



Action, Knowledge, and Will

John Hyman

<https://doi.org/10.1093/acprof:oso/9780198735779.001.0001>

Published: 2015

Online ISBN: 9780191799754

Print ISBN: 9780198735779

Search in this book

CHAPTER

7 Knowledge as an Ability

John Hyman

<https://doi.org/10.1093/acprof:oso/9780198735779.003.0007> Pages 159–190

Published: May 2015

Abstract

Modern epistemology was born in the shadow of scepticism, but we need to separate the task of explaining what knowledge is from any attempt to prove that it exists. In order to decide whether we have or are capable of having knowledge of some kind—knowledge about the past, the future, good and evil, and so on—we need to ask how knowledge of that kind can be obtained. But knowledge is an ability. So if we want to explain what knowledge is, we need to ask what it is an ability to do, in other words, how it gets exercised or expressed. We need to look forward to how it is applied, rather than backward to how it is acquired, turning the investigation round by 180 degrees.

Keywords: [knowledge](#), [belief](#), [ability](#), [scepticism](#), [reason](#), [fact](#)

Subject: [Metaphysics](#)

Collection: [Oxford Scholarship Online](#)

7.1 Introduction

The question, ‘What is individual factual knowledge?’ is quite narrowly focused in three respects. First, there is communal knowledge as well as individual knowledge. Second, a distinction is often drawn between factual knowledge and practical knowledge, or between knowledge that something is the case and knowledge how to do something. Third, as well as asking what knowledge is, it is also possible to ask whether or how knowledge of one kind or another can be acquired. Before attempting to answer the question, I shall dwell briefly on each of these three points.

First, we can distinguish between the possession of knowledge by an individual and the possession of knowledge by a community or group. For example, if we speak or enquire about the state of knowledge in a particular field of biology or history, we are not concerned with what anyone in particular knows about, say, the genetics of fruit flies or the Hundred Years War, but with what the scientific or academic *community* knows. We commonly describe the state of scientific or historical research with the impersonal construction

'It is known that ...', but although there is a close connection between individual and communal knowledge, this does not simply mean 'Someone or other knows that ...'. If we imagine that it does, we are ignoring the role of documents, archives, and libraries in the economy of knowledge. For the invention of writing means that each of us can cease to remember without all of us, collectively, ceasing to know. That is why Thamus, the mythical king of Egypt whom Socrates talks about in the *Phaedrus*, was right to say that writing is a recipe for reminder, and not for memory, but wrong to infer that it cannot extend knowledge, but only a semblance of it.

p. 160 Second, there is sometimes said to be a distinction between practical knowledge and factual knowledge, a distinction that Ryle introduced under the rubric 'knowing how and knowing that'. As many commentators have pointed out, these labels are misleading, since 'that' is not the only pronoun we combine with 'know' to attribute factual knowledge to someone. For example, the statements that Tom knows whether it will rain tomorrow, when and where Alexander was born, what time the match begins, and how sponges reproduce all attribute items of factual knowledge to Tom. The moot questions are, first, whether knowing how to do something is the same as being able to do it or being skilled at doing it, as Ryle is widely thought to have claimed; and second, whether knowing how to do something and knowing that something is the case are essentially different kinds of knowledge. These questions have been studied intensively in recent years, but they require more patient treatment than I can give them here, and I shall not address them.¹

Third, and most importantly for present purposes, the theory of knowledge is concerned not only with the question of what knowledge is, but also with the questions of whether and how knowledge of various kinds can be acquired. This quick formulation embraces many difficult and contentious matters, but the point I should like to emphasize is that the question of what knowledge is and the question of how it can be acquired are quite different, even though they are connected. A recent article by a distinguished British philosopher includes the extraordinary claim that we can 'explain what it is to know that A by identifying different possible means of knowing [i.e. coming to know] it'. At a minimum, he says, 'to understand what it is to know that A, all one needs is an open-ended list of means of knowing that A, perhaps together with some indication of whether some means of knowing that A are more basic than others'.² But this cannot be right.

Take the fact that Tunis is west of Pisa. An open-ended list of the means of knowing this might be: *being told by someone who knows; consulting a map; measuring the longitude of each city with a sextant; and so on*. How could anyone imagine that this list explains *what it is* to know that Tunis is west of Pisa, whether or not it is supplemented by an indication of which of these means, if any, is more basic than the others? Surely it is one thing to explain how knowledge can or cannot be acquired, and hence how knowledge-claims can be tested, and quite another to explain what knowledge is, just as it is one thing to explain how a right or a duty can be acquired, and hence how it is possible to decide whether someone has a certain right or duty, and quite another to explain what a right or a duty is.

p. 161 Consider the ownership of a house. An open-ended list of the means of acquiring a house might be: *buying it, either by private treaty or in an auction; inheriting it; occupying it for a sufficient period, if it has been abandoned; and so on*. But the list does not begin to explain what owning a house is. On the contrary, buying and inheriting are, by definition, ways of acquiring property, ways in which title can be transferred, and we cannot hope to understand what it is to acquire something unless we already know what it is to own it, or what it is to transfer title unless we already know what it is to have title. The same is true of knowledge. We cannot explain what it is to know something by explaining how we can learn it. On the contrary, we cannot understand what it is to learn something unless we already know what it is to know it. A list of ways of learning is no more a theory of knowledge than a list of ways of forgetting.

Why do philosophers confuse or amalgamate the question of what knowledge is and the question of how knowledge can be acquired? There seem to be two main reasons. First, the principal aim of modern epistemology was to refute the sceptical doctrine that we do not in fact know various things that we commonly take ourselves to know. But the claim that we *do* know these things—about causes, about the past or the future, about other people’s thoughts and feelings, and so on—can only be vindicated by showing that there is a satisfactory answer to the question of *how* we know them.* Consequently, when, in the 1960s, defining knowledge assumed the same importance in philosophy as refuting scepticism, there was a tendency for philosophers to fold the question about how knowledge can be acquired into the question of what it is. But it is a mistake to imagine that a definition of knowledge can refute scepticism. If it could, it would be tantamount to an ontological proof of the existence of knowledge, which is absurd. The only concept whose definition proves that it has instances is the concept of definition itself.

p. 162 This is not to say that there are no limits on how someone can learn something. Clearly there are such limits. That is why we can test someone’s claim to know something by asking how he found out, or decide that someone does not know something because he wasn’t in the right place at the right time, because he reasoned from a mistaken premise, or because the man who informed him didn’t know it himself. Equally, we can conclude that someone does not own something because he never had the opportunity to acquire it, because the man he paid did not own it in the first place, or because the document was not properly drawn up. But to repeat, it does not follow, and it is not true, that understanding how property or knowledge of various kinds can be acquired is the same as understanding what it is.

The second main reason why philosophers have confused the question of what knowledge is and the question of how knowledge can be acquired is that it is still commonly assumed that knowledge is a species of belief. But whereas there are limits on how someone can acquire knowledge, there are no similar limits on how someone can acquire belief. For example, one can acquire belief but not knowledge about the outcome of tomorrow’s battle by examining the entrails of a sheep. So if we assume that knowledge is a species of belief, it is tempting to think that deciding whether A’s belief that *p* qualifies as knowledge is a matter of deciding whether it was acquired in the right way. And so the question of how knowledge can be acquired is commonly regarded as part of the larger question of what it is.

For these reasons, and perhaps others, there is a tendency in epistemology for mission creep. We start out wanting to say what knowledge is, but we quickly get embroiled in the question of how it can be acquired. I shall focus exclusively on the question of what knowledge is in this chapter, and on the question of why we value knowledge in the next. Questions about the limits and sources of knowledge have bulked larger in the last four centuries, but they are best devolved to the branches of philosophy that study different kinds of knowledge—philosophy of science, philosophy of mathematics, moral philosophy, and so on—and I shall mostly ignore them here. There is no reason a priori to expect that the ways of acquiring knowledge, or its limits, can be defined informatively in a topic-neutral way, and there is no encouragement to believe this from the record of general epistemology during the last fifty years.

7.2 How knowledge gets expressed

I shall defend the theory that individual factual knowledge (from now on, I shall refer simply to knowledge) is the ability to be guided by the facts, in other words, to take the facts into consideration or account, in what we think or feel or do.

p. 163 Most philosophers still think that knowledge is, in Ryle’s sardonic phrase, an ‘*élite suburb*’ of belief.³ But it is doubtful whether this can be right.

First, the simple fact that philosophers have failed to explain what kind of belief knowledge is, despite (one would guess) more hours of work being devoted to the problem during the past sixty years than were spent philosophizing in the whole history of ancient Greece, suggests that it is not any kind of belief at all.

