

Enquiry and the Value of Knowledge

What, if anything, can we learn about knowledge by reflecting on enquiry? Some philosophers think we can learn important truths. In two articles originally published in the early 1970s, Bernard Williams argued that philosophical theorising about the nature of knowledge had been led astray by a tendency to focus on what he called the ‘examiner situation’, and that the way to correct these errors was to consider instead the position of an enquirer. One aim of this paper is to identify a fundamental flaw in Williams’ arguments. My deeper purpose is constructive rather than critical, however. Even if reflection on enquiry doesn’t have a significant contribution to make to philosophical understanding of the *nature* of knowledge, I shall argue that it nevertheless has a significant contribution to make to philosophical understanding of the *value* of knowledge. The effect of this contribution is not so much to settle this or that issue within the existing debate but rather to reconfigure the framework within which discussion of the value of knowledge takes place. An important upshot of this reconfiguration is that the Platonic question that is normally the focus of discussions of the value of knowledge—‘How can knowledge be more valuable than mere true belief?’¹—may be far less important than it is usually taken to be.

My remarks fall into two parts. In the first part I evaluate the proposal that considerations about enquiry can be used to establish philosophically significant conclusions about the nature of knowledge, by criticising the arguments Williams presents in ‘Deciding to Believe’ (hereafter, DB) and ‘Knowledge and Reasons’ (hereafter, KR). I contend that Williams’ arguments aren’t compelling once we distinguish knowing the answer to a question from merely being a reliable source of information about the answer to a question. In the second part of the paper I develop the proposal that reflection on enquiry nevertheless has a crucial contribution to make to debates about the value of knowledge. This is because the Platonic question that normally initiates discussion of the value of knowledge is underwritten by a more fundamental question about the aim of enquiry: namely, why should an enquirer who wants the truth about whether *p* aim to acquire knowledge of whether *p*, and not merely a true belief about whether *p*?

1. Williams on Enquiry and the Nature of Knowledge

In DB and KR Williams maintains that certain widely endorsed putative necessary conditions for knowing that *p* are spurious. In DB, Williams denies that believing that *p* is necessary for knowing that *p*; in KR, Williams claims that ‘it is possible for A to know that *q* without its being the case that A can rehearse reasons, or at least adequate reasons, for *q*’ (2006: 50). In each instance, Williams argues that the view he regards as mistaken has been encouraged by a tendency on the part of philosophers to concentrate on what he calls ‘the examiner situation’, and that this view is rapidly exposed as erroneous once we concentrate instead on the position of an enquirer—more specifically, on the conditions under which someone (or something) would meet the needs of an enquirer attempting to find out the answer to a question. In each paper, then, the overall structure or form of Williams’ argument is the same.

¹ Millar, Haddock and Pritchard describe this question as ‘[t]he starting point’ (2009: 1) for philosophical discussion of the value of knowledge.

Note that the two parts of Williams' argument are independent of one another, in the following sense: even if a tendency to concentrate on the examiner situation is not the reason why philosophers have endorsed the putative necessary condition for knowledge (e.g. that knowledge requires belief), it might still be the case that reflection on enquiry discredits that condition, and even if reflection on enquiry does not discredit the putative necessary condition on knowledge, it might still be the case that the reason why philosophers have endorsed that condition lies in their tendency to concentrate on the examiner situation. Since my concern in this paper is with the question of what contribution (if any) reflection on enquiry has to make to epistemology, I shall restrict my attention to the second part of Williams' argument in each case. In relation to DB, then, I shall be asking whether considerations about enquiry show that knowledge does not require belief, and, in relation to KR, I shall be asking whether considerations about enquiry show that knowledge does not require the subject to be able to rehearse (adequate) reasons. I begin with the argument of DB.

In DB Williams' discussion of these issues is prompted by the example of a machine which, he claims, has knowledge but not beliefs. Having made this claim about the machine, Williams provides the following in the way of comment:

This [i.e. the claim that the machine knows without believing] goes against what is a rather deep prejudice in philosophy, that knowledge must be at least as grand as belief, that what knowledge is, is belief plus quite a lot; in particular, belief together with truth and good reasons. This approach seems to me largely mistaken. It is encouraged by concentrating on a very particular situation which academic writings about knowledge are notably fond of, that which might be called the *examiner* situation: the situation in which I know that p is true, this other man has asserted that p is true, and I ask the question whether this other man really knows it, or merely believes it. I am represented as checking on someone else's credentials for something about which I know already. That of course encourages the idea that knowledge is belief plus reasons and so forth. But this is far from our standard situation with regard to knowledge; our standard situation with regard to knowledge (in relation to other persons) is rather that of trying to find somebody who knows what we don't know; that is, to find somebody who is a source of reliable information about something. In this sense the machine could certainly know something. Our standard question is not 'Does Jones know that p ?' Our standard question is rather 'Who knows whether p ?' (1973: 146)

When we try to find somebody who knows what we don't know so that we can find it out from them, we are seeking the truth about the answer to that question, and so are engaged in enquiry. Thus, although Williams doesn't put the point in these terms, the proposal he is making here is that reflection on the position of an enquirer—in particular, on what an enquirer wants of an informant—acts as a corrective to the mistaken approach to knowledge, encouraged by concentrating on the examiner situation, according to which 'knowledge must be at least as grand as belief, that what knowledge is, is belief plus quite a lot; in particular, belief together with truth and good reasons'.

Is there anything in Williams' proposal? In line with my present concerns, I shall restrict myself to considering whether the 'mistaken approach' is exposed as such by reflection on enquiry.² Does reflection on enquiry show that the machine 'could certainly know something'?

