

CHAPTER ONE

THE VALUE PROBLEM FOR KNOWLEDGE

§1.1. THE VALUE PROBLEM

It is widely held that knowledge is of distinctive value. Presumably, this is the reason why knowledge—and not, say, justified true belief—has been the principal focus of generations of epistemological theorising. Understanding just why knowledge is distinctively valuable, however, has proved elusive and this has led some to question whether it is distinctively valuable at all. Call this the *value problem*. The goal of this chapter is to set-out what this problem involves. As we will see, this is far from being a straightforward matter, and part of the difficulty here is that there is a related problem—the so-called *swamping problem*—which, while being superficially similar to the value problem, in fact presents a more specific challenge.

§1.2. UNPACKING THE VALUE PROBLEM

Part of the difficulty posed by the value problem involves getting clear about just what it means to say that knowledge is distinctively valuable. One minimal reading of this claim is that knowledge is more valuable than mere true belief.¹ Call the challenge to explain why knowledge is more valuable than mere true belief the *primary value problem*.² Clearly there is more to showing that knowledge is distinctively valuable than answering this problem, and we will consider what additional demands a response to the value problem needs to satisfy in a moment. What ought to be clear, however, is that if we are unable to account even for why knowledge is more valuable than mere true belief then the very project of answering the value problem is a lost cause.

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

Indeed, this was precisely the way that Socrates answers the primary value problem in the *Meno*. Why should you prefer knowledge of the correct way to Larissa rather than mere true belief, given that both will, on the face of it, ensure that you get to your destination? Socrates's answer is that knowledge has a 'stability' which mere true belief lacks. Mere true belief, argues Socrates, is like one of Daedalus's statues untethered in that it is liable to be lost. Knowledge, in contrast, is like one of those statues tethered. For while a mere true belief may well enable you to achieve your goals as well as knowledge, one will be far more insulated from failure by possessing knowledge.

Suppose, for example, that the road to Larissa takes an unexpected course. Someone with mere true belief—where the belief is based on just a hunch, say—may well at this point lose all faith that she is on the right tracks and turn back. Someone who knows that this is the right way to go, however—perhaps because she consulted a reliable map before her departure—will not be so shaken by this turn of events.

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

In order to see this, suppose that one answered the primary value problem by, for example, pointing to a necessary condition for knowledge which in general added practical value (the justification condition, say), but suppose further that the satisfaction of this condition, in conjunction with true belief, was not sufficient for knowledge. Perhaps, for example, when one knows that p it is the fact that one's belief that p is thereby justified that ensures that knowledge has a greater practical value than mere true belief that p alone. One

would thereby have answered the primary value problem while leaving the secondary value problem unanswered. Moreover, let us take it as given that there is no further feature of knowledge which is value-conferring, such that the secondary value problem is regarded not just as unanswered, but as unanswerable.

On the face of it, this lacuna might not seem that problematic, since just so long as one can show that knowledge is more valuable than mere true belief then that would seem to satisfy our intuition that knowledge is of some special value to us (on this view it is, after all, the kind of thing that we should prefer to mere true belief, all other things being equal). The problem, however, is that if the distinctive value of knowledge is due to some feature of knowledge which, with true belief, falls short of knowledge, then it seems that what we should seek is not knowledge as such, but rather that which falls short of knowledge (i.e., true belief plus the value-conferring property X, in this case justification). But if that's right, then why do we regard knowledge as distinctively valuable at all?³

The primary value problem thus naturally leads to the secondary value problem, and it seems that both will need to be answered if we are to account for the distinctive value of knowledge. Even if we can offer a response to the secondary value problem, however, it is still not clear that we have accounted for the distinctive value of knowledge.

This is because the secondary value problem leaves open the possibility that the difference of value at issue is merely one of degree rather than kind. To say that knowledge is of *distinctive* value, however, appears to suggest that the difference in value between knowledge and that which falls short of knowledge is not just a matter of degree, but of *kind*. After all, if one regards knowledge as being more valuable than that which falls short of knowledge merely as a matter of degree rather than kind, then this has the effect of putting knowledge on a kind of continuum of value with regard to the epistemic, 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 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 our intuitions about the value of knowledge requires us to offer an explanation of why knowledge has not just a greater *degree* but also a different *kind* of value than whatever falls short of knowledge. Call this the *tertiary value problem*.

Further support for the tertiary value problem comes from the fact that we often treat knowledge as being, unlike lesser epistemic standings, *precious*, in the sense that its value is not merely a function of its practical import. But if that's correct, then knowledge must be the kind of thing that, unlike that which falls short of knowledge, is valuable for its own sake. That is, it must be non-instrumentally—i.e., *finally*—valuable.⁴

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.

In essence, then, the challenge we face in trying to account for the distinctive value of knowledge is to find a way of demonstrating that knowledge, unlike that which falls short of knowledge, is finally valuable. This is precisely the possibility that I will be exploring in the next chapter.

