

10 Skepticism

I think that we all know many things. I believe I know many facts about my immediate surroundings, much about myself, something about the past, and a little about the future. I believe that we also have scientific knowledge, that we know some general moral truths, and that it is certainly possible that many of us know some religious truths. But there are reasons to doubt much of this. There are reasons to think that at best we know very little, perhaps just self-evident truths, for instance that if no vixens are males then no males are vixens, and a few propositions about our present consciousness, say that I am now thinking about the scope of human knowledge.

The possibility of pervasive error

As I consider these matters, I look back at the green field. I reassure myself that I see it vividly. I certainly cannot *help* believing it is there. But an inescapable belief need not be knowledge, or even justified. Suppose I am hallucinating. Then I would not know (through vision, at least) that the field is there.

Perfectly realistic hallucination

I find it impossible to believe that I am hallucinating. But I might find that impossible even if I were, provided the hallucination was as vivid and steady as my present visual experience. I begin to wonder, then, whether I really know that I am *not* hallucinating. If I do not know this, then even if I am in fact not hallucinating, can I know that there is a green field before me? Similarly, if I do not know that I am not simply having a vivid dream in which it seems to me that there is a green field before me, can I know that there is one there?¹

Remembering that we can justifiedly believe something even if we do not know it, I think that at least I may justifiedly believe that there is a green field before me, even if I do not know that I am not hallucinating one (or merely “seeing” one in a dream). Moreover, if I justifiedly believe that there is a green field before me, how much does it matter whether I also know this? It matters whether the belief is true. But the *likelihood* that it is true, so far as that likelihood is something I can discern, depends on how probable the presence of the field is, given the sensory experience on which my belief is based; and in my attentiveness and caution as an observer, I have contributed all I can to that probability. Despite the possibility of hallucina-

tion, then, it appears that my belief remains justified, and it is as likely to be true as I can make it by any steps in my power, such as more carefully observing the texture of what I see. Internally, in my own consciousness, I am being perfectly reasonable in continuing to believe that there is a green field there. So far as justification is concerned, I am beyond reproach.

These points about justification are plausible, but they give false comfort. Doubtless, we can have beliefs which, though they do not constitute knowledge, are justified, and we can have such a belief even if its basis is hallucinatory. But it is now not merely possible that I am hallucinating: I am also quite *aware* that I could be. Given this awareness, am I still justified in believing that there is a green field there? Should I not regard this belief as unjustified, suspend judgment on whether the field is there, and merely hope that it is?

Two competing epistemic ideals: believing truth and avoiding falsehood

These questions produce a tension. I want to believe that the field is there if it truly is, for I have a deep-seated desire to believe as many significant truths as I can. But I also want to avoid believing that it is there if it is not, for I have a deep-seated desire to avoid believing falsehoods. For most of us, these two desires are important; and they represent ideals that govern much of our thinking. But the two ideals pull against each other. The former inclines us to believe readily, since we may otherwise miss believing a truth; the latter inclines us to suspend judgment, lest we err by believing a falsehood.

The former ideal, calling on us to believe truths, pushes us toward *credulity*: believing on grounds that evidentially are too thin – or without grounds at all – and thereby believing too much. The latter ideal, calling on us to avoid believing falsehoods, pushes us toward *a kind of skepticism*: believing only on absolutely conclusive grounds, and thereby – if common sense is right about the matter – believing too little.

How can we balance these ideals with each other? So far in this book, I have spoken more about how we fulfill the ideal of believing as many significant truths as we can, than about how we might fulfill the ideal of avoiding belief of falsehoods. Clearly, the easiest way to fulfill the latter would be to suspend judgment on every proposition one entertains, or at least on all those which, unlike certain self-evident truths, do not have a luminous certainty that tends to compel assent. This is the kind of response characteristic of *Pyrrhonian skepticism*, an ancient variety tracing to Pyrrho of Elis (c. 360–275 BC).²

These reflections about possible error through hallucination, about the apparent vulnerability of justification in the face of such possibilities, and about the ideal of avoiding error suggest why philosophers have been so concerned with skepticism. In very broad terms, skepticism is most

commonly conceived by philosophers roughly as the view that there is little if any knowledge. Call this *knowledge skepticism*.

A related kind of skepticism is constituted by an attitude or feature of temperament, such as a disapproval of believing without conclusive grounds. This is not our direct concern. But if philosophical skepticism is not justified, then some common skeptical attitudes are not either, and some people who go through life with a skeptical attitude lack the kind of intellectual balance that goes with epistemic virtue. One reason, then, for studying skepticism is to approach a mean between two cognitive traits – *intellectual vices*, in the language of virtue epistemology. One is excessive credulity, which is too weak a disposition to doubt or withhold belief; the other is excessive skepticism, which is too strong a disposition to doubt or withhold belief.

Skepticism may also concern justification. Typically, skeptics do not take our justified beliefs to be significantly more numerous than our beliefs constituting knowledge. Call the view that we have little if any justification for belief *justification skepticism*. How far-reaching might a plausible skepticism of either kind be, and how is skepticism to be assessed? I want to pursue these questions in that order and at some length.

It may seem that skepticism offends so blatantly against common sense, and so lopsidedly prefers the ideal of avoiding falsehood over that of believing truths, that it should be dismissed as ridiculous. But it will soon be evident that skepticism is a serious, perhaps even irrefutable, challenge to common sense. Moreover, even if skepticism turns out, as phenomenalism apparently does, to be quite implausible, we learn a great deal about knowledge and justification from studying it.

A serious exploration of skepticism, whether or not we finally accept some form of it, also tends to help us to avoid dogmatism about our own personal views and a self-satisfied assurance that our collective outlook as rational observers of the world embodies knowledge of the sorts of things we think it does: facts about ourselves, our surroundings, and the ways of nature.

Some dimensions and varieties of skepticism

To understand a skeptical view we should locate it in relation to at least four dimensions: (1) subject matter, say the past or the future or physical objects or other minds; (2) epistemic attitude, such as knowledge, suspended judgment, and justified belief; (3) modality, above all contingency or necessity, or the empirical versus the a priori; and (4) the kind of being it purports to limit, say human, subhuman, or superhuman. Regarding subject matter, my concern is wide-ranging. As for (2)–(4), my concern is with human beings and mainly with knowledge and justification regarding contingent empirical propositions.

Much skepticism, whether about knowledge or about justification, is restricted to a given kind of subject, for instance to propositions about the

world outside oneself, or about the past, or about the future, or about ethics, religion, or science. Skeptical views also differ markedly in the status of the knowledge, and in the degree of the justification, they concern. A strong skepticism regarding propositions about the past, for instance, might hold that there is no knowledge, or even partially justified belief, about the past. A weaker skepticism might hold that although some beliefs about the past are justified to some degree, there is neither certain knowledge of the past nor any beliefs about it that are *sufficiently* justified to make it more reasonable to hold them than to suspend judgment on them.

Still another difference between skeptical views concerns their order. The usual skepticism is first-order, in the sense that it concerns the sorts of beliefs or knowledge we have discussed as typical of the kinds grounded in experience or reason, and not beliefs or knowledge *about* such beliefs or knowledge, say beliefs that ordinary perceptual beliefs often do constitute knowledge. First-order skepticism might deny, then, that I know there is a cold glass in my hand, even when I have the seemingly familiar experience I would describe as smelling the mint in my iced tea and feeling the cold glass in my hand. Second-order skepticism might say that even if I do know this, I do not know that I know it.

It is natural, however, for a first-order skeptic also to maintain second-order skepticism, holding, for instance, that there is no second-order knowledge to the effect that there is (first-order) knowledge, for example that no one knows that there is knowledge of people, places, and things. This second-order skepticism is obviously true if there is in fact no first-order knowledge – since from that it would *follow* that then no one knows there is. But a second-order skeptic can hold that even if there is first-order knowledge, no one knows this.

It is, moreover, natural for skeptics to hold their main views as necessary truths, since, for one thing, they commonly believe that for fallible creatures like us there *cannot* be knowledge or justification of certain kinds. I do not intend to discuss skepticism in detail in each of the many forms described, but what follows will apply to a very wide range of cases.

Skepticism generalized

The skeptical challenges I have brought forward can be directed against all our beliefs about the external world, all our memory beliefs, all our beliefs about the future, and indeed all our beliefs about any subject provided they depend on our memory for their justification or for their status as knowledge. Memory is, after all, at least as liable to error as vision.

Skepticism about direct knowledge and justification

Plainly, if all of the senses can deceive through hallucination, then beliefs grounded in any of the senses may be justificational or epistemically

undermined in the same way my belief that there is a green field before me may be undermined by a realization that I might have been hallucinating. Quite apart from whether perceptual beliefs are true, skeptics tend to claim that either the possibility of such hallucinations prevents them from being justified or, even if they are justified, it precludes their constituting knowledge.

Suppose, for instance, that I might be having an auditory hallucination of bird songs. Then my present experience of (apparently) hearing them may not justify my believing that there are birds nearby and is certainly not a sufficient basis for my knowing there are, even if it is true that there are. Similarly, there is a counterpart of hallucination for memory beliefs: *memorial hallucination*, we might call it. I may have the memorial impression that when I was four I saw my parents kissing under the mistletoe, but this could be just a romantic fantasy masquerading as a memory.

Beliefs about the future are rather different from memory beliefs. The former concern future events and hence are not grounded in experiential states that we think of as in some way causally deriving from things about which we have knowledge (as with past events). But even if there is no counterpart of memorial hallucination here, there are equally undermining possibilities. For instance, a confident belief that I will talk with Jane could be a product of wishful thinking, even when in fact it is grounded in my long-standing intention to talk with her. Perhaps the belief could be an anticipatory delusion. Even my belief that I will live to discuss skepticism could be mistaken owing to many sorts of reasons, including dangers to me of which I am now unaware.

Now consider what we take to be our general knowledge, whether a priori or scientific, say of arithmetic truths or scientific laws. Since it is possible to misremember propositions, or to seem to remember them when one does not, or to have a kind of memorial hallucination that gives rise to a completely groundless belief, it would seem that the only secure beliefs of general propositions are of the relatively few that we can know directly without needing any evidence. This apparently leaves none of our general scientific beliefs, and only our a priori knowledge of self-evident propositions, epistemically unscathed.