The machine in question is introduced by Williams as having three basic features. Firstly, it produces 'assertions'. More precisely, 'it produces messages which express propositions [e.g. on print-outs or through a speech synthesiser], and it has a device which distinguishes between propositions which it asserts, as opposed, for instance, to propositions which it is prepared to hypothesise in the course of an argument' (1973: 145). The second feature is that these 'assertions' issue from states of the machine, and some of these states are based on other states of the same kind in a manner analogous to the way in which my belief that my neighbour is out might be based on my belief that the lights in his house are off. The machine 'goes through a causal process which is at least something like inference' (ibid.). Finally, the machine 'gathers information from the environment; it has various sensory bits and pieces which acquire information about the environment and these get represented in the inner states of the machine, which may eventually issue in appropriate assertions' (ibid.). Having so described the machine, Williams claims that the machine 'would not manifest belief; its states would not be beliefs. They would be instances of a much impoverished notion which I shall call a '*B* state'' (ibid.) According to Williams, the 'essential reason' why the *B* states of the machine would not amount to beliefs is that the machine is incapable of making insincere assertions. Williams is not explicit about why he regards the presence of this capacity as essential for 'full-blooded' belief.³ Interesting as these somewhat inchoate remarks are, the point I wish to press in response to Williams' argument does not concern his claim that the states of the machine would not be beliefs, but rather his further claim that the states of the machine could amount to knowledge. Accordingly, from now on I shall simply grant for the sake of argument that the machine is not a believer; the question I shall focus on is whether Williams succeeds in marshalling convincing considerations in support of the claim that the machine could nevertheless be a knower.

Williams introduces the proposal that the machine could know things in the paragraph immediately preceding the one I quoted earlier:

² Attentive readers will notice that Williams elides two importantly different claims in his description of the 'mistaken approach'. The claim that 'knowledge must be at least as grand as belief' is naturally understood to express the idea that believing that *p* is necessary for knowing that *p*. By contrast, the claim that 'what knowledge is, is belief plus quite a lot; in particular, belief together with truth and good reasons' is much stronger: it says not only that believing that *p* is necessary for knowing that *p*, but that it is possible to give a non-circular analysis of knowing that *p* as a conjunction of believing that *p* with further factors—in particular, truth and good reasons. It is possible to accept the first claim without accepting the second; this, of course, is the position of Williamson (2000). It is clear from the context that Williams' real interest is in the first of these claims.

³ A hint to his thinking may lie in the following earlier remark from DB, from a passage discussing insincere assertion: 'Belief lies at the level of what makes my acceptance sincere or insincere; it does not lie at the level of those acceptances themselves. So we have a picture of belief as of the internal state which my overt assertions may represent or misrepresent. Similarly, statements of intention or promises can be sincere or insincere because they do or do not represent an internal state of intention' (1973: 140–141).

It is, however, a notable feature of this machine that it could have true *B*-states which were non-accidentally arrived at, that is, which were not randomly turned out but were the product of the environment, the programming and so on; and these might be called ‘knowledge’. We could say of this machine that it knows whether so and so is the case, that it knows when the aircraft leaves, that it knows where somebody left something, that it knows where it is itself, that it knows where certain buildings are in the town. I would claim in fact that that the use of the word ‘know’ about this machine is not far away from a lot of uses of the word ‘know’ about human beings, whereas a use of ‘believe’ to refer to the *B*-states that the machine is in, is quite a long way away from the way in which we use the word ‘believe’. (1973: 146)

Does either of these considerations provide a compelling reason to judge that the machine is capable of knowledge, even if it is incapable of belief? The point that it is natural to *say* that the machine ‘knows’ things should not be taken too seriously. After all, it is natural to say to someone who wants to find out when Jane Austen was born that Wikipedia or Google will know when she was, but few, I think, would want to contend that such utterances are intended to be taken literally. Rather, what one means is that the person could *find out* when Jane Austen was born by consulting Wikipedia or Google.⁴ Someone who wished to deny that the machine can literally know things can plausibly maintain that it is natural to say that the machine knows only because we sometimes use ‘knows’ in this manner, to identify ways of finding things out: consulting the machine could certainly be a way of finding out (e.g.) when the aircraft leaves. What of Williams’ other point, that the machine could have true *B*-states which were non-accidentally arrived at? I take it that the point Williams means to be making here would be more accurately expressed by saying that the machine could have *B*-states which were non-accidentally true. He is certainly right about this, and it is also right that many epistemologists have taken such an ‘anti-luck’ condition to be a key requirement for knowledge.⁵ However, objects that are clearly incapable of knowledge can also satisfy this requirement: the statements in a book might be non-accidentally true. Thus, the fact that states of the machine can satisfy this requirement is no reason in itself to think that states of the machine can be states of knowledge.

In fairness to Williams, it should be said that he does not present these two considerations as good reasons, in themselves, for thinking that the machine could know things. He is clearly aware of the fact, for example, that the term ‘knows’ is used in different ways, so the mere fact that we could say of the machine that it ‘knows’ does not establish, in itself, that the

⁴ A quite different metaphorical use of the term ‘knows’ occurs when one uses a knowledge-ascription to convey the fact that something responds in appropriate ways to changes in its environment—e.g. when one says, of glasses with photochromic lenses, that they know when it’s sunny. (Thanks to Joe Cunningham for the example.)

⁵ Including Williams himself, as the following passage from KR demonstrates: ‘If [A] has come to know, then at least it is the case that he has not merely guessed—there must be a route by which the information has come to him, and the fact that what he believes is true must have contributed in some appropriate way to his having come to believe it. In this last consideration, obscure though it is, we shall see a central condition on the concept of knowledge’ (2006: 49).

machine could know things in the sense that human beings can.⁶ The purpose of my critical remarks is not to score cheap points off Williams, but rather to bring out the fact that the argument for the claim that the machine could be a knower, could know things in the same sense that human beings can, consists entirely in the assertion that the machine could service the needs of an enquirer. It is this argument we must now consider.

The argument occurs at the end of the passage I quoted earlier from DB. To repeat: ‘our standard situation with regard to knowledge (in relation to other persons) is rather that of trying to find somebody who knows what we don’t know; that is, to find somebody who is a source of reliable information about something. In this sense the machine could certainly know something’. The problem with the argument in this passage is that being a source of reliable information about whether *p* is not at all the same thing as knowing whether *p*. A tree is a source of reliable information about how old it is—just count the rings—but, as Craig points out, ‘[w]e don’t speak, even metaphorically, of a tree as knowing how old it is (1990: 35).⁷ It seems to me that the plausibility of Williams’ argument relies on us overlooking the distinction between knowers and mere sources of information. It is undeniable that a machine of the kind he describes could be a source of reliable information, and, provided we are not mindful of this distinction, this might persuade us that the machine is capable of knowing things. Once the distinction is drawn, however, it is not clear that the machine can be anything more than a reliable source of information. If this is all the machine can be, the machine cannot literally know anything.