§1.3. THE SWAMPING PROBLEM

Before we can get into the nitty-gritty of exploring how one might account for the final value of knowledge, however, it is first necessary to consider a problem that on the face of it is a variant of (if not simply a sub-species of) the primary value problem. This difficulty is the so-called 'swamping problem', as defended most prominently by Jonathan Kvanvig (e.g., 2003), but also put forward in various forms by Ward Jones (1997), Richard Swinburne (1999; 2000), Wayne Riggs (2002), Linda Zagzebski (2003) and John Greco (*forthcoming*), amongst others. Interestingly, the standard way of expressing this problem is not as posing a general epistemological difficulty which faces all theories of knowledge—which is how the value problem is usually understood—but rather as posing a problem for particular epistemological proposals, such as reliabilism.⁵ With the argument so directed, here is how it goes.⁶

Imagine two great cups of coffee identical in every relevant respect—they look the same, taste the same, smell the same, are of the same quantity, and so on. Clearly, we value great cups of coffee. Moreover, given that we value great cups of coffee, it follows that we also value reliable coffee-making machines—i.e., machines which regularly produce good coffee. Notice, however, that once we've got the great coffee, then we don't then care whether it was produced by a reliable coffee-making machine. That is, that the great coffee was produced by a reliable coffee-making machine doesn't contribute any additional value to it. In order to see this, note that if one were told that only one of the great identical cups of coffee before one had been produced by a reliable coffee-making machine, this would have no bearing at all on the issue of which cup one preferred; one would still be indifferent on this score. In short, whatever value is conferred on a cup of coffee through being produced by a reliable coffee-making machine, this value is 'swamped' by the value conferred on that coffee in virtue of it being a great cup of coffee.

The supposed import of this example to reliabilist theories of knowledge—theories which hold that knowledge is reliably formed true belief—is that it follows, by analogy, that if we are faced with two identical true beliefs, one of them reliably formed and one not, it shouldn't make any difference to us which one we have. After all, we only value reliable belief-forming processes as a means to true belief, just as we only value reliable coffee-making machines as a means to great coffee, and so once we have the good in question—true belief or great coffee—then it shouldn't matter to us whether that good was in addition acquired in a reliable fashion. Intuitively, though, it *is* better to have reliable true belief rather than just mere true belief. Moreover, if, as the reliabilist maintains, reliable true belief is knowledge, then this intuition is stronger still. For surely many of us have the intuition that it would be better to have knowledge rather than mere true belief.

More generally, the axiological claim implicit in the coffee cup case is that if a property (like being reliably formed, when it comes to beliefs, or being reliably produced, when it comes to coffee) is only instrumentally valuable relative to some further good (e.g., true belief or great coffee), then in cases in which the further good in question is already present, no further value is conferred by the presence of the instrumentally valuable property. This is, of course, a general thesis about value. Moreover, as the coffee cup example illustrates, there seems every reason to think that it is true.

Interestingly, once we recognise that what is driving the swamping argument is this general thesis about value, then it ceases to become plausible that it should only affect reliabilist views. After all, the general value thesis just described will have an impact on *any* epistemological proposal which has the same relevant features as reliabilism—*viz.*, which treats the epistemic standing in question as instrumentally valuable only relative to the good of true belief. Moreover, once the point is put in this way it also becomes clear that what is at issue in the swamping argument is not the greater value *simpliciter* of knowledge over mere true belief, but rather the greater *epistemic* value of knowledge over mere true belief. Let us take these points in turn.

In order to see the first point, suppose, for example, that we only value justification because it is a means to true belief—i.e., we only value it instrumentally relative to the good of true belief. If that's right, then we could just as well run the swamping argument for justification as we can for reliability, since if the swamping argument works at all then any value that is conferred on a belief in virtue of its being justified will be swamped by the value conferred by the belief's being true. The problem, then, is not specific to reliabilism but instead applies to any view which treats the value of an epistemic standing as being instrumental value relative only to the good of true belief.⁷

In order to see the second point, notice that there is on the face of it very little to be gained by responding to the swamping problem by arguing that the epistemic standing in question generates a practical value that mere true belief lacks.⁸ For example, suppose one responded to the swamping problem by arguing that knowledge is more valuable than mere true belief because knowledge entails justification and justification is practically valuable. Justified true belief, we might say—in a broadly Socratic fashion (see §1.2 above)—has a 'stability' that mere true belief lacks, and this means that it is more practically useful to us in attaining our goals. The problem with this response, however, is that it doesn't appear to engage with the swamping problem at all. After all, the difficulty that the swamping problem poses concerns how to make sense of the idea that knowledge is more valuable than mere true belief because it involves an epistemic standing which better serves our specifically *epistemic* goals—in particular, the epistemic goal of true belief. Thus, the kind of value that is at issue is specifically an epistemic value. Accordingly, even if it is true that knowledge has more all things considered value because it entails an epistemic standing which adds

practical value to true belief, the problem would still remain that, on the face of it, knowledge is not *epistemically* more valuable than mere true belief.