Second, knowledge does not seem to have the same kind of object as belief. A fact, in the relevant sense of the word, is simply a truth; the facts about something are simply the truth about it, and knowledge of a fact is knowledge of a truth. But what is a truth? One might think *either* that a truth is simply a true proposition *or* that it is the truth *of* a proposition, and these are evidently not the same. For the truth of a proposition is not a true proposition, any more than the beauty of a place is a beautiful place, or the goodness of a man is a good man.⁴

There are three reasons for preferring the view that a fact is the truth of a proposition:

- Facts and propositions have different counterfactual conditions for their existence. For example, if France had won the 2006 World Cup instead of Italy, the proposition that Italy won would still have existed, but the fact that Italy won would not. Hence the fact cannot be identical with the proposition.⁵ (It cannot be identical with the proposition that the proposition that Italy won is true for the same reason.)
- Facts are discovered, learned, known, communicated, forgotten, and ignored. But the discovery (etc.) of a fact is the discovery (etc.) of the truth of a proposition, and not the discovery (etc.) of a proposition. For example, Eratosthenes is credited with discovering the fact that the Earth is round. But he did not discover the proposition, the hypothesis, that the Earth is round. That had been in circulation for some time. He discovered the truth of the proposition that the Earth is round.
- Facts have various properties relating to the effect of learning or recalling them, such as being surprising, shocking, or sad. For example, it is a sad fact that more than fifty percent of cases of sexual abuse of children in the US are by family members. But to be saddened by a fact is to be saddened by the truth of a proposition, even if there are mere propositions that are sad as well.

p. 164 If these points are right, what a person knows, when she knows that *p*, is the truth of the proposition that *p*. But what a person believes, when she believes that *p*, is the proposition that *p* itself. Hence, the object of knowledge is not the same as the object of belief. It might be objected that if this were true, statements like 'I always believed Philby was a traitor; now I know it' would be false. But in fact this does not follow. For 'it' can be used simply to avoid the repetition of a clause or phrase, as in 'Tom is longing for the war to end and Lucy is praying for it'. 'I always believed Philby was a traitor; now I know it' is true if, and only if, I always believed Philby was a traitor, and now I know Philby was a traitor, i.e. if, and only if, I now know the truth of the proposition that Philby was a traitor, and this is a proposition I always believed.

The third reason for denying that knowledge is an 'élite suburb' of belief, in other words, a species of belief that has a special status or value, is the most important for our purposes. As Ryle himself points out, there are several ways in which knowledge seems like an ability or skill, while belief seems like a tendency or disposition.⁶ Among them: belief, like love or trust, can be foolish, passionate, obstinate, fanatical, or whole-hearted, whereas knowledge, like the ability to solve quadratic equations or cast a fly, cannot be any of these things. We can ask or urge someone to believe or not believe something, but we cannot ask or urge someone to know or not know something, but only to remember or forget something, or to find something out. We ask *why* rather than *how* someone believes something, wanting a justification; whereas we ask *how* rather than *why* someone knows something, wanting to be told the means by which the knowledge was acquired. ('How *can* you believe ...' and 'How are you *able* to believe ...' are of course perfectly in order.)

As these remarks suggest, the main alternative to the 'élite suburb' picture of knowledge is that it is an ability. This general idea originates with Plato, but in the twentieth century it was particularly associated

with Wittgenstein and Ryle. Wittgenstein claims that ‘the grammar of the word “knows” is evidently closely related to that of “can”, “is able to”, and Ryle asserts that “know” is of the same family as skill words’.⁷

p. 165 The idea that knowledge is an ability or skill is important, because it transforms the task of defining knowledge. For instead of asking what we need to add to belief to get knowledge, we are forced to ask how knowledge gets exercised or expressed, since this is invariably how abilities and skills are defined. What is the ability to multiply, and how does it differ from the ability to divide? What is the ability to play tennis, and how does it differ from the ability to play squash? The whole answer is given by explaining what multiplying or playing tennis is, and how it differs from dividing or playing squash.

But Wittgenstein’s and Ryle’s remarks are studiously imprecise, and neither of them pursues the idea that knowledge is an ability or skill far enough to make it seem convincing, because they do not attempt to explain what someone who knows something—e.g. that Tunis is west of Pisa—is able to do that someone who does not know it is unable to do, or what skill or skill-like trait he has that the other lacks. Nor did their followers succeed in developing the idea that knowledge is an ability into a convincing theory of knowledge, partly because of a tendency to associate knowledge too closely with the use of language to express knowledge or to communicate information.⁸

For example, White argues that knowledge is the ability to produce the correct answer to a possible question.⁹ But on the most natural interpretation, to produce the correct answer to a question is to state it, and, in this sense, it is possible to know that *p* without being able to produce the correct answer to the question whether *p*. For example, there are circumstances in which we might want to say that a dog knows that it is time for a walk, but it cannot say so. White explicitly acknowledges that ‘there is no reason why young children and animals should not be said to know many things [e.g.] that it is time for a walk’, and he denies that having the ability to produce the right answer to a possible question ‘implies [...] manifesting it in any verbal way’, but he does not explain how, exactly, the exercise of this ability is to be defined.¹⁰

Perhaps what White had in mind is that to produce the correct answer to a possible question is merely to enable someone to state it; for a dog certainly can make its owner aware of the fact that it is time for a walk, and thereby enable him or her to state it. But this definition would be too liberal, because in this sense, a piece of litmus paper, which cannot know anything, can produce the correct answer to the question whether a solution is acidic. Alternatively, it may be that to produce a correct answer is to *intentionally* enable someone to state it.¹¹ But this would not be liberal enough. For even if it is plausible that a dog intends, by leaping towards the door, to get a walk underway, or to get someone to get a walk underway, it is surely *not* plausible that it intends to enable anyone to state that it is time for a walk. White’s definition of knowledge therefore seems to be unhappily poised between tautology and falsehood: to know whether *p* is, p. 166 tautologically, to know the answer to the question whether *p*, but it is not to be able to answer the question whether *p*.

It might be objected that what we are inclined to say about animals is not a convincing reason for rejecting a theory of knowledge. Animals without language have an attenuated awareness of facts, and we should philosophize about knowledge with competent speakers in mind, and allow our conception of the cognitive powers of animals to be decided by our epistemology, and not vice versa. This objection has some force, despite the fact, noted above, that White’s theory of knowledge was meant to be consistent with attributing knowledge to animals. But in any case, knowledge has a far more extensive role in our lives than the ability to provide answers to questions does, and a satisfactory definition will have to reflect this. Bernard Williams points out that the traditional definition of knowledge as justified true belief is too narrowly tailored to what he calls the *examiner’s situation*, where ‘informed questioners are concerned with someone’s credentials with regard to a piece of knowledge’, which is far from typical in practice.¹² White’s theory of knowledge is tailored to what might be called the *expert witness situation*, where uninformed questioners are concerned to establish whether *p* from someone who is presumed to know. Of course, both of these situations arise, but it

is difficult to see why either of them should be thought capable of providing the key to explaining what knowledge is.

Kenny proposes a different definition of knowledge as an ability, which avoids the pitfall of thinking that knowledge is primarily expressed in speech. Knowledge, Kenny says, is 'an ability of a unique kind', but he doubts whether it is possible to say precisely what it is an ability to do:

[T]here is no simple way of specifying how knowledge gets expressed in behaviour, and indeed some pieces of knowledge may never affect behaviour at all. The most that we can say is that to know is to have the ability to modify one's behaviour in indefinite ways relevant to the pursuit of one's goals.¹³

Two things are right about this definition, and two things are wrong with it. The first point in its favour is that if knowledge is an ability, then defining knowledge will, as Kenny implies, be a matter of specifying how it is exercised, manifested, or expressed. This is a perfectly general point about powers of every kind. Dispositions, abilities, and liabilities are all powers; and, as I pointed out earlier, powers are always defined in terms of their exercise, in other words, what they are powers to do. Second, Kenny is right to emphasize the peculiar flexibility of knowledge, in other words, that it can be expressed in indefinitely many ways. For example, Tom's awareness of the fact that the rouble has collapsed can be expressed in his ↪ booking a holiday in Russia, buying shares in Gazprom, or sending dollars to a friend in Moscow.

On the other hand, Kenny seems to equate the question of how knowledge gets expressed in behaviour and the question of what knowledge is an ability to do. But this is a mistake. For if knowledge is an ability, it can be exercised in thought and feeling no less than in behaviour. Performing a calculation in one's head, making an inference, conceiving a desire or forming an intention, believing, doubting, hoping, fearing: all these things can express knowledge just as much as behaviour can. For example, Tom's awareness of the fact that the rouble has collapsed can be expressed in his *forming an intention* to take a holiday in Russia, whether or not he takes any practical steps to do so; or *hoping* that his friend in Moscow has his savings in dollars, instead of supplementing them. Furthermore, while performing a calculation in one's head is something we do in pursuit of our goals, at least for the most part, believing, doubting, hoping, and fearing are not. We do not do these things intentionally or on purpose, and so the scope of knowledge, the extent to which it affects human life, is much larger than the scope of purposes or aims. If we recall the distinction between doing things for reasons and doing things with intentions, knowledge matches reason in its scope (see above, 6.2).