Inferential knowledge and justification: the problem of induction

Even if we leave aside problems about perceptual and memory beliefs, there is a difficulty for the commonsense view that justification or knowledge grounded (directly or indirectly) in a basic source can be transmitted inductively. The classical statement of this *problem of induction* – the problem of how to justify such inductive inferences – comes from David Hume.³ Hume showed that one cannot know a priori that if the premises of a specific piece of inductive reasoning are true, then its conclusion is also true. One can clearly conceive the former being true while the latter is false, whereas one

cannot conceive its being true that (1) all human beings are mortal and Socrates is one of them, and yet false that (2) Socrates is mortal. Thus, no matter how good the inductive reasoning is, it is always (deductively) invalid.

Consider the inductive reasoning from the premise that the sun has always risen each twenty-four hours to the conclusion that it will rise tomorrow. Of all such reasoning – reasoning “concerning matter of fact and existence” – Hume says “That there are no demonstrative [roughly, valid, evidentially conclusive] arguments in the case seems evident, since it implies no contradiction that the course of nature may change and that an object, seemingly like those which we have experienced, may be attended with different or contrary effects.” Hence, even if I do know that the sun has risen every day since time immemorial, and even allowing that I have such an extensive basis for believing that it will rise tomorrow, I *could* be mistaken in believing this, and it seems questionable whether I am justified in believing it.

More generally, Hume’s arguments lead us to ask whether, if our premises could be true yet our conclusion false, we have any reason at all, on the basis of the premises, for believing the conclusion. And how can we ever *know* the conclusion on the basis of such premises? Indeed, how can we even be minimally justified in believing the conclusion on the basis of such premises? The problem of induction, as most often understood, is largely the difficulty of adequately answering these questions.

The problem can also be put in terms of *probability*. We normally operate on the commonsense presumption that when one thing is associated with another, say a sunrise with the passage of twenty-four hours, and the two have never failed to be associated in the same way, then the greater the number of cases of association, the greater the probability that the association will occur in a new case – for instance that the sun will rise tomorrow.⁴ We also operate on the related commonsense presumption that for natural phenomena, such an association can occur sufficiently often to yield justification for believing, and even knowledge, that the association will occur in a new case.

From a Humean perspective, it will not do to argue as follows: I am justified in believing my conclusion on the basis of inductive support for it, such as the past regular behavior of the sun, since past experience has shown that reasoning like this, which has had true premises, has also had true conclusions. For this way of defending an inductively based conclusion simply relies on yet another inductive argument – it gives a kind of inductive reasoning to support the view that certain kinds of inductive arguments justify one in believing their conclusions. It just inductively generalizes about inductive arguments themselves, using as a guide past experience in which we seem to have found out that by and large their conclusions turned out true when their premises were true.

That reasoning, then, apparently *begs the question* against Hume.⁵ For it assumes, without independent evidence, part of what he regards as false, namely, that inductive inference constitutes reasoning that either can ground

knowledge of its conclusion, or can at least *justify* its conclusion, in the sense of providing good reason for it. We have taken the battle to a different field – that of inductive argumentation rather than sunrises – but we have added no new weapons nor enhanced our forces.

The problem of other minds

One of the major points that Hume so powerfully defended – roughly that non-deductive inferences are fallible – is by no means restricted to beliefs about the future. Such beliefs are, however, so prominent in his discussion of inductive inference that sometimes the problem of induction is narrowly conceived as that of how we can show that we have any reason to believe the future will be like the past. This conception is unduly narrow. Recall my observing Jim briskly shuffling papers and angrily mumbling curses. I cannot help believing, on this basis, that he is angry. But this reasoning leaves my belief clearly fallible: even if I know my premises (through perception), it does not follow that he is angry, and that could be false. He could be pretending.

The case of Jim's anger is alarmingly representative. Everything I believe about what is occurring in the inner lives of others seems to rest on grounds that are inductive in this way: what I observe – above all, their behavior – does not *entail* anything about their minds. They could be pretending, or psychologically abnormal; or some other source of error could occur. So if I cannot have knowledge of people's inner lives from their behavior, apparently I can never have it.

Worse still, if I cannot know anything about the inner lives of others, can I even know that there *are* others, as opposed to mere bodies controlled externally, or by hidden microscopic machinery, rather than directed through beliefs and intentions of the kind that I take to animate me?

There is, then, a *problem of other minds*. Can we know, or even justifiedly believe, that there are any? If our experiences would be just as they are if the human bodies we interact with are controlled from outer space and have no inner life of their own, how can we know that those bodies are, as most of us cannot help thinking, animated by minds like ours?⁶

The problem is compounded when we realize that we can never directly verify, as we introspectively can in our own case, what is occurring in someone else's consciousness. Thus, all I can do to check on my inductively grounded beliefs about the inner lives of others is obtain further inductive evidence, for instance by observing whether they behave as one would expect if I am right in thinking them to be, say, angry. I cannot, as in my own case, introspectively focus on the events in their consciousness. How can I know anything about the mental and emotional life of others if I am in principle debarred from decisively verifying my beliefs about the contents and events of their consciousness? Even if I am sometimes right, I can never tell when.

It may be replied that by far my best explanation of why other bodies behave as if they were animated by a mind is that they *are* so animated. The

other hypotheses, such as control from outer space or by a machine, are far-fetched. The suggested reasoning sounds plausible, but notice that it is still a kind of induction: inference to the best explanation (abduction).

The indicated inference to propositions about other minds as best explaining observed behavior is, however, supported by a strong *argument from analogy*: again and again, when my body behaves in a certain way under certain conditions, I am in a certain mental state, say in pain when I am burned and cry out; so (surely) the same behavioral pattern in another body is accompanied by a similar mental state. Other bodies are so much like mine in structure and behavior that they are very probably animated by minds like mine.

There is no need to deny either that positing other minds best explains what we seem to know about other bodies or that the analogical argument just sketched is strong. Still, from one proposition's best explaining another it does not *follow* that the first is true; and the analogies between my body and others at most render *probable*, rather than *entailing*, that some other body is associated with mental states as mine is with my mental states. For one thing, there are other possible explanations (such as the hypothesis of control of other bodies from outer space, or by a powerful and clever evil genius); these explanatory alternatives, if true, would leave my experience exactly as it is. For another thing, some of these alternative hypotheses can well explain the analogies that otherwise seem compelling.

Another way to see the power of these skeptical hypotheses is to note that our experience does not *discriminate* between the skeptical scenario and the commonsense one. Our experience would be just what it is if we were steadfastly hallucinating the external world, including even the human bodies we seem to see. The same holds if we are not hallucinating but the human bodies are externally controlled. How, then, can our experience justify us in believing that there is an external world or that there are other minds?

Putting the problem somewhat differently, if our experience *underdetermines* the truth of propositions we commonly believe about the external world, roughly in the sense that it does not decisively indicate their truth as opposed to the truth of skeptical (or other) alternative hypotheses that can explain our experience, how can our experience justify our believing such commonsense propositions? If it cannot, and if, as Hume plausibly argued, we also cannot know that proposition, how can we be justified in believing anything at all about the external world?⁷

It is only a short step from this full-scale attack on inductive inference to a *problem of the body*. If, as a skeptic might well hold, our apparent knowledge of our own bodies is inductively grounded, being based on perceptions and bodily sensations somewhat as beliefs about external objects are, then can we know, or even justifiedly believe, that we have a body? Could we not be steadily hallucinating even our own flesh?

It might be replied that thoughts, including my reflections on skepticism, necessarily require an embodied thinker. But that point would only

imply that I *have* some kind of body, not that I can know anything about it. The point is also far from self-evident and is indeed denied by philosophers in the powerful tradition of Descartes. They hold that we (persons) are essentially mental (or spiritual), hence non-physical, substances.

In any case, even if it should be true that thoughts can occur only in an embodied thinker, the only embodiment necessary might be a brain. Hence, on the skeptical view imagined, the most one could know is that one is embodied in some way, say in a brain. Whether that brain is itself embodied, or ever interacts with anything else, would be beyond one's knowledge. Why, then, could I not be alone in the world, and perhaps even a "brain in a vat," a brain kept alive in a nurturing liquid and subjected to hallucinations that falsely convey the impression of a normal life?⁸ We could dub this *the envatment problem*.

The egocentric predicament

In this way, skepticism can drive us into an *egocentric predicament*: a position that makes it seem clear that all we can (empirically) know about the world, perhaps all we can justifiably believe about it as well, concerns our own present experience. Perhaps, for all I know, I am a lone conscious ego vividly hallucinating a physical world that has no external reality.

Most skeptics have tended to push no further, or at least not to express very much doubt about our capacity to know propositions of two specific kinds: those about what is currently going on in our minds and at least those a priori propositions that are luminously self-evident. But skeptics *can* push further. Descartes, in the first of his *Meditations*, raised the possibility that there was nothing of which he could (justifiably) be certain. Recall introspectively grounded beliefs, such as that I am thinking about skepticism. It seems possible that this belief is mistaken. If that is possible, how can I know that I am thinking about skepticism? If I know, I cannot be wrong. But here error is possible. Perhaps I do not even have knowledge of my own conscious states.

To make this sort of argument work with beliefs of self-evident propositions we must, I think, strain. Descartes may perhaps be read as seriously entertaining the question of whether it is possible that God, being omnipotent, could have brought it about that even propositions of the sort I am calling self-evident might be false. Could an omnipotent being bring it about that while some dogs are pets, no pets are dogs? I see no reason to think so. As St Thomas Aquinas and many other philosophers have maintained, omnipotence is simply not the power to "do" things that are absolutely impossible.⁹ Power is exercised within the realm of the possible: impossible "deeds" are not candidates for any being to do.

If one accepts this point, one might argue that there is no act of bringing it about that while some dogs are pets, no pets are dogs. Calling this an act misuses the vocabulary of action. Hence, the impossibility that an omni-

potent being can bring it about does not imply that there is any act which that being cannot perform. This point, in turn, deprives the skeptic of a way to argue that beliefs of necessary truths could be false.

This reasoning may not settle the matter, but it is sufficiently plausible to warrant leaving aside skepticism concerning beliefs of luminously self-evident propositions. These propositions seem not only incapable of falsehood, but also incapable of even being believed without justification, at least when carefully and comprehendingly considered. Leaving such skepticism aside takes little from the skeptic in any case. If these are the only propositions we can know, then we can know nothing about our world, not even anything about our innermost consciousness. We are at best in an egocentric predicament.