Of course, one response to this could be to argue that it is so much the worse for the swamping problem that it turns specifically on epistemic value in this way. After all, if the swamping problem is consistent with the greater all things considered value of knowledge over true belief, then how seriously should we take it? Call this the *practical response* to the swamping problem. On this view, the right reaction to the swamping problem is to recognise that its conclusion is in fact harmless. We will explore this response to the swamping problem in more detail below. First, though, we need to do more to pick apart the different moving parts of the argument.

§1.4. FUNDAMENTAL AND NON-FUNDAMENTAL EPISTEMIC GOODS

If the swamping argument is indeed concerned with specifically epistemic value—i.e., a value which specifically concerns epistemic goods—then that prompts the question of what particular view of epistemic value it challenges.⁹ In order to formulate the value thesis in question, we first need to introduce a distinction between *fundamental* and *non-fundamental* epistemic goods. Call a fundamental epistemic good any epistemic good whose epistemic value is at least sometimes not simply instrumental value relative to a further *epistemic* good. That is, such an epistemic good is at least sometimes epistemically valuable entirely for its own sake. Call a non-fundamental epistemic good any good which is not a fundamental epistemic good—i.e., any good whose epistemic value is always instrumental value relative to a further epistemic good (and which is thus never epistemically valuable entirely for its own sake).

Clearly, the very idea of there being such a thing as epistemic value—of assessing the value of something relative to specifically epistemic goals—presupposes that some epistemic goods are fundamental. Something, after all, needs to act as the terminus for the instrumental regress of epistemic value. Nonetheless, one might be puzzled as to why this distinction is being put in the way that it is, since why can't we just talk about instrumental and non-instrumental—i.e., *final*—value?

The problem is that something can be a fundamentally epistemic good and yet not be finally valuable. For while a fundamental epistemic good can act as the terminus of the instrumental regress of epistemic value, this is entirely compatible with that good not being finally valuable *simpliciter*. After all, it could be that the value of the fundamental epistemic good in question is only instrumental value relative to some further non-epistemic goods (e.g., practical goods). For example, suppose that true belief is a fundamental epistemic good—a possibility which we will explore in more detail in a moment. If that's right, then we won't instrumentally value this good relative to any further epistemic good, such as, say, justified belief. Nevertheless, it could be that true belief is not itself finally valuable *simpliciter*, but is instead only instrumentally valuable relative to a further non-epistemic good—e.g., it serves some practical purpose.

There is also a further difference between final and fundamental epistemic value, in that when we describe a good as finally valuable we are claiming that it is in the *nature* of the good in question to be valuable in this way. In contrast, in calling an epistemic good a fundamental epistemic good we are merely saying that at least *sometimes* it has non-instrumental epistemic value. On the face of it at least, this is entirely compatible with the good in question not being in its nature of fundamental epistemic value at all.

So fundamental epistemic value does not entail final value *simpliciter*. Interestingly, however, final value *simpliciter* *does* entail fundamental epistemic value. After all, if an epistemic good has a value which is not an instrumental value then, *a fortiori*, it has a value which is not an instrumental epistemic value either, and hence the reverse direction of fit holds. This might initially seem surprising. After all, one might think that final value is *sui generis* and hence the sort of thing that by definition cannot be thought of as specifically epistemic. I think this worry subsides, however, once one clarifies the claim in play here.

For suppose for the sake of argument that knowledge has final value. This means that knowledge has a value which is independent of all other considerations. But if this is right then, *a fortiori*, knowledge has a value which is independent of all *epistemic* considerations too, and hence it is a fundamental epistemic good. That is, from a purely epistemic point of view where only epistemic considerations are taken into account, knowledge has a value which is not an instrumental value relative to a further epistemic good. It is only in this sense that the final value of an epistemic standing entails that this epistemic standing is a

fundamental epistemic good. Note, however, that there is no conflict between this claim and the idea of final value as being *sui generis*. For to say that the final value of an epistemic standing entails that this epistemic standing is a fundamental epistemic good is not to qualify the final value in question, but merely to note its axiological implications for a particular domain.

One last point is in order regarding our distinction between fundamental and non-fundamental epistemic goods and its relation to final value. Given how weakly the notion of a fundamental epistemic good is understood, such that an epistemic good need only *sometimes* have an epistemic value which is non-instrumental in order to qualify, it ought to be clear that an epistemic standing need only be sometimes of final value in order to ensure that it is a fundamental epistemic good. In particular, it need not, for example, be in the nature of the epistemic standing in question that it is finally valuable. For suppose, for instance, that knowledge only sometimes has final value—i.e., that it is not in the nature of knowledge to be valuable in this way, though it sometimes is. Even then, it will follow that there are occasions on which knowledge retains a value independently of all other considerations, and hence independently of all other epistemic considerations too. Thus, even on this relatively weak supposition about the value of knowledge, it will still count as a fundamental epistemic good.¹⁰

§1.5. THE RELEVANCE OF EPISTEMIC VALUE MONISM

With this distinction in mind, we can now formulate the thesis about epistemic value that is really at issue in the swamping argument, which is a variety of epistemic value monism where the fundamental epistemic value is true belief:

