

Chapter Two

The Structure of Skeptical Arguments and Its Metaepistemological Implications

Kinds of Skepticism

My primary concern in this chapter is to sketch what I take to be the most interesting form of skeptical argument and to examine in a preliminary way the metaepistemological presuppositions, if any, of skeptical arguments. I also attempt to address here some important challenges to the skeptic's method that do not explicitly rely on metaepistemological positions.

I hasten to emphasize that my concern in setting forth what I take to be the most interesting form of skeptical argument is not primarily historical. I do think that one can find the kind of skeptical argument I discuss developed in some detail by the Modern philosophers, and that its paradigm expression can be found in the writings of David Hume, but I will not defend this claim. Much of my view about the structure of skeptical argument also parallels the excellent discussion of this issue in chapter 2 of Ayer's classic book *The Problem of Knowledge*. I also remind the reader that the kind of skepticism in which I am primarily interested involves claims not about knowledge but about the epistemic rationality of belief. I have already observed that the historical discussion of skepticism is couched mostly in terms of what one can or cannot know. I am interested in what one can or cannot rationally believe.

Perhaps we should begin by making some familiar distinctions between kinds of skepticism. One distinction we have just marked—the distinction between skeptical claims about knowledge and skeptical claims about epistemically justified or rational belief. Let us refer to these two kinds of skepticism as *weak* and *strong* skepticism,

respectively. As we noted earlier, the claim that one cannot *know* something (particularly if one adds the appropriate dramatic emphasis to “know”) is relatively weak from the philosophical perspective in that one can easily concede this sort of skepticism with a kind of resigned respect for the fallibility of human consciousness. The claim that one cannot even rationally believe that something is true is much stronger, much more extreme. It seems to give the person who wants to continue believing what common sense requires nothing on which to fall back.

In addition to distinguishing strong and weak skepticism, we can also distinguish global and local skepticism. The global skeptic makes a claim about our epistemic access to *all* truth. Specifically, weak global skepticism maintains that one has no knowledge of anything. Strong global skepticism maintains that one has no epistemically rational beliefs about anything. Local skepticism is skepticism (weak or strong) with respect to a given class of propositions. Thus, we can be a skeptic with respect to propositions about the physical world, the past, other minds, the future, theoretical entities in physics, the existence of God, or any other subclass of propositions.

Skeptical claims can be made with or without modal operators. Thus the skeptic can claim only that we do not have knowledge or rational belief or that we *cannot* have knowledge or rational belief. When we examine more closely the kinds of arguments skeptics advance in support of their claims, we shall see that they typically support the stronger modal claims.

One must also distinguish first-level skepticism from second-level skepticism, or skepticism from what we might call metaskepticism.¹ Second-level skepticism involves skeptical claims about whether or not we have knowledge or rational belief. It is argued by some (even some externalists) that if certain versions of externalism are true, it may make first-level knowledge or rational beliefs *possible* only to invite skepticism about whether or not one ever has such knowledge or rational belief. Some might hope to concede the externalists’ claims at the first level but allow for the legitimacy of traditional skeptical concerns at the next level. I attempt to show later on that from the philosophical perspective, nothing of any interest should change when one moves up a level and that the widespread feeling that something *does* change when one is operating from within an externalist framework has enormously significant implications.

So again, the kind of skepticism with which I am primarily concerned is strong skepticism. Historically, I think there have been very few *global* strong skeptics. It is, of course, the most paradoxical of skepti-

cisms because it entails that one has no epistemic justification for believing it. Whether we should conclude that it is therefore of no philosophical interest is something that we discuss toward the end of this chapter when we examine various charges of self-refutation leveled against the skeptic. The vast majority of skeptics, I argue, have actually presupposed knowledge or justified belief with respect to some class of propositions. Skeptics in the empiricist tradition almost all seemed to presuppose unproblematic access to occurrent mental states. Indeed, the presupposition was so complete that one rarely even finds the Modern philosophers *raising* the question of whether or not one can know that one is in a certain subjective mental state. Furthermore, almost all skeptics seemed to presuppose knowledge of at least logical relations. They seemed to presuppose that one can recognize or “see” contradiction, at least some simple necessary truths, and at least some simple entailments.² As we shall see, the question of whether or not skeptics can “contain” their local skepticism is a matter of some controversy.