Fallibility

In appraising skepticism, I want to formulate some of the main principles that underlie it in what seem its most plausible forms. If they can be shown to be unreasonable, then the skeptical threat to the commonsense view that we have a great deal of knowledge and justification can at least be blunted. In formulating and assessing these principles, it is well to distinguish skeptical threats to the generation of knowledge (or of justification) from skeptical threats to its transmission. It is natural to start with questions about its generation. If no knowledge is generated, there is none to be transmitted.

Three kinds of epistemic infallibility

Is there really any reason to doubt that, normally, introspectively grounded beliefs constitute knowledge? It may be true that such beliefs could be mistaken, but what is a skeptic entitled to make of this? The skeptical argument that comes to mind here is based on what I will call *the infallibility claim about knowledge*: if you know, you cannot be wrong. If we simply add the premise that you can be wrong in holding a given introspective belief, say that you are thinking about skepticism, it would seem to follow that such beliefs do not represent knowledge. This kind of *argument from fallibility*, as we might call it, can be applied to just about every sort of proposition we tend to think we know.

If, however, we look closely, we find that the infallibility claim is multiply ambiguous. There are at least three quite different things it might mean, and hence really three different infallibility principles.

The claim, 'If you know, you can't be wrong,' might have the meaning of

- (1) It must be the case that if you know that something is true, then it is true (i.e., you cannot know something that is false).

Call (1) the *verity principle*, since it says simply that knowledge must be of truths (verities). Knowledge can never have a falsehood as its object. The claim might, on the other hand, have the meaning of

- (2) If you know that something is true, then it must be true, that is, the proposition you know is necessarily true (i.e., you can know only necessary truths).

Call (2) the *necessity principle*, since it says simply that knowledge is of necessary truths. Knowledge never has among its objects any propositions that could possibly fail to hold.¹⁰ The claim 'If you know, you can't be wrong' might also have the meaning of

- (3) If you know that something is true, then your *belief* of it must be true, in the sense that your believing it entails or guarantees its truth (i.e., only beliefs that cannot be false constitute knowledge).

Call (3) the *infallibility principle proper*, since in saying that only infallible beliefs constitute knowledge it connects with skepticism more closely than (1) or (2). Knowledge, it says, is never constituted by fallible beliefs, those that can have falsehoods among their objects.

Unlike (2), (3) implies nothing about the propositional or other *objects* of knowledge; instead, it restricts the kind of *belief* that can constitute knowledge. And by contrast with (2), (3) also allows for knowledge of contingent (non-necessary) truths, such as that I exist. This proposition can be false (that I exist is not a necessary truth); but my *belief* of it is infallible and therefore cannot be false. If I now believe that I exist, then it follows that I do now exist.

Knowledge and fallibility

We can now assess the skeptical reasoning that employs the infallibility claim in one or another interpretation. I will be quite brief in discussing the first two; the third is the most controversial and most important for skepticism.

The verity principle, (1), is plainly true: one cannot know something that is false. In *this* sense, knowledge is infallible. If it is false that the maple is taller than the spruce, then I do not know it is. But if this is all the infallibility claim comes to, it provides no reason to conclude that I do not know that I am thinking (or that anything else I believe is not genuine knowledge). Granted, it must be true that *if* I know I am thinking, then I am. But that tells us nothing about whether I *do* know I am. The verity principle is itself a verity, but it does not advance the skeptical cause.

The necessity principle, on the other hand, principle (2), seems mistaken. Surely I know some propositions that are not necessary, such as that I exist

(as noted earlier, it is not a necessary truth that I exist, as it is that vixens are female). Even the skeptic would grant that I cannot falsely believe this, since my believing it entails that I exist (non-existent things cannot have beliefs at all, true or false). It may indeed be impossible for me even to be unjustified in believing the proposition that I exist when I comprehendingly consider it (Descartes seems to maintain in *Meditation II* that this case is impossible). The same holds, of course, for you in relation to your belief that you exist.¹¹

It might seem that we may grant the skeptic that the only kinds of *proposition* that cannot be falsely believed are necessary. But that would also be a mistake: any proposition entailed by there being at least one belief is incapable of being falsely believed. Anyone's *believing* it would entail that it is true – though it might be possible, given certain logical deficiencies, to believe such a proposition without having justification for it, as is certainly possible for necessary truths of mathematics. Such cases suggest that there is no simple relationship between the kind of proposition believable with infallibility and the conditions for knowing or justifiedly believing it.

Even if the necessity principle were true, however, a skeptic could not reasonably *use* it, without first defending it by adequate argument, against the commonsense view that introspective or even perceptual beliefs normally constitute knowledge. For clearly they are not beliefs of necessary truths, and defenders of common sense do not take them to be; hence, invoking the necessity principle against common sense, without first arguing for the principle, would be in effect a flat denial that such beliefs constitute knowledge. That would beg the question against the commonsense view.

Suppose, for instance, that a skeptic says that if you know, you cannot be wrong, where this means (2), then notes that introspective and perceptual beliefs (which are of propositions that are not necessary) can be false, and concludes that such beliefs do not constitute knowledge. This would not be presenting a good reason to believe the conclusion, but just asserting, disguisedly given (2), that the commonsense view that we have introspective and perceptual knowledge is mistaken. There may seem to be a good argument here, because it is so easy to take 'If you know, you can't be wrong' as asserting the verity principle. But that principle is acceptable to common sense, whereas the necessity principle is not. To argue for the latter by allowing the plausibility of the former to serve as support for it is to trade on an ambiguity. It masks poor reasoning – or the absence of any argument or support at all.

The infallibility principle proper, (3), in effect says that only infallible beliefs can be knowledge. Now as we have seen some beliefs of contingent propositions are infallible. Consider my belief that I now exist, and my more specific belief that I have a belief. Just as my believing I exist entails that I do exist, if I *believe* I have a belief, it follows that I have one: I have at least that very belief even if I have no other beliefs. Beliefs like these might be called *self-grounding*, since their possession constitutes a sufficient ground for their truth.

The infallibility of these two contingently true, but self-grounding, beliefs shows that despite appearances, (3) is not equivalent to (2), since (3), the infallibility principle, but not (2), the necessity principle, allows knowledge of propositions that are *not* necessary (contingent propositions). But why should we accept (3)? What reason can the skeptic give for it? Not that if you know, you cannot be wrong; for when we look closely, we find that when plausibly interpreted as meaning (1), that is no help to the skeptic, and when interpreted as (2) or (3) it just asserts the skeptical position against common sense.

What makes the infallibility claim seem to give the skeptic an argument against common sense is the way skepticism can trade on the ambiguity of that formulation: one finds the argument from fallibility attractive because its main premise, conceived as equivalent to (1), is so plausible; yet the argument succeeds against common sense only if (2) or (3) are legitimate premises, and I think it is doubtful that the skeptic has a cogent argument for either of them. It will help to consider first the bearing of the concept of uncertainty, one closely related to the notion of infallibility.

Uncertainty

Like fallibility, uncertainty has seemed to many skeptics to leave us with little, if any, knowledge. Recall the possibility that I am hallucinating a green field before me when there is none there. Can it possibly be certain, then, that there is one there? And can I ever *tell* for certain whether or not I am hallucinating a field? Skeptics tend to answer these questions negatively and to contend that if we cannot tell for certain whether we are hallucinating, we do not know we are not hallucinating. They also tend to argue that if one does not know that one is not hallucinating, surely one does not know that there is a field there.

Moreover, in a way uncertainty cuts deeper than fallibility: for even if I believe a theorem of logic that cannot be false and so have an infallible belief, I may not be justified in taking my proof to be sound and hence cannot be justifiedly certain. Uncertainty arises where one's grounds are not conclusive, and it can arise, as with beliefs of theorems, even when one's belief is infallible. Thus, even infallibility is not enough to render a belief knowledge. At least two important principles are suggested here.

One principle suggested by reflection on these questions about possible error is *the certainty principle*: if one cannot tell for certain whether something is so, then one does not know it is so. This principle is plausible in part because, typically, 'How can you tell?' and 'How can you be certain?' are appropriate challenges to a claim to know something. Moreover, 'I know, but I am not certain' sounds self-defeating, in a way that might encourage a skeptic to consider it contradictory. Further support for the certainty principle can be derived from the idea that if our grounds for a belief

underdetermine its truth – as where a skeptical possibility such as the Cartesian demon hypothesis can also explain our having those grounds – then one cannot tell for certain that the belief is true.

Another principle suggested by our questions about the possibility of hallucination is *the back-up principle*: if one believes something, say that there is a field before one, which is inconsistent with a further proposition – such as that one is merely hallucinating a field where none exists – then one's belief constitutes knowledge only if it is backed up by one's knowing, or at least being in a position to know, that the further (undermining) proposition is false.

The back-up principle is plausible in part because one is in a sense responsible for the implications of what one claims to know. If, for instance, I claim to know that there is a green field before me, and that proposition implies that the field is not a pavement textured and painted to look just like a green field, it would seem that I had better know that it is not such a pavement. This, in turn, is commonly taken to imply that I must at least be justified in rejecting that strange possibility.

The upshot of this skeptical reasoning is that if I know that there is a green field before me, I apparently must be prepared to back that up by justifiedly rejecting exactly the sorts of possibilities that the skeptic reminds us are always there, in abundance. But must I be thus prepared? Let us consider the certainty and back-up principles in turn.

Knowing, knowing for certain, and telling for certain

Chapter 8 argued that knowing does not imply knowing for certain. This conclusion suggests that the kind of certainty in question, epistemic certainty, is not required for knowledge, and that having such certainty may be something quite different from simply knowing. Still, from the point that knowing need not be knowing for certain, it does not follow that one can know without being *able to tell* for certain. Thus, the skeptic may still maintain that the certainty principle undermines the commonsense view that we have perceptual knowledge.

Let us first ask what it is to tell for certain. A skeptic may mean by this acquiring knowledge, in the form of an infallible belief, of a proposition that entails the truth of what one can tell is so. Thus, to tell (for certain) that one is not hallucinating a green field one might, like Descartes in the *Meditations*, prove that there is a God of such goodness and power that – since it would be evil for God to allow it – one could not be mistaken in a belief properly based on such a vivid and steadfast perception as one now has of a green field. We can tell for certain that there is an object before us because we can prove that God would not allow us to believe this under the present conditions unless it were true.