The second weakness in Kenny's account of knowledge is that he fails to explain what unites the heterogeneous variety of things that can express the knowledge that p as opposed to the knowledge that q . To know, he says, is to have the ability to modify one's behaviour in indefinite ways relevant to the pursuit of one's goals. But if this is the most we can say about how knowledge is exercised, the difference between knowing one thing and knowing another is bound to remain obscure.

But is this a weakness? Naturally, the answer depends on what form an explanation of the difference between knowing that p and knowing that q is supposed to take. But if what is being sought is a formula connecting specific kinds of thought and action with the possession of specific pieces of knowledge, it may be objected that this is unfeasible. There is a trivial sense in which the various modifications to a person's thought or conduct that express her knowledge that p have something in common, for they are all informed by her knowledge that p . Hence, if Tom's knowledge that the rouble has collapsed is an ability, there can be nothing wrong with saying that its exercise will consist in thought and conduct informed by the knowledge that the rouble has collapsed. But it is not very illuminating. And it becomes positively misleading if it makes us think that we can say specifically what acts Tom is able to perform if and only if Tom knows that the ↪ rouble has collapsed. For if A knows that p , it does not follow that A is able to do *all* the things that

could be informed by the knowledge that *p*. For example, if Tom knows that the rouble has collapsed, it does not follow that he is able to arrange a loan from the World Bank. And if A is able to do *at least one* of the things that can be informed by the knowledge that *p*, it does not follow that A knows that *p*. For example, an investment in Gazprom can be informed by the knowledge that the rouble has collapsed, but there is no need for Tom to know that the rouble has collapsed in order to buy shares.

The reply to this objection is that the difference between knowing one thing and knowing another *cannot* be explained by saying specifically what thoughts and acts can express the knowledge that *p*. But it does not follow that we cannot explain what knowing that *p* is an ability to do. Compare the question ‘What is enthusiasm?’ Enthusiasm isn’t a tendency to do specific things. For instance, enthusiastic sportsmen do not invariably hop from foot to foot and punch the air. But it is, in part, a tendency to do a range of things in an enthusiastic manner. We can call it an *adverbial* tendency, if the phrase is acceptable: a tendency to do things enthusiastically. Similarly, unpunctuality is a tendency to do things such as arriving for meetings, submitting tax returns, and answering correspondence later than the appointed time. But there is no other way of classifying unpunctual acts except by saying that they are done late, which is about how they are related to clocks, not about which acts they are. So being unpunctual can also be described as an adverbial tendency, although in this case the adverb is an adjunct of time, rather than of manner.

The lesson of these examples is that while enthusiasm and unpunctuality are not tendencies to perform specific acts, this does not prevent us from explaining what they are tendencies to do. But is it possible to say what knowing that *p* is an ability to do in a similar way, by specifying an adverbial clause which we can introduce into a sentence to indicate that the act or thought it mentions was informed by the agent’s or thinker’s knowledge of a particular fact?

In fact this is a relatively simple thing to do. For we can convey exactly the same information by saying either that an act was informed by the agent’s knowledge that *p*, or that the agent did the act in view of, or in the light of, the fact that *p*. And a fact in view of which, or in the light of which, an agent does an act is a fact *because of which* she does it, not in the general sense of a reason why, but in the particular sense of a fact she was guided by, or took into consideration or account. For example, ‘James went to church \downarrow *because it would please his mother*’, as this would normally be meant and understood, implies that in going to church James was guided by the fact that doing so would please his mother. The adverbial clause ‘because it would please his mother’ is not of course an adjunct of time or manner, but what is called an adjunct of contingency, like ‘He called *in order to warn her*’, or ‘He bought a magazine *in case she was delayed*’.

So the idea that knowledge is an ability not only transforms the task of defining knowledge, it makes it tractable. For while philosophers have been unable to say what kind of belief knowledge is, or what needs to be added to belief to get knowledge, short of knowledge itself, it is quite easy to say how knowledge gets expressed. To take another example, my knowledge of the fact that Hampstead is on the Northern Line gets expressed whenever I am guided by the fact that Hampstead is on the Northern Line, in what I think or do. This happens when I head for the Northern Line at King’s Cross to get home, or when I merely deduce from the fact that Balham is on the Northern Line that one can travel from Hampstead to Balham without changing lines.

Talk of being guided by a fact is metaphorical of course, but the phrase is familiar to English-speakers—the OED records it in the fifteenth century—and it is not hard to explain. To be guided by a fact is to take it into consideration or account.¹⁴ So the metaphor is dispensable, but it is helpful, because it draws attention to the similarity between being aware of facts and being aware of things. Think of a cat stalking a bird. The cat expresses its awareness of the bird when it modifies its behaviour in response to way the bird modifies *its* behaviour. The bird hops this way, the cat turns this way; the bird flutters across the courtyard, the cat advances a few paces; and so on. The cat’s movements are responsive to, are guided by, the bird. Or think of a traveller following a guide. The guide takes the left fork, so the traveller takes the left fork; the guide

pauses, so the traveller pauses; and so on. Being guided by facts is similar. Whether one is guided by facts or by things, one is responsive to what one is guided by. But, as Wittgenstein pointed out, this is not like a train being guided by the rails, because the way knowledge gets expressed depends on one's aims or purposes, by one's desires and values, and because being guided by the facts is not passive or constraining, any more than reading is, although the reader's thought (if she is reading silently) or speech (if she is reading out loud) is guided by the words she sees on the page.

p. 170 Hence, if we conceive of knowledge as an ability, and approach the problem of defining knowledge by asking how it gets expressed in our mental lives and in our conduct, instead of asking how it can be certified or acquired, we can define it as the ability to be guided by the facts. But the word 'fact' must be understood in the broadest possible sense, so that does not involve a contrast between facts and values, past and future, contingent and necessary, or a posteriori and a priori. For example, if Sheldon believes that the circumference of a dime is more than an inch and a half because its diameter is more than half an inch and π is greater than three, the view defended here implies that he knows that π is greater than three in just the same way as it implies that he knows that the diameter of a dime is more than half an inch. Even if mathematical truths are rules, as Wittgenstein argued, it is still a fact that π is greater than three, in this broad sense of 'fact', just as it is a fact that a driver must stop at a red light.

How does this definition compare with Kenny's and White's? First, it is flexible, in the sense that it allows knowledge to be expressed in indefinitely many ways, and it is precise, in the sense that it reveals precisely what the difference is between knowing one fact and knowing another. For one knows that p if and only if one's thought and conduct can be guided by the fact that p , and one knows that q if and only if one's thought and conduct can be guided by the fact that q . White's definition is precise but inflexible, whereas Kenny's is flexible but imprecise. The definition proposed here has the advantage of combining both virtues.

Second, it also subsumes both Kenny's and White's definitions: Kenny's, because the ability to be guided by the facts in general includes the ability to modify one's behaviour in the pursuit of one's goals, and White's, because the ability to be guided by the fact that p in particular includes, in the case of competent and uninhibited communicators, the ability to give the correct answer to the question whether p . Kenny's definition is, in effect, the result we obtain if we apply the idea of being guided by the facts to the particular case of goal-directed behaviour, as opposed to believing, wanting, hoping, fearing, or any other kind of thought or feeling that knowledge can inform. As for White's definition, we saw earlier that giving the right answer to a question is only one of the ways in which knowledge can be expressed, but of course it need not be an expression of knowledge at all. For it would be absurd to suggest that someone knows that p if he happens to give the right answer to the question whether p for whatever reason, for instance, because it is the answer he thinks the person asking will be glad to hear. What
p. 171 competent and uninhibited communicators who know whether p can do that others cannot do is give the right answer to the question whether p because it is the right answer—that is, either give an affirmative answer because p or give a negative answer because not- p . So White's definition is the result we obtain if we apply the idea of being guided by the facts to the particular case of answering the question whether or not something is the case.

In sum, if we pursue Wittgenstein's idea that the verb 'knows' is closely related to 'can' and 'is able to', or Ryle's idea that 'know' is a skill word, we can define knowledge as the ability to be guided by the facts. This explains why competent and uninhibited communicators who know that p can give the right answer to the question whether p , and more generally why knowledge allows us to modify our behaviour in the pursuit of our goals. But our beliefs and desires and doubts and hopes and fears can be guided by facts, no less than our deeds. So defining knowledge in this way also reflects the fact that knowledge does not only get expressed in our conduct, but also in the elements of our mental lives from which it flows.