Epistemic Value T-Monism

True belief is the sole fundamental epistemic good.¹¹

The driving idea behind epistemic value T-monism is that there is only one fundamental epistemic good (this is the ‘monism’ part) and that this good is true belief (this is the ‘T’ part). Accordingly, the epistemic value of all other epistemic goods is instrumental value relative to this epistemic good. So, for example, on this picture the epistemic value of an

epistemic standing like justification is to be understood purely instrumentally relative to the fundamental epistemic good of true belief. That is, justification is epistemically valuable, but only because it is a means to true belief; it is not, unlike true belief, epistemically valuable in its own right.¹²

I take it that many epistemologists would be attracted to epistemic value T-monism, at least pre-theoretically (i.e., independently of considering any problems that a commitment to this thesis would generate). Indeed, one might go so far as to hold that it is a mere truism that the ultimate epistemic good just is the good of true belief, and hence that all other epistemic goods must be understood relative to this ultimate epistemic good. Epistemic value T-monism is clearly an immediate consequence of a picture of epistemic value of this sort.

So why does the swamping problem presuppose epistemic value T-monism? The key point in this regard is that the setting up of the problem demands that whatever epistemic value is contributed by the epistemic standing in question—reliability, justification, etc.—that value is only instrumental epistemic value relative to the further good of true belief. Moreover, since the swamping argument is meant to be a decisive strike against the greater epistemic value of knowledge over mere true belief, it is clearly being presupposed here that the additional value cannot be generated in any further fashion, and thus it is being taken for granted that the conception of epistemic value at issue is a monistic one regarding the value of true belief. In effect, then, the claim made by the swamping problem is that with this conception of epistemic value in play, it follows that knowledge cannot be more valuable than mere true belief.

More formally, we can express the swamping argument in terms of an inconsistent triad of claims. The first claim is just the thesis of epistemic value T-monism, which we will for convenience re-phrase as (1):

- (1) The epistemic value conferred on a belief by that belief having an epistemic property is instrumental epistemic value relative to the further epistemic good of true belief.¹³

The second claim is the general thesis about value that we noted in §1.3:

- (2) If the value of X is only instrumental value relative to a further good and that good is already present, then it can confer no additional value.

The difficulty posed by the swamping problem, however, is that (1) and (2) entail that there cannot be an epistemic standing which is epistemically more valuable than mere true belief. After all, whatever epistemic value an epistemic standing contributes could only, given (1), be instrumental epistemic value relative to the epistemic good of true belief. But (2) makes clear that this value gets swamped by the value of true belief. Thus, there can never be a true belief which, in virtue of possessing an epistemic property, is epistemically more valuable than a corresponding mere true belief.

In order to get our inconsistent triad, then, all we require is a thesis to the effect that there is an epistemic standing that is at least sometimes epistemically more valuable than mere true belief. For the sake of simplicity, we will here focus on the epistemic standing of knowledge, since as I noted above we have the strong intuition that knowledge is more valuable than mere true belief. We thus get (3):

(3) Knowledge that p is sometimes more epistemically valuable than mere true belief that p .

In short, then, (1), (2) and (3) are jointly inconsistent because it straightforwardly follows from (1) and (2) that (3) is false. Thus, if we accept (1) and (2), we are thereby committed to the denial of (3). Given that this is an inconsistent triad, it follows that one must reject one of the claims that make up this triad.

Now that we have more clearly set-out the swamping problem it should be easier to see that the challenge it poses differs from the superficially similar challenge at issue in the primary value problem as we formulated it above. For while the latter poses a general problem about how knowledge is more valuable than mere true belief, the swamping problem instead issues in a far more restricted difficulty, one that is focussed specifically on epistemic value and which, as a result, should not be considered, at least not without qualification, to simply be a variant on the primary value problem. As we noted in §1.3, one could answer the primary value problem by appealing to the greater practical value of knowledge over mere true belief. Given that the swamping problem is explicitly expressed in terms of epistemic value, however, the issue of whether this approach to the primary value problem can be extended to the swamping problem—i.e., the practical response—is clearly moot. At the very least, one would need to supplement this proposal in some way to make it effective for both value problems.

With this point in mind, we can now consider the question of how one should go about responding to this problem.

§1.6. RESPONDING TO THE SWAMPING PROBLEM I: THE PRACTICAL RESPONSE

One option, of course, is just to reject the general claim about value expressed as (2). Unfortunately, I can see no way of objecting to this claim, nor am I aware of any good objections to this thesis in the literature (typically, responses to the swamping argument in effect focus on other elements of the argument than this claim). Accordingly, in what follows I am going to take this thesis for granted and focus instead on the status of the other two members of this triad, since this is where I think the real promise of responding to this problem lies.