Let us turn to KR. As I said earlier, Williams’ concern in that paper is not to dispute the claim that knowing *p* requires believing *p*, but rather to establish that ‘it is possible for A to know that *q* without its being the case that A can rehearse reasons, or at least adequate reasons, for *q*’ (2006: 50). This formulation is inadequate, because substituting an indicative sentence for the propositional variable ‘*q*’ produces a sentence that is ungrammatical: it does not make sense to say ‘it is possible for A to know that Fergus smokes without its being the case that A can rehearse reasons, or at least adequate reasons, for Fergus smokes’. A better formulation, consonant with Williams’ intentions, is that it is possible for A to know that *q* without its being the case that A can rehearse reasons, or at least adequate reasons, in favour of the truth of his belief that *q*.⁸ This claim provides one way of understanding the dispute between

⁶ This is indicated by his next remark in DB: ‘I would claim in fact that the use of the word ‘know’ about this machine is not far away from a lot of uses we make of the word ‘know’ about human beings, whereas a use of ‘believe’ to refer to the *B*-states that this machine is in, is quite a long way away from the way in which we use the word ‘believe’ (1973: 146).

⁷ Doesn’t this contradict the point I made earlier, that we sometimes use sentences of the form ‘X knows whether *p*’ to convey that someone could find out whether *p* from X? After all, if one can find out how old a tree is by counting its rings, surely it ought to be quite natural to say that the tree knows how old it is, but, as Craig observes, this is a very unnatural thing to say. We can explain why this metaphorical use of ‘knows’ is unnatural in relation to the tree by appealing to the thought that the metaphor seems apt only when there is something like a conversational exchange of asking and being told. In some sense, one can ‘ask’ the machine when the aircraft leaves, and be ‘told’ when it does in reply. The same point applies in relation to internet search engines—hence the existence of ‘Ask Jeeves’. By contrast, it is strange to think of counting tree rings as a way of ‘asking’ a tree how old it is.

⁸ Isn’t it a problem with this formulation that it implies that believing that *p* is necessary for knowing that *p*, something Williams denies? No. In KR Williams explicitly assumes this for the sake of argument, as we shall see shortly.

internalists and externalists in epistemology. On this understanding, the externalist denies that the ability to rehearse reasons in favour of the truth of one's belief that *p* is required for knowing (and perhaps even justifiably believing) that *p*. The internalist, by contrast, insists on its being necessary. Using the terminology in this way, Williams sides with the externalist.⁹

Williams argues for externalism on the basis of considerations about enquiry:

Let us return to the cases in which our interest in another's possession of knowledge concerns a matter about which we lack knowledge: our concern (unlike the examiner's) is to find out, not so much about this person, as about the matter in hand. If "p" represents some sentence which may be used to make a statement of fact, let "wh-p" represent in general direct and (with any necessary grammatical modification) indirect questions that may be formed out of that sentence...Now the following statements all seem to me to be true:

- a. In many standard situations, all that is necessary for it to be the case that A knows wh-p—besides his actually being right in this case (see (b) below)—are such things as that A is almost always right about matters of this kind, because e.g., in the matter of the train, he has learned up the time-table; or that we know, what perhaps A himself does not know, that A has come by his beliefs on this subject by reliable means. Thus it may be that A has come by his beliefs from having been told by B, and we know that B is a reliable authority on these matters, though A himself may not, never (for instance) having reflected on that question.
- b. If conditions of this kind are satisfied, then all that is further necessary for it to be the case that A knows wh-p is that the beliefs he has with regard to the question be true, i.e. that he actually be right. (I shall assume throughout this paper that *belief*, involving a fair degree of conviction, is a necessary condition of knowledge. I doubt whether that is true without qualification, but the assumption will serve for the present discussion.)
- c. If A knows wh-p, and P is the class of correct answers to the question "wh-p?," then for some member of P, q, A knows that q. This rather cumbersome formulation is intended to allow for such facts as that if "wh-p?" is a "who"-question—say, of the form "who did X?"—A may know of a certain person under some descriptions that he did X, but not under other descriptions.

If statements (a)–(c) are all true, then it is possible for A to know that q without its being the case that A can rehearse reasons, or at least adequate reasons, for q. For the sorts of conditions mentioned in (a) do not necessarily imply anything, or anything very substantial, about A's *consciousness* of reasons for q, or of his own relation to the truth of q. (Williams 2006: 49–50)

⁹ Although I will use the internalism/externalism terminology in this way throughout this paper, it is not at all important for my purposes that this is 'the' correct way of understanding the debate—supposing that there is such a thing. Readers who dislike this way of understanding the debate should regard my way of understanding the terminology as purely stipulative.

As I understand it, Williams' argument has three basic premises. The first is that S can know *wh-p* without being able to rehearse (adequate)¹⁰ reasons in favour of the truth of his belief(s) about *wh-p*. Call this the *enquiry premise*, because it is supposed to be established by reflection on the conditions under which we would judge someone to know *wh-p* when we are in the position of enquirers seeking the truth about *wh-p*. The second premise is that, if A knows *wh-p*, then, for at least one *q* in the set of correct answers to the question of *wh-p*, A knows that *q*. Knowledge *wh-p* entails knowledge *that*. The final premise is that if S knows that *p* then S believes that *p*. We can now reason as follows to arrive at the conclusion that it is possible for a subject to know that *p* without being able to rehearse reasons in favour of the truth of his belief that *p*. Suppose that A knows *wh-p* and, furthermore, that A is unable to rehearse reasons in favour of the truth of his belief(s) about *wh-p*. (The possibility of such a case is a consequence of the enquiry premise.) It follows from the second premise that, for some *q* in the set of correct answers to the question *wh-p*, A knows that *q*, and, from the third premise, that A believes that *q*. Trivially, where *q* is a correct answer to the question of *wh-p*, the belief that *q* qualifies as a belief about *wh-p*. Given that A cannot rehearse reasons in favour of any of his beliefs about *wh-p*, it follows that A cannot rehearse reasons in favour of his belief that *q*. But A knows that *q*. So A knows that *q* even though he cannot rehearse reasons in favour of his belief that *q*; so it is possible for a subject to know that *p* even if he is unable to rehearse reasons in favour of his belief that *p*.