7.3 Objections and replies

Philosophers have objected to the claim that knowledge of a fact is the ability to be guided by it on various grounds, which may be summarized as follows: knowing a fact is not (1) necessary, or (2) sufficient, for being able to be guided by it; (3) there is no such ability as the ability to be guided by a fact; (4) knowledge is not an ability but a state; (5) a person's knowledge of a fact explains her ability to be guided by it, so the knowledge and the ability cannot be identical. I shall comment on these objections in turn.

7.3.1 Knowing a fact is not necessary for being able to be guided by it

I argued in 6.4 that one cannot be guided by a fact one does not know. Readers may find it helpful to review the argument there now. Instead of recapitulating it in detail, I shall add some comments about the relationship between explanations of thought and action that refer to knowledge, and ones that refer to belief.

p. 172 The argument in 6.4 made use of several cases of true belief without knowledge, including ones that were invented when the project of defining \hookrightarrow knowledge in terms of belief was in full swing. For the same examples that were used then to show that knowledge cannot be equated with justified true belief, or justified true belief with no false lemmas, or true belief caused by the right event, can be used to show that if one has the justified true belief that p , or the justified true belief with no false lemmas, or the true belief caused by the right event, it does not follow that one can be guided by the fact that p . Something one believes but does not know can be the ground on which one does something, but one cannot be guided by a fact one does not know, any more than one can follow a guide one cannot see. In fact, the relationship between knowledge and true belief is like the relationship between perception and veridical hallucination. For if a traveller sees the guide taking the left fork and follows him, then he is guided by the guide; whereas if he cannot see the guide, but he hallucinates him taking the left fork, and takes it himself for that reason, he is not guided by the guide. A fact that guides one is not merely a ground one acts on, and one can only be guided by a fact one knows.*

When philosophers assumed that knowledge is a species of belief, they naturally also assumed that explaining thought and action in terms of knowledge is essentially the same as explaining it in terms of belief. For whatever part belief plays in causing thought and action, it surely plays the same part whether or not it has the characteristics that were thought to make it qualify as knowledge, such as having the right kind of justification or the right kind of cause. But it turns out that they were wrong. Explanations involving knowledge are quite different from ones involving belief, because an agent can only be guided by the facts he knows.

p. 173 Of course, it can seem to someone who merely believes something to be the case just as if he knows it to be the case, and it can be impossible to tell belief from knowledge in another. For believing something tends to influence thought and behaviour in the same way as knowing it for a fact. In some cases, the influence of a belief can be diminished by the subject's knowledge that his evidence is weak or that he is affected by prejudice. In this sort of case, the extent to which he is influenced by the belief will depend on the extent to which he controls what he thinks and does by \hookrightarrow reminding himself of such facts. But even so, if someone would do X or believe Y or feel Z if he knew that p , then he will be disposed to do X or believe Y or feel Z if he believes that p . It is sometimes claimed that to believe that p is to be disposed to act (think, feel) 'as if p ' or 'as if p were true'. But this must mean that to believe that p is to be disposed to act (think, feel) *as one would if one knew that p , or as one would if one were guided by the fact that p* , since the mere fact that something is the case need not have any influence on us at all. How would I be disposed to act if there were an odd number of stars in the galaxy? Exactly as I *am* disposed to act. Yet I do not believe this, or believe the opposite.*

Hence, one can do or think or feel something that the facts would justify one in doing or thinking or feeling if one knew them, without being guided by the facts. And this is exactly what conceiving of knowledge as an ability rather than a species of belief leads us to expect. For many performances that normally display an ability or skill can also be done occasionally by luck. For instance, a bad speller might spell ‘accommodate’ correctly from time to time, and a novice fly fisherman might cast a perfect fly. If knowledge is the ability to be guided by the facts, this is exactly what happens when someone who does not know a certain fact acts in a way that seems to express knowledge of it, because he happens to have the right opinion: e.g. when a guide who believes that the left fork is the road to Larissa and happens to be right, but has never been there and does not really know, points a traveller in the right direction, or gives him the right advice.

p. 174 In sum, belief mimics knowledge, so it can be difficult to tell them apart. But the fact that it can be difficult to distinguish between belief and knowledge in practice is not an obstacle to explaining the difference between them in theory, and the root of the difference is that one cannot be guided by a fact one does not know. Of course it does not follow that an act that expresses belief and not knowledge is done aimlessly, or for no reason. For example, if Sybil feeds James oysters because she believes they \hookrightarrow are aphrodisiac, it is fairly obvious what her aim is, and her reason for doing it—that is, her *ground*—is the content of her belief. But she is not guided by the facts about the effect of eating oysters, and as the analogy between false belief and veridical hallucination suggests, and the argument in 6.4 confirms in detail, this would remain the case if her belief happened to be true.

Hence, explanations that involve knowledge, such as ‘He took the left fork because he knew that it led to Larissa’, are quite different from ones that involve belief, such as ‘He took the left fork because he believed that it led to Larissa’. The *ground* on which the man is said to have acted is the same in both cases, but the knowledge-involving explanation refers to a feature of his situation he was aware of and took into account, whereas the belief-involving explanation merely refers to his state of mind.

7.3.2 Knowing a fact is not sufficient for being able to be guided by it

It might be objected that knowledge cannot be the ability to be guided by the facts, because knowing a fact is not sufficient for being able to be guided by it, either on the grounds that some facts it is possible to know are too recondite to guide anyone in doing anything at all, or on the grounds that knowledge requires less intellectual sophistication than being guided by facts.

The first argument does not merit extensive comment. Every fact, however insignificant or obscure, is potentially someone’s reason for doing or for not doing something, for believing or wanting something, or for hoping or fearing something. It would be tedious to illustrate this with examples, and besides, it is easily proved. For at a minimum, the fact that p is self-evidently a reason for believing that the proposition that p is true, or that the proposition that not- p is false, or that if p implies q , then q . The second argument may seem more convincing. For example, I mentioned earlier that there are circumstances in which we might want to say that a dog knows it is time for a walk. But it could be argued that since dogs cannot give reasons for their actions, they cannot act for reasons, and so the fact that it is time for a walk cannot be a dog’s reason for doing anything. If that is right, then knowledge requires less intellectual sophistication than the ability to be guided by reasons, and so they cannot be identified or equated.

p. 175 Evidently, dogs cannot give reasons for their actions, at least if giving a reason means stating it, because they cannot speak. But the objection is a weak one, for two reasons. First, it is not obvious that animals cannot do \hookrightarrow something for a reason unless they can give reasons for doing something, and, as a matter of fact, there are well-known examples of animals doing things which seem to depend on reasoning, such as the story about Chrysippus’s hunting dog. In hot pursuit, it reached a place where the path branched into

three; it sniffed one path, but didn't detect its quarry's scent; it sniffed the second, but drew a blank again; and then it took the third path *without* sniffing. Annas and Barnes comment on the story as follows:

How is its action to be explained? Chrysippus argues that it must have engaged in some simple reasoning: it said to itself, in effect: 'Either A or B or C; but not A, and not B: therefore C.' [...] It is the fact that the dog selects the third track without further experiment which requires explanation —and the explanation which Chrysippus offers is highly plausible.¹⁵

If Annas and Barnes are right, and Chrysippus's explanation is plausible, then the syllogism sets out the dog's reason for taking the third path, despite the fact that the dog could not have stated this or any other reason, for this or any other act.

Annas and Barnes are unusual for philosophers. For although the ethological literature is widely believed to provide copious evidence of reasoning by dogs, primates, and some birds, philosophers have tended to be more sceptical than non-philosophers about the intellectual abilities of animals without language. But — and this is the second reason why the objection is weak — suppose there are compelling arguments to show that the evidence from studies of animal behaviour has been misinterpreted, and Chrysippus's dog only engaged in some *quasi*-reasoning, and only *quasi*-acted for the reason set out in the syllogism. Still, the objection will not have been made out if these arguments (or other ones) also prove that it only *quasi*-knew that its quarry had taken the third path. Certainly, the terms we use to describe the exercise of rational powers by human beings apply only in an attenuated or analogical sense to many animals. Which animals these are is a matter of dispute, but the general proposition is not. But the same is true where cognition is concerned. So if the objection is to bite, it has to be shown that when guidance by facts has become attrited, knowledge remains intact. But the arguments which have persuaded philosophers that only human beings can act for reasons do not show this. They do not drive a wedge between knowledge and the ability to be guided by the facts. On the contrary. Davidson, for example, argues that only members of a speech community can act for reasons, because a creature cannot act for reasons if it cannot have beliefs, and only a member of a speech community can have beliefs.¹⁶ But a creature cannot be incapable of having beliefs and yet be capable of knowledge. (This is true whether or not knowledge implies belief, simply because we are fallible.)

p. 176

Davidson's reason for claiming that only a member of a speech community can have beliefs is that an animal cannot have a belief 'unless it understands the possibility of being mistaken'; but this 'requires grasping the contrast between truth and error', a contrast Davidson claims can only emerge within a speech community.¹⁷ But it is not self-evident that an animal cannot have a belief unless it understands the possibility of being mistaken. A dog can be hungry or lustful or want to bury a bone, whether or not it can understand the possibility of being frustrated or disappointed, and whether or not it can grasp the contrast between success and failure, and if the analogy between belief and desire is misjudged, Davidson does not say why this is so. But the cogency of Davidson's argument is beside the point. What matters for present purposes is that his sceptical conclusion about animals does not cast doubt on the claim that knowledge that *p* is sufficient for being able to be guided by the fact that *p*. Indeed, far from weakening the link between cognition and reason, Davidson reinforces it, for he also claims that attributing beliefs and desires to a creature only makes sense to the extent that doing so reveals a broadly rational pattern in its behaviour.