The Structure of Skeptical Arguments for Strong Local Skepticism

If one examines classic arguments for strong local skepticism, one can discover, I think, a recurring pattern. First the skeptic indicates the class of propositions under skeptical attack. Then the skeptic attempts to exhaustively characterize the most plausible candidate for something that could conceivably justify, or make rational, belief in this kind of proposition. Next the skeptic attempts to drive a logical wedge between the available justification and the proposition it is supposed to justify. The wedge is logical. The claim at this point is only that the justification available for that belief does not logically guarantee the truth of the proposition believed. It is conceivable that someone has precisely that sort of justification even though the belief in question is false. At this point, the Cartesian skeptic might end the argument with the weak skeptical conclusion that it is not possible to know *with certainty* the proposition believed. But this conclusion does not get one strong skepticism. The strong skeptic goes on to argue that the logical gap cannot be bridged using any legitimate nondeductive reasoning.

Let us try to illustrate the kind of skeptical argument discussed above with a few examples, and let us begin with familiar epistemological problems concerning our access to the external world. Our strong skeptic with respect to the physical world argues that it is not epistemically

rational for us to believe any proposition asserting the existence of a physical object. To what evidence might we appeal in trying to justify our belief in the existence of some object? The best evidence we could possibly get (according to common sense) is the testimony of our senses. The presumption is that if I cannot rationally believe that there is a table in front of me now when I seem to see and feel a table, there is no proposition describing a physical object that it would be epistemically rational for me to believe. But does any number of truths about the phenomenological character of my subjective and fleeting sensations ever logically guarantee the truth of any proposition describing the physical world? The answer, the skeptic argues, is clearly no.

To support this answer the skeptic will often appeal to the famous *skeptical scenarios*. A skeptical scenario is simply a description of a perfectly intelligible hypothetical situation in which someone has the best possible justification for believing a proposition about the physical world, even though that proposition is false. In the case of beliefs about physical objects, the skeptical scenarios describe hypothetical situations in which one has the best possible evidence in the form of sensation that, for instance, the table exists, even though it does not. The great fascination of skeptical arguments no doubt is owed in large part to their exotic appeals to the possibilities of dreams, hallucinations, malevolent demons, brains in a vat, telepathic powers, and the like. No matter how vivid my visual and tactile sensations may seem to me right now, who could deny that it is at least conceivable that I have these qualitatively same sensations in a vivid dream or in a drug-induced hallucination? And if I am dreaming or hallucinating, it would be mere chance that the table I take to be there exists.

Mad neurophysiologists with futuristic knowledge of the workings of the brain provide particularly useful grist for the skeptic's mill. Many of the antiskeptics have an almost slavish devotion to the dictates of science, and science seems to tell us that it is brain events that are the immediate causes of (or, on some views, are identical to) sensations. By stimulating the relevant part of the brain in the appropriate way, it seems in principle possible to produce the very electrical discharge that will cause me to seem to see a table. If we "tickle" another region of the brain in just the right way, we can get a tactile "table" sensation. Indeed, if our neurophysiologist of the future has enough skill, sophisticated instruments, and knowledge of the brain, there is no reason to suppose that the brain could not be played like a piano to produce the extraordinarily complex set of sensations associated with visiting the Grand Canyon. And if it can be done, it can be done surreptitiously

without the knowledge (or memory, if the knowledge once existed) of the subject. The intelligibility of the hypotheses seems hardly in question. It is the stuff of some extremely good and utterly *intelligible* literature and cinema.³ To be sure, the skeptic's appeals to such possibilities are not considered unproblematic in the context of a skeptical argument, and we shall examine some of the complaints shortly. Still, it seems to be almost obvious that one can distinguish sensations from truths about the physical world and that no conjunction of truths about sensations will ever entail a truth about the physical world.

But so what? Epistemically justified or rational belief does not require the inconceivability of error, and even weaker concepts of knowledge defined in terms of justified true belief seem to make knowledge perfectly compatible with the conceivability of error. It is here that Hume took the skeptical concern to its natural conclusion. If sensations can occur in the absence of the physical objects we take them to indicate, what reason do we have for supposing that it is even likely that when we have certain sensations, certain physical objects exist? Well, how do we establish one thing as evidence for the existence of another? Perhaps the most familiar pattern of inference we employ to answer such questions is inductive argument. We take dark clouds to be a good indication of an imminent storm because in the past we have *observed* a correlation between the presence of such clouds and subsequent storms. We take the sound of barking to be a reliable (if not infallible) indicator of the presence of a dog because in the past we have *observed* a correlation between the occurrence of such sounds and the presence of a dog.