Is Williams' argument compelling? One might dispute the premise that knowing *p* entails believing *p*. This would miss Williams' point, however. As we have seen, Williams was doubtful himself that knowledge requires belief, and in KR he only assumes that it does for the sake of argument. His point is that even if knowledge does require belief, we should still reject the stronger claim that knowledge requires the subject to be able to rehearse reasons in favour of the truth of his belief. A quite different response to Williams' argument would be to question the assumption that knowledge *wh-p* entails knowledge *that*. Recently, Jonathan Schaffer (2007) has attacked what he calls the 'reductive view' of knowledge-*wh* ascriptions, but closer inspection of Schaffer's argument reveals that he means to challenge only the proposal that knowing that *q*, where *q* is a correct answer to the question *wh-p*, is sufficient for knowing *wh-p*. It is consistent with Schaffer's argument that knowing that *q*, where *q* is a correct answer to the question *wh-p*, is necessary for knowing *wh-p* (indeed, Schaffer endorses this claim), and this is all that is needed for Williams' argument to go through. Since I do not wish to question the validity of the reasoning, or the apparently obvious principle that, where *q* is a correct answer to the question *wh-p*, the belief that *q* qualifies as a belief about *wh-p*, the only issue that remains is whether Williams supplies compelling grounds for the enquiry premise. I shall argue that he does not. Once again, the considerations he adduces are undermined by the distinction between knowing and merely being a reliable source of information.¹¹

¹⁰ For smoother exposition, I will drop Williams' hedge from now on.

¹¹ Note that I will not concern myself with the question of what it is to be able to rehearse reasons, or adequate reasons, in favour of one's beliefs, because the objection I wish to make to Williams' argument does not depend on niceties about this.

First, let me identify what is not in dispute here. I do not want to raise doubts about the claim that an enquirer could find out *wh-p* from A, provided he knows that A arrived at his belief about *wh-p* by reliable means. This ‘concession’ is perfectly consistent with an internalist stance on knowledge and justification. Even a dyed in the wool internalist can accept that an enquirer could come to know *wh-p* by gauging A’s opinion on the matter, provided he knows that A arrived at his belief about *wh-p* by consulting B, and that B is a reliable authority. One can deny that Norman knows or justifiably believes that the President is in New York whilst still allowing that an enquirer who knew of Norman’s clairvoyant powers could come to know of the President’s whereabouts by consulting Norman.¹² Indeed, I take it that most internalists would accept this. The crucial point, however, is that this does not settle the question of whether A and Norman *know* the things that others are able to learn from them. One can learn *wh-p* from someone who is a reliable source of information about *wh-p*, but being a reliable source of information about *wh-p* is insufficient for knowing *wh-p*. I suggested above that Williams’ *B*-state machine can plausibly be regarded as a mere source of information, and I think that the internalist can make the same move with respect to alleged instances of ‘unreasoned knowledge’. That is to say, they should hold that, strictly speaking, although A is a reliable source of information about (e.g.) the train times, A does not literally know when the trains go. Even in cases in which one learns that *p* from a sincere assertion that *p* by someone who truly believes that *p* via a reliable method, the person’s belief that *p* may be neither justified nor knowledge, and they may be functioning as no more than a reliable source of information; or so the internalist should say, at any rate.

Let me clarify what is being claimed here. I am not asserting that A and Norman do not really know, that they are only reliable sources of information. Rather, the claim is that once the distinction between knowing and being a reliable source of information is on the table, *it is not obvious* that subjects like A and Norman can serve as anything more than sources of information. It helps to think in terms of a spectrum of cases here. At one end we have clear cases of mere sources of information—tree rings, books, DVDs, and the like. At the other end of the spectrum we have clear cases of knowledge—the knowledge of mature human beings like you and me. Between lie cases where things are not so clear, such as Williams’ *B*-state machine and the ‘unreasoned knowledge’ of subjects like A. It might turn out that we should count these cases as genuine instances of knowledge. Recall, however, that our present concern is with whether the considerations about enquiry Williams adduces establish significant conclusions about the nature of knowledge. It should be clear that these considerations are powerless when it comes to determining whether or not a case on the spectrum is one of knowledge. The fact that an enquirer could use A to get what he wants does not speak in favour of the idea that A knows. When we are in the position of enquirers, Williams says, ‘our concern (unlike the examiner’s) is to find out, not so much about this person, as about the matter in hand’. An enquirer can find out *wh-p* from a mere source of information about *wh-p*, so the fact that an enquirer can find out *wh-p* from A establishes only that A is a source of information about *wh-p*; it does not speak to the question of whether A knows. Whether he does is not relevant to the enquirer’s concerns.

¹² The example is, of course, due to Bonjour (1980).

So far I have argued that the argumentative strategy employed by Williams, that of appealing to the thought that a person who does not satisfy putative necessary conditions for knowledge could nevertheless satisfy the needs of an enquirer, cannot be used to establish significant results about what is required for knowledge. The reason is simple: a mere source of information can give an enquirer what he is after, but someone (or something) can be a source of information on some matter without knowing about it. It is therefore open to the proponent of the relevant necessary condition to argue that this applies in the cases appealed to by Williams.

2. Enquiry and the Value of Knowledge

Let us turn, at last, to the place where I think that reflection on enquiry does have a significant contribution to make to philosophical understanding of knowledge. This is to the resolution of puzzles about the value of knowledge. To see why this is so, it is necessary to begin at the beginning, with Plato's seminal discussion of the value of knowledge in the dialogue *Meno*.

In *Meno*, Socrates famously compares a guide who knows the way to Larissa with a guide who merely has a true belief about the way. He claims that the latter guide is just as good as the former, and, more generally, that true belief 'is just as good a guide as knowledge, when it comes to guaranteeing correctness of action' (97b). Furthermore, true belief is obviously just as true as knowledge. These observations provoke puzzlement in *Meno*:

All this is making me wonder, Socrates, why, if this is so, knowledge is so much more highly valued than true belief and on what grounds one can distinguish between them. (97c–97d)