The argument here neither supports nor contradicts the sceptical view about animal intelligence which Davidson and, for different reasons, Aristotle, Aquinas, and Descartes all recommend. But it does support the proposition, on which their arguments converge, that knowledge and the ability to be guided by facts are either co-present or co-absent. If the theory of knowledge defended here is right, the reason for this could not be simpler: they are identical.

7.3.3 There is no such ability as the ability to be guided by a fact

Hacker argues that knowledge of a fact cannot be the ability to be guided by it because there is no such ability. Abilities, he claims, correspond to ‘act descriptions’:

p. 177

‘To φ ’, one might say, is the general form of an act description. But ‘to φ for the reason that p ’ is no more a different act description from ‘to φ for the reason that q ’ than ‘to φ for A’s sake’ is a different act description from ‘to φ for B’s sake’.¹⁸

For example, ‘to marry for the sake of money’ and ‘to marry for the sake of love’ are not different act descriptions, they are the same act description, namely ‘to marry’, combined with different adverbial clauses. Hence, there is no such ability as the ability to marry for money, or the ability to marry for love: there is simply the ability to marry, which one can exercise for different reasons, or in the light of different facts.

This argument has three main weaknesses. First, Hacker does not say what verb phrases he regards as admissible substitutions for the variable φ in an ‘act description’; second, he does not explain how we are to decide whether or not two phrases are the same ‘act description’, and correspond to the same ability; and third, the comparison between doing an act for one reason rather than another and doing an act for A’s sake rather than B’s sake does not support his case. I shall comment on these points in turn.

(A)

Abilities are generally attributed in English with the auxiliary verb ‘can’ or the auxiliary construction ‘is able to’, followed in both cases by an infinitive verb or verb phrase. But the verb or verb phrase need not signify an act in the sense of a deed. Thus, ‘walk’, ‘run’, ‘play tennis’, and ‘play squash’ signify activities as opposed to acts, while ‘fall asleep’ and ‘stay awake’ signify neither. Furthermore, a verb phrase used to attribute an ability will commonly include an adverbial phrase as a constituent. For example, a multi-talented man may be able to run a mile *in four minutes*, fall asleep *on a crowded train*, remain calm *under fire*, and sing *in tune*. Hence, one cannot argue that there is no such ability as the ability to do an act for a particular reason on the ground that ‘for the reason that p ’ is an adverbial phrase. But Hacker does not say which adverbial phrases can be used to specify an ‘act description’, and which cannot.

(B)

How are we to decide whether the ability to φ and the ability to ψ are identical or distinct? For example, is the ability to run a mile in four minutes a different ability from the ability to run a mile in five minutes? Or are they the same ability, i.e. the ability to run a mile, which can be exercised at different speeds? What about the ability to fall asleep on a crowded train? Does this specific ability exist, or is there just the ability to fall asleep, which can be exercised in different places?

p. 178

We can assume that for all φ and for all ψ , the ability to φ and the ability to ψ are different abilities if it is possible to have one without having the other. \hookrightarrow It follows that the ability to run a mile in four minutes and the ability to run a mile in five minutes are different abilities, since it is possible to have the second without having the first, even though it is not possible to have the first without having the second. Again, the ability to fall asleep on a crowded train cannot be the same as the ability to fall asleep on a crowded bus, because someone could have one, in this case either one, without having the other. (The different noises and rhythms could be the reason, or just a hatred of buses and a love of trains.)

Similarly, if p and q are different reasons, the ability to φ for the reason that p cannot be the same as the ability to φ for the reason that q , regardless of whether the theory of knowledge defended here is right,

because it is always possible that someone should be able to do a certain act for one reason but not for another. For instance, it is possible that a sports fan should be able to stay awake until three in the morning because a match is being broadcast on TV, but not because she needs to finish a chapter of *Syntactic Structures*, or that a man should be able to marry for the sake of money and unable to marry for the sake of love, because he is unable to love.

This is not to say that the abilities to marry for money and to marry for love are unrelated. It is obvious that they *are* related, since someone who marries for money and someone who marries for love both marry, just as someone who runs a mile in five minutes and someone who runs a mile in four minutes both run a mile. But this is not because marrying for money (or running a mile in five minutes) and marrying for love (or running a mile in four minutes) are the same act, and correspond to the same ability. It is because they are both species of the generic act of marrying (or running a mile), and the corresponding abilities are related in the same way. Compare dispositions. Tin melts at 232°C, silver at 961°C. Are 'to melt at 232°C' and 'to melt at 961°C' different 'act descriptions', in Hacker's terminology, or are they the same 'act description' with different adverbial clauses attached? I would not venture to say. But the disposition to melt at 232°C and the disposition to melt at 961°C are evidently different dispositions, although of course they are both species of the generic disposition to melt.

(C)

Finally, the example of marrying for the sake of money and marrying for the sake of love shows that Hacker's comparison between doing an act for one reason rather than another and doing an act for A's sake rather than B's sake does not support his case, for here too we have distinct species of the same generic act, and, corresponding to them, distinct species of the same generic ability.

p. 179 7.3.4 Knowledge is not an ability but a state

The use of the term 'state' is partly stipulative, and so we find different conceptions of a state in philosophical writings. Some philosophers contrast states with abilities or dispositions, while others include an object's abilities and dispositions among its states, and of course this difference can lead to apparently conflicting claims. For example, some philosophers prefer to say that being gaseous is a state, whereas being volatile is not; or that being depressed is a state, but being prone to depression is not. Others prefer to say that all these conditions are states and to distinguish between states of different kinds. But these apparently conflicting claims may only reflect a terminological disagreement, a disagreement in the use of the term 'state'.

One common approach to defining a state is to contrast states with processes, and to explain the distinction between them by distinguishing between verbs that have progressive tenses, corresponding to processes, and verbs that do not have progressive tenses, corresponding to states. For example, 'kill', 'knit', 'learn', and 'prove' have progressive tenses: 'is killing', 'was knitting', etc.; whereas 'know', 'believe', 'be able', and 'be blue' do not. We can say that a child is learning how to multiply or that a cat is killing a bird; but not that a child is knowing how to multiply or that a cat is being able to kill a bird. So on this approach, killing and knitting are held to be processes, whereas knowing and believing are held to be states.

In fact, this way of drawing the distinction between states and processes needs fine-tuning, because many verbs that are generally acknowledged to signify states as opposed to processes have an idiomatic progressive tense. This is true especially of psychological verbs. Thus we have, 'I am hoping', 'I am intending', and so on. What distinguishes these verbs from ones that signify processes is that for them the rule holds that A is ϕ ing if and only if A ϕ s. Thus, I am hoping that Lucy will arrive soon if and only if I hope that Lucy will arrive soon, I am intending to give up whisky for Lent if and only if I intend to give up whisky

for Lent, and so on. Hence, hoping and intending are states. By contrast, memorizing is a process, since I can spend all morning memorizing *Kubla Khan* without completing the task.¹⁹

p. 180 With or without this qualification, knowledge is certainly an intellectual state, in this sense of the term 'state', but it is not a state as opposed to an ability, because by these criteria abilities are states. But what distinguishes abilities from other kinds of states? I shall not attempt to define an ability here, I shall confine myself to the following observations.

An ability is a power, and therefore an actual property of the substance that possesses it, and not merely the possibility that the kind of act or event that manifests it should occur. For example, if *aqua regia* is able to dissolve gold, this does not mean merely that it is possible for gold coming into contact with it to dissolve. The point applies to every ability, but there is an additional reason for refusing to classify powers as possibilities in those cases where their exercise involves some physical or intellectual skill. For, as I pointed out earlier, many tasks that can be performed by agents with the relevant skills can also be performed occasionally by luck by those without them, as when a novice casts a perfect fly. If he does so, it follows that it is possible for him to do so, but not that he has the ability to do so. So the possibility and the ability are not the same.

