

PART I  
Skepticism

---



# Introduction

---

Like Rene Descartes, we have all asked ourselves at one time or another “Couldn’t everything I seem to see, hear, etc. be illusory? Might I in fact be dreaming all this? If so, what do I really know of the outside world?” The skeptic’s answers are pessimistic: yes, you could be dreaming, and so you know nothing of the outside world. The conclusion is outlandish, and yet the reasoning behind it hardly seems strained at all. We feel the pressure towards skepticism in the movement from the question about the trustworthiness of our senses to the question of our ability to know. Given that the bulk of our knowledge of the outside world derives from the senses, how can we know anything about the world unless we first show that our senses can be trusted? The core of the skeptical strategy is more general: how can one gain knowledge using a source of belief unless one first shows that the source is trustworthy?

In his selection, Barry Stroud presents the skeptic’s argument in its most favorable light. The skeptic does not hold us up to an uncommonly high standard of knowledge only to make the obvious point that we fail to meet it. The skeptic invokes only the standards presupposed in everyday knowledge attributions. To use an example of Stroud’s, if no goldfinch could possibly be a canary, then if one is to know that the bird one sees is a goldfinch, one must be able to rule out its being a canary. More generally, to know that  $p$ , one must be able to rule out every possibility one knows to be incompatible with one’s knowing that  $p$ . The skeptic then has her wedge: to know that you’re sitting beside a warm fire, you must be able to rule out any possibility which excludes this knowledge, including innumerable “skeptical possibilities,” such as that you’re dreaming, that you’re being deceived by a malicious demon, and that you’re a brain in a vat stimulated to have the experiences and apparent memories you now have. But it’s hard to see how you can rule these out.

In each of the selections from the work of G. E. Moore, the tables are turned on the skeptic. Moore provides a counter-argument in “Proof of an External World.” A good proof, he explains, proceeds from known premises to a distinct conclusion to which they can be seen to lead. He then produces an example: raising his hands, one after the other, he exclaims “Here is a hand. Here is another hand,” and he concludes “There are

hands.” If asked to prove his premises, he would reject the demand, for not everything that is known can be proved.

Moore nevertheless takes the skeptic seriously. In “Certainty,” he grants that if he doesn’t know he is not dreaming, he doesn’t know he is standing up giving a lecture. Still he asks why there is any more plausibility in using this premise as part of a *modus ponens* argument to conclude that he doesn’t know he is standing up than in using it as a part of the corresponding *modus tollens* argument to conclude that he does know after all that he is not dreaming.

In “Four Forms of Scepticism,” Moore fully admits that skeptical scenarios are logically possible, but he finds it more certain that something has gone awful in the skeptical argument than that he lacks knowledge that he has hands (or is holding a pencil). Moreover, he concludes that since the only way he *could* know this is through some inductive or analogical argument from the character of his experience, such an argument *must* exist.

The selections from Stroud and Moore concern our knowledge of the external world. One might hope that, even if it is hard to answer the skeptical challenge for *knowledge*, at least it could be satisfactorily answered for *justification*. Peter Klein calls the view that we cannot be justified in our beliefs about how things *are* (as opposed to how they seem) “Academic skepticism” and contrasts it with an older form of skepticism: Pyrrhonism. Pyrrhonism, in Klein’s view, is a more moderate skepticism than its Academic cousin, for Pyrrhonism allows that our beliefs can be conditionally or provisionally justified. But it is still a form of skepticism, because it denies that our beliefs can be completely justified. Only if reasoning could settle the matter of whether a belief is true could that belief be completely justified. But how can reasoning settle anything? If it were legitimate to end reasoning with a proposition for which we could not provide a further reason, then it seems reasoning could settle some matter. But this is not legitimate. Nor is it legitimate to reason in a circle. Therefore, the only way for reasoning to settle matters would be to complete an infinite regress of non-repeating reasons (a view Klein refers to as “infinetism,” discussed in more detail in his contribution to Part II). While this would be a legitimate way to settle some matter, it cannot, in fact, be done.

The lesson for Academic skepticism is that the arguments invoked in favor of Academic skepticism are themselves fallacious in that they either rely on arbitrary premises or beg the question in favor of their conclusion. Thus, consider the Academic skeptic’s claim that we cannot know whether we are dreaming or deceived by a malicious demon. This claim is central to the argument for Academic skepticism. If it is unsupported, it is arbitrary. To support the claim, the Academic skeptic must first demonstrate that we cannot know, say, that there is a table in front of us. But “I cannot know there is a table in front of me” is the ostensible conclusion of the skeptical argument. Therefore, Academic skepticism, like the inadequate models of reasoning, must either rely on arbitrary premises or beg the question.

Michael Williams argues that if there is such a thing as *knowledge of the external world*, the kind of knowledge the Cartesian skeptic questions, it seems impossible for us to see ourselves as having it. That is, the skeptic would carry the day. But he asks: *is* there such a thing as knowledge of the world? His answer is *no*. The concept of knowledge of the external world is a theoretical concept, and so, unlike practical concepts such as the concept of a chair, it lacks application entirely unless there is an appropriate unified

domain of reality whose contours are there for it to match. But there is no such epistemic domain. There could be only if (empirical) beliefs divided into two classes: those that could only be known on the basis of beliefs about immediate experience, i.e., beliefs about the external world, and those that could be known directly from immediate experience. Yet an examination of our practices in attributing knowledge and justification suggests that beliefs do not divide into these epistemic categories nor into any objective epistemic categories.

Williams describes his view as a form of contextualism. But it is a version of contextualism quite different from those appearing in Part VIII of this volume. The contextualist theories of DeRose and Cohen, and to a lesser extent Lewis, presuppose the existence of a unified range of objective characteristics which, given a speech context, comprise the truth-conditions for knowledge attributions in that context. For DeRose, there are the objective (context-invariant) notions of sensitivity and strength of epistemic position, and for Cohen objective notions of strength of evidence or justification. For Lewis, there are the objective factors of one's evidence and which possibilities it rules out. For all three of these epistemologists, the function of context is to set the bar on which (or what degree) of a relatively unified range of objective factors count. Thus, for them, there is an independent place for epistemological inquiry into the nature of these objective factors as well as into how they feed into the semantics of knowledge attribution. According to Williams, by contrast, there is no range of objective factors, with the result that there is nothing at all to serve as an object of theoretical investigation for the epistemologist.

Part and parcel of repudiating skepticism, then, is repudiating traditional epistemology. Both rely for their livelihood on the assumption that Williams calls "epistemological realism," viz. that there are objective relations of epistemic priority waiting to be described.

### Further Reading

- Annas, Julia and Jonathan Barnes, *The Modes of Scepticism* (Cambridge: Cambridge University Press, 1985).
- Burnyeat, M. F. (ed.), *The Skeptical Tradition* (Berkeley: University of California Press, 1983).
- Clarke, Thompson, "The Legacy of Skepticism," *Journal of Philosophy* 69 (1972), pp. 754–69.
- DeRose, K. and T. Warfield (eds), *Skepticism: A Contemporary Reader* (Oxford: Oxford University Press, 1999).
- Fumerton, R., *Metaepistemology and Skepticism* (Langham, MD: Rowman and Littlefield, 1995).
- Hadot, Pierre, *Philosophy as a Way of Life* (Oxford: Blackwell, 1995).
- Hankinson, R. J., *The Sceptics* (London: Routledge, 1995).
- Huemer, Michael, *Skepticism and the Veil of Perception* (Langham, MD: Rowman and Littlefield, 2001).
- Landesman, Charles and Roblin Meeks (eds), *Philosophical Skepticism: From Plato to Rorty* (Oxford: Blackwell, 2002).
- Klein, P., *Certainty: A Refutation of Scepticism* (Minneapolis: University of Minnesota Press, 1991).
- Nozick, Robert, *Philosophical Explanations* (Oxford: Oxford University Press, 1981).
- Popkin, Richard, *Skepticism from Erasmus to Spinoza* (Berkeley: University of California Press, 1979).
- Roth, Michael D. and Glenn Ross (eds), *Doubting: Contemporary Perspectives on Skepticism* (Dordrecht: Kluwer Academic Publishers, 1990).
- Sosa, Ernest, "The Sceptic's Appeal," in Marjorie Clay and Keith Lehrer (eds), *Knowledge and Skepticism* (Boulder, CO: Westview Press, 1989), pp. 51–68.

- Strawson, P. F., *Skepticism and Naturalism: Some Varieties* (London: Methuen, 1985).
- Stroud, Barry, "Understanding Human Knowledge in General," in Marjorie Clay and Keith Lehrer (eds), *Knowledge and Skepticism* (Boulder, CO: Westview Press, 1989), pp. 31–50.
- , *The Significance of Philosophical Scepticism* (Oxford: Oxford University Press, 1984).
- Unger, Peter, *Ignorance: A Case for Scepticism* (Oxford: Clarendon Press, 1975).
- Williams, Michael, *Unnatural Doubts* (Princeton: Princeton University Press, 1996).