What does *Meno* mean in asking why 'knowledge is so much more highly valued than true belief'? There are at least two different ways of understanding this part of *Meno*'s question. On one interpretation, *Meno* is here alluding to the fact that human beings desire and seek knowledge, and not merely true beliefs. (As Aristotle later put it, 'All men by nature desire to know' (*Metaphysics*: I.980a20).) Following this line, we can summarize *Meno*'s question by asking, why do we desire and seek knowledge, and not merely true beliefs, given that true belief is just as true, and just as useful, as knowledge? Call this question *the enquiry question*.¹³ The enquiry question is one that concerns us in the role of enquirers, i.e. as subjects who desire and seek the truth. *Meno*'s problem is often formulated in a way that does not bring in the idea of a subject engaged in enquiry, however. According to this second understanding of *Meno*'s question, the problem is to explain why it is better to know that *p* than to have a mere true belief that *p*. Call this question *the comparative state question*. When the problem is expressed in this way, the predicament of an enquirer who neither knows nor truly believes that *p*, but wants to know whether *p*, disappears from view. The exercise becomes one of comparing and contrasting two subjects, one of whom knows that *p* and the other of whom merely truly believes that *p*, and trying to identify respects in which the former subject is better off than the latter. This second interpretation of the question is

¹³ In other work, I have also referred to this question as 'the activity question'.

encouraged by the subsequent direction of Plato's discussion, which proceeds in precisely this way.¹⁴

The first claim I wish to emphasise is that we really do have two different questions here. The fact that the second question no longer mentions us in the role of enquirers, as subjects engaged in the pursuit of truth, is not merely a cosmetic difference. There are various ways of substantiating this point, but perhaps the most fundamental is to point out that the two questions presuppose different things. The enquiry question asks why we desire and seek knowledge, and not merely true beliefs, given that true belief is just as true, and apparently just as useful, as knowledge. This question presupposes that when we are interested in having the truth on some matter we are at least sometimes interested in having the truth, more specifically, in the manner of knowing the truth. This presupposition is difficult to deny. We frequently profess that we want to know things, from the trivial ('I want to know the solution to this crossword clue') to the momentous ('I want to know who my real father is'). We also speak of wanting to work out/find out/discover/determine etc., whether..., or who..., or when..., or what..., or why..., or where..., or how..., and working out/finding out/discovering/determining the answer to a question of any of these forms is plausibly a matter of acquiring knowledge of the answer to that question. Given that we often express our desires and aims in more or less explicitly epistemic terms, it seems undeniable that we do in fact desire and seek knowledge of the answers to many questions. By contrast, the comparative state question presupposes something quite different. That question asks why a subject who knows that *p* is in a superior position, other things being equal, to a subject who merely truly believes that *p*. Obviously, this presupposes that a subject who knows that *p* is in a superior position, other things being equal, to a subject who merely truly believes that *p*. Is this presupposition correct? It is not clear that we have any commitment to this claim prior to engaging in philosophy. Nor is there any obvious respect in which knowledge is superior to mere true belief, as the Platonic Socrates shows. The question this raises, of course, is why so many philosophers have nevertheless assumed that knowing that *p must be* in some way superior to merely truly believing that *p*. Why have they taken it that the comparative state question raises a serious philosophical problem about the value of knowledge?

I think philosophers have assumed this because they have assumed that the right way to answer the enquiry question is by answering the comparative state question. That is to say, they have assumed that the correct way to explain why enquirers are interested in acquiring knowledge, and not merely true beliefs, is to identify respects in which knowing that *p* is superior to merely truly believing that *p*. This is not a stupid assumption: in general, a good way of explaining why an agent has X as his aim rather than Y is to identify respects in which X is superior to Y relative to the agent's interests. I suspect that philosophers have assumed that the explanation of why we are interested in knowledge, and not merely in true belief, must conform to this pattern. With this assumption in place, the question of why we are interested in knowledge and not just true belief becomes the question of why it is

¹⁴ Note that here my primary concern is not with the interpretation of Plato, but rather with the philosophical basis of the Meno problem and how that basis has been understood by recent philosophers. It would be interesting to know if either interpretation of Meno's question is favoured by the original Greek.

advantageous to know that p rather than to merely truly believe that p . Hence, I conjecture, the widespread acceptance of the presupposition of the comparative state question, and also the failure to distinguish the comparative state question from the enquiry question. At least in principle, however, there might be ways of explaining why an agent seeks X rather than Y other than by appealing to the idea that X and not Y. This idea is important, and I will come back to it towards the end of this paper. Before I do so, however, I want to offer some evidence in support of my hypothesis about why philosophers have tended, when reflecting on the value of knowledge, to focus on the comparative state question. There is not space here to attempt anything close to an exhaustive survey of the literature. All I can do is identify a few passages, and a few features of the wider debate, which seem to me to betray the fact that the fundamental concern here is with the interests of enquirers rather than the comparative value of states.

Let me begin with a passage from Duncan Pritchard's contribution to Haddock, Millar and Pritchard (2010). In the passage Pritchard argues that even if we are able to solve what he calls the 'primary value problem'—i.e. the comparative state question—by identifying some feature of knowledge that mere true belief lacks, this still may not 'secure our intuition that knowledge is distinctively valuable' (2010: 6), for we may still be left with what Pritchard calls the 'secondary value problem'. The secondary value arises in the following way. Suppose that the feature, F, in virtue of which knowledge is more valuable than mere true belief is insufficient for knowledge—i.e. that 'S has a true belief that p with feature F' does not entail 'S knows that p '. Then we would still appear to be left with the problem of explaining why it is better to know that p than to have a mere true belief that p with feature F. To bring out the force of this problem, Pritchard asks us to imagine 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' (2010: 7). What he says in response to this possibility is revealing:

On the face of it, this lacuna [i.e. our inability to solve the secondary value problem, despite solving the primary value problem] 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). (2010: 7)

According to Pritchard, the problem with leaving the secondary value problem unanswered is that it seems to imply that enquirers should be interested only in the state, true belief + F, which, though superior to mere true belief, is nevertheless insufficient for knowledge. In other words, in leaving the secondary value problem unanswered we forfeit the possibility of giving an adequate response to what I have been calling the enquiry question. But if this is 'the problem' with failing to answer the secondary value problem, then presumably the

reason why Pritchard is interested in that problem is that he thinks we need to explain why enquirers are (rightly) interested in knowledge and not just in that which falls short of knowledge, and assumes that the explanation of why this is so must conform to the pattern identified above, that of identifying respects in which knowledge is superior to anything that falls short of knowledge. Furthermore, if this is why Pritchard takes it that we need to solve the secondary value problem, presumably it is also at least part of the reason why he takes it that we need to solve the primary value problem.¹⁵ My diagnosis of why philosophers have tended to focus on the comparative state question (or versions thereof), and to accept its presupposition without argument, seems to fit Pritchard's discussion like a glove.