However, abilities are sometimes confused with certain kinds of possibilities for which we also use the auxiliary verb 'can', specifically, circumstantial possibilities (opportunities) and epistemic possibilities (cases of consistency with known fact). As Kenny points out, both these distinctions are marked linguistically in English.²⁰ We refer to epistemic possibilities but not abilities with the word 'may'. For instance, 'Peter may still be in Rome' does not mean that he is able to remain there, it means that for all we know he is still there. And we refer to abilities and opportunities differently in the future tense. Kenny's own example is that 'I can speak Russian tomorrow, we have guests coming from Moscow' is correct, whereas 'I can speak Russian next spring; I shall be taking a beginner's course this fall' is incorrect. We refer to future opportunities either with 'can' or with 'will be able', but we refer to future abilities only with 'will be able'.

p. 181 Both Hume and Ryle confuse abilities with epistemic possibilities, but the distinction between abilities and opportunities is the easier one to lose sight of in this context.²¹ But it is important not to lose sight of it, for, if we do, we may imagine that a person is able to take a fact into account without knowing it, as long as it is a fact he can easily ascertain. 'Surely,' we may say, 'Tom could easily have taken the fact that it was going to rain into account: all he had to do was check the forecast.' Fair enough. But the phrase 'could have' here refers to an opportunity: it means that Tom's circumstances allowed him to do the thing in question, i.e. to take the fact that it was going to rain into account. That is why we can also say that Tom could have taken (i.e. would have been able to take) the fact that it was going to rain into account *if only* he had checked the forecast.

The distinction between an ability and an opportunity is not affected by how easily the ability is acquired. Some abilities can be acquired easily and in moments—for example, the ability to pronounce the name 'Luigi' or to spell the word 'fruit'. Even in these cases, possessing an ability is not the same as being in circumstances in which it is easy to acquire it. An ability is internal to the agent, and a positive factor in accounting for an act, whereas an opportunity is external, and may amount to no more than the absence of circumstances that would prevent the act from being done or interfere with it. As Kenny points out, it is difficult to make the intuitive truth that abilities are internal whereas opportunities are external precise.²² But it does not seem necessary to argue that acquiring knowledge enlarges our abilities rather than altering our circumstances.

7.3.5 A person's knowledge of a fact explains her ability to be guided by it, so the knowledge and the ability cannot be identical

However, it is possible to agree that knowledge enlarges our abilities because it enables us to be guided by facts, without agreeing that knowledge of a fact is the ability to be guided by it. Williamson and Setiya both take this position. They argue, independently, that a person's knowledge of a fact explains her ability to be guided by it, so the knowledge and the ability cannot be the same state. I shall comment on their arguments in turn.

Williamson agrees that someone who believes truly that p without knowing that p cannot do X because p , cannot do Y because p , cannot do Z because p , and so on:

But a single failure to know explains all these incapacities. If the incapacities constituted the failure to know, the correlation between the incapacities would be an unexplained coincidence.²³

This objection is surely unconvincing. Of course, explanations vary depending on the context in which they are given and the purpose for which they are required, with what is assumed to be common knowledge, and so on. But it simply is not true that if someone could not do X because p , could not do Y because p , etc., and these incapacities constituted her failure to know, the correlation between them would be an unexplained coincidence. Clearly, it would be explained by whatever explains her failure to know, such as the fact that she did not see this morning's paper. Similarly, if the inability to do crawl, or breaststroke, or doggy paddle, ... constitutes the inability to swim, it does not follow that someone's inability to do crawl, or breaststroke, or doggy paddle, ... is an unexplained coincidence. It is explained by whatever explains her inability to swim. Williamson says that 'a single failure to know explains all these incapacities'. And in a sense it does. It explains them by including them, in the way that a person's inability to ride a bicycle explains why she cannot cycle from X to Y , from Y to Z , from Z to W , and so on.

Setiya's argument is similar. He writes as follows:

[The theory that knowledge of a fact is the ability to be guided by it] cannot be right. For the relation of knowledge to this capacity is not symmetric. When someone lacks knowledge, they cannot act because p because they do not know that p . Compare the converse claim: S does not know that p because he cannot do things because p . That explanation is false. [...] On Hyman's view, there should be no asymmetry: the explanations should be on a par.²⁴

The argument is not quite accurately expressed. First, 'the relation of knowledge to this capacity is not symmetric' cannot be exactly what Setiya means. For things are not related by a single relation, but by many, and it is possible for things to stand in both symmetric and asymmetric relations, e.g. two siblings, one taller than the other. Perhaps he means that nothing can stand in an asymmetric relation to itself. But that brings us to the second point, which is that the asymmetry Setiya detects is not between a person's knowledge of a fact and her ability to be guided by it, it is between two facts: the fact that a person does not know a fact, and the fact that she cannot be guided by it. Setiya holds that the first fact explains the second, but the second does not explain the first. But if he is right about this, it merely follows that these are different facts. It does not follow that knowledge of a fact and the ability to be guided by it are different states. Similarly, if you park on a double yellow line, you contravene section 238 of the Highway Code because you do so, and not the other way around: you do not park on a double yellow line because you contravene section 238 of the Highway Code. So these are two facts asymmetrically related to each other. Still, parking on a double yellow line and contravening section 238 are one and the same act.²⁵

Hence, if 'S cannot be guided by the fact p because he does not know it' sounds right, while 'S does not know the fact that p because he cannot be guided by it' sounds wrong, this does not refute the theory that

knowledge of a fact is the ability to be guided by it. If it did refute it, many philosophical theories would be easier to dismiss than in fact they are. For example, some philosophers hold that to have a certain colour is to look that colour to normal observers in normal conditions. I do not myself believe this theory is correct, but it would be too quick to dismiss it on the grounds that ‘Her hair looks blue because it is blue (and not because of a trick of the light)’ is perfectly intelligible and may be true, whereas ‘Her hair is blue because it looks blue (and not because of the shampoo she uses)’ scarcely makes sense.²⁶

Presumably, the fact that one statement sounds right and the other sounds wrong reflects the fact that there are different ways of interpreting the ‘because’—i.e., there are different kinds of explanations—and if we interpret the ‘because’ in the wrong way, a valid explanation can seem false.²⁷ Thus, we can expect a philosopher who accepts this theory of colour to say that ‘Her hair looks blue because it is blue’ is potentially a valid causal explanation of why someone’s hair looks blue, whereas ‘Her hair is blue because it looks blue’ is part of a valid explanation of the kind of property being blue is. It is likely to sound baffling if we are expecting a causal explanation, such as ‘... because of the shampoo she uses’, because it is not that kind of explanation. But we cannot infer from this that the theory of colours it illustrates is false.

Similarly, ‘S does not know the fact that *p* because he cannot be guided by it’ will sound baffling if we are expecting a causal explanation—e.g. ‘Tom does not know that I’m here because I didn’t tell him’, or ‘Anne does not know who won because she didn’t hear the news’—for the same reason. But this has no bearing on whether it is part of a valid explanation of the kind of mental state knowing a fact is, i.e. on whether knowledge is the ability to be guided by the facts.

p. 184 But if neither Williamson’s argument nor Setiya’s is convincing, it does not follow that knowledge is the ability to be guided by the facts, as opposed to a state that explains why we have the ability, as they suggest. Consider the soporific power—the *virtus dormitiva*—of a drug. The drug has a particular internal structure, which explains why it has the ability to cause sleep. Having this internal structure is a matter of how it is, specifically, how its parts are arranged in relation to each other, whereas being able to cause sleep is a matter of what it *does*, when it is ingested. The drug has the power because it has the structure, the structure is the ‘ground’ or ‘basis’ of the power. Might knowledge be analogous to the drug’s structure rather than its power?

To answer this question we need to recall the distinction between possessing knowledge and containing information. We can assume that the state that ‘grounds’ a person’s ability to be guided by a fact is an informational state, which is realized in the synaptic connections and neural firing patterns in her brain. Philosophers and neuroscientists have sometimes assumed that knowledge itself is an informational state, but containing information is not the same as possessing knowledge. For instance, an old vinyl record does not know the songs recorded on it, but it contains the information that enables a record-player to play them, because it was stamped with the right shaped groove (the information is realized in the groove). So it would be a mistake to think that knowledge is analogous to the informational state of the record, as opposed to the record-player’s ability to play the song. It is the ability rather than the state that grounds it. That is why the acid test of what a person knows—as of any ability or skill—is what she says and does, for example, when she modifies her behaviour in pursuit of her goals.*

Furthermore, it is a basic rational principle that simpler empirically adequate hypotheses are to be preferred. There are certainly other theories of knowledge which imply that knowing that *p* is necessary and sufficient for being able to be guided by the fact that *p*, apart from the theory that they are the same intellectual state, such as the functionalist theory that *knowing a fact is having the second-order property of having some first-order property that grounds the ability to be guided by it*. But the theory that knowing a fact is being able to be guided by it is obviously the simplest explanation, hence the onus of proof must fall on those who prefer a different one. Besides, if the functionalist theory turned out to be preferable, I would regard this as more of a vindication than a defeat. For the principle I have been at pains to defend in this

chapter, that the right way to define knowledge is to explain how it gets exercised or expressed in rational thought and behaviour, instead of asking how it can be certified or acquired, would be confirmed.