Some thinkers might embrace Descartes's theistic solution here. But one might also reject the skeptical principle in question, the infallibility principle.

To require that a belief can be knowledge only if – whether in Descartes's way or a similar fashion – it can be conclusively shown to be true would again beg the question against the commonsense view that a belief can constitute knowledge without being infallible (a belief that can be absolutely conclusively shown to be true *is* infallible). Thus, if skeptics have no good argument for the principle of infallibility proper, they should not assume that principle in defending the view that we can know only what we can “tell for certain” in this strong sense of the phrase.

Perhaps, on the other hand, telling for certain is simply a matter of ascertaining the truth in question by some means that justifies one in being (psychologically) certain of what one can tell, even if not maximally certain (if there is a maximum here). If so, perhaps we normally *can* tell for certain that we are not hallucinating, for instance by seeing whether the senses of touch and smell confirm our visual impression. To be sure, the confirming experiences do not *entail* that there is a green field before me. But we still have no good argument that certainty (or knowledge) may arise only from entailing grounds (another controversial view, shortly to be discussed). Thus, this point does not establish that confirming experiences cannot enable us to tell for certain that we are not hallucinating.

Moreover, suppose that we interpret telling for certain in the modest way just suggested, and that we can tell for certain in this sense that what we know is true. In that case, perhaps there is a weak sense in which beliefs constituting knowledge are infallible. They need not be such that it is absolutely impossible (logically impossible, in a broad sense) that they be false, as in the case of my belief that I exist. There need only be something about our grounds for them in virtue of which they (empirically) cannot be false, say because it would violate the laws of nature. Water cannot flow (as opposed to being pumped) uphill, but this is empirically impossible, not absolutely so, as it is impossible for some pets to be dogs without any dogs being pets.

It may be true that grounds of what is commonly considered to be knowledge *are* typically such that, *given* those grounds, the belief constituting that knowledge cannot be false (at least cannot be false within the laws of nature). Suppose this is true. Should we now say to the skeptic that the beliefs commonsensically considered knowledge, such as many perceptual ones, are empirically certain? We may say this only if we keep in mind what was wrong with inferring the necessity principle from ‘If you know, you can't be wrong’. There surely might be causal laws of nature which guarantee that *if* one is situated before a field in good light, as I am, and one has visual experiences like mine caused by the field as mine are, *then* one sees it, and hence cannot falsely believe that it is there. But this does not imply that my belief is empirically necessary, as a law of nature at least commonly is, any more than the “logical” law that it is necessary that *if* one knows that *p*, then *p* is true, implies that *p* itself is necessary. A guarantee of truth is not a guarantee of even empirically necessary truth,

much less of epistemic certainty, any more than a guarantee of payment is a guarantee of payment in gold or in some medium that *cannot* be devalued.

If the existence of causal laws and associated causes of many of our beliefs implies the truth of many of our beliefs grounded in experience, such as my belief that there is a green field before me, it does not follow that those true ones *cannot* be mistaken, in the sense of being epistemically certain, or conclusively justified, or of any other epistemic status high enough to satisfy a skeptic. What follows is only that *given* the laws of nature and their causal grounding, they are *true*. This seems more than enough for common sense. The skeptic gives us no good argument to show that there are no such laws or that such law-based truth is not quite sufficient to render a belief knowledge whether it represents certainty or not.¹²

Entailment as a requirement for inferential justification

The back-up principle fares no better than the infallibility principle proper. For one thing, it depends on the assumption, which defenders of common sense stoutly reject, that in order to know that something is true, one must have grounds that entail its truth. To see that the back-up principle depends on this, consider first a very simple case. Take the proposition that it is false that there is a green field before me. This is inconsistent with what I believe, namely, that there *is* one before me. Hence, the back-up principle requires that I at least be in a position to know that this is false.

The back-up principle may seem true because one may think: How else, besides being able to know the falsity of propositions incompatible with what I believe, can I be adequately armed against the threat of falsehood? If I am not in a position to know that propositions plainly incompatible with what I believe are false, I cannot properly back up what I believe. But the falsity of the negative proposition that there is not a green field before me entails that there *is* one before me; for if it is false that it is false that there is one, then it is true that there is one. Thus, if, by virtue of how I must be able to back up my original claim, I do know that this negative proposition is false, then I thereby have (and know) an entailing ground for the truth of what I originally believed – that there is a green field before me.

Now take a case in which backing up what I think I know is more complicated. Consider the proposition that what I take to be a green field is really a pavement with such a realistic-seeming grassy green texture that I cannot tell (perceptually) that it is really not a field. Must I be in a position to know that this is false in order to know that there is a green field before me? The very description of the case suggests that I cannot know, at least by using the senses unaided by experimentation or specialized knowledge, that the field is not a textured pavement. But why must I be able to tell this at all? Is there any reason to think that the field might actually be dyed pavement? (Is that a relevant alternative, some philosophers would ask?) Must I,

in order to know, not only have a well-grounded true belief but also grounds for knowing, of every possible explanation of how my belief could be false, that this explanation is incorrect? I do not see that I must.

One might object that in order to know a proposition I must be in a position to know whatever follows from it (or at least obviously follows from it). After all, if something does follow from what I know, I could infer it by valid steps from what I initially know, and thereby come to know it, too.

This is an important objection. But in discussing the transmission of knowledge and justification, we considered cases that apparently undercut the objection. I can apparently know the sum of a column of figures even if I cannot, without further checking, know something which obviously follows from it: that if my wife (whom I justifiably believe to be a better arithmetician) says this is not the sum, then she is wrong. If this can be true of me, then neither knowledge nor justification is automatically transmitted across valid deductive inference.

It apparently will not do, then, to say that we can always count on the transmission of knowledge from propositions we know or justifiably believe to those they entail, even when the entailment is, as in our example, obvious. Thus, even though my seeing a green field plainly entails (for instance) that I am not seeing a pavement with a textured grassy-looking surface, I presumably do not have to be in a position to know or justifiably believe (by inferring it) that this proposition is false.

Suppose, however, that this view is mistaken, and that knowledge and justification *are* always transmitted across valid deductive inference. It may be plausibly argued that I *do* have justification for rejecting the skeptical hypothesis that there is a pavement before me textured to look just like a green field. It is not just that it appears to me that there is a green field before me; I also have no reason to think there is anything abnormal in the situation, and some reason to think that, in cases like this, large, nearby familiar kinds of things are as they appear to me in such vivid and careful observation. Thus, I may reject the skeptical hypothesis and I do know or justifiably believe that there is a green field before me. We could also stress that the *kinds* of grounds I have for believing there is a green field before me are plainly sufficient for knowing this proposition and then that very proposition is my premise for the entailed conclusion that there is not a pavement before me textured to look like grass. On this view, the point is that by virtue of perceptual justification we gain (commonsense) knowledge of a conclusive ground for rejecting the skeptical hypothesis.¹³

There are other factors one might cite, indeed, too many to discuss here. My point is simply this. Since the skeptic has not provided good reasons for the principles I have already rejected (or for comparably strong principles), even if knowledge and justification are always transmitted across valid inference, there may be good reason to say that skeptical hypotheses, such as that the "field" is a cleverly painted and textured pavement, may be justifiably rejected.¹⁴

Knowing and showing

There is something we may grant to the skeptic that will help to justify my rejection of the certainty and back-up principles. Admittedly, in order to *show* the skeptic that my original belief is knowledge, in the face of the *suggestion* that one of those explanations of its falsity holds, I may have to know that, and perhaps why, this explanation does not hold. Showing something, after all, commonly requires invoking premises for it, and one must presumably know or justifiably believe those premises if one is to show a conclusion from them.¹⁵ The question 'Do you know?' tends to move discussion to a second-order context in which one seeks not simply to offer grounds for what one takes oneself to know, but grounds for the *second-order* proposition that one knows it. After all, a direct answer to, for instance, 'Do you know that she missed the train?' is something like, 'Yes, I know because I just checked the station,' rather than 'I just checked the station'. The latter reply provides a *ground* on which one may know that she missed the train and only indirectly implies that I do know it. A skeptic would not grant this evidential power to such a ground, however, and would indeed not take my citing the schedule to provide an adequate answer to 'Do you know that she missed the train?'

Still, we may ask, why, in the absence of the need to *show* that I know, must I, in order simply to *have* knowledge, have the capacity to show that I have it, as the back-up principle would require? Surely I need not. I can know that if some dogs are pets then some pets are dogs, even if I cannot show this self-evident truth – perhaps simply because I can think of nothing more obvious to use as a reasonable premise from which to show it. And if my wife raises no question of whether my arithmetic answer is correct, I can know that answer even if I cannot show – without obtaining further grounds for the answer – that I do know it. (If my original justification were good enough to enable me to show that if she says the sum is wrong, then she is wrong, perhaps it would also enable me to know, even without showing it, that if she says this, she is wrong.)

The point that one can know without being able to show that one does drastically weakens the case for the back-up principle. Moreover, if, as seems quite possible, I can know the sum on the basis of my calculations without being able to show that I do – apart from gaining new evidence – then I can know it without being able, given my evidence from careful calculation, to *tell* for certain whether it is true. That would require new calculations and hence new evidence. This second point directly cuts against the certainty principle as well as against the back-up principle.

Examining the relation between knowing something and being able to show it also indicates that the converse of the certainty principle – the *show-know principle*, we might call it – should also be rejected: being able to show something one believes, even being able to prove it, does not entail knowing it. This can be seen from our example. Suppose I can now show that if she says the sum is wrong, she is mistaken, by doing a more careful calculation

twice over. Does it follow that I now know this proposition for certain? I do not see that it does. From the fact that I now have the ability to show something I believe, it does not follow that I now know it at all. Having the raw materials to create something – here grounds for knowledge – does not entail already having it. Moreover, suppose that, as sometimes happens, I am lucky in a mathematical hunch; I might still be capable of constructing a proof I myself would not have expected to discover. We may by good fortune have raw materials to create a foundation for something we have fabricated by a stroke of luck.