If we use this model to understand our reasons for relying on sensation as an indicator of physical objects, then to avoid strong skepticism we would need to make plausible the claim that we have in some sense *observed* a constant or near constant correlation between the occurrence of certain sensations and the existence of certain objects. But this, of course, we cannot do. We cannot step outside sensation to compare the sensation with the physical object it is supposed to represent. To use an inductive argument for the conclusion that sensation is a reliable indicator of the presence of physical objects, we would need access to physical objects that is independent of sensation. Without such access we could never discover the necessary constant conjunction of sensation and object. But we have no access to the physical world except through our sensations. Assuming we have unproblematic access to the past, we can perhaps correlate sensations. We can discover all sorts of interesting connections between visual, kinesthetic, tactile, auditory, gustatory, and

olfactory sensations. But we can never step outside our mental states in order to correlate a mental state with something other than a mental state. And in the absence of our ability to discover correlations between the mental and the physical, we will never be able to rationally believe that there is a connection between the two.

In developing the above skeptical argument, one does need to establish the intelligibility of sensations occurring in the absence of the physical objects we take them to represent. Notice, however, that one can take the argument seriously and remain neutral on many of the controversies concerning the metaphysical analysis of sensation. One can think that the visual sensations that occur in the absence of physical objects involve our being related to another kind of object, a sense datum, or one can think that visual sensation should be understood as a nonrelational property of the mind or self (the so-called appearing or adverbial theory of sensation). All that one needs to take the skeptical challenge seriously is some understanding of sensation that allows us to speak meaningfully of the occurrence of sensations in the absence of physical objects.

Although epistemological problems of perception have occupied center stage in the history of skeptical challenges, epistemological problems concerning our access to the past through memory almost certainly have a more fundamental logical place in the ordering of skeptical issues to be resolved. The problem has received far less attention from philosophers partly because it is much less obvious how to characterize the nature of the available justification. If we suppose for a moment that there is such a thing as a memory "experience," an experience that can be veridical or nonveridical with respect to the past, the skeptical argument involving the past will closely parallel the argument for strong skepticism with respect to the physical world. *Ultimately*, in reaching conclusions about the past one must rely on what one seems to remember. But one scarcely needs an argument (particularly when one gets to be my age) that memory is fallible. One can have a *vivid* apparent recollection that one did something even if one did not. If an argument is needed, we can return to the apparent causal dependency of experience on brain events. If one can in principle produce nonveridical sensations by stimulating the brain, one can presumably find that region of the brain responsible for memory "experiences" and produce them at will. But if the occurrence of memory experiences is logically compatible with the events we seem to remember not having occurred, then what reason do we have for thinking that such memory experiences are reliable indicators of past events? It is tempting to rely again on induc-

tive reasoning, but it would seem that such reasoning is once more unavailable. An inductive justification for the reliability of memory would proceed from a premise describing correlations between *past* memory experiences and the events we took them to correctly represent. But the skeptic wants to know what your reason is for supposing that in the past memory has typically been reliable. And, of course, the skeptic wants you to answer that question without begging the question, that is, without relying on memory. But if one cannot rely on the fact that one remembers having veridical memory experiences to justify one's belief that a memory experience is a reliable indicator of a past event, how could one ever get away from the present to gain access to the past?

Notice how much more fundamental the problem of memory is than other epistemological problems. Almost all candidates for resolving the problem of perception minimally presuppose that we have access at least to past sensations. Discussion of the notorious problem of induction again almost always presupposes that we have access to past correlations between properties in order to ask how one can justifiably project these past correlations into the future. In the context of worrying about how we can justify our belief in other minds relying only on observational knowledge of physical behavior, epistemologists typically "give" one knowledge of past correlations between one's own behavior and one's own mental states, that is to say, they presuppose that there is some solution to both the epistemological problems of perception and of memory. And philosophers of science who worry about the possibility of justifying belief in hypotheses that deal with microphenomena that are in principle unobservable (compared with the way in which ordinary macro-objects are presumed to be observable) typically assume the legitimacy of our conclusions concerning the macro-sized objects of the physical world, the past, and projectibility of observed correlations.

I suggested that the relative lack of concern with epistemological problems concerning memory might be due in part to the difficulty one has formulating the problem. As long as one supposes that there are such things as memory experiences, the skeptical argument goes relatively straightforwardly. But the existence of memory experience is far from uncontroversial. To be sure, philosophers do not agree much on how to understand visual, kinesthetic, tactile, gustatory, and olfactory sensations, or even the intelligibility of talk about the occurrence of sensations without physical objects. But that something new and in some sense *occurrent* comes into existence when I open my eyes, and ceases to exist when I close them, seems hardly problematic. It is far

less obvious that when I remember putting my car keys on the desk something came into existence at a certain time, which can be meaningfully described as an occurrent mental state of seeming to remember having done something. As we shall see, the problematic nature of the presupposition that there are memory experiences can contribute to the attraction that externalist analyses of knowledge and justification have for so many philosophers.

