For although this is a low probability event, very little about the actual world needs to change in order for this event to obtain—in most lotteries, just a few coloured balls need to fall in a different configuration—and so there will be scenarios just like the scenario the agent is actually in in which she continues to form her belief on the same basis and yet ends up with a false belief as a result.⁴

Finally, the ability condition enables the proposal to deal with a range of cases like ‘Temp’ in which the agent has a non-lucky true belief and yet does not count as knowing because the true belief does not arise out of the cognitive abilities of the agent. Since such cases do not satisfy the ability condition, they pose no problem for a view of this sort.

Indeed, it is worth noting that the way that the ability condition deals with the Temp case also enables it to deal with cases involving necessary propositions which robust anti-luck epistemology has traditionally struggled with. Suppose one has formed one’s true belief in a necessary proposition in an inappropriate way. For example, suppose that one forms a true belief in a mathematical proposition by writing out a series of calculations in one’s notebook, but that one has some radically false beliefs about how to do mathematical calculations such that one is generally unreliable in these matters. Now clearly we do not want to regard this as knowledge, but the problem is that on the face of it one’s belief does satisfy the anti-luck condition. After all, since the proposition believed is necessarily true, this can hardly be a true belief that could have easily been false.

The proponent of robust anti-luck epistemology might try to get around this problem by talking more generally of a doxastic output of a belief-forming method, rather than specifically of belief in the target proposition. The issue, then, would not be whether there is a near-by possible world where one believes the target proposition on the same basis as in the actual world but where one’s belief is false, but rather whether there is a near-by possible world in which one employs the same belief-forming method as in the actual world and thereby comes to form a false belief as a result. Given that the target proposition is necessarily true, it would be impossible for any belief in the target proposition to be false. Even so, however, it is possible that the belief-forming process might issue in a *different* belief in a near-by possible world which is false. Employing a faulty system of mathematical calculation will almost certainly lead you to form a false belief in a mathematical proposition in some near-by possible worlds.

The problem, however, is that one can simply change the set-up of the example to re-introduce the original problem. All one needs to do, as in the Temp case, is to describe the case so that there is some factor present which always ‘cancels-out’ one’s mistakes. For example, suppose that there is a wizard in the room who is observing your calculations and whose goal it is to ensure that you end up with a correct answer. Using his magic, the wizard makes sure that any mistakes you make in your calculations get cancelled-out elsewhere so

that your belief-forming method always results in a true belief.

Just as in the Temp case, the problem here is that one is forming one's true belief in such a way that this belief has the required modal properties to ensure that the anti-luck condition is met. Nevertheless, this isn't knowledge, and the natural explanation of why is that the truth of the agent's belief is not appropriately connected to the agent's cognitive abilities. In particular, the truth of the agent's belief is not due even in minor part to the agent's cognitive abilities, but is rather the result of the intervention of the wizard. Adding an ability condition to one's theory of knowledge thus enables one to deal with such cases. Again, then, we find that we need both an anti-luck and an ability condition on knowledge.

Anti-luck virtue epistemology thus handles the standard range of epistemology examples very well. That anti-luck virtue epistemology is able to straightforwardly deal with these cases lends the view a great deal of *prima facie* support.

§3.5. INTERLUDE: IS ANTI-LUCK VIRTUE EPISTEMOLOGY A REDUCTIVE THEORY OF KNOWLEDGE?

In the last chapter we examined the issue of whether robust virtue epistemology is a reductive or non-reductive account of knowledge (i.e., whether it involves a 'decomposition' of knowledge into its component parts where these parts can be specified in a knowledge-independent fashion). This issue was important there because of how we were evaluating the account's credentials at dealing with the swamping problem where this question looms large. In particular, we noted that non-reductive accounts of knowledge which hold that knowledge is sometimes more valuable than mere true belief are naturally allied to a certain kind of monistic response to the swamping problem which treats knowledge as the sole fundamental epistemic good. In contrast, non-reductive accounts which hold that knowledge is sometimes more valuable than mere true belief are more naturally allied to a pluralist response to the swamping problem which treats knowledge as one fundamental epistemic good amongst others. Insofar as anti-luck virtue epistemology also attempts to answer the swamping problem—we shall examine what it says about this problem, and the value problem more generally as well, shortly—this issue will be important here to. Moreover, even if this issue

didn't have a bearing on the swamping problem, this would still be a natural question to ask about the view.