Of course, Pritchard is only one participant (albeit an influential one) in an ongoing debate about the value of knowledge, and it would be premature in the extreme to assume that his thinking on this matter is representative of everyone else's. To strengthen the case for the hypothesis that philosophers have focussed on the comparative state question because they have assumed that the correct response to the enquiry question is to appeal to the idea that knowledge is superior to mere true belief, I shall show how that hypothesis can be used to explain otherwise puzzling features of the 'intuition' that knowledge is superior to mere true belief. To do so, I shall draw on the critical analysis of the value problem offered by Jason Baehr (2009).

Baehr argues that what he regards as the standard conception of the value problem is unmotivated.¹⁶ Although he accepts the existence of 'a widespread, reasonably compelling, pretheoretical intuition to the effect that knowledge is more valuable than true belief' (2009: 45, emphasis removed), he contends that the content of this intuition has been misconstrued, and that once it is correctly understood the value problem no longer arises in its standard form. He draws attention to two features of the intuition in particular. The first is that it is supposed to be entirely formal, in the sense that it does not specify, even implicitly, any features in virtue of which knowledge is more valuable than mere true belief. The intuition is supposed to be just that there are such features, without containing any hints about *what* they are. The second feature of the alleged intuition is that it is entirely general, in the sense that any instance of knowledge is supposed to be more valuable than the corresponding instance of true belief.¹⁷ Baehr proceeds to argue that our pretheoretical intuition that knowledge is more valuable than true belief has neither of these features, and consequently that the intuition is incapable of motivating a formal constraint on the analysis of knowledge—something he takes to be central to the standard conception of the value problem. He suggests

¹⁵ Pritchard cites another supposed reason for thinking that knowledge must be more valuable than mere true belief: '[p]resumably, this is the reason why knowledge—and not, say, justified true belief—has been the principal focus of generations of epistemological theorising' (2010: 5). The fact that knowledge has been the focus of epistemology suggests only that knowledge is a central concern in our cognitive lives, however. It remains to be seen whether this is to be accounted for in terms of the idea that knowledge is a distinctively valuable state.

¹⁶ See Baehr (2009: 42) for his account of the problem as involving a formal constraint on the analysis of knowledge.

¹⁷ Clearly, an 'other things being equal' clause is required here. Suppose you tell me, truthfully, that if I have a mere true belief about the way to Larissa you will not harm me, but if I know the way you will have me murdered by bandits. In these circumstances it is obviously better for me to have a mere true belief about the way than to know it. I do not want to get sidetracked by this kind of issue here.

that the reason why philosophers have mistakenly taken the intuition to have these features is that ‘those who have addressed the [value] problem appear not to have thought very carefully about the specific character and content of the guiding intuition; instead they have tended to treat it as a kind of necessary, self-evident truth’ (2009: 54).

I agree with Baehr that the alleged intuition is normally taken to be both formal and general, in something like the senses he identifies. What I want to take issue with is his explanation of why philosophers have taken the supposed intuition to have these features. I think that there is more going on here than the mere absence of careful thought. Both features can be elegantly explained on the assumption that the fundamental concern is with the enquiry question.

It should be obvious how my hypothesis can explain the formality of the ‘intuition’ that knowledge is more valuable than mere true belief. On my understanding, the ‘intuition’ arises from the perceived need to explain why enquirers care about knowledge, and not just true belief. The thought is that knowledge must have some feature (or features) that are capable of playing this explanatory role. But simply specifying that the feature must be capable of playing this explanatory role leaves it wide open what the feature is; hence the formality of the alleged intuition.

There is a twist in the tale here worth noting. Despite the formality of the intuition that knowledge is more valuable than mere true belief, certain accounts of the value of knowledge have been objected to on the grounds that the distinctive value-conferring features of knowledge that they identify do not provide a satisfying response to the Meno problem.¹⁸ I have in mind what are sometimes called ‘virtue theoretic’ accounts of knowledge and its value. On the face of it, the cogency of this style of objection is puzzling: if the intuition that knowledge is more valuable than true belief is entirely formal, as Baehr puts it, how can one object to an account of the value of knowledge on the grounds that it identifies the wrong kind of (admittedly value-conferring) feature? My diagnosis of why philosophers have been interested in the comparative state question provides an answer. Although, in principle, any number of features might explain why enquirers rightly prefer knowledge to mere true belief (hence the formality of the ‘intuition’ that knowledge is more valuable than true belief), it may nevertheless be clear of certain allegedly distinctive features of knowledge that, although they would confer additional value on knowledge if genuine, the kind of value they confer is not one that enquirers are necessarily interested in. If my diagnosis is correct, when this occurs it should be felt that the proposed solution to the comparative state question somehow misses the point; and my diagnosis will be additionally confirmed if this is connected with the idea that the solution fails to engage with the concerns of an enquirer. Both of these expectations are met in discussions of virtue theoretic accounts of the value of knowledge, as the following passage from Haddock, Millar and Pritchard illustrates:

Virtue theorists attempt to show that knowledge is valuable for its own sake since there is value in having come by the truth through the exercise of particular virtues or

¹⁸ Again, whether or not we should identify the Meno problem with what I have been calling the comparative state question is something I wish to remain neutral on. Certainly, this is how it has been standardly understood.

competences...As Ernest Sosa views the matter (2007: lecture 4), we rightly think that credit is due to performances, for instance, the shot of archer, which achieve the aim of relevant activity not merely by accident or luck, but due to the exercise of competence on the part of the performer. This interesting approach raises a host of interesting issues...It also faces a number of challenges. For instance, one might think that an adequate answer to the *Meno* problem should explain why the tourist seeking the way to the cathedral should prefer to know rather than have a merely true, or merely true and justified, belief on the matter. From this perspective it is not clear that the consideration that knowing would be worthy of credit gives the right kind of answer. We were led by the problem to wonder what good it would be to an enquirer to know, not how good the enquirer would be to know. (2009: 3)

Not only, then, is my diagnosis able to explain the formality of the 'intuition' that knowledge is more valuable than mere true belief: it can also explain why, in spite of this formality, accounts of the value of knowledge can nevertheless fail to engage with the concern driving the comparative state question, even if it is allowed that they identify distinctive, value-conferring features of knowledge (e.g. that knowledge, unlike mere true belief, is the product of the exercise of intellectual virtues or competences).