Metaepistemological Presuppositions of the Skeptical Arguments

It seems to me that reflection on the above examples of skeptical reasoning strongly suggests that the skeptic relies implicitly on a principle that I call the principle of inferential justification:

To be justified in believing one proposition *P* on the basis of another proposition *E*, one must be (1) justified in believing *E* and (2) justified in believing that *E* makes probable *P*.

Arguments for strong local skepticism typically invoke clause 2 of the principle first and then often counter proposed attempts to satisfy clause 2 by relying on clause 1. Thus the argument for strong skepticism with respect to the physical world relies on clause 2 by insisting that a belief in physical objects *inferred* from what we know about the character of our sensations is rational only if we have some reason to suppose that there is a connection between the occurrence of certain sensations and the existence of certain objects. When one attempts to inductively infer the existence of such a connection from a premise describing past correlations between sensations and objects, the skeptic invokes clause 1 of the principle to challenge our justification for believing the premise of that inductive argument. The argument for strong skepticism with respect to the past also relies on clause 2 insisting that any justified conclusion about the past inferred from what we seem to remember must include justification for believing that memory experiences are a reliable indicator of past events. Again, when an inductive justification of such reliability is attempted, clause 1 of the principle is invoked to challenge our ability to rationally believe the premise of the inductive argument, a premise that will describe past events and again require an inference based on memory.

It is tempting to think that the question of whether or not one accepts

the principle of inferential justification determines whether or not one is an internalist or an externalist in epistemology. It is certainly true, as we shall see, that paradigm externalists reject at least clause 2 of the principle. It is *not* true, however, that paradigm externalists reject clause 1. Nevertheless, the above skeptical arguments do not get off the ground unless it is presumed that a reasonable conclusion based on premises requires a reasonable belief in the proposition that the premises make probable the conclusion. And on one reading, externalists avoid traditional skeptical problems by refusing to accept clause 2 of the principle of inferential justification.

The issue is, however, complicated. First, it is not clear that the principle of inferential justification constitutes a metaepistemological principle. On the face of it, one could accept the principle as a very general normative principle of epistemology. Second, it is not clear that one should simply define the internalism/externalism controversy in terms of whether one does or does not accept both clauses of the principle of inferential justification. In the next chapter I examine more carefully the controversy between internalists and externalists in epistemology. I try to define precisely the controversy, or more accurately a number of controversies, and we return to the question of which version of internalism, if any, classic skepticism presupposes. In chapters 3 through 7 we attempt to arrive at some conclusions about the philosophical plausibility of various metaepistemological views and the position they leave us in with respect to answering the skeptic. But before we leave this preliminary discussion of skeptical arguments, I want to emphasize and clarify certain features of the arguments. I also want to reply to a number of objections that are leveled against skepticism, objections that do not explicitly focus on the principle of inferential justification and any metaepistemological implications that acceptance of that principle might be thought to have. My goal is to leave the skeptic in as strong a position as possible when we consider the question of whether the externalist revolution in epistemology is the only effective way to circumvent the skeptical challenge.

Clarification of the Skeptical Arguments and Charges of Self-Refutation

Inference and Inferential Justification

If the skeptic does rely on the principle of *inferential* justification in presenting skeptical arguments with respect to our epistemic access to

the physical world, the past, other minds, the future, and so on, he is obviously committed to the conclusion that the only justification available in all of these cases is inferential in character. This is a feature of the classical skeptical argument that has understandably come under considerable attack. In what sense do I actually *infer* the existence of the familiar objects around me now? Is it not even more strained to talk as if my belief about what I had for breakfast this morning involved some inference from propositions describing the phenomenological character of my present memory states? Many philosophers would even suggest that our beliefs about the conscious states of those around us are misleadingly described as involving inference.⁴ When someone is writhing on the ground before me, I “see” the suffering directly. I hardly notice first a pattern of behavior, think about correlations between behavior and pain, and then reach the conclusion that the person before me is in pain. Of course, if we are relying on phenomenological evidence in order to determine whether it is plausible to maintain that an inference has taken place, it is equally problematic to suppose even that beliefs about the future typically involve inference. When I drink the water I expect it to quench my thirst. This expectation is probably (although not necessarily) *caused* by past associations of drinking water and subsequent diminution of thirst, but I can confidently assert that I never recall having attempted to list the occasions on which drinking water quenched my thirst in order to generate the premises of an inductive argument. My memory is so bad, I would be hard pressed to come up with more than a dozen such occasions, hardly enough to get me the kind of impressive correlation one should have for a strong inductive argument.