I think the right answer to give here is the same as the answer given when we discussed robust virtue epistemology—*viz.*, that the default reading of this view is as offering a reductive account, but that ultimately the devil will lie in the detail. That is, on the face of it there seems no obvious reason why one cannot specify the anti-luck and the ability conditions in knowledge-independent terms, and if that's right then this will constitute a reductive account of knowledge. Nonetheless, it is entirely possible that on closer inspection it could turn out that this is not in fact possible at all. In particular, one might find that one cannot specify the ability condition without making appeal to knowledge—e.g., one might end up defining cognitive abilities in terms of how they are knowledge-conducive. If that turns out to be the case, then the proposal would be non-reductive.

For our present purposes, we can take a liberal view of this issue. The crux of the matter is just that knowledge has the structure set out by anti-luck virtue epistemology. Whether this structure presents us with a reductive account of knowledge is a further question, as is the issue of what bearing this question has on the response that the proponent of anti-luck virtue epistemology should offer to the swamping problem (we will set-out the options on this score more fully in a moment).

§3.6. DIAGNOSING THE STRUCTURE OF KNOWLEDGE

One question we might naturally ask about anti-luck virtue epistemology is why knowledge has this bipartite structure. A relevant contrast here is with robust virtue epistemology, since on this view there is a very straightforward account of why knowledge has the alleged structure which adverts to how knowledge is part of a more general phenomenon of achievements. The category of achievement is important to us, and hence it is no surprise that we have a separate sub-category of knowledge which picks out specifically cognitive achievements. This straightforward kind of story is not available to the proponent of anti-luck virtue epistemology, however. Similarly, the kind of anti-luck story told by the proponent of robust anti-luck epistemology—such that the overarching purpose of knowledge is to signal

non-lucky true belief—is similarly unavailable here.

It may be helpful in this regard to undertake a thought-experiment that is suggested by Edward Craig (1991) which offers a very plausible genealogical account of our concept of knowledge. Craig tries to cast light on the nature of knowledge by considering the question of why an imaginary society which lacked this concept may feel the need to introduce it. Craig's answer is to appeal to the fact that we clearly care about being able to identify reliable informants. He thus argues that a proto-concept of knowledge would develop very quickly as a means of picking-out these informants. This concept is just a proto-concept of knowledge, however, in that it is not yet similar in relevant respects to our own concept of knowledge. In particular, what is (in essence) required is that this concept, over time, should be used not just to classify actual reliable informants, but also to classify the potential reliability of informants, including oneself. This process, what Craig refers to as 'objectification', has the effect of making the concept more demanding in its application, since it means that non-actual (but modally close) error-possibilities can now be relevant to whether or not an agent knows.

The details of Craig's view, interesting though they are, are not important for our purposes here. What is salient are rather the general features of this proposal, since they highlight a possible avenue of defence for the proponent of anti-luck virtue epistemology. In particular, what this proposal suggests is that the central importance of the concept of knowledge resides in the practical need to pick-out reliable informants—informants that one can rely on.

Now one might think that this element of the view is actually contrary to the spirit of anti-luck virtue epistemology, in that it favours the virtue-theoretic account of knowledge with its focus on reliable cognitive abilities. Nevertheless, there is room for the proponent of anti-luck virtue epistemology to incorporate this motivation. The reason for this lies in an ambiguity in the very notion of a reliable informant. In one sense, it means an informant who possesses a reliable cognitive ability with regard to the target subject matter (and who is willing to sincerely communicate what she believes, something that we will take for granted in what follows). In another sense, it means an informant that one can rely on.

Now one might naturally think that this is a distinction without a difference, in that informants who possess reliable cognitive abilities in the sense just specified are thereby

informants that one can rely on, and *vice versa*. Closer inspection, however, reveals that first appearances are deceptive on this score. In order to see this, we just need to notice that it can be appropriate to rely on an informant who is forming her true belief via an unreliable cognitive ability, and also that it can be inappropriate to rely on an informant who nevertheless is forming a true belief via a reliable cognitive ability.