Let me start with the option of denying (3).¹⁴ This will be the strategy favoured by the proponent of the practical response to the swamping problem that we described at the end of §1.3. We have just noted that merely arguing for the greater practical value of knowledge over mere true belief would not in itself engage with the swamping problem, since the issue at hand specifically concerns epistemic value. However, although this point by itself does not gain any purchase on the swamping problem, it may be possible to supplement this claim in such a way that it does have an impact. In particular, if one wished to respond to the swamping problem by denying (3), then one might use this claim as part of a diagnostic story about why we tend to suppose that (3) is true. On this view, the greater practical value of knowledge over mere true belief—a non-epistemic value which, accordingly, cannot be swamped by the value of true belief—is meant to account for our initial conviction that (3) is true. The idea is that once we make the distinction between practical and specifically epistemic value clear—and once, in addition, we realise that knowledge is of greater practical value when compared with mere true belief—we are meant to see that there is no harm in denying (3). Moreover, there is a key advantage of this approach, in that by denying (3) the swamping problem is completely blocked and hence one is free, on this score at least, to endorse epistemic value T-monism.

We thus get the practical response to the swamping problem whereby one responds to the swamping problem by arguing that if this problem is indeed specifically focussed on epistemic value then that in itself should serve to undermine its plausibility, in that so long as knowledge is generally practically more valuable than mere true belief then it is harmless to concede that knowledge is not epistemically more valuable than mere true belief. Indeed, one might go so far as to maintain that to insist on epistemic value in the absence of practical value in this way is to succumb to a kind of theoretical fetishism. Is this view sustainable?

In order to evaluate the merits of this position, it is worth first reflecting on just how weak (3) is—it merely holds that knowledge is *sometimes* of greater epistemic value than mere true belief. Given the weakness of this claim, to deny it is to endorse a very strong thesis: that knowledge is *never* of greater epistemic value than mere true belief. Indeed, recall that our focus on knowledge was arbitrary, in that we could have easily focussed on *any* epistemic standing. More generally, then, the detractor of (3) has to argue that *no* epistemic standing is *ever* of greater epistemic value relative to true belief. This is strong stuff indeed.

Nevertheless, the proponent of this strategy will no doubt be prepared to accept this consequence of their position. Of course, in order for their view to be plausible it will be essential that they are able to offer an adequate response to the more general primary value problem which we outlined in §1.2. After all, it is only if knowledge is indeed of greater practical value when compared with mere true belief that the proponent of this response to the swamping problem can offer her crucial diagnostic story of why we might (wrongly) be initially inclined to regard (3) as true. Still, as we saw in §1.1, it is plausible to hold that such a story is forthcoming.

Interestingly, note that the revisionism of the proposal does not end with the denial of (3), since the proponent of this view must also argue that the tertiary value problem is a pseudo-problem. Recall that the tertiary value problem demands an explanation of why knowledge is more valuable, not merely as a matter of degree but of kind, than that which falls short of knowledge, and we noted in §1.2 that this meant that knowledge, unlike that which falls short of knowledge, must be of final value. Recall, however, that we also noted above that if knowledge has final value then it follows that it is a fundamental epistemic good. Hence the thesis that knowledge has final value is inconsistent with epistemic value T-monism. It follows that the defender of the practical response to the swamping problem must

regard knowledge as lacking in final value, and hence must reject the intuitions that led to the tertiary value problem.

Will the proponent of the practical response be troubled by this? I suspect not. For their view is meant to be a ‘de-mystifying’ stance which can extend even to the tertiary value problem. That is, they will argue that just as we can be tempted into thinking that knowledge has greater epistemic value than mere true belief because we confuse practical value with epistemic value, so we can be tempted into thinking that knowledge has final value by reifying the practical value that it possesses.

Rather than engage with this proposal directly, I want to instead explore the relative plausibility of alternative proposals. For if an alternative response to the swamping problem is available—one that is compatible with a positive answer to the tertiary value problem—then this would go a long way towards mitigating the attraction of this view. Given that one accepts (2) and (3), the only remaining alternative is to deny (1).

§1.7. RESPONDING TO THE SWAMPING PROBLEM II: THE MONISTIC RESPONSE

Interestingly, there are two very different ways of denying (1). In particular, while any response to the swamping problem which denied (1) would be committed to denying epistemic value T-monism, it does not follow that in denying this thesis one is thereby endorsing epistemic pluralism.

The reason for this is that there are different ways of being an epistemic value monist than being a T-monist. Consider, for example, someone like Timothy Williamson (e.g., 2000), who has explicitly argued for what he calls a ‘knowledge-first’ epistemology. Although (as far as I am aware), he has not endorsed such a view in print, presumably he would be very attracted, given his wider epistemological views, to an epistemic value monism of the following form:

Epistemic Value K-Monism

Knowledge is the sole fundamental epistemic good.

According to this view, there is only one fundamental epistemic good and that's knowledge, which means that the value of all other epistemic goods is to be understood along instrumental lines relative to this fundamental good.