7.4 Our knowledge of our own mental states

The theory of knowledge defended in this chapter is partly inspired by Wittgenstein's claim that 'the grammar of the word "knows" is evidently closely related to that of "can", "is able to"'. But it contradicts his famous remark: 'It can't be said of me at all (except perhaps as a joke) that I *know* I am in pain.'²⁸ For the fact that I am in pain certainly *can* be my reason for doing, wanting, or believing something, say, for taking an aspirin, for wanting to lie down, or for believing I have spent too long in the sun. So if the theory of knowledge defended here is right, it *can* be said of me (not merely as a joke) that I know I am in pain. Is this a reason for disputing Wittgenstein's remark? Or is it a reason for rejecting or modifying the theory of knowledge?

Two preliminary points should be noted. First, Wittgenstein made less famous remarks that express a different view. For instance, in a remark about my knowledge of the disposition of my limbs, e.g. whether my arm is extended or not, he says: 'I simply *know* how it is—without knowing it *because* ... Just as I also know where I feel pain—but do not know it *because*.'²⁹ (I have no quarrel with this.) Second, the claim that I cannot be said to know that I am in pain is an ingenious gambit to play against the philosopher who is tempted by the thought that *only* I can know whether I am in pain. It turns the tables on him in a dramatic way, and puts him on the defensive. But it is not an essential part of Wittgenstein's attack on the idea that we cannot ever really know what others think and feel, or even whether they have thoughts and feelings at all.

Wittgenstein supports the claim that I cannot be said to know that I am in pain with two arguments, or rather, with a challenge and an argument. Here is the challenge:

It can't be said of me at all (except perhaps as a joke) that I *know* I am in pain. What is it supposed to mean—except perhaps that I *am* in pain?

And here is the argument: our use of the verb 'know' is coordinated with the use of a number of other verbs, and can only occur in circumstances which might also permit the use of these other verbs—verbs such as 'learn', 'doubt', 'believe', 'suspect', 'find out', and so on. But I cannot be said to learn, doubt, believe, suspect, or find out that I am in pain; hence, I cannot be said to know I am in pain either. The relevant remarks are these:

Other people cannot be said to learn of my sensations *only* from my behaviour,—for I cannot be said to learn of them. I *have* them.

The truth is: it makes sense to say about other people that they doubt whether I am in pain; but not to say it about myself.

'I know ...' may mean 'I do not doubt ...' but does not mean that the words 'I doubt ...' are *senseless*, that doubt is logically excluded.

One says 'I know' where one can also say 'I believe' or 'I suspect'; where one can find out.³⁰

I shall comment on the challenge and the argument in turn.

(A)

So far as the challenge is concerned, there appears to be a simple answer. The difference between ‘Katy knows she is in pain’ and ‘Katy is in pain’ is that ‘Katy knows she is in pain’ implies that Katy can be guided by the fact that she is in pain, whereas ‘Katy is in pain’ does not. A fish or a new-born baby cannot be aware of the fact that it is in pain—although it can of course be in pain—because a fish or a baby cannot be guided by reasons in the way it acts. Hence, if Katy is a new-born baby, ‘Katy is in pain’ may be true, but ‘Katy knows that she is in pain’ will certainly be false. A new-born baby can *be* in pain, but she cannot *know* whether she is in pain, either that she is or that she is not.

(B)

What about the argument? First, there are cases other than the contested ones about my current conscious state where some but not all of the verbs Wittgenstein mentions seem to apply happily, such as the most elementary truths of arithmetic. For example, I know, and I once learned, that $3 + 2 = 5$, but is this something I could suspect or surmise? It seems not. So the coordination of these verbs appears to be looser and more variable than Wittgenstein’s remarks imply. Second, it invites the following question: why can’t we concede that the use of the verb ‘know’ is *normally* or *generally* coordinated with the use of the other verbs Wittgenstein mentions, while insisting that facts about my present conscious experience, and possibly other facts that I can know, are anomalous cases? After all, there are plenty of anomalies of this sort in our use of words, and the verb ‘know’ itself provides several examples.

For example, Ryle says: ‘To know a truth I must have discovered or established it’, but the ‘must’ is rash, because there are anomalous cases.³¹ I know that I live in London and that I speak French, but these are not things that I have discovered or established. In a similar vein, Wittgenstein says: ‘If someone [...] knows something, then the question “how does he know?” must be capable of being answered.’³² But must it? Elsewhere, he denies that it must:

‘But how could I know that I should have reacted like this if you had asked me?’—How? There is no How. But there are indications that I am right in saying this.³³

Again, Wittgenstein says: ‘Whether I *know* something depends on whether the evidence backs me up or contradicts me.’³⁴ But what is the evidence that $1 + 1 = 2$ or that identity is symmetric? And yet I know these things. So Wittgenstein may have been right to deny that I can suspect or surmise that I am in pain, and to insist that the question ‘How do you know that you are in pain?’ is absurd. But why should we infer that I cannot know that I am in pain? He does not say.

Hence, Wittgenstein’s position appears to be a weak one, unless he had a plausible conception of knowledge in mind, which draws the boundaries of what one can be said to know more narrowly than I wish to, and which supports his idea about the coordination of verbs. But it is difficult to be certain whether he did have such a conception of knowledge in mind. As far as I know, the only commentator who has addressed this question directly is Kenny. He writes as follows:

Throughout his life [Wittgenstein] thought of knowledge as involving the possession of a true description of a state of affairs. That was why in the *Tractatus* he declared knowledge of tautologies impossible, since tautologies were not pictures of states of affairs. That is why, since ‘I am in pain’ is not, when I am in pain, a true description in the normal sense, ‘I know that I am in pain’ cannot be in order if ‘know’ is being used in the normal sense.³⁵

Wittgenstein certainly suggests that the sentence ‘I am in pain’ is not a description: when you say that you are in pain, he argues, you do not ‘read off what you say from the facts’.³⁶ If he thought of knowledge as involving the possession of a true description of a state of affairs, the claim that I cannot be said to know that I am in pain would have looked like a corollary of this doctrine. But is Kenny’s interpretation correct?

On the face of it, it is plausible. It explains Wittgenstein's remarks about knowledge of one's own conscious state, and it also accounts for the doctrine that we ought not to speak of knowing mathematical propositions either, which some commentators have attributed to Wittgenstein, presumably on the basis of the following lines from the *Remarks on the Foundations of Mathematics*: 'If \hookrightarrow you know a mathematical proposition, that's not to say you yet know *anything*. // I.e., the mathematical proposition is only supposed to supply a framework for a description.'³⁷ But there are other remarks, notably in *On Certainty*, which suggest that even though he distinguished sharply between mathematical propositions and true descriptions of states of affairs, he was happy to speak of knowledge in mathematics, and which therefore tell against Kenny's interpretation. For example: 'If I say "I know" in mathematics, then the justification for this is a proof.'³⁸

But whether or not Kenny's interpretation is right, the conception of knowledge he attributes to Wittgenstein is mistaken. If Rebecca knows that the angles of a triangle are equal to two right-angles, that is something Rebecca knows, and there is nothing wrong with saying so. We can acknowledge the differences between the ways in which mathematical knowledge and empirical knowledge are acquired, used, and tested, without drawing such narrow limits around the field of knowledge that mathematics is excluded from it. And the same goes for our own conscious states. The sort of knowledge that is expressed in what Wittgenstein thought of as a true description of an empirical state of affairs is of course one pre-eminent paradigm of knowledge. But we should not allow ourselves to become so captivated by a paradigm that no other sort of knowledge seems possible. My knowledge extends to every fact I can be guided by, be it a fact about my perceptible environment, a mathematical fact or a fact about my own conscious state. In every case, I can respond to the fact rationally, or take it into consideration or account, and that is all we need to talk of knowledge.

In sum, Wittgenstein's famous remark, 'It can't be said of me at all (except perhaps as a joke) that I know I am in pain', is a dramatic reversal of philosophical orthodoxy. But the arguments he offers to support it are inconclusive, and, if Kenny's interpretation is correct, its ultimate source was a mistaken conception of knowledge itself. It would be generally agreed that in normal circumstances an adult human being knows whether or not he is in pain, whereas a man in a delirium, a new-born baby, or a fish does not. The theory of knowledge defended here explains why.

7.5 Conclusion

The idea that knowledge is an ability was occasionally pursued by philosophers influenced by Wittgenstein and Ryle between the 1960s and 1990s, \hookrightarrow but it was almost completely eclipsed by the idea that knowledge is a species of belief. Gettier's influential article 'Is justified true belief knowledge?' persuaded most philosophers that knowledge is not simply justified true belief, but several decades passed before scepticism about the idea that knowledge is *some* kind of belief started to gain ground.