First, consider an agent who possesses a reliable cognitive ability as regards a certain subject matter but who is in an environment in which there exists a misleading defeater, one which you know about, but which the agent (and prospective informant) does not, and one which moreover, you are unable to defeat. An example might be an agent who is a reliable barn-detector but where you have been given a misleading ground for supposing that she is in barn façade county (e.g., false testimony from a good source). Such an agent is in fact a *reliable* informant about the relevant subject matter. But given that you know about the misleading defeater, and are aware that you are unable to defeat that defeater, would you be able to *rely on* this informant? Surely not.

The converse point also holds. After all, an agent might have a cognitive ability which is unreliable and yet be such that there are compensating factors known to us that mean that we can rely on this informant even though she is unreliable. Consider again, for example, the case described above of an agent with poor mathematical skills who is trying to work-out a series of mathematical problems, but who is unbeknownst to him being helped by a wizard who ensures that all his beliefs formed on this basis are true. If we know that this compensating factor is in play, then we can rely on what this informant tells us even though she is not reliable about this subject matter.

So one can have an informant that one can rely on but who lacks the relevant cognitive abilities, and one can have an informant who possesses the relevant cognitive abilities but whom one cannot rely on. If we accept this distinction, and we also buy into Craig's more general idea that the core motivation for introducing the proto-concept of knowledge is to pick out reliable informants, then we are just a few steps away from giving a diagnostic rationale for anti-luck virtue epistemology. For given this ambiguity in the idea of a reliable informant, it is natural that the concept of knowledge that evolves from the proto-concept of knowledge will generate *both* the anti-luck and the ability intuition.

That is, that as the range of cases which the concept is meant to apply to widens, so the distinction will open up between informants who are reliable and informants that we can rely

on, and we would expect the concept of knowledge that results to respect both sides of this distinction. In particular, examples where an agent exhibits a reliable cognitive ability but where the presence of epistemic luck means that we would not be able to rely on this agent *qua* informant would not be counted as cases of knowledge. Similarly, cases in which an agent forms a true belief in an epistemically friendly environment—such that any true belief so formed would not be subject to epistemic luck—would not be counted as cases of knowledge so long as the agent concerned failed to exhibit the relevant cognitive ability (even though we could rely on this agent *qua* informant). In short, the concept of knowledge that results will both (i) disallow cases of true belief as knowledge where the belief isn't appropriately due to a substantive degree to the relevant cognitive ability, and (ii) disallow cases of true belief as knowledge where the truth of the belief is substantively due to luck and hence unsafe.

Craig's very plausible story about the genealogy of the concept of knowledge thus lends support to anti-luck virtue epistemology after all, despite first appearances. In fact, if I am right that the goal of picking out reliable informants is ambiguous in the way just described, then contrary to the prevailing wisdom on this score, Craig's genealogical account of the concept of knowledge actually *favours* anti-luck virtue epistemology over rival proposals.

§3.7. BACK TO THE VALUE PROBLEM

But what import, ultimately, does anti-luck virtue epistemology have for the value problem? We noted in the last chapter that the best hope of accounting for the final value of knowledge lay in the robust virtue epistemic proposal to identify knowledge with cognitive achievements. This proposal, however, was shown to be untenable, and hence we concluded that a positive response to the tertiary value problem was unavailable. The best we can hope for then is a positive response to the primary and secondary value problems and an appropriate *diagnostic* response to the tertiary problem—i.e., an explanation of why we might naturally suppose knowledge to be distinctively valuable in the manner that the tertiary value problem demands even though it is not in fact valuable in this way. Interestingly, anti-luck

virtue epistemology does seem to be in a very good position to offer a diagnostic story of this sort, and also to answer the primary and secondary value problems.

To begin with, notice that while we have denied that knowledge is finally valuable, we have granted that the separate epistemic category of cognitive achievements are distinctively valuable in the very sense at issue in the tertiary value problem. That is, they are more valuable, not merely as a matter of degree but of kind, than lesser epistemic standings, where this entails that they are finally valuable (though note that this claim will be qualified somewhat in the next chapter). This in itself gives anti-luck virtue epistemology some purchase when it comes to diagnosing the intuition that knowledge is distinctively valuable. After all, when we think of paradigm cases of knowledge we tend to think of cases where there is in addition the corresponding cognitive achievement present. In particular, when asked to describe paradigm cases of knowledge it is not, say, Jenny-style cases that spring to mind. This fact in itself goes some way towards explaining why we might initially suppose knowledge to be (distinctively, and hence) finally valuable, since if the paradigm cases of knowledge coincide with cognitive achievement, and the latter is (distinctively, and hence) finally valuable, then it is only natural to ascribe this property to knowledge more generally too.