On this view, there is no reason to assent to (1), and thus the argument cannot even get started. Moreover, notice that one cannot run a swamping argument against this view. For what is crucial to the swamping argument described above is that we are comparing mere true belief with a more elevated epistemic standing, and asking the question of why the latter epistemic standing should be epistemically more valuable. With epistemic value K-monism in play, however, the question of why knowledge is more valuable than mere true belief has a straightforward answer: because knowledge, unlike mere true belief, is a fundamental epistemic good. Call this response to the swamping problem the *monistic response*.

In order to motivate such a position one would need to give it the further theoretical support that it requires. That is, one would need to re-examine the intuitions that drive epistemologists to commit themselves to epistemic value T-monism and offer an alternative diagnosis of what is going on here. More generally, insofar as epistemic value K-monism is wedded to a knowledge-first epistemology then it will represent a far more revisionary approach to our contemporary thinking about epistemology. That said, of course, perhaps such revisionism is required, in which case one should not shirk from it.

Indeed, notice that taking this line might also enable one to make some progress with the value problem more generally. For one thing, if knowledge is the sole fundamental epistemic value then one might plausibly contend that one can resolve the secondary value problem, and thereby the primary value problem, on account of how on this view it follows immediately that knowledge is of greater epistemic value than any epistemic standing that falls short of knowledge. Given the plausible additional claim that no lesser epistemic standing is of greater non-epistemic (e.g., practical) value than knowledge, then the secondary problem is completely neutralised.

What about the tertiary value problem? Here matters are not so clear. We noted earlier that it does not follow from the fact that an epistemic good is fundamental that it is thereby of final value. Given that demonstrating that knowledge has final value is key to answering the tertiary value problem, it follows that one will need to do more to answer the tertiary value problem than simply defend epistemic value K-monism. If this further argumentation is not

forthcoming, then it would be incumbent upon the proponent of the monistic response to explain why the tertiary value problem need not be answered.

§1.8. RESPONDING TO THE SWAMPING PROBLEM III: THE PLURALIST RESPONSE

The only remaining option available to us is to deny claim (1) not by advancing a distinct form of epistemic value monism but rather by advocating an appropriate form of epistemic value pluralism.

To begin with, notice that simply opting for a form of epistemic value pluralism might not be enough by itself to rescue our intuition that knowledge is at least sometimes epistemically more valuable than mere true belief. Consider, for example, a form of epistemic value pluralism which maintained that both true belief and a further epistemic property—let's call it 'X'—are fundamental epistemic goods. Simply arguing for such a position would not, however, be enough to ensure the truth of (3). After all, suppose that it were the case that, necessarily, when one knows a proposition, one's belief *cannot* exhibit the target epistemic property, X. Although this may be unlikely, this doesn't seem to be a possibility that we can rule out in advance *a priori*, since it depends upon one's wider epistemological commitments (in particular, what theory of knowledge one endorses). If this were the case, however, then clearly that X is an epistemic property which, in addition to true belief, is fundamentally epistemically valuable would not suffice to ensure that knowledge is sometimes epistemically more valuable than mere true belief, and hence (3) would still be in question. The whole point of denying (1) in the face of this inconsistent triad, however, is surely so that one is able to consistently endorse (2) and (3).

Accordingly, what is required of this strategy is some way of denying (1) so that (3) gets the support that it requires. In particular, what is needed is a form of epistemic value pluralism which can account for why knowledge is sometimes of greater epistemic value than mere true belief by appeal to the fact that knowledge, in addition to true belief, is of fundamental epistemic value.¹⁵ Call this the *pluralist response* to the swamping problem.

Consider, for example, the kind of epistemic value pluralism endorsed by Kvanvig (e.g., 2003). Kvanvig argues that there are several fundamental epistemic goods other than true belief, the chief example that he offers in this respect being understanding. Interestingly, however, Kvanvig does not offer his epistemic value pluralism as a means of rescuing the intuition that knowledge is epistemically more valuable than mere true belief, since he instead uses the swamping argument to demonstrate that we are mistaken in thinking that knowledge is a particularly valuable epistemic standing, as opposed to understanding. Accordingly, it would seem that Kvanvig responds to our inconsistent triad by rejecting *both* (1) and (3).¹⁶

I don't think that this can be the full story, however, for while Kvanvig argues that understanding is not a species of knowledge, he does not argue for the stronger claim that it is *never* the case that when one knows a proposition one's true belief also has the relevant epistemic property of understanding. Indeed, although the details of his view are too involved to be usefully expounded here—we will be discussing understanding more fully in chapter four—his conception of understanding and its relationship to knowledge in fact strongly suggests that it is quite often the case that when one has what he calls 'propositional' understanding, one will also have knowledge of the corresponding proposition.¹⁷ If this is right, though, then Kvanvig ought to be willing to accept (3) after all, and thereby accept that knowledge is at least sometimes epistemically more valuable than mere true belief, since in those cases where knowledge and understanding coincide, the known true belief in question will have an epistemic value which is greater than the epistemic value of the corresponding mere true belief. The ultimate problem posed by the swamping argument on Kvanvig's view, then, is thus not for (3) at all, but rather for epistemic value T-monism (i.e., (1)).