My examples of showing something are cases of deductive demonstration. But it would be a mistake to think that all instances of showing require deductive demonstration or proceed from absolutely conclusive grounds. Both points were discredited by our discussion of how propositions are established through scientific reasoning.

Negative versus positive defenses of common sense

In the context of thinking about skepticism, it is easy to forget that knowing something does not require being able to show that one knows it. For in thinking about skepticism we are likely to be trying to defend, against a skeptical onslaught, the commonsense view that there is much knowledge, and we easily think of defending this view as requiring us to show that there *is* knowledge. There is, however, more than one kind of defense. The two kinds I have in mind are analogous to standing firm as opposed to attacking.

A *negative defense of common sense*, one that seeks to show that skeptical arguments do not justify the skeptic's conclusion, does not require accomplishing the second-order task of showing that there is knowledge or justified belief. That achievement is required by a *positive defense of common sense*, one that seeks to show that we have the kinds of knowledge and justified beliefs common sense takes us to have. A negative defense requires only contending that skepticism provides no good argument against common sense.

It does not appear that skepticism as so far examined provides a good argument against common sense. Why, for instance, should the skeptic's merely suggesting a possible explanation of how there could be no green field before me, without giving any reason for thinking the explanation is correct, require me to know, or be in a position to know, that it is not correct? All things considered, then, I reject the skeptical case as so far described.

Deducibility, evidential transmission, and induction

When we come to the problem of induction, it seems clear that one assumption the skeptic is making is that if we believe something on the basis of one or more premises, then we can know it on the basis of those premises only if

it follows from them, in the sense that they entail it. Call this *the entailment principle*. It says in effect that *knowledge can be transmitted only deductively*.

Why should we accept this principle? Not simply because inductive reasoning is "invalid"; for that term may be held to be improperly applied to it: inductive reasoning is strong or weak, probable or otherwise, but it does not even "aim" at (deductive) validity. Even if it may be properly said to be (deductively) invalid, however, that may be considered an uncontroversial technical point about its logical classification. It is a point of logic, not of epistemology. So conceived, the point does not imply either that knowledge of the premises of inductive reasoning cannot ground knowledge of its conclusions, or that justified beliefs of those premises cannot ground justified beliefs of their conclusions.

One might, on the other hand, accept the entailment principle and argue that when properly spelled out inductive reasoning can be replaced by valid deductive reasoning. For instance, suppose we add, as an overarching premise in inductive reasoning, *the uniformity of nature principle*, which says that nature is a domain of regular patterns, in the sense of patterns that do not change over time. From this together with the premise that the sun has always risen each day it apparently does follow that it will rise tomorrow.¹⁶

But what entitles us to the premise that nature is uniform? Hume would reply that it is not knowable a priori, and that to say that we know it through experience – a way of knowing it that would depend on inductive reasoning – would beg the question against him. (On the Humean view, if our belief of the uniformity principle is grounded wholly in premises that only inductively support it, we do not know it.) I believe that this Humean response is highly plausible. The problem of induction must be approached differently.

Epistemic and logical possibility

What perhaps above all makes the entailment principle plausible is the thought that if our premises could be true and yet our conclusion *might* be false, then we cannot know (or even justifiably believe) the conclusion on the basis of those premises. At first, this thought may sound like just another formulation of the entailment principle. It is not; it is different and considerably more plausible. That is partly why, when it is conflated with the entailment principle, it seems to support that principle. The 'might' in question is epistemic; it is like a farmer's in 'That wood dust might mean carpenter ants' or a physician's in 'Those abdominal pains might mean appendicitis'. This 'might' suggests not only that for all we know (or may take ourselves to know) the pains do mean appendicitis, but also that there is *reason* for at least some degree of suspicion that there is appendicitis and perhaps some need to rule it out.

The statement that certain abdominal pains might mean appendicitis is not merely an expression of a bare logical possibility of appendicitis – a

statement that appendicitis is possible without contradiction – based, say, on no one's being absolutely invulnerable to it. If that very weak and general statement represents all we know about the case, we are not entitled to say that the pains might mean appendicitis. It is also not a logical impossibility that the Tower of London levitate above the City; but we would be quite unjustified in saying that it might.

This distinction between *epistemic possibility* – what is expressed by the epistemic 'might' just illustrated – and mere *logical possibility* – what can be the case without contradiction or some other kind of necessary falsehood¹⁷ – bears importantly on the problem of induction. It is true that if, no matter how good inductive reasoning is, its premises could be true and yet its conclusion might, in the *epistemic* sense, be false, perhaps we cannot know the conclusion on the basis of them. But is this generally the case with inductive reasoning? I cannot see that it is.

Moreover, suppose it could be true that, relative to its premises, the conclusion of inductive reasoning might, in the epistemic sense, be false, what reason is there to think that this really is true? Skeptics cannot justifiably argue for this claim as they sometimes do, maintaining, simply on the ground that the premises do not entail the conclusion, that the conclusion might be false. Arguing in this way is rather like saying, of just any stomach ache a child gets after eating too much Halloween candy, that it might mean appendicitis.

It is barely possible that, relative to all we know or are justified in believing about the child, the stomach ache means appendicitis. But from that bare possibility we may not automatically conclude that appendicitis is *epistemically possible* – roughly, that relative to all we know or are justified in believing, we are unjustified in disbelieving that the stomach ache might mean appendicitis. Nor does this bare possibility rule out our knowing, on inductive grounds, that overeating is the cause.¹⁸

Entailment, certainty, and fallibility

There are other reasons for the attractiveness of the entailment principle, at least from a skeptical point of view. If one embraces the infallibility principle, one is in fact committed to the entailment principle. For suppose that, from known – and hence on this view infallibly believed – premises, one inductively derives a belief which is not itself infallible, as (empirical) beliefs which are inferentially grounded typically are not. Since inductive transmission allows inference of a false conclusion from true premises, the belief one derives could, as far as sheer logic goes, be false despite the truth of its inductive premises and one's infallibly believing them. True premises, even if infallibly believed, simply do not absolutely guarantee the truth of a conclusion inductively inferred from them. Hence, beliefs of such inductively inferred conclusions would be fallible (unless they happened to be self-grounding or to have necessary truths as objects). But then, being fallible, these beliefs would be capable of falsehood and hence would not

constitute knowledge. Thus, knowledge can be inferentially transmitted only by deductive inference. *Only valid deduction inferentially preserves infallibility.*¹⁹

If one thinks of knowledge as entailing absolute certainty, one might again be drawn to the entailment principle. For even if a fallible belief can be absolutely certain, a belief that is only inductively based on it will presumably be at least a bit less certain and thus not absolutely certain. For the proposition believed – the conclusion belief – is supported by the original belief only with some (perhaps high) degree of probability, rather than with absolute certainty, as where the conclusion is entailed by the premises. This would allow that the premise belief be certain and the conclusion belief not certain (or less so), since it would not inherit from the premise belief the same degree of protection against falsehood.

To see this, suppose that the premise belief only minimally meets the standard for absolute certainty. Then a belief inductively grounded on it can fall below that standard and thereby fail to be knowledge. Putting the point in terms of probability, we might imagine a case in which our premise meets the bare minimum conditions for absolute certainty, which we might represent by a probability of 1. Then, any conclusion that follows only inductively from this premise will thereby inherit from it only some lower probability and hence fall below the minimum level for absolute certainty. Thus, again the skeptic will argue that only deduction is sufficient to transmit knowledge.

But we have already seen reason to doubt both the infallibility principle and the view that a belief constitutes knowledge only if its status is absolutely certainty. Indeed, I do not see that skeptics give us good reason to believe either these principles or the entailment principle. It does not follow from the absence of good arguments for the principles that they are, as they appear to be, false; but if there is no good reason to believe them, even skeptics would approve of our refusing to accept them.

Absolute certainty is a high, and in some ways beautiful, ideal; but it is neither adequate to the concept of knowledge nor appropriate to the human condition.

The authority of knowledge and the cogency of its grounds

There is one further principle we should consider, one rather different from those examined so far and apparently more modest. It derives in part from the idea that if you know something, you have a certain authority regarding it, an authority presumably due to your being in a position to see the truth which you know. This authority is in part what accounts for the possibility of knowledge through testimony: if you know something, you have an authority about it such that normally I can come to know it, as well as to acquire justification for believing it, from your testimony.

Epistemic authority and cogent grounds

Indeed, if you tell someone that you know something – especially where you are asked whether you *really* know it – you put yourself on the line. It is as if you gave your firmest assurance – an epistemic promise, as it were – that it is true. If it turns out to be false, your position is somewhat like that of a person who has broken a promise. You are open to a kind of criticism and may have to make amends. A good theory of knowledge should account for this *epistemic authority* that seems to go with knowing.

A stronger, but closely associated view is that if you know that a proposition is true, then you must be able to say something on behalf of it. After all, the question how one knows is always intelligible, at least for beliefs that are not of self-evident truths or propositions about one's current consciousness (two kinds not in question for the most important kinds of skepticism); and if one really does know, one should be able to give more than a dogmatic answer, such as 'I can just see that it is true'.

The associated principle might be expressed in what I shall call *the cogency principle*: with the possible exception of beliefs of certain self-evident propositions and propositions about one's current consciousness, one knows that something is so only if one has grounds for it from which one can (in principle) argue cogently for it on the basis of those grounds.

Since the cogency principle requires only that one *can* argue cogently for what one knows, temporary inability to mount an argument would not prevent one's knowing. Even little children might have knowledge, for perhaps if they could just find a way to express their grounds they could provide cogent arguments. And since self-evident propositions and propositions about one's current consciousness are knowable even according to most skeptics, and may be objects of directly justified belief, there is a stopping place in epistemic chains and no regress need result when one produces a series of arguments to support a claim. What is known must simply be either traceable to those secure foundations or otherwise defensible by appeal to adequate grounds.

If the cogency principle is combined with the entailment principle, it will immediately preclude anyone's having knowledge on inductive grounds; for the entailment principle implies that inductive grounds are never cogent. But it need not be combined with the entailment principle. If it is not, it can allow for inductive reasoning of certain kinds to be cogent and thereby to transmit knowledge.