Indeed, in the next chapter we will see that there is more that the proponent of anti-luck virtue epistemology can do on this score to diagnose where the intuition that knowledge is finally valuable comes from. In particular, there is a kind of epistemic standing—a species of understanding—which is finally valuable and which also tends to coincide with knowledge. This fact can thus help to explain why when we initially think about knowledge we tend to suppose that it is finally valuable. Moreover, it turns out that it should be no surprise that understanding of this sort is finally valuable since it is also by its nature a cognitive achievement. This is an important result, since while a number of commentators have noted the distinctive value of understanding, they have not offered an explanation of why it is finally valuable. If this species of understanding constitutes a cognitive achievement, however, then we have a straightforward explanation available to us of why it is finally valuable.

That leaves the primary and secondary formulations of the value problem. Here anti-luck virtue epistemology seems to be on strong ground when it comes to offering a positive

resolution of these difficulties. After all, if the Craig-style story about the concept of knowledge is right, then we have the beginnings of a plausible answer to the secondary value problem and, thereby, the primary value problem too. For it seems that knowledge, on this view, marks-out a distinctive epistemic standing which is of particular instrumental value to us. We would thus expect knowledge to be of more instrumental value than that which falls short of knowledge, even though knowledge is not finally valuable.

Finally, there is the swamping problem to deal with. Given that anti-luck virtue epistemology denies that knowledge is in its nature finally valuable, and given that a defence of the final value of knowledge would constitute the most direct way of defending the claim that knowledge is a fundamental epistemic good (and thus that true belief is not the sole fundamental epistemic good), then on the face of it it would appear that anti-luck epistemology ought not be allied to those responses to the swamping problem—i.e., the monistic and pluralist responses—which deal with the problem by rejecting epistemic value T-monism. In this vein, one might regard anti-luck virtue epistemology as most naturally allied to the practical response to the swamping problem which maintains that knowledge is never epistemically more valuable than mere true belief, even though it is typically of greater practical value. Indeed, insofar as the practical response is viewed as a ‘de-mystifying’ response to the swamping problem, then it seems only natural that it should go hand-in-hand with a response to the value problem which treats the tertiary value problem as a pseudo-problem.

The matter becomes more complex, however, once one remembers that one it only has to be the case that *sometimes* knowledge is of final value for it to be a fundamental epistemic good (see §1.4). And since we have so far granted the claim that cognitive achievements are finally valuable—we shall be exploring this thesis in more detail in the next chapter—it is open to the proponent of anti-luck virtue epistemology to argue that because sometimes knowledge constitutes a cognitive achievement, so sometimes it is of final value and hence knowledge is a fundamental epistemic good.

Of course, which non-practical response to the swamping problem anti-luck virtue epistemology should offer in the light of this move—whether the monistic response which treats knowledge as the sole fundamental epistemic good, or the pluralist response which allows other epistemic goods to be fundamental too—is a further matter, and will depend

upon one's wider theoretical commitments (such as whether one regards one's view as a reductive theory of knowledge). We do not need to take a stand on this point here, however, since all that matters for our present concerns is that there is a response to the swamping problem available on this view.⁵

§3.8. CONCLUDING REMARKS

In this chapter we have seen that out of the ashes of robust virtue epistemology we can delineate a new account of knowledge—anti-luck virtue epistemology—which is able to avoid the problems facing both virtue-theoretic and anti-luck epistemologies. Moreover, we have also seen that while this view is unable to offer a positive response to the tertiary value problem, it is in a position to offer a diagnosis of why we might initially find the claim that knowledge is finally valuable intuitive. In essence, this diagnostic story appeals to the fact that on this view knowledge is at least sometimes a cognitive achievement, and cognitive achievements are plausibly of final value. Moreover, anti-luck virtue epistemology is in a position to offer a positive response to the secondary (and hence primary) value problem and also the swamping problem.