Nevertheless, I think that Kvanvig is onto something here when he seems to in effect treat the swamping argument as ultimately undermining both (1) and (3). For suppose that one argues for the negation of (1) by appeal to a form of epistemic value pluralism which treats knowledge as sometimes epistemically more valuable than mere true belief because of the fundamental epistemic value of an epistemic property which is only sometimes present when one has knowledge (such as understanding). Indeed, just to make this point particularly vivid, suppose that one argues for the negation of (1) by appeal to a form of epistemic value pluralism which treats knowledge as *always* epistemically more valuable than mere true belief because of the fundamental epistemic value of an epistemic property which is *always*

present when one has knowledge, even though it is not sufficient, with true belief, for knowledge (i.e., the epistemic property in question is merely necessary for knowledge). Would we regard such a stance as demonstrating that knowledge has the kind of epistemic value that we typically suppose it to have?

I think not. In order to see this, notice that part of what is at issue in the wider debate about epistemic value is the central role that knowledge plays in epistemological inquiry. If it were to turn out, however, that knowledge is only of greater epistemic value than mere true belief because of the greater epistemic value of a necessary component of knowledge (still less, a non-necessary component), then that would surely threaten the central role that knowledge plays in epistemological theorizing almost as much as the claim that knowledge is *never* of greater epistemic value than mere true belief. After all, why should we now care whether we have knowledge, specifically, rather than just true belief plus the extra fundamentally valuable epistemic property? (The observant reader will spot that this is essentially a variant of the secondary value problem outlined in §1.2 as applied to the specific issue of epistemic value raised by the swamping problem).

Ideally, then, what we want is a defence of (3) which is able to support our wider intuition that knowledge is worthy of the central focus that it enjoys in epistemological theorising. The most straightforward way of doing this within the model of epistemic pluralism would simply be to argue that knowledge is itself a fundamental epistemic value, in addition to the fundamental epistemic good of true belief. The trouble with this suggestion, however, is that once one grants that true belief is a fundamental epistemic good then it is hard to see why knowledge should in addition be a fundamental epistemic value also, rather than being an epistemic standing which is instrumentally epistemically valuable relative to the epistemic good of true belief.

Moreover, the point raised earlier against epistemic value K-monism that the fundamental epistemic good in question needs to have final value if this response to the swamping problem is to extend to the value problem more generally is equally applicable here. For if the greater value of knowledge turns out to ultimately just be a greater non-epistemic instrumental value, then the pluralist response to the swamping problem will not extend to the tertiary value problem. Conversely, however, if one could show that knowledge is finally valuable, then one would thereby have shown that knowledge is fundamentally epistemically valuable as well.¹⁸

So if one wishes to extend the pluralist response to the swamping problem to the tertiary value problem it is essential that one is able to demonstrate that knowledge is not only fundamentally epistemically valuable but also finally valuable. Alternatively, if such additional argumentation is not forthcoming, then it will be incumbent upon the proponent of the pluralist response to explain away the tertiary value problem—i.e., demonstrate that this is a pseudo-problem that does not require a response.¹⁹

§1.9. CONCLUDING REMARKS

In this chapter we have examined the value problem for knowledge: the problem of accounting for the distinctive value of knowledge. We have found that there are three ‘levels’ of response required in this regard. The primary value problem demands that one show that knowledge is more valuable than mere true belief. The secondary value problem demands that one show that knowledge is more valuable than that which falls short of knowledge. And finally the tertiary value problem demands that one show that knowledge is more valuable than that which falls short of knowledge not merely as a matter of degree but of kind. We argued that any adequate response to the tertiary value problem must be able to defend the claim that knowledge has final value.

In addition, we also examined a problem concerning the value of knowledge which is superficially similar to the primary value problem—*viz.*, the swamping problem. As we saw, this problem in effect arises from the fact that three highly intuitive claims are collectively inconsistent. These claims are: (i) epistemic value monism regarding true belief (epistemic value T-monism); (ii) a general claim about value; and (iii) the thesis that knowledge is at least sometimes of greater epistemic value than mere true belief.

We found that there are three *prima facie* plausible responses to the swamping problem. The first—the practical response—appeals to the practical value of knowledge in order to explain why knowledge is not epistemically more valuable than mere true belief after all. This view is entirely compatible with epistemic value T-monism. The other two responses, in contrast, maintain that knowledge is a fundamental epistemic good and hence reject epistemic value T-monism. The first approach—the monistic response—argues that

knowledge is the only fundamental epistemic good, and hence blocks the swamping problem by appeal to a different form of epistemic value monism (epistemic value K-monism). The second approach—the pluralist response—instead puts forward a form of epistemic value pluralism which holds that both knowledge and true belief are fundamental epistemic goods.

We noted that the swamping problem poses a distinct problem to the value problem for knowledge. Indeed, it poses a distinct problem even when compared to the primary value problem for knowledge. Since ideally we want an account of the value of knowledge that can deal with both difficulties, it is thus incumbent upon us to consider how a response to the one problem bears on the other problem. In particular, since any response to the value problem for knowledge requires us to defend the thesis that knowledge is finally valuable, we need to consider how the three responses to the swamping problem would fit with this claim.