# CHAPTER 1

## The Problem of the External World

---

Barry Stroud

Since at least the time of Descartes in the seventeenth century there has been a philosophical problem about our knowledge of the world around us.<sup>1</sup> Put most simply, the problem is to show how we can have any knowledge of the world at all. The conclusion that we cannot, that no one knows anything about the world around us, is what I call “scepticism about the external world”, so we could also say that the problem is to show how or why scepticism about the external world is not correct. My aim is not to solve the problem but to understand it. I believe the problem has no solution; or rather that the only answer to the question as it is meant to be understood is that we can know nothing about the world around us. But how is the question meant to be understood? It can be expressed in a few English words familiar to all of us, but I hope to show that an understanding of the special philosophical character of the question, and of the inevitability of an unsatisfactory answer to it, cannot be guaranteed by our understanding of those words alone. To see how the problem is meant to be understood we must therefore examine what is perhaps best described as its source – how the problem arises and how it acquires that special character that makes an unsatisfactory negative answer inevitable. We must try to understand

the *philosophical* problem of our knowledge of the external world.

The problem arose for Descartes in the course of reflecting on everything he knows. He reached a point in his life at which he tried to sit back and reflect on everything he had ever been taught or told, everything he had learned or discovered or believed since he was old enough to know or believe anything.<sup>2</sup> We might say that he was reflecting on his knowledge, but putting it that way could suggest that what he was directing his attention to was indeed knowledge, and whether it was knowledge or not is precisely what he wanted to determine. “Among all the things I believe or take to be true, what amounts to knowledge and what does not?”; that is the question Descartes asks himself. It is obviously a very general question, since it asks about everything he believes or takes to be true, but in other respects it sounds just like the sort of question we are perfectly familiar with in everyday life and often know how to answer.

For example, I have come to accept over the years a great many things about the common cold. I have always been told that one can catch cold by getting wet feet, or from sitting in a draught, or from not drying one’s hair before going outdoors in cold weather. I have also learned that the common cold is the effect of a virus transmitted by an already infected person. And I also believe that one is more vulnerable to colds when over-tired, under stress, or otherwise in less than the best of health. Some of these beliefs seem

---

Originally published in B. Stroud, *The Significance of Philosophical Scepticism* (Oxford: Clarendon Press, 1984), ch. 1.

to me on reflection to be inconsistent with some others; I see that it is very unlikely that all of them could be true. Perhaps they could be, but I acknowledge that there is much I do not understand. If I sit back and try to think about all my “knowledge” of the common cold, then, I might easily come to wonder how much of it really amounts to knowledge and how much does not. What do I really know about the common cold? If I were sufficiently interested in pursuing the matter it would be natural to look into the source of my beliefs. Has there ever been any good reason for thinking that colds are even correlated with wet hair in cold weather, for example, or with sitting in a draught? Are the people from whom I learned such things likely to have believed them for good reasons? Are those beliefs just old wives’ tales, or are they really true, and perhaps even known to be true by some people? These are questions I might ask myself, and I have at least a general idea of how to go about answering them.

Apart from my impression of the implausibility of all my beliefs about the common cold being true together, I have not mentioned any other reason for being interested in investigating the state of my knowledge on that subject. But for the moment that does not seem to affect the intelligibility or the feasibility of the reflective project. There is nothing mysterious about it. It is the sort of task we can be led to undertake for a number of reasons, and often very good reasons, in so far as we have very good reasons for preferring knowledge and firm belief to guesswork or wishful thinking or simply taking things for granted.

Reflection on or investigation of our putative knowledge need not always extend to a wide area of interest. It might be important to ask whether some quite specific and particular thing I believe or have been taking for granted is really something I know. As a member of a jury I might find that I have been ruling out one suspect in my mind because he was a thousand miles away, in Cleveland, at the time of the crime. But I might then begin to ask myself whether that is really something that I know. I would reflect on the source of my belief, but reflection in this case need not involve a general scrutiny of everything I take myself to know about the case. Re-examining the man’s alibi and the credentials of its supporting witnesses might be enough to satisfy me. Indeed I might find that its reliability on those counts is precisely what I had been going on all along.

In pointing out that we are perfectly familiar with the idea of investigating or reviewing our knowledge on some particular matter or in some general area I do not mean to suggest that it is always easy to settle the question. Depending on the nature of the case, it might be very difficult, perhaps even impossible at the time, to reach a firm conclusion. For example, it would probably be very difficult if not impossible for me to trace and assess the origins of many of those things I believe about the common cold. But it is equally true that sometimes it is not impossible or even especially difficult to answer the question. We do sometimes discover that we do not really know what we previously thought we knew. I might find that what I had previously believed is not even true – that sitting in draughts is not even correlated with catching a cold, for example. Or I might find that there is not or perhaps never was any good reason to believe what I believed – that the man’s alibi was concocted and then falsely testified to by his friends. I could reasonably conclude in each case that I, and everyone else for that matter, never did know what I had previously thought I knew. We are all familiar with the ordinary activity of reviewing our knowledge, and with the experience of reaching a positive verdict in some cases and a negative verdict in others.

Descartes’s own interest in what he knows and how he knows it is part of his search for what he calls a general method for “rightly conducting reason and seeking truth in the sciences”.<sup>3</sup> He wants a method of inquiry that he can be assured in advance will lead only to the truth if properly followed. I think we do not need to endorse the wisdom of that search or the feasibility of that programme in order to try to go along with Descartes in his general assessment of the position he is in with respect to the things he believes. He comes to find his putative knowledge wanting in certain general respects, and it is in the course of that original negative assessment that the problem I am interested in arises. I call the assessment “negative” because by the end of his *First Meditation* Descartes finds that he has no good reason to believe anything about the world around him and therefore that he can know nothing of the external world.

How is that assessment conducted, and how closely does it parallel the familiar kind of review of our knowledge that we all know how to conduct



in everyday life? The question in one form or another will be with us for the rest of this book. It is the question of what exactly the problem of our knowledge of the external world amounts to, and how it arises with its special philosophical character. The source of the problem is to be found somewhere within or behind the kind of thinking Descartes engages in.

One way Descartes's question about his knowledge differs from the everyday examples I considered is in being concerned with *everything* he believes or takes to be true. How does one go about assessing all of one's knowledge all at once? I was able to list a few of the things I believe about the common cold and then to ask about each of them whether I really know it, and if so how. But although I can certainly list a number of the things I believe, and I would assent to many more of them as soon as they were put to me, there obviously is no hope of assessing everything I believe in this piecemeal way. For one thing, it probably makes no sense, strictly speaking, to talk of the number of things one believes. If I am asked whether it is one of my beliefs that I went to see a film last night I can truly answer "Yes". If I were asked whether it is one of my beliefs that I went to the movies last night I would give the same answer. Have I thereby identified two, or only one, of my beliefs? How is that question ever to be settled? If we say that I identified only one of my beliefs, it would seem that I must also be said to hold the further belief that going to see a film and going to the movies are one and the same thing. So we would have more than one belief after all. The prospects of arriving even at a principle for counting beliefs, let alone at an actual number of them, seem dim.

Even if it did make sense to count the things we believe it is pretty clear that the number would be indefinitely large and so an assessment of our beliefs one by one could never be completed anyway. This is easily seen by considering only some of the simplest things one knows, for example in arithmetic. One thing I know is that one plus one equals two. Another thing I know is that one plus two is three, and another, that one plus three is four. Obviously there could be no end to the task of assessing my knowledge if I had to investigate separately the source of each one of my beliefs in that series. And even if I succeeded I would only have assessed the things I know about

the addition of the number one to a given number; I would still have to do the same for the addition of two, and then the addition of three, and so on. And even that would exhaust only my beliefs about addition; all my other mathematical beliefs, not to mention all the rest of my knowledge, would remain so far unexamined. Obviously the job cannot be done piecemeal, one by one. Some method must be found for assessing large classes of beliefs all at once.

One way to do this would be to look for common sources or channels or bases of our beliefs, and then to examine the reliability of those sources or bases, just as I examined the source or basis of my belief that the suspect was in Cleveland. Descartes describes such a search as a search for "principles" of human knowledge, "principles" whose general credentials he can then investigate (HR, 145). If some "principles" are found to be involved in all or even most of our knowledge, an assessment of the reliability of those "principles" could be an assessment of all or most of our knowledge. If I found good reason to doubt the reliability of the suspect's alibi, for example, and that was all I had to go on in my belief that he was in Cleveland, then what I earlier took to be my knowledge that he was in Cleveland would have been found wanting or called into question. Its source or basis would have been undermined. Similarly, if one of the "principles" or bases on which all my knowledge of the world depends were found to be unreliable, my knowledge of the world would to that extent have been found wanting or called into question as well.