Even a moderate skeptic, however, is likely to accept at most a restricted kind of induction, a kind whose premises make its conclusion at least close to certain. This kind meets a higher standard than is usually applied to inductive inference. Thus, even though the cogency principle is separable from the entailment principle, it need not be combined with the entailment principle to be very hostile to the commonsense view that we can know the sorts of things I have been suggesting we can know, at least if this view is understood in a foundationalist framework. For this principle strikes at some

of the main sources of knowledge as they are plausibly understood, and it threatens to undermine our claim to knowledge of the past, the future, and the external world. Let us pursue this.

It is true that some of our beliefs that constitute direct knowledge (and are directly justified) can be supported by apparently more secure premises. For instance, my belief that I see a green field before me can be supported by premises about how things appear to me, which concern only my present consciousness. After all, that this is so seems to be the best explanation of why my visual field contains a green field. But this supportability by premises need not hold for *everything* it seems reasonable to regard as directly known. It may not hold in the case of apparently memorial knowledge. As we saw in Chapter 2, one might know something through the success of one's sheer retentive powers even when the only premises one knows or is justified in using to support it fail to justify it.

A proponent of the cogency principle would certainly tend to deny that my memory can be trusted as a source of direct knowledge or direct justification, in part because memory seems far more liable to error than perception. Moreover, I might be unable to provide good inductive reasoning to support the reliability of memory even in cases where it is very vivid, if only because such reasoning would require my depending on my memory for my justification in believing its premises, say premises about how often my past memory beliefs have been confirmed. To summarize their track record, I must remember how things turned out in the past – or at least remember that I wrote the results down as they occurred. I would thus be relying on memory in order to vindicate it. Still, even if I could give no cogent argument to justify my memorial beliefs, it does not follow that they are not justified, or that they do not constitute knowledge.

Grounds of knowledge as conferring epistemic authority

Must we accept even the apparently modest cogency principle, which requires that in order to know something, one needs grounds for it from which one can (in principle) argue cogently for it? I do not see why. Certainly one can have a kind of authority without being able to defend it by premises or exhibit it in argumentation. Consider, for instance, someone who can always tell "identical" twins apart but cannot say how. Moreover, saying 'I see it' need not be a dogmatic answer to 'How do you know?' It may simply specify one's grounds, as where one says, 'I see it' in answer to 'How do you know there is still ice on the road?' It says *how* one knows; it need not (though it may) *show* that one does, particularly if showing this requires more than exhibiting an appropriate source of the challenged belief.

There is a general lesson here. When skeptics ask how we know something, this is typically a challenge to show it. I have already argued that

knowing something does not require being able to show it, so this challenge is not always appropriate. What I now want to stress is that the skeptical emphasis on 'How do you know?' as a request for a *demonstration* must not be allowed to obscure the possibility of taking it as just a request to specify a *ground* of one's knowledge and of fulfilling that request simply by giving an adequate ground. In doing this successfully, one shows how one knows.

In saying how I know something by citing my ground I may also be doing something further: *exhibiting* my knowledge of what the ground supports and perhaps even of the fact that it does so. But even exhibiting knowledge need not be showing that one has it; and even if I know that, say, seeing ice on the roads shows that they are icy, I may not know how to use this connection to show that I know they are. Still, if my knowledge that they are carries an authority that can be confirmed by my citing my ground, why need I be able to go on to the sophisticated task of showing that I have knowledge? Again the analogy to virtue is pertinent: having it requires being able to manifest it in appropriate circumstances, but not to be able to show – in any further way – that one has it.

Exhibiting knowledge versus dogmatically claiming it

One might think this approach licenses dogmatism. Granted, saying 'I see it' could be dogmatic if intended to show conclusively that I know, for instance by serving as absolutely proving that there is green grass before me. But the same words can simply indicate the basis for my knowledge. This is different from flatly claiming that I have knowledge. Indeed, saying it exhibits a ground for my belief which, if adequate, suggests that I am not being dogmatic in taking myself to know. Perhaps it is precisely because the skeptic's 'How do you know?' is commonly *meant* as a challenge to be shown conclusively that one knows, and not as a request to specify a source or a ground of the knowledge, that saying 'I see it' *seems* dogmatic in the context of discussing skepticism, even when the function of saying this is mainly explanatory.

If the issue raised by skepticism is whether we can show that we have knowledge, the point that an appeal to visual experience does not conclusively establish visual knowledge is an important concession. But the issue here is whether the skeptic succeeds in showing that we do not have perceptual knowledge. In that context, the point is not a concession. Once again, we can see how skepticism can gain credibility because skeptics make it sound as if their case against the existence of one or another kind of knowledge succeeds if we cannot show that there is such knowledge. In fact, we need not be able to show that there is knowledge in order to have it; and the skeptic must give us good reason not to believe that there is knowledge.

Refutation and rebuttal

Have I, then, refuted skepticism, even in the few forms considered here? I have not tried to. *Refutation* would require showing that those forms of skepticism are false, which in turn would entail showing that there is knowledge (and justified belief). What I have tried to do is to *rebut* skepticism in certain plausible forms, to show that the arguments for those skeptical views do not establish that we do not have knowledge (and justified belief). Refutation of skepticism suffices for its rebuttal; but rebuttal does not require refutation. Now suppose I have succeeded in rebutting skepticism. Where do we stand? May we believe that we have knowledge, or may we only suspend judgment both on this and on skeptical claims that we do not?

I have already argued, by implication, that one can know something without knowing that one knows it. For instance, in arguing that much of our knowledge is not self-conscious, I indicated how I can know that there is a green field before me without even believing that I know this. I do not even form such self-conscious beliefs in most everyday situations. Moreover, toddlers, who do not understand what knowledge is – and so are not in a position to believe they know anything – can apparently know such simple things as that Mama is before them.

Even if I did have the second-order knowledge that I know the field is there, I surely would not possess (if it is even possible for me to possess) the infinite series of beliefs required by the view that knowing entails knowing that one knows – *the KK thesis*, as it has often been called – the series that continues with my knowing that I know that I know; knowing that I know, that I know that I know; and so forth. There is no plausibility in thinking that if I know that (for example) the field is there, I must know that I know that I know ... this, up to the limit of my capacity. I never have such a repetitive thought. Moreover, I do not think that I believe (or disbelieve) the proposition in question (I have not tested my memory here); and I cannot imagine a good use for it.²⁰ Given these points (among others), it would be a mistake to think, as some skeptics might like us to, that if we do not know that we have knowledge, then we do not.

For similar reasons, it seems possible that we might be justified in believing that we have knowledge even if we are properly unwilling to claim that we know we do, and perhaps even if we are properly unwilling to claim justification for believing that we do. Let us explore these possibilities.

If foundationalism is correct, then if one can know anything, one can know at least something directly. Moreover, some of the sorts of things that a plausible foundationalism says we know directly – for instance, self-evident truths and some propositions about our present consciousness – are the kinds of things which, simply on the basis of reflection on the examples involved, it is plausible to think we know. Perhaps, of course, this reflection, even if it does not involve arguing from premises, shows that we have knowledge. In any case, I think that we are justified in believing that we

have some knowledge even if we cannot show that we do; and I am aware of no good argument against the view that we have some knowledge.

Might there be a way, however, to give a cogent, positive defense of common sense: to show that we have knowledge, even of the external world? And could we establish this second-order thesis even to the satisfaction of some skeptics? There is no satisfying a *radical skeptic*, one who denies that there can be *any* knowledge or justified belief (including justification of that very claim, which the skeptic simply asserts as a challenge). For a radical skeptic, nothing one presents as a reason for asserting something will count as justifying it.

Could anything be said, however, to show that there is knowledge that might be plausible to a *moderate skeptic*: one who holds, say, that although transmission of justification and knowledge must be deductive, we may justifiably believe, and perhaps know, at least self-evident propositions and propositions about our present consciousness? Even if the answer is negative, perhaps we can show that there is knowledge, or at least justified belief, whether any skeptics would find our argument plausible or not.

Prospects for a positive defense of common sense

How might an argument for a positive defense of common sense go? Let us consider justified belief first, since showing that certain of our beliefs are justified, unlike showing that some of them constitute knowledge, does not require showing that the beliefs in question are true.

A case for justified belief

One might view the issue this way: if we are to show that there are justified beliefs, then one result of our argument will itself be producing justification, specifically justification for the second-order belief that there are justified beliefs. For to show something by argument is at least to produce justification for believing it.

If we are to provide such second-order justification, we apparently need at least two things: a general premise expressing a sufficient condition for justification, and one or more specific premises saying that a particular belief meets that condition. For instance, the general premise might be the justification principle that

- (1) An attentive belief to the effect that one is now in an occurrent mental state, such as thinking, is (prima facie) justified,

where an *attentive* belief is one based on careful attention to the matter in question, and where the justification is not absolute but prima facie; it must simply be strong enough to make it appropriate for a rational person to hold the belief.²¹ The particular premise might be that

- (2) I have an attentive belief that I am now in such a state, namely thinking.

If I am justified in believing these premises, I would surely be justified in inferring deductively from them, and thereby in believing on the basis of them, what they self-evidently entail, namely, that

- (3) My belief that I am thinking is (prima facie) justified.

Here I would be inferentially justified in believing (3), at least if I can hold all three propositions before my mind in a way that avoids dependence on memory of my premises; and this seems possible for me. (If I needed so many premises, or such complicated premises, that I could not hold them in mind at once, then my justification for my conclusion would depend on that of my memory belief(s) of my premises.) Premises (1) and (2) self-evidently entail (3), and a moderate skeptic will very likely grant that if, from premises that I am justified in believing, I infer (without dependence on memory) a conclusion that self-evidently follows from them, I am justified in believing that too.

But how am I now justified in believing premises (1) and (2), if I am? There is some plausibility in holding that the general principle, (1), is justifiable directly (non-inferentially) by reflection, and so my belief of it might itself be directly justified. This is not to deny that it *could* be justified by prior premises; the point is only that it is arguably justified by reflection not dependent on one's appealing to such premises. As for the particular premise, (2), I might be directly justified in holding it by virtue of a justification principle similar to the general one, but applying to beliefs, a principle to the effect that if, on careful introspection, one believes that one attentively believes a proposition, *p*, then one is justified in believing one does (presumably directly justified, if one has introspected carefully).