What about the generality of the supposed intuition? Is this explained by the hypothesis that philosophers have focused on the comparative state question because they have assumed that answering that question is the way to answer the enquiry question? Recall that the claim that the intuition is entirely general is the claim that, according to the intuition, any instance of knowledge is more valuable than the corresponding instance of true belief. One way of explaining this feature of the intuition in terms of my diagnosis would be to appeal to the idea that whenever enquirers desire or seek the truth about the answers to questions, what they desire or seek is knowledge of the answers to those questions. If one then thought that the way to tackle the enquiry question was by tackling the comparative state question, it would seem that, in order to explain why knowledge was all that interested enquirers, one would have to identify value-conferring features of knowledge in virtue of which knowledge was of greater value to enquirers in more or less any circumstances, regardless of the subject matter they were concerned with. Unfortunately, it is false that whenever we are interested in the truth in some matter what we are interested in is knowledge. Sometimes we aspire only to guess at the truth, or to arrive at correct assumptions for practical purposes, and when this is the case our aspiration clearly falls short of acquiring knowledge of the answers to our questions. Thus, we cannot appeal to the claim that whenever enquirers are interested in truth they are interested in knowledge in order to explain the generality of the supposed intuition that knowledge is more valuable than mere true belief. However, I think that there is a related claim that can do the required explanatory work. This is that whenever enquirers are at least interested in arriving at true beliefs about the answers to questions, they are interested in arriving at true beliefs more fundamentally because they are interested in acquiring

knowledge of the answers to questions.¹⁹ Although enquirers sometimes aim for objectives that fall short of knowledge (e.g. correct guesses), they never ordinarily do so in virtue of merely aiming to arrive at true beliefs. Intuitively, there is something odd about the idea of merely aiming to acquire a true belief about the answer to some question without aiming to find out or determine the answer. (The oddity is not to do with the bare possibility of doing this, but rather to do with *why* one would ever do this.) If one displays some sensitivity to this point, and one also thinks that the correct way to tackle the enquiry question is by tackling the comparative state question, then it becomes comprehensible why the intuition that knowledge is more valuable than mere true belief should strike one as entirely general in something like the sense identified by Baehr. For one will think that we need to explain why, on any occasion on which a subject is at least interested in acquiring a true belief on some matter, he is interested not only in acquiring a true belief but in acquiring knowledge. If we are to do this by tackling the comparative state question, then we will have to explain why any instance of knowledge is more valuable than the corresponding instance of true belief. In committing ourselves to explaining this, we commit ourselves to the generality of the ‘intuition’ that knowledge is more valuable than mere true belief.

To summarise, I have presented three main pieces of evidence for the hypothesis that philosophers have accepted the presupposition of the comparative state question, that it is better for the subject to know that *p* than to have a mere true belief that *p*, because they assumed that the way to answer the enquiry question is by appealing to this claim. The first piece of evidence was textual, taken from Pritchard’s discussion of the secondary value problem. To further confirm the hypothesis, I showed how it could be used to explain certain features of the ‘intuition’ that knowledge is more valuable than mere true belief identified by Baehr. I indicated how it could be used to explain the ‘formality’ of the alleged intuition, whilst also accounting for the otherwise puzzling cogency of an objection to virtue theoretic accounts of the value of knowledge. I also showed how it can be used to explain the generality of the supposed intuition. Taken together, this seems to me impressive evidence in favour of taking the hypothesis seriously.

In any case, my primary interest is not in why philosophers have tended to accept the presupposition of the comparative state question, but rather in what reason there actually *is* think that knowledge is superior to mere true belief. It seems to me that the only good general reason to accept this claim, in advance of identifying specific respects in which knowledge is superior to mere true belief, is that we need to appeal to it in order to explain why knowledge, and not just true belief, is the aim of enquiry. To see this, imagine that we had an alternative explanation of why knowledge and not just true belief is of interest to enquirers, one that did not appeal to the idea that knowledge is superior to mere true belief. I think it would be odd at this point if someone continued to press the question of why someone who knows that *p* is in a superior state, other things being equal, to someone who merely truly believes that *p*. We already have a perfectly adequate understanding of why knowledge, and not just true belief,

¹⁹ A closely related idea is expressed by Matthew Soteriou when he writes, ‘In cases of theoretical reasoning in which it *is* a requirement on fulfilling one’s aim that one acquire a belief as to whether or not *p*, this is because the more fundamental aim is that of acquiring knowledge of *whether p*’ (2013: 350–351).

is of interest to enquirers; why think it will be possible, in addition, to identify respects in which knowledge is superior to mere true belief? What is it that is meant to be justified or explained by reference to such features? There is no obvious answer to this question. If this right, then the basic problem in this area is really that presented by the enquiry question, and tackling the comparative state question is no more than *one* way of attacking that problem. Furthermore, it should also be clear that the line of attack represented by the comparative state question is not a promising one. I refer here not just to the conspicuous absence of an agreed solution to the comparative state question, despite repeated attempts, but also to the formidable obstacles that stand in the way of giving an adequate answer to the enquiry question by way of the comparative state question. For one thing, the account has to explain not just why knowledge is often or usually of greater value to an enquirer than mere true belief, but why its value exceeds that of mere true belief in such a way that on any occasion in which an enquirer is interested in acquiring at least a true belief on some matter, this is because he wishes to know about it. For another, it also has to explain not just why knowledge is more valuable than mere true belief, but why knowledge is more valuable than anything that falls short of knowledge; that is to say, it has to be capable of resolving Pritchard's secondary value problem. In the face of these obstacles, and the record of failure, it is surely worth considering whether we have not got off on the wrong track in pursuing the comparative state question. But how else might we go about explaining why enquirers desire and seek knowledge, and not just true beliefs, other than by identifying respects in which knowledge is superior to mere true belief?