Some philosophers challenged the assumption that knowledge is a species of belief. In particular, Ryle pointed out several ways in which knowledge resembles other abilities and skills, while belief resembles tendencies and dispositions, and Vendler argued that belief and knowledge have different objects: what is believed is a proposition, whereas what is known is a fact.³⁹ (See 7.2.) But these arguments had little effect on opinion, and when the assumption that knowledge is a species of belief began to be more widely questioned, the main reason was exhaustion. Gettier instigated a vast amount of work, but no definition of knowledge in terms of belief commanded general assent. Counterexamples were invented that seemed to disprove every published proposal. And inevitably, the more arcane the definitions proposed in the literature became, the more likely it seemed that even if one was devised to which no counterexample could be found, it would not be a plausible explanation of what knowledge is, because there could be no point in

using such a complicated concept. In other words, it could not matter whether someone met the conditions comprising it or not.

For these reasons, a sceptical view of the prospects for a definition of knowledge in terms of belief has become quite common, and the view that knowledge—as opposed to justified belief—does not matter has also been aired.⁴⁰ But the idea that knowledge is a kind of belief has cast such a long shadow that even those who reject it tend to equate the question of whether it is possible to define knowledge in terms of belief with the question of whether it is possible to define knowledge at all, and so the conviction that knowledge is indefinable has increased in popularity at the same time.⁴¹

p. 190 If knowledge is the ability to be guided by the facts, the project of defining knowledge did not run into the sands because knowledge is inherently resistant to definition, but because it is an ability and not a species of belief, and so the question we need to ask in order to define it is how it gets exercised or expressed, rather than how it can be certified or acquired. We need to think about knowledge applied, employed, in the infinitely varied circumstances that arise in the course of human life. We need to think about knowledge in action, knowledge in use. Why did so few philosophers see this during the last fifty years? The main reason is probably that until quite recently, hardly anyone saw that the kind of explanation of thought and behaviour that refers to knowledge cannot be subsumed into the kind that refers to belief⁴² (see 6.4).

A theory of factual knowledge is first and last a testable definition of the kind of intellectual state an individual's knowledge of a fact is. But a convincing theory will not only explain what knowledge is: it will also explain why we value it. I shall turn to this problem in the next chapter.

Notes

- * That is why sceptical claims (e.g. that I do not know that I am not a brain in a vat) are commonly supported by claiming that there is no satisfactory answer to the question of how we know what the sceptic denies that we know (viz. that I am not a brain in a vat). See for example, DeRose, 'Solving the Sceptical Problem', p. 2.
- * Another alternative, proposed by Williamson, is that 'knowing is the most general factive stative attitude' (*Knowledge and its Limits*, p. 34). I discuss this proposal, and its relationship with the definition of knowledge defended here, in my article, "The most general factive stative attitude".
- * Guidance by a rule is similar, for, as Quine points out, behaviour *fits* any rule it conforms to, but it can only be *guided* by a rule the behavior knows ('Methodological Reflections on Current Linguistic Theory', p. 386). Interestingly, Wittgenstein draws the contrast between following and conforming to rules in terms of reasons (*BB*, pp. 13–15).
- * The phrase 'disposition to act as if *p*' is widely used without the gloss Braithwaite gave it when he introduced it. He explains that it is to be treated as an abbreviation for 'disposition to act in ways that would be conducive to the satisfaction of one's needs if *p* were true' ('The Nature of Believing', p. 134). But this is not a plausible theory of belief, since plants and primitive animals have the disposition to act in ways that are conducive to satisfying their needs without having beliefs. The phrase is now generally treated as an abbreviation for 'disposition to act in ways that would be conducive to the satisfaction of one's desires if *p* were true'. Velleman is rightly critical of the idea that belief can be defined purely in terms of its effect on behaviour (*The Possibility of Practical Reason*, pp. 255f.).
- * In Wittgenstein's terms, a person's behaviour is not a *symptom* of knowledge, but a *criterion*. This is not a reversion to behaviourism; it is perfectly compatible with the point made earlier, that knowledge gets expressed in thought and feeling as well as in behaviour.

Notes

1. See for example Stanley & Williamson, 'Knowing How'; Rumfitt, 'Savoir Faire'; Snowdon, 'Knowing How and Knowing That: A Distinction Reconsidered'; Stanley, *Know How*; Wiggins, 'Practical Knowledge: Knowing How To and Knowing That'; Bengson & Moffett (eds), *Knowing How: Essays on Knowledge, Mind, and Action*.
2. Cassam, 'Can the Concept of Knowledge be Analysed?', p. 27.

3. Ryle, 'Mowgli in Babel', p. 5.
4. The comparison is due to Fine, 'First-Order Modal Theories III: Facts', p. 52. See also Slote, *Metaphysics and Essence*, ch. 6. For the view that facts are simply true propositions, see Carnap, *Meaning and Necessity*, p. 28; Prior, *Objects of Thought*, pp. 4–6; Williamson, *Knowledge and its Limits*, pp. 42f.
5. This argument is expounded in detail in Fine, 'First-Order Modal Theories III: Facts', pp. 46–50.
6. Ryle, *The Concept of Mind*, pp. 116ff.
7. Plato, *Republic* 477d; Wittgenstein, *Philosophical Investigations*, §150; Ryle, *The Concept of Mind*, p. 129.
8. See, for example, Margolis, 'Knowledge, Belief and Thought'; White, *The Nature of Knowledge*, ch. 6; Craig, *Knowledge and the State of Nature*, pp. 11f and passim.
9. White, *The Nature of Knowledge*, p. 119; cf. Wittgenstein, *On Certainty*, §586.
10. White, *The Nature of Knowledge*, p. 121.
11. White, *The Nature of Knowledge*, p. 120.
12. Bernard Williams, 'Knowledge and Reasons', p. 3.
13. Kenny, *The Metaphysics of Mind*, pp. 108f.
14. In 'Knowledge, Explanation and Motivating Reasons', Dustin Locke denies that if one takes a fact into consideration one is thereby guided by it, on the dubious ground that a decision to do an act of a certain kind can take reasons *against* doing it into consideration, but cannot be guided by those reasons.
15. Annas & Barnes, *The Modes of Scepticism*, pp. 47–8; cf. Sorabji, *Animal Minds and Human Morals*, ch.7. The story is told by Philo, *On Animals*, 45–6, and by Sextus Empiricus, *Outlines of Pyrrhonism*, I §69.
16. Davidson, 'Thought and Talk', repr. in *Inquiries into Truth and Interpretation*, pp. 155–70.
17. Davidson, 'Thought and Talk', repr. in *Inquiries into Truth and Interpretation*, p. 170.
18. Hacker, *The Intellectual Powers: A Study of Human Nature*, p. 183.
19. Kenny, *Action, Emotion, and Will*, p. 175.
20. Kenny, *Will, Freedom and Power*, p. 132.
21. Hume, *A Treatise of Human Nature*, 2.1.10; Ryle, *The Concept of Mind*, pp. 114f.
22. Kenny, *Will, Freedom and Power*, p. 133.
23. Williamson, *Knowledge and its Limits*, p. 64.
24. Setiya, 'Causality in Action', p. 511.
25. I discuss the identity of acts in 3.2.
26. I discuss colours and appearances in Hyman, *The Objective Eye*, ch.1, and p. 240, n. 6.
27. See Wiggins, *Needs, Values, Truth*, pp. 106f; and cf. Hyman, *The Objective Eye*, pp. 53–6.
28. Wittgenstein, *Philosophical Investigations*, §246.
29. Wittgenstein, *Zettel*, §481.
30. Wittgenstein, *Philosophical Investigations*, §246 & II, p. 221.
31. Ryle, 'Knowing How and Knowing That', in *Collected Essays 1929–1968*, vol.2, p. 224.
32. Wittgenstein, *On Certainty*, §550.
33. Wittgenstein, *Remarks on the Philosophy of Psychology I*, §428.
34. Wittgenstein, *On Certainty*, §504; but cf. *Remarks on the Foundations of Mathematics*, pp. 336–7.
35. Kenny, *Wittgenstein*, p. 201.
36. Wittgenstein, *Philosophical Investigations*, §292; cf. II, p. 189.
37. Wittgenstein, *Remarks on the Foundations of Mathematics*, p. 356.
38. Wittgenstein, *On Certainty*, §563.
39. Vendler, *Res Cogitans*, ch. 5. The argument in 7.2 supports this conclusion, but the conception of facts proposed there is different from Vendler's.
40. Kaplan, 'It's Not What You Know That Counts', pp. 354–6.
41. E.g. Williamson, *Knowledge and Its Limits*, ch.1, s.3; cf. Millar, 'Knowledge in Recent Epistemology: Some Problems', 5.2.
42. I discuss this point in detail in 6.4.