Are there any important "principles" of human knowledge in Descartes's sense? It takes very little reflection on the human organism to convince us of the importance of the senses – sight, hearing, touch, taste, and smell. Descartes puts the point most strongly when he says that "all that up to the present time I have accepted as most true and certain I have learned either from the senses or through the senses" (HR, 145). Exactly what he would include under "the senses" here is perhaps somewhat indeterminate, but even if it is left vague many philosophers would deny what Descartes appears to be saying. They would hold that, for example, the mathematical knowledge I mentioned earlier is not and could not be acquired from the senses or through the senses, so not *everything* I know is known in that way. Whether

Descartes is really denying the views of those who believe in the non-sensory character of mathematical knowledge, and whether, if he were, he would be right, are issues we can set aside for the moment. It is clear that the senses are at least very important for human knowledge. Even restricting ourselves to the traditional five senses we can begin to appreciate their importance by reflecting on how little someone would ever come to know without them. A person blind and deaf from birth who also lacked taste buds and a sense of smell would know very little about anything, no matter how long he lived. To imagine him also anaesthetized or without a sense of touch is perhaps to stretch altogether too far one's conception of a human organism, or at least a human organism from whom we can hope to learn something about human knowledge. The importance of the senses as a source or channel of knowledge seems undeniable. It seems possible, then, to acknowledge their importance and to assess the reliability of that source, quite independently of the difficult question of whether *all* our knowledge comes to us in that way. We would then be assessing the credentials of what is often called our "sensory" or "experiential" or "empirical" knowledge, and that, as we shall see, is quite enough to be going on with.

Having found an extremely important "principle" or source of our knowledge, how can we investigate or assess *all* the knowledge we get from that source? As before, we are faced with the problem of the inexhaustibility of the things we believe on that basis, so no piecemeal, one-by-one procedure will do. But perhaps we can make a sweeping negative assessment. It might seem that as soon as we have found that the senses are one of the sources of our beliefs we are immediately in a position to condemn all putative knowledge derived from them. Some philosophers appear to have reasoned in this way, and many have even supposed that Descartes is among them. The idea is that if I am assessing the reliability of my beliefs and asking whether I really know what I take myself to know, and I come across a large class of beliefs which have come to me through the senses, I can immediately dismiss all those beliefs as unreliable or as not amounting to knowledge because of the obvious fact that I can sometimes be wrong in my beliefs based on the senses. Things are not always as they appear, so if on the basis of the way they appear to me I believe that they

really are a certain way, I might still be wrong. We have all found at one time or another that we have been misled by appearances; we know that the senses are not always reliable. Should we not conclude, then, that as a general source of knowledge the senses are not to be trusted? As Descartes puts it, is it not wiser never "to trust entirely to any thing by which we have once been deceived" (HR, 145)? Don't we have here a quite general way of condemning as not fully reliable *all* of our beliefs acquired by means of the senses?

I think the answer to that question is "No, we do not", and I think Descartes would agree with that answer. It is true that he does talk of the senses "deceiving" us on particular occasions, and he does ask whether that is not enough to condemn the senses in general as a source of knowledge, but he immediately reminds us of the obvious fact that the circumstances in which the senses "deceive" us might be special in certain ascertainable ways, and so their occasional failures would not support a blanket condemnation of their reliability.

Sometimes, to give an ancient example, a tower looks round from a distance when it is actually square. If we relied only on the appearances of the moment we might say that the distant tower is round, and we would be wrong. We also know that there are many small organisms invisible to the naked eye. If the table before me is covered with such organisms at the moment but I look at it and say there is nothing on the table at all, once again I will be wrong. But all that follows from these familiar facts, as Descartes points out, is that there are things about which we can be wrong, or there are situations in which we can get false beliefs, if we rely entirely on our senses at that moment. So sometimes we should be careful about what we believe on the basis of the senses, or sometimes perhaps we should withhold our assent from any statement about how things are – when things are too far away to be seen properly, for example, or too small to be seen at all. But that obviously is not enough to support the policy of never trusting one's senses, or never believing anything based on them. Nor does it show that I can never know anything by means of the senses. If my car starts promptly every morning for two years in temperate weather at sea level but then fails to start one morning in freezing weather at the top of a high mountain, that does not support

the policy of never trusting my car to start again once I return to the temperate lower altitude from which I so foolishly took it. Nor does it show that I can never know whether my car will ever start again. It shows only that there are certain circumstances in which my otherwise fully reliable car might not start. So the fact that we are sometimes wrong or “deceived” in our judgements based on the senses is not enough in itself to show that the senses are never to be trusted and are therefore never reliable as a source of knowledge.

Descartes’s negative assessment of all of his sensory knowledge does not depend on any such reasoning. He starts his investigation, rather, in what would seem to be the most favourable conditions for the reliable operation of the senses as a source of knowledge. While engaging in the very philosophical reflections he is writing about in his *First Meditation* Descartes is sitting in a warm room, by the fire, in a dressing gown, with a piece of paper in his hand. He finds that although he might be able to doubt that a distant tower that looks round really is round, it seems impossible to doubt that he really is sitting there by the fire in his dressing gown with a piece of paper in his hand. The fire and the piece of paper are not too small or too far away to be seen properly, they are right there before his eyes; it seems to be the best kind of position someone could be in for getting reliable beliefs or knowledge by means of the senses about what is going on around him. That is just how Descartes regards it. Its being a best-possible case of that kind is precisely what he thinks enables him to investigate or assess at one fell swoop all our sensory knowledge of the world around us. The verdict he arrives at about his putative knowledge that he is sitting by the fire with a piece of paper in his hand in that particular situation serves as the basis for a completely general assessment of the senses as a source of knowledge about the world around us.

How can that be so? How can he so easily reach a general verdict about all his sensory knowledge on the basis of a single example? Obviously not simply by generalizing from one particular example to all cases of sensory knowledge, as one might wildly leap to a conclusion about all red-haired men on the basis of one or two individuals. Rather, he takes the particular example of his conviction that he is sitting by the fire with a piece of paper in his hand as representative of the best

position any of us can ever be in for knowing things about the world around us on the basis of the senses. What is true of a representative case, if it is truly representative and does not depend on special peculiarities of its own, can legitimately support a general conclusion. A demonstration that a particular isosceles triangle has a certain property, for example, can be taken as a demonstration that all isosceles triangles have that property, as long as the original instance was typical or representative of the whole class. Whether Descartes’s investigation of the general reliability of the senses really does follow that familiar pattern is a difficult question. Whether, or in precisely what sense, the example he considers can be treated as representative of our relation to the world around us is, I believe, the key to understanding the problem of our knowledge of the external world. But if it turns out that there is nothing illegitimate about the way his negative conclusion is reached, the problem will be properly posed.

For the moment I think at least this much can be said about Descartes’s reasoning. He chooses the situation in which he finds himself as representative of the best position we can be in for knowing things about the world in the sense that, if it is impossible for him in that position to know that he is sitting by the fire with a piece of paper in his hand then it is also impossible for him in other situations to know anything about the world around him on the basis of his senses. A negative verdict in the chosen case would support a negative verdict everywhere else. The example Descartes considers is in that sense meant to be the *best* kind of case there could be of sensory knowledge about the world around us. I think we must admit that it is very difficult to see how Descartes or anyone else could be any better off with respect to knowing something about the world around him on the basis of the senses than he is in the case he considers. But if no one could be in any better position for knowing, it seems natural to conclude that any negative verdict arrived at about this example, any discovery that Descartes’s beliefs in this case are not reliable or do not amount to knowledge, could safely be generalized into a negative conclusion about all of our sensory “knowledge” of the world. If candidates with the best possible credentials are found wanting, all those with less impressive credentials must fall short as well.

It will seem at first sight that in conceding that the whole question turns on whether Descartes knows in this particular case we are conceding very little; it seems obvious that Descartes on that occasion does know what he thinks he knows about the world around him. But in fact Descartes finds that he cannot know in this case that he is sitting by the fire with a piece of paper in his hand. If the case is truly representative of our sensory knowledge in general, that will show that no one can know anything about the world around us. But how could he ever arrive at that negative verdict in the particular case he considers? How could anyone possibly doubt in such a case that the fire and the piece of paper are there? The paper is in Descartes's hand, the fire is right there before his open eyes, and he feels its warmth. Wouldn't anyone have to be mad to deny that he can know something about what is going on around him in those circumstances? Descartes first answers "Yes". He says that if he were to doubt or deny on that occasion that he is sitting by the fire with a piece of paper in his hand he would be no less mad than those paupers who say they are kings or those madmen who think they are pumpkins or are made of glass. But his reflections continue:

At the same time I must remember that I am a man, and that consequently I am in the habit of sleeping, and in my dreams representing to myself the same things or sometimes even less probable things, than do those who are insane in their waking moments. How often has it happened to me that in the night I dreamt that I found myself in this particular place, that I was dressed and seated near the fire, whilst in reality I was lying undressed in bed! At this moment it does indeed seem to me that it is with eyes awake that I am looking at this paper; that this head which I move is not asleep, that it is deliberately and of set purpose that I extend my hand and perceive it; what happens in sleep does not appear so clear nor so distinct as does all this. But in thinking over this I remind myself that on many occasions I have in sleep been deceived by similar illusions, and in dwelling carefully on this reflection I see so manifestly that there are no certain indications by which we may clearly distinguish wakefulness from sleep that I am lost in astonishment. And my astonishment is such that it is almost capable of persuading me that I now dream. (HR, 145–6)

With this thought, if he is right, Descartes has lost the whole world. He knows what he is experiencing, he knows how things appear to him, but he does not know whether he is in fact sitting by the fire with a piece of paper in his hand. It is, for him, exactly as if he were sitting by the fire with a piece of paper in his hand, but he does not know whether there really is a fire or a piece of paper there or not; he does not know what is really happening in the world around him. He realizes that if everything he can ever learn about what is happening in the world around him comes to him through the senses, but he cannot tell by means of the senses whether or not he is dreaming, then all the sensory experiences he is having are compatible with his merely dreaming of a world around him while in fact that world is very different from the way he takes it to be. That is why he thinks he must find some way to tell that he is not dreaming. Far from its being mad to deny that he knows in this case, he thinks his recognition of the possibility that he might be dreaming gives him "very powerful and maturely considered" (HR, 148) reasons for withholding his judgement about how things are in the world around him. He thinks it is eminently reasonable to insist that if he is to know that he is sitting by the fire he must know that he is not dreaming that he is sitting by the fire. That is seen as a necessary condition of knowing something about the world around him. And he finds that that condition cannot be fulfilled. On careful reflection he discovers that "there are no certain indications by which we may clearly distinguish wakefulness from sleep". He concludes that he knows nothing about the world around him because he cannot tell that he is not dreaming; he cannot fulfil one of the conditions necessary for knowing something about the world.