The purpose of the next chapter is to extend this diagnostic story by showing how a certain kind of understanding is finally valuable, and finally valuable precisely because, unlike knowledge, it constitutes a cognitive achievement. Along the way we will also acquire a more nuanced conception of what constitutes a cognitive achievement and a greater understanding of why achievements—or at least a certain type of achievement anyway—are of final value.

NOTES

¹ 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 Unger (1968) and Pritchard (2005; 2007*a*).

² The general idea that knowledge entails a true belief that could not have easily been false is often referred to as the safety principle, and there are various more fine-grained formulations of this principle that have been offered in the recent literature. For our present purposes, however, we can stick with the more general

formulation of safety. For some key defences of safety, see Sosa (1999; 2000) and Pritchard (2002; 2003; 2005; 2007a).

³ The anti-luck condition will also deal with Harman's (1973, 142-54) assassination case as well, at least so long as it is read in a way that the resulting belief is unsafe (see Pritchard (2005, ch. 6) for a discussion of how the assassination case is ambiguous in this respect). Cases involving misleading defeaters—which are often discussed in the same context as Gettier-style cases—can also be dealt with by this account, though not simply by appeal to the anti-luck condition. On the one hand, I argue in Pritchard (2005, ch. 5) that cases where an agent's belief satisfies the anti-luck and ability conditions but where she could so very easily have been misled by a misleading defeater (but wasn't) are entirely compatible with knowledge, and so do not pose any kind of Gettier-style challenge to one's theory of knowledge. On the other hand, cases where the agent ought to have been aware of the misleading defeater, but wasn't—e.g., because she was inattentive to some aspect of the environment that she should have been attentive to—are dealt with by the ability condition, in that given this feature of the environment it follows that she wasn't appropriately exercising her cognitive ability (this is just a manifestation of the point made earlier that the more epistemically unfriendly the environment is, the more cognitive ability one needs to exhibit in order to know). The same point goes for cases in which the agent groundlessly disregards the misleading defeater and so continues to believe the target proposition nonetheless. Finally, it is worth noting that since the anti-luck condition is formulated in terms of safety rather than sensitivity, then it avoids the problems cases that have beset sensitivity, such as Sosa's (2000) garbage chute case. For further discussion of the relative merits of safety and sensitivity, see Pritchard (2005, ch. 6; 2008f).

⁴ There are some issues regarding the proper formulation of safety that are raised by the lottery problem, though it would take us too far afield to explore these here. For more on this point, see the exchange between Greco (2007c) and Pritchard (2007a). Note that anti-luck virtue epistemology is also able to deal with a related problem which is sometimes discussed—e.g., by Hawthorne (2004)—as part of the lottery problem just described. This problem arises out of the fact that one can, it seems, competently deduce from certain propositions that one appears to know that one is in possession of a losing lottery ticket, and thereby gain knowledge of this proposition in that way. For example, one might take it as obvious that one knows that one will not have sufficient funds to buy a sports car next week. One might know, however, that one would only have such funds if one's lottery ticket were a winner. But doesn't that mean that one can competently deduce from the fact that one knows that one will not have sufficient funds to buy a sports car next week that one's lottery ticket must be a loser, and so thereby gain knowledge of this proposition? I think the key to understanding these cases is to realise that one *doesn't* know that one will not have sufficient funds to buy a sports car next week so long as one is in possession of a lottery ticket where one doesn't know the outcome of the lottery. This is not as counterintuitive as it may at first appear, for the whole point of buying a lottery ticket is to make scenarios that would ordinarily be far-fetched—such as having sufficient funds to buy a sports car—modally close. Accordingly, while one can normally know that one will not have sufficient funds to buy a sports car next week, one can't know this when one is in possession of a lottery ticket in a lottery that one does not know the outcome of, since in that case one does not satisfy the anti-luck condition: even where one truly believes the target proposition one's belief could have very easily have been false. For more on this point, see Pritchard (2007c).

⁵ That said, as I've noted at a number of junctures I think that the default stance is to treat a theory of knowledge as reductive until there are grounds for regarding it otherwise. Accordingly, insofar as a reductive anti-luck virtue epistemology is more naturally allied to a pluralist response to the swamping problem, then the pluralist response is the default response to take in this regard. I explore anti-luck virtue epistemology at greater length in Pritchard (2008c).