Whether or not the justification available for a given belief involves inference will be a difficult question to answer and will inevitably involve complex philosophical controversies. In characterizing a kind of justification as inferential, we seem to be implicitly contrasting that justification with some other kind of justification—noninferential justification. Obviously one needs an analysis of the distinction, and that we attempt to give in chapter 3. For now, let us make some relatively innocuous comments about ways of thinking about the distinction that bear on the reader’s rather natural concern that the skeptic is talking about inference where there appears to be no inference.

The first distinction one might make when facing this sort of objection to talk about inference is the familiar distinction in philosophy of science between the context of discovery and the context of justification. It is not obvious that in order for one’s justification to be inferen-

tial one must actually go through some process of conscious inference. It is one thing to arrive at a conclusion. It is another to justify that conclusion. It is almost certainly the case, it seems to me, that we do *not* usually formulate, and perhaps *never* have formulated, premises describing the phenomenological character of sensation as part of an attempt to gather evidence in support of our beliefs about the world around us. Indeed, it is probably seriously misleading to talk about my *belief* that there is a table before me and a wall behind me now. Philosophers have a tendency to oversimplify the range of intentional states that the conscious human mind can exemplify, and our language may even be inadequate to capture the subtle differences between them. There does seem to me to be a difference between belief and what one might call expectation. I do not so much believe that the wall is behind me as I expect it to be there. The expectation may consist in nothing more than the disposition to believe the proposition were I to entertain it, and the disposition to be extremely surprised if I were to turn around and not have the familiar range of experiences associated with there being a wall there. In any event, these beliefs (occurrent or dispositional) and expectations are still the kinds of things that can be justified or unjustified and we can still raise questions about the justification *available* to support such beliefs.

But surely, one might object, this is a little too slick. The skeptics and their opponents are raising questions about whether our *actual* beliefs are *actually* rational or not. And if we accept the framework within which the skeptics are asking this question, it would seem that their opponents run the risk of getting caught in the presupposition that if we are actually justified in thinking that there is table before us, we must have actually gone through some process of inference. The *availability* of a legitimate inference that we might have engaged in seems neither here nor there when it comes to the question of whether a belief is justified. This is perhaps a point well taken, and I argue later that a complete answer to it requires that we distinguish a number of different senses in which we can talk about a belief’s being inferentially justified. In particular, it will be useful to distinguish an *ideal* sort of justification from derivative concepts of justification. For now, let us be content with these observations. First, nearly everyone who talks about inferential justification wants to allow that the beliefs that are involved in inferentially justifying another belief might be merely dispositional. Second, engaging in conscious consideration of some set of premises on the way to reaching a conclusion that is consciously thought of as following from those premises is not a necessary condition for the justification supporting that belief to be inferential.

The Role of Skeptical Scenarios

Although it is perhaps obvious already from the earlier discussion of the form of classical skeptical arguments, I want to emphasize the limited role that the familiar skeptical scenarios play in generating the skeptical conclusion.⁵ As I presented those skeptical arguments, appeal to the possibility of dreams, hallucinations, artificially stimulated brains, and the like is primarily designed to support the conclusion that one has no *direct*, unproblematic access to the truth of the propositions under skeptical attack, and one cannot *deduce* the propositions under skeptical attack from the available evidence. The skeptic I am interested in is *not* presenting the following superficially similar argument for skepticism concerning the external world: The hypothesis that there is no table before you now but you are dreaming (hallucinating, having your brain artificially stimulated by a neurophysiologist, being deceived by an evil demon) contradicts your commonsense belief, and you do not know or have reason to believe that this alternative hypothesis is false. Therefore, you do not know or have reason to believe that you are really seeing a table. This argument simply invites the kind of response that Peter Klein developed so plausibly in his excellent book *Certainty*. That disarmingly straightforward response is to announce that we do have epistemic reason to reject the skeptical scenarios, a reason that consists in our being justified in accepting commonsense conclusions about the physical world. One might just as well argue, after all, that since the skeptical hypotheses are incompatible with the dictates of common sense and the dictates of common sense are epistemically rational, then the skeptical hypotheses are irrational. We can deduce their falsehood from the commonsense premises rationally believed. There is nothing wrong with Klein's *strategy* for defeating this form of skeptical argument, provided that he can establish the crucial conclusion that our commonplace beliefs are epistemically rational.