Interestingly, we noted that the practical response to the swamping problem is incompatible with the thesis that knowledge is finally valuable, and hence must regard the tertiary value problem as a pseudo-problem. In contrast, the monistic and pluralist responses to the swamping problem are compatible with the final value of knowledge. Thus, if one can show that knowledge is finally valuable, then this would put these two views at an advantage when compared with the practical response. Conversely, if it were to turn out that knowledge is not finally valuable, then the practical response to the swamping problem could not be faulted on account of its inability to deal with the tertiary value problem.

The central question raised by this chapter is thus whether knowledge is finally valuable, and this question will be the concern of the next chapter.²⁰

NOTES

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

² Of course, 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).

³ Elsewhere, I've called this 'Kaplan's problem'. Kaplan (1985) famously argues 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 in effect answers

the primary value problem while arguing that the secondary value problem does not need to be answered. 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.

⁴ Final value, and in particular its relation to intrinsic value, is discussed in more detail in the next chapter.

⁵ Though see Riggs (2002).

⁶ The following discussion closely follows that offered by Zagzebski (2003).

⁷ Indeed, the problem isn't specific to externalist epistemological proposals more generally either. The argument that was just ran as regards justification would work with equal force whether the notion of justification in question were cashed-out in an externalist or internalist fashion. All that matters is that justification has the target property in question (i.e., that it is only instrumentally valuable relative to the epistemic good of true belief).

⁸ Since something could be both practically and epistemically valuable, note that in what follows unless I specify otherwise when I talk of something being practically valuable I have in mind a practical *non-epistemic* value.

⁹ The observant reader will have spotted a kind of circularity in play here. On the one hand, one's account of epistemic value will have a bearing on what one counts as an epistemic good. On the other hand, one will evaluate a theory of epistemic value by considering how well it accords with the intuitive class of epistemic goods. Such circularity is unavoidable, but it is not vicious. Here, as elsewhere in philosophy, the solution is to seek a reflective equilibrium, where that might involve, for example, regarding some goods which one intuitively supposed as epistemic as being non-epistemic.

¹⁰ Sosa (2007, ch. 4) draws a similar distinction between fundamental and non-fundamental epistemic goods to that marked out here, though (if I've read him correctly anyway) our presentation of this distinction diverges in some important details.

¹¹ Note that sometimes this view is expressed in terms of the thesis that it is *truth* (rather than true belief) which is the only ultimate fundamental epistemic good (see, for example, Kvanvig 2003, ch. 2). The difference between these two views is complex and raises some subtle questions. However, since nothing is lost for our purposes by focussing on the version of epistemic value T-monism set out here, I will ignore this complication in what follows.

¹² Given that this is a monistic proposal, the epistemic value of true belief must *always* be non-instrumental.

¹³ Clearly I am here taking truth not to be an epistemic property of one's beliefs. If one wishes to think of truth as an epistemic property of belief, then one should reformulate this claim accordingly.

¹⁴ This view has been expressed to me in conversation by both Michael Lynch and Mike Ridge. See also Petersen (2008). Kvanvig (2003) could also be read as responding to the swamping problem by denying (3). As I note below, however, I think that ultimately he is best read as rejecting (1).

¹⁵ Note that one might further claim that there are other epistemic standings which are of fundamental epistemic value. Since what is crucial to the view (*qua* response to the swamping problem at least) is only that knowledge and true belief are fundamental epistemic goods, I will set this complication to one side in what follows.

¹⁶ Of course, evaluations of this sort are inevitably moot given that Kvanvig wasn't responding to the swamping problem in the specific form that it is presented here. What follows should be read with this *caveat* in mind.

¹⁷ For example, consider one's true belief that one's house burned down because of faulty wiring. It seems entirely plausible that it could be the case that this belief both has the epistemic property of knowledge and the epistemic property of understanding (i.e., that one both knows and understands that one's house burned down because of faulty wiring). As just noted, we will be discussing understanding more fully in chapter four. In the meantime, for more on Kvanvig's conception of understanding, see Grimm (2006), Brogaard (2007b), Elgin (*forthcoming*), Kvanvig (*forthcoming*) and Pritchard (*forthcomingb*); for more on the notion of understanding more generally, see Zagzebski (2001) and Riggs (2003).

¹⁸ For further discussion of the idea that final value is relevant to our understanding of the swamping problem, see Brogaard (2007a) and Pritchard (2007b; 2007d; 2008c; *forthcomingc*; *forthcomingd*). A precursor to this approach can be found in Percival (2003).

¹⁹ For more on the debate between epistemic value pluralists and epistemic value T-monists, see the exchange between David (2005) and Kvanvig (2005).

²⁰ For a recent survey of research on both the value problem for knowledge and the swamping problem, see Pritchard (2007b; cf. Pritchard 2007d). See also Baehr (*forthcoming*) and Pritchard (2008g; *forthcomingd*).