Now if my belief of my general premise is justified, and if I may justifiably hold the particular premise, then surely I may justifiably conclude that I am justified in my belief that I am thinking. I may justifiably conclude this even if my justification in believing my premises is not direct, as I am tentatively assuming it is. Moreover, if my beliefs of (1) to (3) are true, they may also constitute knowledge: my justification for each seems strong enough, and apart from this matter of degree there appears to be no other kind of bar to knowledge.²²

Given the plausibility of the premises just used to try to show that I am justified in holding a belief about my own mental life, I am inclined to think that it can be shown that there are some justified beliefs, even some justified empirical ones. But even if the line of argument I have used is successful, one might question whether it extends to any beliefs about the external world. What would be our general principle for, say, visual perceptual beliefs?

In answer, perhaps we might begin with an instance of a justification principle stated in Chapter 1. Applied to the green field example, this would say that:

- (a) If one believes, on the basis of a vivid and steady visual experience in which one has the impression of something green before one, that there is something green before one, then one is (*prima facie*) justified in so believing.

Surely we may say this, particularly since the justification in question is admittedly defeasible. (It could, for instance, be undermined by my knowing that I have frequently been hallucinating greens lately.) Suppose this premise may be believed with direct justification, and we may also believe (possibly with direct justification) that:

- (b) I have a belief (that there is something green before me) grounded in the way the premise – principle (a) – requires.

Then I may, much as before, justifiedly conclude that:

- (c) I am (*prima facie*) justified in believing that there is something green before me.

To be sure, my perceptual belief is only *prima facie* justified – roughly, justified in the absence of defeating factors. But this is still a significant conclusion, even if (as seems possible) I could not, by reflection alone, rule out all of those defeaters.

The regress of demonstration

Supposing this line of argument against the skeptic is sound, have I shown anything? If showing something is producing a good argument for it from true premises that one is justified in believing, presumably I have. It is easy, however, to think that the skeptic would be correct in denying that I have shown anything. For there is a subtlety here that is easily missed. Even if I have shown my conclusion, I might not be justified in *saying*, to the skeptic or anyone else, that I *have* shown it, or even in believing I have shown it. For justification for asserting or believing that second-order proposition about my beliefs would ordinarily require holding (or at least having justification for holding) *third-order* beliefs, such as the belief that my second-order belief that I believe I am thinking is justified and true (since this second-order belief has been shown by good argument). And what in my situation would give me that still higher-order justification?

The general point is that whatever one's justificational or epistemic achievement, justifiedly *saying* or even justifiedly believing that one has

succeeded in it requires justification or knowledge at the *next* higher level. This higher-order justification or knowledge may or may not be forthcoming. Initially, this point may seem to doom my original attempt to show that I have a justified belief. But I do not think it does; it brings out only that one can show this without automatically showing the higher-order proposition that one *has* shown it. Plainly, we can achieve something even if we are in no way entitled to credit ourselves with achieving it.

What we have encountered here is a counterpart of the regress of justification, a *regress of demonstration*: if one shows anything at all, there will be some unshown shower (at least one unshown premise). In a sense, there will always be some point at which self-congratulation – or a final dismissal of the skeptical challenges – is inappropriate. This is another reason why a moderate foundationalist perspective should not be thought to lead to dogmatism.

Still, even if I do not know that I have shown that my belief that I am thinking is justified, I may yet *have* shown this; and if I have, then I may well know the proposition that I have shown: that my belief that I am thinking is justified. Perhaps, moreover, a similar procedure may be repeated, with equal success, at each higher level to which one can ascend without losing track of the progressively more complex issue. Then, with sufficient patience, one could show any given justification-ascribing proposition in the hierarchy – that one has shown that one has a justified belief, that one has shown that one has shown this, and so on, to the limit of one's comprehension.

A case for knowledge

The reasoning we have been exploring in connection with justification also bears on skepticism about knowledge. If the premise beliefs, (a) and (b), are true, they may constitute knowledge. I may, then, not only be showing that I am justified in holding a belief about the external world; I also may, as a result of my reasoning, know that I am justified in holding it. There would then be at least justificational knowledge: self-knowledge to the effect that one has justification for some beliefs.

In reasoning from (a) and (b) to (c), however, I do not automatically know that I am showing that I am justified in believing something about the external world. Suppose I do not know this. Perhaps I only hope that I am showing it. Then, even if I do have second-order knowledge that I have a justified belief about the external world, I may not be justified in holding the *third-order* belief that I have (second-order) knowledge that I have this (first-order) justified belief, the belief that there is something green before me. I have as yet no principle that would justify me in concluding that I know or justifiedly believe that I have a justified first-order belief. I lack a principle stating conditions that generate second-order knowledge or second-order justified belief.

It seems, however, that the sort of justification I apparently have for all the relevant beliefs, including the belief that I have a justified belief about the external world, is the kind whose possession by *true* beliefs is sufficient for their constituting knowledge. Thus, through reasoning using premises like (a) and (b) I may well know that I have justified beliefs about the external world. Certainly I have reason to think that the skeptic does not know, or justifiedly believe, that I *lack* justified beliefs about the external world.

One assumption it is natural to make in using this strategy against skepticism deserves emphasis: the assumption that the crucial principles of justification are a priori, and believing them is justified by reflection (directly, or at least on the basis of self-evident steps from directly justified beliefs of a priori premises). Suppose the principles are empirical. Then our justification for believing them would presumably be broadly inductive. A skeptic could plausibly deny that, on an inductive basis, we can justifiedly believe them. Let us pursue this possibility.

A circularity problem

There would apparently be a circularity problem if we had to justify our crucial principles inductively. For justifying them by inductive reasoning would seem to presuppose using just such principles, principles that specify, for instance, under what conditions inductive inference can transmit justification or knowledge. We would have to use induction to develop a track record for such inferences, say by determining, observationally, how often their premises are found true and their conclusions determined to be false. We would need to rely on perception and memory to do this – in addition to using induction to infer from a good track record on the part of a source to its general reliability. To acquire justified beliefs about the reliability of perception, moreover, we would need to use perception, for instance in looking at objects again to see if our initial color perceptions were accurate.

Are the kinds of principles of justification I have been using a priori? That is certainly arguable; but it is also controversial. On the most plausible kind of reliability theory of justification, for instance, a belief is justified by virtue of being grounded in reliable belief-producing processes such as perceptual ones; and it is apparently not an a priori matter what processes are *reliable*, that is, actually produce a suitably large proportion of true beliefs. This is the sort of thing that must be determined largely by observation.²³

Thus, for reliabilism, in order to know what principles account for justification, one must know what processes tend to generate true beliefs. One could determine that only through considerable experience. Hence, if these principles are empirical, the circularity problem just mentioned would beset the attempt, within a reliabilist framework, to justify them.

On the other hand, I argued above that reliability theories are less plausible for justification than for knowledge, and I believe that it is more reasonable, though by no means obviously correct, to suppose that at least some principles about the conditions for justification are a priori. I would include various principles expressing ways in which – as described by Chapters 1 through 4 – justification is produced by its basic sources.

Skepticism and common sense

Where, then, does this chapter leave us with respect to appraising skepticism? To begin with, there are forms of skepticism I have not mentioned, and I have also not discussed every plausible argument for the skeptical principles I have addressed: chiefly the infallibility, certainty, back-up, entailment, and cogency principles. But these principles are in some important ways representative of those on which even moderate skepticism rests. I have offered reasons for rejecting them, and on that basis I have maintained that skepticism, at least insofar as it depends on these and similar principles, can be rebutted. It can be shown to be rationally resistible.

We are, then, warranted in refusing to accept skepticism concerning justification and knowledge of propositions other than those that are self-evident or attribute to the believer a present occurrent mental property. If it is not false, it is at least not justified by what seem the main arguments for it. It is not clear, however, that anything said above refutes the kinds of skepticism we have considered. For refuting those views entails showing them to be false, and it is not altogether clear what that requires.

Positively, I have suggested that on one plausible notion of showing something, namely, validly and justifiably deducing it from true premises which one justifiedly believes and are good grounds for it, we *can* show that there are some beliefs we are justified in holding, probably even some justified beliefs about the external world (these may perhaps include some about the inner lives of others).²⁴ I am less inclined to say that we can – by this strategy – show that there is knowledge, particularly knowledge of the external world. Much depends on the kind of grounding required for such premises as that I see a green field before me, which, because simple seeing *entails* the existence of the object seen, in turn entails that there is something external. Much also depends on how rigorous a standard of showing is appropriate.

I have argued, however, for the commonsense view that we can know that there is both justified belief and knowledge about the external world, and can know this even if we cannot show that there is. I also maintain that there is justified belief and knowledge about one's own consciousness and about certain a priori matters. Sceptics certainly do not seem to have shown that we do not have knowledge and justified belief of these kinds. I believe that we have both.

Moreover, if, as argued in Chapter 9, it is true that rationality is a more permissive notion than justification, then whatever the anti-skeptical case for our having justification, it will count more strongly for the counterpart views concerning the rationality of our beliefs and other epistemic attitudes. Even if rationality, as applied to beliefs, is significantly weaker than justification, it is still the kind of status skeptics tend to deny is ever achieved by our beliefs about the external world, the past, and many other things.

Perhaps viewing knowledge, justification, and rationality in the way I have might be thought to be an article of epistemological faith. I do not think it is; but the difficulty of determining whether it is partly an article of unverifiable faith, or can be established by cogent argument, or is more than the former yet less than the latter, is some testimony to the depth and complexity of skeptical problems.

Notes

- 1 Some writers on skepticism prefer to raise skeptical possibilities using the dream case rather than the hallucination one, perhaps in part because Descartes so famously used a dream argument in his *Meditations on First Philosophy*. For relevant discussion, esp. of the dream argument, see, for example, Barry Stroud, *The Significance of Philosophical Skepticism* (Oxford: Oxford University Press, 1984); Michael Williams, *Unnatural Doubts* (Oxford: Basil Blackwell, 1991); and Robert Fogelin, *Pyrrhonian Reflections on Knowledge and Justification* (Oxford and New York: Oxford University Press, 1994). Another way to raise skeptical possibilities is to imagine that one is a "brain in a vat," i.e., that one's brain is kept alive artificially in a liquid and given just the sorts of stimulation it now has, so that one would presumably experience things just as one now does. This example derives from Hilary Putnam, who uses it to argue, against skepticism, that the very content of such sentences as 'I am a brain in a vat' prevents their being intelligibly thought in the way Descartes apparently believed they could be. Owing to how language and conceptualization work,

although the people in that possible world [in which they are brains in a vat over their whole existence] can think and 'say' any words we can think and say, they cannot (I claim) refer to what we can refer to. In particular they cannot think or say that they are brains in a vat (even by thinking 'we are brains in a vat').