The Cartesian problem of our knowledge of the external world therefore becomes: how can we know anything about the world around us on the basis of the senses if the senses give us only what Descartes says they give us? What we gain through the senses is on Descartes's view only information that is compatible with our dreaming things about the world around us and not knowing anything about the world. How then can we know anything about the world by means of the senses? The Cartesian argument presents a challenge to our knowledge, and the problem of our knowledge of the external world is to show how that challenge can be met.

When I speak here of the Cartesian argument or of Descartes's sceptical conclusion or of his negative verdict about his knowledge I refer of course only to the position he finds himself in by the end of his *First Meditation*. Having at that point discovered and stated the problem of the external world, Descartes goes on in the rest of his *Meditations* to try to solve it, and by the end of the *Sixth Meditation* he thinks he has explained how he knows almost all those familiar things he began by putting in question. So when I ascribe to Descartes the view that we can know nothing about the world around us I do not mean to suggest that that is his final and considered view; it is nothing more than a conclusion he feels almost inevitably driven to at the early stages of his reflections. But those are the only stages of his thinking I am interested in here. That is where the philosophical problem of our knowledge of the external world gets posed, and before we can consider possible solutions we must be sure we understand exactly what the problem is.

I have described it as that of showing or explaining how knowledge of the world around us is possible by means of the senses. It is important to keep in mind that that demand for an explanation arises in the face of a challenge or apparent obstacle to our knowledge of the world. The possibility that he is dreaming is seen as an obstacle to Descartes's knowing that he is sitting by the fire, and it must be explained how that obstacle can either be avoided or overcome. It must be shown or explained *how* it is possible for us to know things about the world, given that the sense-experiences we get are compatible with our merely dreaming. Explaining how something is nevertheless possible, despite what looks like an obstacle to it, requires more than showing merely that there is no impossibility involved in the thing – that it is consistent with the principles of logic and the laws of nature and so in that sense *could* exist. The mere possibility of the state of affairs is not enough to settle the question of how our knowledge of the world is possible; we must understand how the apparent obstacle is to be got round.

Descartes's reasoning can be examined and criticized at many different points, and has been closely scrutinized by many philosophers for centuries. It has also been accepted by many, perhaps by more than would admit or even realize that

they accept it. There seems to me no doubt about the force and the fascination – I would say the almost overwhelming persuasiveness – of his reflections. That alone is something that needs accounting for. I cannot possibly do justice to all reasonable reactions to them here. In the rest of this chapter I want to concentrate on deepening and strengthening the problem and trying to locate more precisely the source of its power.

There are at least three distinct questions that could be pressed. Is the possibility that Descartes might be dreaming really a threat to his knowledge of the world around him? Is he right in thinking that he must know that he is not dreaming if he is to know something about the world around him? And is he right in his "discovery" that he can never know that he is not dreaming? If Descartes were wrong on any of these points it might be possible to avoid the problem and perhaps even to explain without difficulty how we know things about the world around us.

On the first question, it certainly seems right to say that if Descartes were dreaming that he is sitting by the fire with a piece of paper in his hand he would not then know that he is sitting by the fire with a piece of paper in his hand. When you dream that something is going on in the world around you you do not thereby know that it is. Most often, of course, what we dream is not even true; no one is actually chasing us when we are lying asleep in bed dreaming, nor are we actually climbing stairs. But although usually what we dream is not really so, that is not the real reason for our lack of knowledge. Even if Descartes were in fact sitting by the fire and actually had a piece of paper in his hand at the very time he was dreaming that he is sitting by the fire with a piece of paper in his hand, he would not thereby know he was sitting there with that paper. He would be like a certain Duke of Devonshire who, according to G. E. Moore, once dreamt he was speaking in the House of Lords and woke up to find that he *was* speaking in the House of Lords.<sup>4</sup> What he was dreaming was in fact so. But even if what you are dreaming is in fact so you do not thereby know that it is. Even if we allow that when you are dreaming that something is so you can be said, at least for the time being, to think or to believe that it is so, there is still no real connection between your thinking or believing what you do and its being so. At best you have a thought or a belief

which just happens to be true, but that is no more than coincidence and not knowledge. So Descartes's first step relies on what seems to be an undeniable fact about dreams: if you are dreaming that something is so you do not thereby know that it is so.

This bald claim needs to be qualified and more carefully explained, but I do not think that will diminish the force of the point for Descartes's purposes. Sometimes what is going on in the world around us has an effect on what we dream; for example, a banging shutter might actually cause me to dream, among other things, that a shutter is banging. If my environment affects me in that way, and if in dreams I can be said to think or believe that something is so, would I not in that case know that a shutter is banging? It seems to me that I would not, but I confess it is difficult to say exactly why I think so. That is probably because it is difficult to say exactly what is required for knowledge. We use the term "know" confidently, we quite easily distinguish cases of knowledge from cases of its absence, but we are not always in a position to state what we are going on in applying or withholding the term in the ways we do. I think that in the case of the banging shutter it would not be knowledge because I would be *dreaming*, I would not even be awake. At least it can be said, I think, that even if Descartes's sitting by the fire with a piece of paper in his hand (like the banging shutter) is what in fact causes him to dream that he is sitting by the fire with a piece of paper in his hand, that is still no help to him in coming to know what is going on in the world around him. He realizes that he could be dreaming that he is sitting by the fire even if he is in fact sitting there, and that is the possibility he finds he has to rule out.

I have said that if you are dreaming that something is so you do not thereby know that it is so, and it might seem as if that is not always true. Suppose a man and a child are both sleeping. I say of the child that it is so young it does not know what seven times nine is, whereas the grown man does know that. If the man happens at that very moment to be dreaming that seven times nine is sixty-three (perhaps he is dreaming that he is computing his income tax), then he is a man who is dreaming that something is so and also knows that it is so. The same kind of thing is possible for knowledge about the world around him. He

might be a physicist who knows a great deal about the way things are which the child does not know. If the man also dreams that things are that way he can once again be said to be dreaming that something is so and also to know that it is so. There is therefore no incompatibility between dreaming and knowing. That is true, but I do not think it affects Descartes's argument. He is led to consider how he knows he is not dreaming at the moment by reflecting on how he knows at that moment that he is sitting by the fire with a piece of paper in his hand. If he knows that at all, he thinks, he knows it on the basis of the senses. But he realizes that his having the sensory experiences he is now having is compatible with his merely dreaming that he is sitting by the fire with a piece of paper in his hand. So he does not know on the basis of the sensory experiences he is having at the moment that he is sitting by the fire. Nor, of course, did the man in my examples know the things he was said to know on the basis of the sensory experiences he was having at that moment. He knew certain things to be so, and he was dreaming those things to be so, but in dreaming them he did not *thereby* know them to be so.

But as long as we allow that the sleeping man does know certain things about the world around him, even if he does not know them on the basis of the very dreams he is having at the moment, isn't that enough to show that Descartes must nevertheless be wrong in his conclusion that no one can know anything about the world around him? No. It shows at most that we were hasty or were ignoring Descartes's conclusion in conceding that someone could know something about the world around him. If Descartes's reasoning is correct the dreaming physicist, even when he is awake, does not really know any of the things we were uncritically crediting him with knowing about the way things are – or at least he does not know them on the basis of the senses. In order to know them on the basis of the senses there would have to have been at least some time at which he knew something about what was going on around him at that time. But if Descartes is right he could not have known any such thing unless he had established that he was not dreaming at that time; and according to Descartes he could never establish that. So the fact about dreams that Descartes relies on – that one who dreams that something is so does not thereby know that it is so – is enough

to yield his conclusion if the other steps of his reasoning are correct.