Again, however, as I construe the skeptic's appeal to the intelligibility of skeptical hypotheses, they are designed only to show something about the nature of our justification for believing propositions about the physical world, the past, other minds, and so on, namely that such justification involves nondeductive inference. If this conclusion can be reached, then one forgets about skeptical scenarios and invokes straightforwardly the principle of inferential justification. Once we have agreed, for example, that the occurrence of our sensations is perfectly compatible with there being no physical world, the skeptic can invoke the principle of inferential justification in order to request some positive reason

to suppose that there is at least a *probabilistic* connection between the occurrence of certain sensations and the existence of certain objects.

Notice that here there is no burden of proof. Many discussions of skeptical challenges begin with a frantic jockeying for position. The antiskeptical wants the skeptic to give some positive reason for supposing that there is no physical world, and in the absence of such argument proposes that we continue with the beliefs that we are in any event disposed to have. The skeptic, however, armed with the principle of inferential justification, can adopt what seems to me the correct philosophical attitude that the principle of inferential justification plays no favorites, recognizes no special burdens of proof. The astrologer and the astronomer, the gypsy fortune-teller and the economic forecaster, the druid examining entrails and the physicist looking at tracks in cloud chambers are all expected to have reason to believe that their respective evidence makes probable their conclusions if the conclusions are to be rational. And you are expected to have reason to think that your sensations make probable the existence of the objects you take them to indicate if you are to be justified in believing the dictates of common sense.

Before leaving this preliminary discussion of the role played by appeals to skeptical scenarios, I should add that there is at least one attempt to satisfy the second condition of the principle of inferential justification that can give rise to a second purpose served by appeals to the intelligibility of such hypotheses. Whereas earlier in this century philosophers tended to emphasize the prominence of enumerative induction as the most obvious candidate for legitimate nondeductive reasoning, contemporary philosophers who realize the limitations of inductive reasoning as a means of regaining commonsense beliefs about the world have often turned to so-called reasoning to the best explanation.⁶ The physical world, the past, other minds, lawful regularities, and theoretical entities are posited, the argument goes, as the best explanation for the order in which sensations come and go, the existence of memory experiences, the behavior of other bodies, observed regularity, and phenomena in the macroworld. In response to this gambit the skeptic will often request criteria for "best explanation" and turn again to skeptical scenarios to argue that there are always alternative explanations that compete with our "commonsense" hypotheses and that satisfy equally well the criteria of good explanation.⁷ In chapter 7 we will examine in detail the appeal to best explanation as a solution to the skeptical challenge and discuss the way in which the skeptic might invoke skeptical scenarios to counter the use of reasoning to the best explanation.

Epistemological Commonsensism

I have suggested that the skeptic who invokes the principle of inferential justification wants a level playing field. In particular, the skeptic will not give the mere fact that you are inclined to believe a hypothesis any particular weight. There are, however, many philosophers who would argue that one must simply rule out skeptical conclusions from the start. The most common form of argument in analytic philosophy is the *reductio*. One objects to a philosophical position by pointing out that it has absurd consequences and is therefore absurd. But strong skepticism with respect to commonsense beliefs is itself patently absurd, the argument continues, and is a sufficient reason to reject any view, including a metaepistemological view, that leads to it. We might call the view that rules out skepticism from the start and evaluates metaepistemological views in part by the way in which they allow one to avoid skepticism, epistemological commonsensism. I return later to the suggestion that skeptical conclusions are absurd, but for now I want to make at least a preliminary comment on one sort of argument for it. The most obvious question the skeptic will ask is *why* we should assume at the outset that the beliefs we take to be justified are justified. The answer that we must start somewhere will no doubt not please a skeptic who is disinclined to start a careful reexamination of all of our beliefs with the presupposition that most of those we take to be justified are justified.

A somewhat more sophisticated answer involves appeal to science and evolution. We can assume, the argument goes, that most of our beliefs are justified because it is obviously evolutionarily advantageous to have justified beliefs and science tells us that what is evolutionarily advantageous has a high probability of occurring, other things being equal. Such an argument will not impress traditional skeptics, of course, because they will correctly point out that the pronouncements of science can be used to refute skepticism only if they are themselves justified. Their justification, however, presupposes solutions to the various problems the skeptic presents.