See the selection from Hilary Putnam, *Reason, Truth and History* (Cambridge: Cambridge University Press, 1981), entitled 'Brains in a Vat,' in Huemer, op. cit., pp. 524–38. The literature contains much critical discussion, including Huemer's 'Direct Realism and the Brain-in-a-Vat Argument,' reprinted (from *Philosophy and Phenomenological Research* 61, 2, 2000) in Huemer, op. cit., pp. 575–88.

- 2 Pyrrhonian skepticism need not imply that suspending judgment is always psychologically possible; and it may also allow that one can accept a proposition for the sake of argument, and so, without believing it, act on it as one who believes it would act; but there is no need to discuss this position here. If what I say about skepticism in general is sound, it can be readily applied to the Pyrrhonian form.
- 3 See, for instance, Section IV of Hume's *Inquiry Concerning Human Understanding* (first published in 1748).
- 4 This formulation is roughly the one given by Bertrand Russell in *The Problems of Philosophy* (Oxford: Oxford University Press, 1912). The formulation should be understood to apply only to phenomena of a certain kind, such as we explore in scientific inquiry and much everyday investigation: it will not hold for certain special cases. For

instance, with an increase in the number of instances in which I lose a fair sweepstakes in which I hold one of a million coupons, there is *no* change in the probability that I will lose; the probability remains the ratio of the number of tickets I hold – one – to the total number: 1 million. To think my good day is now more likely to come is to commit the gambler's fallacy.

- 5 I cannot take time here to consider begging the question in any detail; it is an important but elusive notion. For a detail treatment see Walter Sinnott-Armstrong, 'Begging the Question,' *Australasian Journal of Philosophy*, 77, 2 (1999), 174–91.
- 6 A remark attributed to Bertrand Russell introduces (humorously) the possibility that we may, perhaps haphazardly, be sometimes right and sometimes wrong about other minds: "There may be other minds," he quipped, "but there certainly aren't many of them." In this spirit I might note another twist to the problem of other minds. As usually conceived, it concerns whether, in effect, there are as *many* minds as there seem to be. But reflection on skepticism can also lead us to a converse worry. How do I know that when I am asleep my body is not taken over by another mind, one connected, perhaps, with a part of the same brain as goes with my mind? And why might there not be several others who control this body when I do not? Granted there could be a conflict with another mind over, say, the movements of my right arm; but I could be built (or programmed) so as never to be conscious when another mind takes over this body. Call this the *problem of too many minds*.
- 7 There is much recent literature on the extent to which skeptical hypotheses undermine commonsense views about the extent of our knowledge. See, for instance, Burchvarov's *Skepticism and the External World* (New York, 1998), Moser's *Philosophy after Objectivity* (Oxford and New York, 1993) and the books by Fogelin, Huemer, Klein, Nozick, Rescher, Sosa, Stroud, and Michael Williams also cited in the bibliography. For a helpful discussion of skepticism focusing on the underdetermination problem, see Jonathan Vogel, 'Dismissing Skeptical Possibilities,' *Philosophical Studies* 70 (1993), 235–50.
- 8 This is the term widely used in recent literature in connection with a debate between, on the non-skeptical side, Hilary Putnam, who (as quoted in note 1) denies that one really could be both a brain in a vat and have the mastery of language needed to raise the question of skepticism, and, closer to the skeptical side in interpreting such examples, a number of other philosophers, including Anthony Brueckner, 'Trying to Get Outside Your Own Skin,' *Philosophical Topics* 23, 1 (1995), 79–111, which contains references to Putnam's original discussions of the brain-in-a-vat problem and a number of more recent discussions. See also the paper by Huemer cited in note 1 and his *Skepticism and the Veil of Perception* (Lanham, MD: Rowman and Littlefield, 2001).
- 9 See, for example, Thomas Aquinas's *Summa Theologica* (written in the thirteenth century), Ia, question 25, a.3.
- 10 As an epistemic principle, (2) – and indeed (1) and (3) as well – would commonly be taken by proponents to be necessarily true. Thus (2) would rule out even the possibility of knowledge of falsehoods, as opposed to the mere occurrence of it. But for our purposes the formulations as more simply stated will serve.
- 11 That my thinking entails my existing does not, of course, entail that my existing entails my thinking. But Descartes's claim (also in *Meditation II*) that his essence *is* to be a thinking thing, led to the following joke (which I recount as I remember it). *Bartender to customer*: Do you want another? *Customer*: I think not. *Outcome*: the customer disappears.
- 12 If there are such laws, then there is empirical grounding that is conclusive in the sense that it implies the proposition it grounds with "natural necessity", the kind appropriate to causal laws. Since those are not absolutely necessary, as are logical laws and necessary truths as described in Chapter 4, it would not follow that the implication is an entailment.
- 13 This is the kind of strategy taken by Perer D. Klein in *Certainty: A Refutation of Skepticism* (Minneapolis: University of Minnesota Press, 1981); he extends it in the paper cited in Chapter 6 and reprinted in Huemer's collection, pp. 552–74. In part, the issue concerns

- whether we may simply take it as obvious that we do know certain things before we have a *criterion* of knowledge, e.g. an account that tells us both what knowledge is and whether beliefs constituting it must be infallible. For detailed discussion of this *problem of the criterion* – the problem of whether cases of knowledge are prior to accounts or vice versa – see R.M. Chisholm, 'The Problem of the Criterion,' in his *The Foundations of Knowing* (Minneapolis: University of Minnesota Press, 1982), reprinted in Huemer, op. cit., 590–601; and Robert Amico, *The Problem of the Criterion* (Lanham, Md.: Rowman and Littlefield, 1994). For an examination of Moorean commonsensism in relation to criteria, dogmatism, and skepticism, see James Pryor, 'The Skeptic and the Dogmatist,' *Noûs* 34, (2000).
- 14 There is a sophisticated and plausible compromise with skepticism that deserves note here. One could argue that knowledge must be understood not unqualifiedly but in terms of relativization to "contrast classes". Thus, relative to the contrast class of near perfect imitations, I do not know that there is a green field before me; relative to the contrast class of ordinary green things encountered in the same visual way, such as green ponds and green canvas laid out for picnicking, I do. For a detailed statement of this view (which may be regarded as a kind of contextualism), see Walter Sinnott-Armstrong's chapter on moral skepticism in Walter Sinnott-Armstrong and Mark Timmons (eds) *Moral Knowledge?* (New York: Oxford University Press, 1996). One reply is that knowledge may be understood without relativization provided we at least (1) recognize that knowledge *attributions* occur in a context and are in some way relative to it, and (2) distinguish between the following kinds of things: (a) knowing there is a green field before one and (b) knowing that there is a green field before one *as opposed to* a pavement textured to look just like one. Knowing the former does not entail knowing the latter, contrastive proposition, but it is easily taken to require that because (i) the latter proposition apparently follows from the former and (ii) *noting* the possibility of ignorance of the latter is an intelligible challenge to a *claim* to know the former.
 - 15 The 'presumably' is meant to allow that there is a sense of 'show' in which one need not have justification for one's premises: if they are true and are cogent grounds for what one wants to show, then invoking them may serve to show it. Here, however, one might not be justified in believing the very thing one shows. This objective, externalist way of showing – call it *de facto showing* – something is not the one of interest here, which might be called *dialectical showing*, since it figures crucially in philosophical disputation.
 - 16 It does not strictly follow unless we define regularity to preclude the following kind of thing: the sun rises every day except every trillionth after the Earth came to be, where tomorrow *is* the trillionth.
 - 17 As noted in Chapter 4, the synthetic a priori might be necessary without being logically necessary in the strict sense; similarly, something can be a synthetically necessary falsehood (thus impossible, as is a round square), but not strictly logically impossible.
 - 18 Epistemic possibility for us is sometimes characterized simply in terms of what is possible given what we *know*, but I think it is appropriate to include justification here, if only because we may be better able to tell by reflection on our current overall cognitive resources what we are justified in believing than what we know. For an account of the notion and some of the difficulties confronting it, see David Chalmers' 'Epistemic Possibility' (MS.).
 - 19 This should be taken to apply to non-formal validity: infallibility would be preserved by the inference from something's being round to its not being square, but this is a case of synthetic a priori, as opposed to formal, entailment.
 - 20 Reasons for doubting that we should posit such beliefs are given in my 'Dispositional Beliefs and Dispositions to Believe,' *Noûs* 28, 4, 1994, 419–34.
 - 21 We might perhaps think of the degree of justification in question here as the kind that is sufficient to render a true belief knowledge, in the absence of the sorts of cases discussed in Chapter 8 showing that justified true belief need not constitute knowledge.
 - 22 I ignore here the point that I might have only situational justification for my conclusion if my believing it is not *based* on my believing my premises. Note 24 comments on this problem.
 - 23 For extensive discussion of the kind of circularity problem in question here see William P. Alston, 'Epistemic Circularity,' *Philosophy and Phenomenological Research* 46 (1986). Cf. Ernest Sosa, 'Philosophical Scepticism and Epistemic Circularity,' *Proceedings of the Aristotelian Society* supplementary volume 68 (1994), followed by a response from Barry Stroud.
 - 24 I have not directly argued that there are *justified beliefs*. For I have not argued for the premise, apparently needed for this conclusion, that we are non-inferentially justified in believing that the relevant beliefs, such as the belief that there is something green before me, are *based* on the visual impressions constituting one's grounds for it. This basing is partly causal, and skeptics are likely to argue that justification for attributing causal propositions requires inductive, hence inferential, grounds. This is not self-evident, and I have challenged it in 'Causalist Internalism,' *American Philosophical Quarterly* 26 (1989), reprinted in my *The Structure of Justification* (Cambridge and New York: Cambridge University Press, 1993). In any case, even if I am not justified in believing that my external world belief is based on a sense impression, but only that I *have* the relevant impression, I am justified in believing that I *have* this justification for the external world belief: we might say that I am entitled to hold it even though I may not hold it on the basis of my entitlement.