When he first introduces the possibility that he might be dreaming Descartes seems to be relying on some knowledge about how things are or were in the world around him. He says “I remind myself that on many occasions I have in sleep been deceived by similar illusions”, so he seems to be relying on some knowledge to the effect that he has actually dreamt in the past and that he remembers having been “deceived” by those dreams. That is more than he actually needs for his reflections about knowledge to have the force he thinks they have. He does not need to support his judgement that he has actually dreamt in the past. The only thought he needs is that it is now *possible* for him to be dreaming that he is sitting by the fire, and that if that possibility were realized he would not know that he is sitting by the fire. Of course it was no doubt true that Descartes had dreamt in the past and that his knowledge that he had done so was partly what he was going on in acknowledging the possibility of his dreaming on this particular occasion. But neither the fact of past dreams nor knowledge of their actual occurrence would seem to be strictly required in order to grant what Descartes relies on – the possibility of dreaming, and the absence of knowledge if that possibility were realized. The thought that he *might* be dreaming that he is sitting by the fire with a piece of paper in his hand, and the fact that if he were he wouldn’t know he was sitting there, is what gives Descartes pause. That would worry him in the way it does even if he had never actually had any dreams exactly like it in the past – if he had never dreamt about fires and pieces of paper at all. In fact, I think he need never have actually dreamt of anything before, and certainly needn’t know that he ever has, in order to be worried in the way he is by the thought that he might be dreaming now.

The fact that the possibility of dreaming is all Descartes needs to appeal to brings out another truth about dreams that his argument depends on – that anything that can be going on or that one can experience in one’s waking life can also be dreamt about. This again is only a statement of possibility – no sensible person would suggest that we *do* at some time dream of everything that actually happens to us, or that everything we dream about does in fact happen sometime. But

it is very plausible to say that there is nothing we *could* not dream about, nothing that could be the case that we *could* not dream to be the case. I say it is very plausible; of course I cannot prove it to be true. But even if it is not true with complete generality, we must surely grant that it is possible to dream that one is sitting by a fire with a piece of paper in one’s hand, and possible to dream of countless other equally obvious and equally mundane states of affairs as well, and those possibilities are what Descartes sees as threatening to his knowledge of the world around him.

There seems little hope, then, of objecting that it is simply not possible for Descartes to dream that he is sitting by the fire with a piece of paper in his hand. Nor is it any more promising to say that even if he were dreaming it would not follow that he did not know that he was sitting there. I think both those steps or assumptions of Descartes’s reasoning are perfectly correct, and further defence of them at this stage is unnecessary. If his argument and the problem to which it gives rise are to be avoided, it might seem that the best hope is therefore to accept his challenge and show that it can be met. That would be in effect to argue that Descartes’s alleged “discovery” is no discovery at all: we *can* sometimes know that we are not dreaming.

This can easily seem to be the most straightforward and most promising strategy. It allows that Descartes is right in thinking that knowing that one is not dreaming is a condition of knowing something about the world around us, but wrong in thinking that that condition can never be met. And that certainly seems plausible. Surely it is not impossible for me to know that I am not dreaming? Isn’t that something I often know, and isn’t it something I can sometimes find out if the question arises? If it is, then the fact that I must know that I am not dreaming if I am to know anything about the world around me will be no threat to my knowledge of the world.

However obvious and undeniable it might be that we often do know that we are not dreaming, I think this straightforward response to Descartes’s challenge is a total failure. In calling it straightforward I mean that it accepts Descartes’s conditions for knowledge of the world and tries to show that they can be fulfilled. That is what I think cannot be done. To put the same point in another way: I think Descartes would be perfectly correct in

saying “there are no certain indications by which we may clearly distinguish wakefulness from sleep”, and so we could never tell we are not dreaming, *if* he were also right that knowing that one is not dreaming is a condition of knowing something about the world around us. That is why I think one cannot accept that condition and then go on to establish that one is not dreaming. I do not mean to be saying simply that Descartes is right – that we can never know that we are not dreaming. But I do want to argue that either we can never know that we are not dreaming or else what Descartes says is a condition of knowing things about the world is not really a condition in general of knowing things about the world. The straightforward strategy denies both alternatives. I will try to explain why I think we must accept one alternative or the other.

When Descartes asks himself how he knows that he is sitting by the fire with a piece of paper in his hand why does he immediately go on to ask himself how he knows he is not dreaming that he is sitting by the fire with a piece of paper in his hand? I have suggested that it is because he recognizes that if he were dreaming he would not know on the basis of his senses at the moment that he is sitting there, and so he thinks he must know that that possibility does not obtain if he is to know that he is in fact sitting there. But this particular example was chosen, not for any peculiarities it might be thought to possess, but because it could be taken as typical of the best position we can ever be in for coming to know things about the world around us on the basis of the senses. What is true of this case that is relevant to Descartes’s investigation of knowledge is supposed to be true of all cases of knowledge of the world by means of the senses; that is why the verdict arrived at here can be taken to be true of our sensory knowledge generally. But what Descartes thinks is true of this particular case of sensory knowledge of the world is that he must know he is not dreaming if he is to know that he is sitting by the fire with a piece of paper in his hand. That is required, not because of any peculiarities of this particular case, but presumably because, according to Descartes, it is a necessary condition of any case – even a best possible case – of knowledge of the world by means of the senses. That is why I ascribed to Descartes the quite general thesis that knowing that one is not dreaming is a condition

of knowing something about the world around us on the basis of the senses. Since he thinks the possibility of his dreaming must be ruled out in the case he considers, and the case he considers is regarded as typical and without special characteristics of its own, he thinks that the possibility that he is dreaming must be ruled out in every case of knowing something about the world by means of the senses.

If that really is a condition of knowing something about the world, I think it can be shown that Descartes is right in holding that it can never be fulfilled. That is what the straightforward response denies, and that is why I think that response must be wrong. We cannot accept the terms of Descartes’s challenge and then hope to meet it.

Suppose Descartes tries to determine that he is not dreaming in order to fulfil what he sees as a necessary condition of knowing that he is sitting by the fire with a piece of paper in his hand. How is he to proceed? He realizes that his seeing his hand and seeing and feeling a piece of paper before him and feeling the warmth of the fire – in fact his getting all the sensory experiences or all the sensory information he is then getting – is something that could be happening even if he were dreaming. To establish that he is not dreaming he would therefore need something more than just those experiences or that information alone. He would also need to know whether those experiences and that information are reliable, not merely dreamt. If he could find some operation or test, or if he could find some circumstance or state of affairs, that indicated to him that he was not dreaming, perhaps he could then fulfil the condition – he could know that he was not dreaming. But how could a test or a circumstance or a state of affairs indicate to him that he is not dreaming *if* a condition of knowing *anything* about the world is that he know he is not dreaming? It could not. He could never fulfil the condition.

Let us suppose that there is in fact some test which a person can perform successfully only if he is not dreaming, or some circumstance or state of affairs which obtains only if that person is not dreaming. Of course for that test or state of affairs to be of any use to him Descartes would have to know of it. He would have to know that there is such a test or that there is a state of affairs that shows that he is not dreaming; without such



information he would be no better off for telling that he is not dreaming than he would be if there were no such test or state of affairs at all. To have acquired that information he would at some time have to have known more than just something about the course of his sensory experience, since the connection between the performance of a certain test, or between a certain state of affairs, and someone's not dreaming is not itself just a fact about the course of that person's sensory experience; it is a fact about the world beyond his sensory experiences. Now strictly speaking if it is a condition of knowing *anything* about the world beyond one's sensory experiences that one know that one is not dreaming, there is an obvious obstacle to Descartes's ever having got the information he needs about that test or state of affairs. He would have to have known at some time that he was not dreaming in order to get the information he needs to tell at *any* time that he is not dreaming – and that cannot be done.

But suppose we forget about this difficulty and concede that Descartes does indeed know (somehow) that there is a test or circumstance or state of affairs that unfailingly indicates that he is not dreaming. Still, there is an obstacle to his ever using that test or state of affairs to tell that he is not dreaming and thereby fulfilling the condition for knowledge of the world. The test would have to be something he could know he had performed successfully, the state of affairs would have to be something he could know obtains. If he completely unwittingly happened to perform the test, or if the state of affairs happened to obtain but he didn't know that it did, he would be in no better position for telling whether he was dreaming than he would be if he had done nothing or did not even know that there was such a test. But how is he to know that the test has been performed successfully or that the state of affairs in question does in fact obtain? Anything one can experience in one's walking life can also be dreamt about; it is possible to dream that one has performed a certain test or dream that one has established that a certain state of affairs obtains. And, as we have seen, to dream that something about the world around you is so is not thereby to know that it is so. In order to know that his test has been performed or that the state of affairs in question obtains Descartes would therefore have to establish that he is not merely dreaming that he performed the test

successfully or that he established that the state of affairs obtains. How could that in turn be known? Obviously the particular test or state of affairs already in question cannot serve as a guarantee of its own authenticity, since it might have been merely dreamt, so some further test or state of affairs would be needed to indicate that the original test was actually performed and not merely dreamt, or that the state of affairs in question was actually ascertained to obtain and not just dreamt to obtain. But this further test or state of affairs is subject to the same general condition in turn. *Every* piece of knowledge that goes beyond one's sensory experiences requires that one know one is not dreaming. This second test or state of affairs will therefore be of use only if Descartes knows that he is not merely dreaming that he is performing or ascertaining it, since merely to dream that he had established the authenticity of the first test is not to have established it. And so on. At no point can he find a test for not dreaming which he can know has been successfully performed or a state of affairs correlated with not dreaming which he can know obtains. He can therefore never fulfil what Descartes says is a necessary condition of knowing something about the world around him. He can never know that he is not dreaming.