Whether or not science can refute traditional skepticism will, as we shall see, depend itself on the plausibility of certain metaepistemological views.⁸ But for now I also want to remind the reader that even if we give ourselves full access to the pronouncements of science, it is not clear that science does tell us that it is evolutionarily advantageous to have justified or rational beliefs. Many of the empirical conclusions of science seem to suggest that much of what we expect or take for granted

is “programmed” into us through evolution. If children had to reason deductively and nondeductively to the various conclusions they take for granted, their chances of survival would no doubt be rather slim. There is no reason to believe that we are not simply programmed to respond to certain stimuli with certain intentional states, just as lower life forms appear to be programmed to respond to certain stimuli with appropriate behavior. Now *given* certain metaepistemological views that we will discuss later, the causal origin of these spontaneous unreflective beliefs might be sufficient to make them justified, but one might also conclude that nature has simply no need to satisfy the *philosopher’s* desire for having fully justified belief. One might, in other words, argue that if what science tells us is true, one might well expect that nature has probably not constructed us to believe only that which we have good reason to believe. Commonsense science might well tell us that through evolution nature has decided that it would be better for us to have true beliefs than justified beliefs.⁹

Of course, if we could *know* this we would again have reason to believe that most of our beliefs are true. My only concern here is to point out that our scientific *beliefs* are perfectly compatible with the conclusion that those beliefs are quite unjustified, and may well even suggest that conclusion, given certain metaepistemological views about what is required for justification. It is probably just this thought that led Hume to observe, with respect to the question of whether or not man should believe in a physical world, “Nature has not left this to his choice, and doubtless esteem’d it an affair of too great importance to be trusted to our uncertain reasonings and speculations.”¹⁰

Charges of Self-Refutation

Epistemological commonsensism maintains that skepticism runs afoul of methodological constraints on epistemological investigation. The charge of self-refutation is a much more specific criticism of skepticism and must be treated very carefully. Let us begin by distinguishing two ways in which an argument might be charged with self-refutation. The first charge of self-refutation consists in the claim that the conclusion of a skeptical argument is inconsistent with the premises that are used to reach that conclusion, or in the claim that the very intelligibility of the skeptical conclusion requires that it be rejected. Let us say that an argument with this feature is formally self-refuting. If the conclusion of an argument really is inconsistent with its premises, then of course the argument either is formally invalid or has necessarily false

premises. In either case an argument that is formally self-refuting in this way is always unsound. A more interesting and probably more common criticism of skeptical arguments is that they are *epistemically* self-refuting. Let us say that a skeptical argument is epistemically self-refuting if the truth of its conclusion implies that one has no justification for accepting its premises. But if one has no justification for accepting its premises, then the very principle of inferential justification upon which the skeptic relies implies that one cannot be justified in believing the conclusion by inferring it from the premises.

Of the two ways in which one might assert that skepticism is self-refuting, the charge of formal self-refutation is going to be the hardest to make stick. The most common allegations of formal self-refutation today involve claims about language and intentionality. Largely under the influence of the later Wittgenstein, a number of philosophers have explicitly or implicitly adopted what one might call a *contrast* theory of meaning. Roughly, the idea seems to be that a predicate expression "X" only has meaning if there are things that are both correctly and incorrectly described as being X. Thus, on my reading of Wittgenstein's private language argument, the fundamental objection to a private language has nothing much to do with memory. The problem is that a private linguist is the sole arbiter of how similar something must be to a paradigm member of a class to count as similar enough to be described in the same way. But as the sole judge it will not be possible to make a mistake, and where there is no possibility of error there is no possibility of getting it right. It is only meaningful to talk about the correct application of a rule if it can be contrasted with an incorrect application of the rule. If one applies this principle to dreams, hallucinations, and more generally, nonveridical experience, it will make sense to speak of nonveridical experience only if we are contrasting such experience with something else.

The contrast theory of meaning as stated is far too crude to evaluate. Suffice it to say, for our purposes, any remotely plausible version of the view would have to stress the modal operators. It will make sense to talk about dreams only if it is *possible* to have an experience that is not part of a dream. It will make sense to say of something that it is not a unicorn only if it is *possible* for something to be a unicorn. But to entertain the possibility that one is always dreaming, always hallucinating, or always being deceived by the senses is not to entertain the hypothesis that veridical experience is in any way *impossible*.