It is impossible for me to provide a full answer to this question here. All I want to do is highlight three points which seem to me of central importance. In Plato's original discussion, Meno is led to raise his problem by Socrates' observation that true belief is just as true, and apparently just as useful, as knowledge. The first very basic point I wish to make is that care needs to be taken in spelling out exactly why this observation makes our interest in knowledge puzzling. *That* it raises a puzzle about why enquirers desire and seek knowledge I do not deny—Meno's question is undoubtedly intuitively forceful. The point I want to insist on is just that philosophers ought to pay far more attention to exactly what the puzzle is about enquiry before they assume that the way to answer it must be by identifying respects in which knowledge is superior to mere true belief. The second, connected point is that we should not assume, in tackling the enquiry question, that enquirers are basically or fundamentally only interested in acquiring true beliefs. Two remarks should be made here. The first is that there is no obvious reason to accept this assumption. The point that mere true belief is just as true, and just as useful, as knowledge, is irrelevant here. It is agreed that our conception of the aim of enquirers should be more specific than 'useful mental state with a true content'; why should it be more specific in the manner of believing truly than in that of knowing?²⁰ The second is that, if we make this assumption, it is very hard to see how we can possibly explain why enquirers like ourselves have any interest in knowledge. For example, the argument that enquirers seeking true beliefs ought to use methods that reliably produce true beliefs, and that true beliefs produced by reliable methods constitute knowledge, cannot be used to show that

²⁰ Williamson makes this objection to the functional account of the concept of knowledge presented by Craig (1990). See Williamson 2000: 31, note 3. (I have worded the objection in much the same way as Williamson.)

enquirers seek not just true beliefs but knowledge; to draw this conclusion is to conflate the means an enquirer ought to use to pursue his end with the end itself.²¹ The final point is that, in thinking about why enquirers are interested in knowledge and not just in true beliefs, we must be sensitive to what is involved in an enquirer arriving at so much as a belief through his investigations. One might think, roughly, that until I take myself to have worked out or determined (i.e. acquired knowledge of) the answer to my question, I won't have arrived at so much as a belief about the answer to my question, as opposed to (say) a mere hunch or suspicion about the answer. If such a claim can be made out, it may shed light on what would be odd about merely setting out to acquire true beliefs. What would be odd about merely adopting this objective is that I will not have achieved it until I take myself to have achieved the further objective of acquiring knowledge. Let me emphasise again, however, that these remarks are merely intended as pointers to an alternative approach to the enquiry question that bypasses the comparative state question. The detailed spelling out and defence of such an approach must be left for another occasion.

Where does all of this leave the comparative state question? Note that I have not rejected the presupposition of that question, that knowledge is superior to mere true belief. It is consistent with what I have said here that there are respects in which knowledge is superior to mere true belief. Furthermore, I am in fact sympathetic to Williamson's proposal that knowledge is more valuable than mere true belief because '[p]resent knowledge is less vulnerable than mere present true belief to *rational* undermining by future evidence' (Williamson 2000: 79),²² and is thus more stable. Notice too, however, that it is very hard to see how this proposal can afford an adequate response to the enquiry question. Enquirers are often interested in knowledge on occasions on which they have little or no interest in arriving at *stable* true beliefs. Such a proposal is therefore unable to explain why, on these occasions, enquirers remain interested in the acquisition of knowledge.²³ (This, of course, is the problem of accounting for the generality of the supposed intuition that knowledge is more valuable than mere true belief.) The correct response to the persistence of such problems is not to try to identify further respects in which knowledge is superior to mere true belief, but rather to jettison the assumption that the right way to tackle the enquiry question is by tackling the comparative state question. Even if considering enquiry does not shed much light on the

²¹ Bernard Williams reasons in this way in considering why knowledge, and not just true belief, is the aim of enquiry. See Williams 2005: 23–31.

²² This formulation of Williamson's point is not quite accurate, and invites the objection that there are ways in which knowledge can be undermined by future evidence in which mere true belief cannot—for example, by future evidence of which the subject is unaware. (For an alleged example of this, see Harman's (1973) assassination case.) Here is Williamson's strictly accurate formulation: 'If your cognitive faculties are in good order, the probability of your believing *p* tomorrow is greater conditional on your knowing *p* today than on your merely believing *p* truly today (that is, believing *p* truly without knowing *p*)' (2000: 79). The point, pressed by Kvanvig (2003), that there are ways in which knowledge can be undermined by future evidence in which mere true belief cannot, is irrelevant to this claim.

²³ This point is eloquently made by Craig (1990: 7): 'Whether the stabilisation of true beliefs is important or not depends on which beliefs we are considering, and the circumstances of the agent—many beliefs are required for the guidance of single, 'one-off' actions under circumstances which will not recur, and once the particular occasion is past there is no obvious value at all in their persistence. (I might now need a true belief about the time; but that this belief should persist, so that tomorrow I will still know what the time was today, at the moment when I wanted to know it, may be of no interest to me whatever.)'

nature of knowledge, the most promising way of understanding the central role of knowledge in our cognitive lives lies in close attention to the position of an enquirer.

Bibliography

Baehr, J. (2009) 'Is There a Value Problem?'. In: A. Millar, A. Haddock and D. Pritchard, *Epistemic Value*. Oxford: Oxford University Press: 42–59.

BonJour, L. (1980) 'Externalist Theories of Empirical Knowledge'. *Midwest Studies in Philosophy*, 5 (1): 53–74.

Craig, E. (1990) *Knowledge and the State of Nature*. Oxford: Oxford University Press.

Harman (1973) *Thought*. Princeton, New Jersey: Princeton University Press.

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

A. Millar, A. Haddock and D. Pritchard (2009) 'Introduction'. In: A. Millar, A. Haddock and D. Pritchard, *Epistemic Value*. Oxford: Oxford University Press: 1–15.

Plato (2005) *Meno and Other Dialogues*. Translated by R. Waterfield. Oxford: Oxford University Press.

Pritchard, D. (2010) 'Knowledge and Understanding'. In: D. Pritchard, A. Millar and A. Haddock, *The Nature and Value of Knowledge: Three Investigations*. Oxford: Oxford University Press: 3–88.

Schaffer, J. (2007) 'Knowing the Answer'. *Philosophy and Phenomenological Research*, 75 (2): 383–403.

Sosa, E. (2007) *A Virtue Epistemology: Apt Belief and Reflective Knowledge*. Oxford: Clarendon Press.

Soteriou, M. (2013) *The Mind's Construction: The Ontology of Mind and Mental Action*. Oxford: Oxford University Press.

Williams, B. (1973) 'Deciding to Believe'. In: B. Williams, *Problems of the Self*. Cambridge: Cambridge University Press: 136–151.

Williams, B. (2005) *Descartes: The Project of Pure Enquiry*. London: Routledge.

Williams, B. (2006) 'Knowledge and Reasons'. In: B. Williams, *Philosophy as a Humanistic Discipline*. Princeton: Princeton University Press: 47–56.

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