I must emphasize that this conclusion is reached *only* on the assumption that it is a condition of knowing anything about the world around us on the basis of the senses that we know we are not dreaming that the thing is so. I think it is his acceptance of that condition that leads Descartes to "see so manifestly that there are no certain indications by which we may clearly distinguish wakefulness from sleep". And I think Descartes is absolutely right to draw that conclusion, *given* what he thinks is a condition of knowledge of the world. But all I have argued on Descartes's behalf (he never spells out his reasoning) is that we cannot both accept that condition of knowledge and hope to fulfil it, as the straightforward response hopes to do. And of course if one of the necessary conditions of knowledge of the world can never be fulfilled, knowledge of the world around us will be impossible.

I think we have now located Descartes's reason for his negative verdict about sensory knowledge in general. If we agree that he must know that he is not dreaming if he is to know in his particular

case that he is sitting by the fire with a piece of paper in his hand, we must also agree that we can know nothing about the world around us.

Once we recognize that the condition Descartes takes as necessary can never be fulfilled if he is right in thinking it is indeed necessary, we are naturally led to the question whether Descartes is right. Is it really a condition of knowing something about the world that one know one is not dreaming? That is the second of the three questions I distinguished. It is the one that has received the least attention. In asking it now I do not mean to be going back on something I said earlier was undeniably true, viz., that if one is dreaming that something about the world is so one does not thereby know that it is so. That still seems to me undeniable, but it is not the same as Descartes's assumption that one must know that one is not dreaming if one is to know something about the world. The undeniable truth says only that you lack knowledge if you are dreaming; Descartes says that you lack knowledge if you don't know that you are not dreaming. Only with the stronger assumption can his sceptical conclusion be reached.

Is that assumption true? In so far as we find Descartes's reasoning convincing, or even plausible, I think it is because we too on reflection find that it is true. I said that not much attention had been paid to that particular part of Descartes's reasoning, and I think that too is because, as he presents it, the step seems perfectly convincing and so only other parts of the argument appear vulnerable. Why is that so? Is it because Descartes's assumption is indeed true? Is there anything we can do that would help us determine whether it is true or not? The question is important because I have argued so far that *if* it is true we can never know anything about the world around us on the basis of the senses, and philosophical scepticism about the external world is correct. We would have to find that conclusion as convincing or as plausible as we find the assumption from which it is derived.

Given our original favourable response to Descartes's reasoning, then, it can scarcely be denied that what I have called his assumption or condition *seems* perfectly natural to insist on. Perhaps it seems like nothing more than an instance of a familiar commonplace about knowledge. We are all aware that, even in the most ordinary

circumstances when nothing very important turns on the outcome, we cannot know a particular thing unless we have ruled out certain possibilities that we recognize are incompatible with our knowing that thing.

Suppose that on looking out the window I announce casually that there is a goldfinch in the garden. If I am asked how I know it is a goldfinch and I reply that it is yellow, we all recognize that in the normal case that is not enough for knowledge. "For all you've said so far," it might be replied, "the thing could be a canary, so how do you know it's a goldfinch?" A certain possibility compatible with everything I have said so far has been raised, and if what I have said so far is all I have got to go on and I don't know that the thing in the garden is not a canary, then I do not know that there is a goldfinch in the garden. I must be able to rule out the possibility that it is a canary if I am to know that it is a goldfinch. Anyone who speaks about knowledge and understands what others say about it will recognize this fact or condition in particular cases.

In this example what is said to be possible is something incompatible with the truth of what I claim to know – if that bird were a canary it would not be a goldfinch in the garden, but a canary. What I believe in believing it is a goldfinch would be false. But that is not the only way a possibility can work against my knowledge. If I come to suspect that all the witnesses have conspired and made up a story about the man's being in Cleveland that night, for example, and their testimony is all I have got to go on in believing that he was in Cleveland, I might find that I no longer know whether he was there or not until I have some reason to rule out my suspicion. If their testimony were all invented I would not know that the man was in Cleveland. But strictly speaking his being in Cleveland is not incompatible with their making up a story saying he was. They might have invented a story to protect him, whereas in fact, unknown to them, he was there all the time. Such a complicated plot is not necessary to bring out the point; Moore's Duke of Devonshire is enough. From the fact that he was dreaming that he was speaking in the House of Lords it did not follow that he was not speaking in the House of Lords. In fact he was. The possibility of dreaming – which was actual in that case – did not imply the falsity of what was believed. A possible deficiency in the

basis of my belief can interfere with my knowledge without itself rendering false the very thing I believe. A hallucinogenic drug might cause me to see my bed covered with a huge pile of leaves, for example.<sup>5</sup> Having taken that drug, I will know the actual state of my bed only if I know that what I see is not just the effect of the drug; I must be able to rule out the possibility that I am hallucinating the bed and the leaves. But however improbable it might be that my bed is actually covered with leaves, its not being covered with leaves does not follow from the fact that I am hallucinating that it is. What I am hallucinating could nevertheless be (unknown to me) true. But a goldfinch simply could not be a canary. So although there are two different ways in which a certain possibility can threaten my knowledge, it remains true that there are always certain possibilities which must be known not to obtain if I am to know what I claim to know.

I think these are just familiar facts about human knowledge, something we all recognize and abide by in our thought and talk about knowing things. We know what would be a valid challenge to a claim to know something, and we can recognize the relevance and force of objections made to our claims to know. The question before us is to what extent Descartes's investigation of his knowledge that he is sitting by the fire with a piece of paper in his hand follows these recognized everyday procedures for assessing claims to know. If it does follow them faithfully, and yet leads to the conclusion that he cannot know where he is or what is happening around him, we seem forced to accept his negative conclusion about knowledge in general just as we are forced to accept the conclusion that I do not know it is a goldfinch or do not know the witness was in Cleveland because I cannot rule out the possibilities which must be ruled out if I am to know such things. Is Descartes's introduction of the possibility that he might be dreaming just like the introduction of the possibility that it might be a canary in the garden or that the alibi might be contrived or that it might be a hallucination of my bed covered with leaves?

Those possibilities were all such that if they obtained I did not know what I claimed to know, and they had to be known not to obtain in order for the original knowledge-claim to be true. Does Descartes's dream-possibility fulfil both of those

conditions? I have already said that it seems undeniable that it fulfils the first. If he *were* dreaming Descartes would not know what he claims to know. Someone who is dreaming does not thereby know anything about the world around him even if the world around him happens to be just the way he dreams or believes it to be. So his dreaming *is* incompatible with his knowing. But does it fulfil the second condition? Is it a possibility which must be known not to obtain if Descartes is to know that he is sitting by the fire with a piece of paper in his hand? I think it is difficult simply to deny that it is. The evident force of Descartes's reasoning when we first encounter it is enough to show that it certainly strikes us as a relevant possibility, as something that he should know not to obtain if he is to know where he is and what is happening around him.

When that possibility strikes us as obviously relevant in Descartes's investigation we might come to think that it is because of a simple and obvious fact about knowledge. In the case of the goldfinch we immediately recognize that I must know that it is not a canary if I am to know it is a goldfinch. And it is very natural to think that that is simply because its being a canary is incompatible with its being a goldfinch. If it were a canary it would not be a goldfinch, and I would therefore be wrong in saying that it is; so if I am to know it is a goldfinch I must rule out the possibility that it is a canary. The idea is that the two conditions I distinguished in the previous paragraph are not really separate after all. As soon as we see that a certain possibility is incompatible with our knowing such-and-such, it is suggested, we immediately recognize that it is a possibility that must be known not to obtain if we are to know the such-and-such in question. We see that the dream-possibility satisfies that first condition in Descartes's case (if he were dreaming, he wouldn't know), and that is why, according to this suggestion, we immediately see that it is relevant and must be ruled out. Something we all recognize about knowledge is what is said to make that obvious to us.

But is the "simple and obvious fact about knowledge" appealed to in this explanation really something that is true of human knowledge even in the most ordinary circumstances? What exactly is the "fact" in question supposed to be? I have described it so far, as applied to the case of the

goldfinch, as the fact that if I know something  $p$  (it's a goldfinch) I must know the falsity of all those things incompatible with  $p$  (e.g., it's a canary). If there were one of those things that I did not know to be false, and it were in fact true, I would not know that  $p$ , since in that case something incompatible with  $p$  would be true and so  $p$  would not be true. But to say that I must know that all those things incompatible with  $p$  are false is the same as saying that I must know that truth of all those things that must be true if  $p$  is true. And it is extremely implausible to say that that is a "simple and obvious fact" we all recognize about human knowledge.

The difficulty is that there are no determinate limits to the number of things that follow from the things I already know. But it cannot be said that I now know all those indeterminately many things, although they all must be true if the things that I already know are true. Even granting that I now know a great deal about a lot of different things, my knowledge obviously does not extend to everything that follows from what I now know. If it did, mathematics, to take only one example, would be a great deal easier than it is – or else impossibly difficult. In knowing the truth of the simple axioms of number theory, for example, I would thereby know the truth of everything that follows from them; every theorem of number theory would already be known. Or, taking the pessimistic side, since obviously no one does know all the theorems of number theory, it would follow that no one even knows that those simple axioms are true.

It is absurd to say that we enjoy or require such virtual omniscience, so it is more plausible to hold that the "simple and obvious fact" we all recognize about knowledge is the weaker requirement that we must know the falsity of all those things that we *know* to be incompatible with the things we know. I know that a bird's being a canary is incompatible with its being a goldfinch; that is not some farflung, unknown consequence of its being a goldfinch, but something that anyone would know who knew anything about goldfinches at all. And the idea is that that is why I must know that it is not a canary if I am to know that it is a goldfinch. Perhaps, in order to know something,  $p$ , I do not need to know the falsity of all those things that are incompatible with  $p$ , but it can seem that at least I must know the falsity of

all those things that I *know* to be incompatible with  $p$ . Since I claim to know that the bird is a goldfinch, and I know that its being a goldfinch implies that it is not a canary, I must for that reason know that it is not a canary if my original claim is true. In claiming to know it is a goldfinch I was, so to speak, committing myself to knowing that it is not a canary, and I must honour my commitments.

This requirement as it stands, even if it does explain why I must know that the bird is not a canary, does not account for the relevance of the other sorts of possibilities I have mentioned. The reason in the goldfinch case was said to be that I know that its being a canary is incompatible with its being a goldfinch. But that will not explain why I must rule out the possibility that the witnesses have invented a story about the man's being in Cleveland, or the possibility that I am hallucinating my bed covered with a pile of leaves. Nor will it explain why Descartes must rule out the possibility that he is dreaming. What I claimed to know in the first case is that the man was in Cleveland that night. But, as we saw earlier, it is not a consequence of his being in Cleveland that no one will invent a story to the effect that he was in Cleveland; they might mistakenly believe he was not there and then tell what they think is a lie. Nor is it a consequence of my bed's being covered with leaves that I am not hallucinating that it is. But we recognize that in order to know in those cases I nevertheless had to rule out those possibilities. Similarly, as the Duke of Devonshire reminds us, it is not a consequence of Descartes's sitting by the fire with a piece of paper in his hand that he is not dreaming that he is. So if it is obvious to us that Descartes must know that he is not dreaming if he is to know that he is sitting by the fire, it cannot be simply because the possibility in question is known to be incompatible with what he claims to know. It is not.

If there is some "simple and obvious fact about knowledge" that we recognize and rely on in responding to Descartes's reasoning it must therefore be more complicated than what has been suggested so far. Reflecting even on the uncontroversial everyday examples alone can easily lead us to suppose that it is something like this: if somebody knows something,  $p$ , he must know the falsity of all those things incompatible with his knowing that  $p$  (or perhaps all those things he

knows to be incompatible with his knowing that *p*). I will not speculate further on the qualifications or emendations needed to make the principle less implausible. The question now is whether it is our adherence to any such principle or requirement that is responsible for our recognition that the possibility that the bird is a canary or the possibility that the witnesses made up a story must be known not to obtain if I am to know the things I said I knew in those cases. What exactly are the procedures or standards we follow in the most ordinary, humdrum cases of putative knowledge? Reflection on the source of Descartes's sceptical reasoning has led to difficulties in describing and therefore in understanding even the most familiar procedures we follow in everyday life. That is one of the rewards of a study of philosophical scepticism.

The main difficulty in understanding our ordinary procedures is that no principle like those I have mentioned could possibly describe the way we proceed in everyday life. Or, to put it less dogmatically, if our adherence to some such requirement were responsible for our reactions in those ordinary cases, Descartes would be perfectly correct, and philosophical scepticism about the external world would be true. Nobody would know anything about the world around us. If, in order to know something, we must rule out a possibility which is known to be incompatible with our knowing it, Descartes is perfectly right to insist that he must know that he is not dreaming if he is to know that he is sitting by the fire with a piece of paper in his hand. He knows his dreaming is incompatible with his knowing. I have already argued that if he is right in insisting that that condition must be fulfilled for knowledge of the world around us he is also right in concluding that it can never be fulfilled; fulfilling it would require knowledge which itself would be possible only if the condition were fulfilled. So both steps of Descartes's reasoning would be valid and his conclusion would be true.

That conclusion can be avoided, it seems to me, only if we can find some way to avoid the requirement that we must know we are not dreaming if we are to know anything about the world around us. But that requirement cannot be avoided if it is nothing more than an instance of a general procedure we recognize and insist on in making and assessing knowledge-claims in every-

day and scientific life. We have no notion of knowledge other than what is embodied in those procedures and practices. So if that requirement is a "fact" of our ordinary conception of knowledge we will have to accept the conclusion that no one knows anything about the world around us.

I now want to say a few more words about the position we would all be in if Descartes's conclusion as he understands it were correct. I described him earlier as having lost the whole world, as knowing at most what he is experiencing or how things appear to him, but knowing nothing about how things really are in the world around him. To show how anyone in that position could come to know anything about the world around him is what I am calling the problem of our knowledge of the external world, and it is worth dwelling for a moment on just how difficult a problem that turns out to be if it has been properly raised.

If we are in the predicament Descartes finds himself in at the end of his *First Meditation* we cannot tell by means of the senses whether we are dreaming or not; all the sensory experiences we are having are compatible with our merely dreaming of a world around us while that world is in fact very different from the way we take it to be. Our knowledge is in that way confined to our sensory experiences. There seems to be no way of going beyond them to know that the world around us really is this way rather than that. Of course we might have very strongly held beliefs about the way things are. We might even be unable to get rid of the conviction that we are sitting by the fire holding a piece of paper, for example. But if we acknowledge that our sensory experiences are all we ever have to go on in gaining knowledge about the world, and we acknowledge, as we must, that given our experiences as they are we could nevertheless be simply dreaming of sitting by the fire, we must concede that we do not know that we are sitting by the fire. Of course, we are in no position to claim the opposite either. We cannot conclude that we are not sitting by the fire; we simply cannot tell which is the case. Our sensory experience gives us no basis for believing one thing about the world around us rather than its opposite, but our sensory experience is all we have got to go on. So whatever unshakeable conviction we might nevertheless retain, that conviction cannot be knowledge. Even if we are in fact

holding a piece of paper by the fire, so that what we are convinced of is in fact true, that true conviction is still not knowledge. The world around us, whatever it might be like, is in that way beyond our grasp. We can know nothing of how it is, no matter what convictions, beliefs, or opinions we continue, perhaps inevitably, to hold about it.

What *can* we know in such a predicament? We can perhaps know what sensory experiences we are having, or how things seem to us to be. At least that much of our knowledge will not be threatened by the kind of attack Descartes makes on our knowledge of the world beyond our experiences. What we can know turns out to be a great deal less than we thought we knew before engaging in that assessment of our knowledge. Our position is much more restricted, much poorer, than we had originally supposed. We are confined at best to what Descartes calls “ideas” of things around us, representations of things or states of affairs which, for all we can know, might or might not have something corresponding to them in reality. We are in a sense imprisoned within those representations, at least with respect to our knowledge. Any attempt to go beyond them to try and tell whether the world really is as they represent it to be can yield only more representations, more deliverances of sense experience which themselves are compatible with reality’s being very different from the way we take it to be on the basis of our sensory experiences. There is a gap, then, between the most that we can ever find out on the basis of our sensory experience and the way things really are. In knowing the one we do not thereby know the other.

This can seem to leave us in the position of finding a barrier between ourselves and the world around us. There would then be a veil of sensory experiences or sensory objects which we could not penetrate but which would be no reliable guide to the world beyond the veil. If we were in such a position, I think it is quite clear that we could not know what is going on beyond the veil. There would be no possibility of our getting reliable sensory information about the world beyond the veil; all such reports would simply be more representations, further ingredients of the evermore-complicated veil. We would know nothing but the veil itself. We would be in the position of someone waking up to find himself locked in a room full of television sets and trying to find out what is going

on in the world outside. For all he can know, whatever is producing the patterns he can see on the screens in front of him might be something other than well-function cameras directed on to the passing show outside the room. The victim might switch on more of the sets in the room to try to get more information, and he might find that some of the sets show events exactly similar or coherently related to those already visible on the screens he can see. But all those pictures will be no help to him without some independent information, some knowledge which does not come to him from the pictures themselves, about how the pictures he does see before him are connected with what is going on outside the room. The problem of the external world is the problem of finding out, or knowing how we could find out, about the world around us if we were in that sort of predicament. It is perhaps enough simply to put the problem this way to convince us that it can never be given a satisfactory solution.

But putting the problem this way, or only this way, has its drawbacks. For one thing, it encourages a facile dismissive response; not a solution to the problem as posed, but a rejection of it. I do not mean that we should not find a way to reject the problem – I think that is our only hope – but this particular response, I believe, is wrong, or at the very least premature. It is derived almost entirely from the perhaps overly dramatic description of the predicament I have just given.

I have described Descartes’s sceptical conclusion as implying that we are permanently sealed off from a world we can never reach. We are restricted to the passing show on the veil of perception, with no possibility of extending our knowledge to the world beyond. We are confined to appearances we can never know to match or to deviate from the imperceptible reality that is forever denied us. This way of putting it naturally encourages us to minimize the seriousness of the predicament, to try to settle for what is undeniably available to us, or perhaps even to argue that nothing that concerns us or makes human life worthwhile has been left out.

If an imperceptible “reality”, as it is called on this picture, is forever inaccessible to us, what concern can it be of ours? How can something we can have no contact with, something from which we are permanently sealed off, even make sense to us at all? Why should we be distressed by an

alleged limitation of our knowledge if it is not even possible for the “limitation” to be overcome? If it makes no sense to aspire to anything beyond what is possible for us, it will seem that we should give no further thought to this allegedly imperceptible “reality”. Our sensory experiences, past, present, and future, will then be thought to be all we are or should be concerned with, and the idea of a “reality” lying beyond them necessarily out of our reach will seem like nothing more than a philosopher’s invention. What a sceptical philosopher would be denying us would then be nothing we could have ordinary commerce with or interest in anyway. Nothing distressing about our ordinary position in the familiar world would have been revealed by a philosopher who simply invents or constructs something he calls “reality” or “the external world” and then demonstrates that we can have no access to it. That would show nothing wrong with the everyday sensory knowledge we seek and think we find in ordinary life and in scientific laboratories, nor would it show that our relation to the ordinary reality that concerns us is different from what we originally thought it to be.

I think this reaction to the picture of our being somehow imprisoned behind the veil of our own sensory experiences is very natural and immediately appealing. It is natural and perhaps always advisable for a prisoner to try to make the best of the restricted life behind bars. But however much more bearable it makes the prospect of life-imprisonment, it should not lead him to deny the greater desirability, let alone the existence, of life outside. In so far as the comfort of this response to philosophical scepticism depends on such a denial it is at the very least premature and is probably based on misunderstanding. It depends on a particular diagnosis or account of how and why the philosophical argument succeeds in reaching its conclusion. The idea is that the “conclusion” is reached only by contrivance. The inaccessible “reality” denied to us is said to be simply an artefact of the philosopher’s investigation and not something that otherwise should concern us. That is partly a claim about how the philosophical investigation of knowledge works; as such, it needs to be explained and argued for. We can draw no consolation from it until we have some reason to think it might be an accurate account of what the philosopher does. So far we have no such

reason. On the contrary; so far we have every reason to think that Descartes has revealed the impossibility of the very knowledge of the world that we are most interested in and which we began by thinking we possess or can easily acquire. In any case, that would be the only conclusion to draw if Descartes’s investigation does indeed parallel the ordinary kinds of assessments we make of our knowledge in everyday life.

We saw that I can ask what I really know about the common cold, or whether I really know that the witness was in Cleveland on the night in question, and that I can go on to discover that I do not really know what I thought I knew. In such ordinary cases there is no suggestion that what I have discovered is that I lack some special, esoteric thing called “real knowledge”, or that I lack knowledge of some exotic, hitherto-unheard-of domain called “reality”. If I ask what I know about the common cold, and I come to realize that I do not really know whether it can be caused by sitting in a draught or not, the kind of knowledge I discover I lack is precisely what I was asking about or taking it for granted I had at the outset. I do not conclude with a shrug that it no longer matters because what I now find I lack is only knowledge about a special domain called “reality” that was somehow invented only to serve as the inaccessible realm of something called “real knowledge”. I simply conclude that I don’t really know whether colds are caused by sitting in draughts or not. If I say in a jury-room on Monday that we can eliminate the suspect because we know he was in Cleveland that night, and I then discover by reflection on Tuesday that I don’t really know he was in Cleveland that night, what I am denying I have on Tuesday is the very thing I said on Monday that I had.

There is no suggestion in these and countless similar everyday cases that somehow in the course of our reflections on whether and how we know something we are inevitably led to change or elevate our conception of knowledge into something else called “real knowledge” which we showed no signs of being interested in at the beginning. Nor is it plausible to suggest that our ordinary assessments of knowledge somehow lead us to postulate a “reality” that is simply an artefact of our inquiries about our knowledge. When we ask whether we really know something we are simply asking whether we know that thing. The “really”

signifies that we have had second thoughts on the matter, or that we are subjecting it to more careful scrutiny, or that knowledge is being contrasted with something else, but not that we believe in something called “real knowledge” which is different from or more elevated than the ordinary knowledge we are interested in. Knowing something differs from merely believing it or assuming it or taking it for granted or simply being under the impression that it is true, and so forth, so asking whether we really know something is asking whether we know it as opposed to, for example, merely believing it or assuming it or taking it for granted or simply being under the impression that it is true.

If that is true of our ordinary assessments of knowledge, and if Descartes’s investigation of his knowledge that he is sitting by the fire with a piece of paper in his hand is just like those ordinary cases, his discovery that he doesn’t know in the case he considers will have the same significance as it has in those ordinary cases. And if that example is indeed representative of our knowledge of the world around us, the kind of knowledge we are shown to lack will be the very kind of knowledge we originally thought we had of things like our sitting by the fire holding a piece of paper. Without a demonstration that Descartes’s philosophical investigation differs from our ordinary assessments in some way that prevents its negative conclusion from having the kind of significance similar conclusions are rightly taken to have in everyday life, we can derive no consolation from the ungrounded idea that the reality from which he shows our knowledge is excluded does not or should not concern us anyway. It is the investigation of his everyday knowledge, and not merely the fanciful picture of a veil of perception, that generates Descartes’s negative verdict.

But even if we did try to console ourselves with the thought that we can settle for what we *can* know on Descartes’s account, how much consolation could it give us? The position Descartes’s argument says we are in is much worse than what is contemplated in the optimistic response of merely shrugging off any concern with an imperceptible “reality”.

For one thing, we would not in fact be left with what we have always taken to be the familiar objects of our everyday experience – tables and chairs, trees and flowers, bread and wine. If

Descartes is right, we know nothing of such things. What we perceive and are in direct sensory contact with is never a physical object or state of affairs, but only a representation – something that could be just the way it is even if there were no objects at all of the sort it represents. So if we were to settle for the realm of things we could have knowledge about even if Descartes’s conclusion were correct, we would not be settling for the comfortable world with which we began. We would have lost all of that, at least as something we can know anything about, and we would be restricted to facts about how things seem to us at the moment rather than how they are.

It might still be felt that after all nothing is certain in this changing world, so we should not aspire to firm truths about how things are. As long as we know that all or most of us agree about how things seem to us, or have seemed to us up till now, we might feel we have enough to give our social, cultural, and intellectual life as much stability as we can reasonably expect or need. But again this reaction does not really acknowledge the poverty or restrictedness of the position Descartes’s sceptical conclusion would leave each of us in. Strictly speaking, there is no community of acting, experiencing and thinking persons I can know anything about if Descartes is correct. Other people, as I understand them, are not simply sensory experiences of mine; they too, if they exist, will therefore inhabit the unreachable world beyond my sensory experiences, along with the tables and chairs and other things about which I can know nothing. So at least with respect to what I can know I could not console myself with thoughts of a like-minded community of perceivers all working together and cheerfully making do with what a communal veil of perception provides. I would have no more reason to believe that there are any other people than I have to believe that I am now sitting in a chair writing. The representations or sensory experiences to which Descartes’s conclusion would restrict my knowledge could be no other than my own sensory experiences; there could be no communal knowledge even of the veil of perception itself. If my own sensory experiences do not make it possible for me to know things about the world around me they do not make it possible for me to know even whether there are any other sensory experiences or any other perceiving beings at all.



The consequences of accepting Descartes's conclusion as it is meant to be understood are truly disastrous. There is no easy way of accommodating oneself to its profound negative implications. But perhaps by now we have come far enough to feel that the whole idea is simply absurd, that ultimately it is not even intelligible,

and that there can be no question of "accepting" Descartes's conclusion at all. I have no wish to discourage such a reaction. I would only insist that the alleged absurdity or unintelligibility must be identified and made out. I think that is the only way we can hope to learn whatever there is to be learned from Descartes's investigation.

### Notes

---

- 1 It has been argued that the problem in the completely general form in which I discuss it here is new in Descartes, and that nothing exactly similar appears in philosophy before that time. See M. F. Burnyeat, "Idealism and Greek Philosophy: What Descartes Saw and Berkeley Missed", *The Philosophical Review* (1982).
- 2 See the beginning of the first of his *Meditations on First Philosophy* in *The Philosophical Works of Descartes*, edited and translated by E. S. Haldane and G. R. T. Ross (2 vols, New York, 1955), vol. I, p. 145. (Hereafter cited as HR.)
- 3 See his *Discourse on the Method of Rightly Conducting Reason and Seeking Truth in the Sciences* in HR, pp. 81 ff.
- 4 See G. E. Moore, "Certainty", this vol., ch. 4.
- 5 A memorable example H. H. Price gave in a lecture in 1962. It is my impression that Price was reporting on an actual hallucination of his.