In response, the antiskeptical pushing this argument might appeal to a still more controversial verificationist theory of meaning. Roughly, the

idea is that for X to be meaningful, not only must it be possible for both something to be X and something not to be X, but we must have *criteria* for distinguishing the Xs from the not-Xs. But for the more vague expression "criteria," this principle is, of course, familiar as at least a relative of the logical positivists' old verifiability criterion of empirical meaningfulness. On the face of it, however, the verifiability criterion of meaning, even in its very weak forms, has little to recommend it. As others have pointed out, we can entertain perfectly meaningful hypotheses that we could not in principle verify or disconfirm. Consider, for example, the proposition that there are things of which no one has ever or will have ever thought. What are our criteria for picking out the things about which we have never thought? Far from being meaningless, the hypothesis is entirely plausible. But if we can entertain and even believe such a hypothesis, why should we be unable to suppose that the causes of our sensations are radically different from any of which we have thought, and different in a way that would make *all* of our ordinary beliefs about the causes of sensations false?

Without focusing as much on language alone as a means of representation, Putnam also seriously questions the intelligibility of the skeptical conclusion given its premises. In the now-famous discussion of "Brains in a Vat,"¹¹ Putnam argues that if we were brains in a vat we couldn't assert that we are brains in a vat. It will not be possible to do justice to the argument here because it rests on a highly sophisticated, controversial, and incomplete theory of what is involved in one thing representing another. Given the attention the argument has received, however, I should at least indicate how I would respond.

Crudely, Putnam's idea is that for a mental state, a thought, to be a thought of X, X must causally interact with the state that represents it. Like most theories of representation, the principle would be qualified to allow for complex thoughts of nonexistent things that are "constructed" by the mind out of simpler thoughts. Hume, for example, thought that all simple ideas had their source in experience but allowed that one could form the idea of a unicorn without seeing one by "putting together" the idea of a horse and a horn. So, too, Putnam could allow that one might form the idea of things that do not exist by putting together the ideas of things with which one has had causal interaction. Within the framework of such a view, for my talk or thought of physical objects to mean what it does, my language and mental states must have interacted in the appropriate way with the physical world. Brains in a vat, by hypothesis, do not interact with the physical world in the way that would be relevant to allowing such brains to represent physical

objects. But if brains in a vat cannot represent physical objects, then they cannot coherently frame the hypothesis (as we understand it) that they might be *brains* in a *vat*. By the same reasoning, they cannot coherently frame the hypothesis (as we understand it) that there might be no *physical world*. For these hypotheses to be genuinely meaningful, they would have to be false.

I have argued elsewhere that the theories of representation on which the argument rests are false, and for that reason, and because an excursion into the metaphysics of intentionality would take us too far afield, perhaps we should content ourselves with the observation that even if something like Putnam's conception of representation were correct, he would not get much ammunition for use against the skeptic. By now I think that almost everyone agrees that Putnam's original argument involved rhetorical "stretch." For obvious reasons; even if we assume as true everything Putnam claimed about representation, it would still be possible for *me* to be a brain in a vat now, for *me* to have always been a brain in a vat, for *all* humans to be brains in a vat now, and for *all* humans to always have been brains in a vat. The reason all of these hypotheses are intelligible on Putnam's view is that our ability to represent can "piggyback" on prior representation. The relevant causal chains that allow us to represent an *X* can be extraordinarily convoluted. And the skeptic can surely get what is needed (particularly given the limited role skeptical scenarios play on my construal of the skeptical argument) from the intelligibility of the above skeptical scenarios.

The other observation I would make is that the ability to construct complex ideas of things that do not exist out of simpler ideas representing existents leaves enormous room for the skeptic to argue that hypotheses about the physical world might all be intelligible while at the same time being false. Put briefly, why should we not suppose that the concept of a physical object just is one of those innumerable many concepts of something that does not exist? On one historically prominent conception of the physical world, physical objects are *thought* of as the causes of certain sensations that stand in certain isomorphic relations with the sense data they cause. If one can construct the thought of phlogiston only to find out that there is nothing corresponding to it, why can one not construct the thought of a physical world only to realize that one has no reason to believe that there is such a thing? The ideas of sensation and causation might be traced to actual phenomena. But a complex idea formed out of these might not. I am not arguing for such an analysis of the concept of a physical object. I am only pointing out that in the absence of considerable argument to the contrary, there is

every reason to suppose that if one can have ideas of nonexistent things, one can have the idea of a nonexistent physical world.

A more whimsical argument still for the incoherence of skeptical challenges to common beliefs is suggested by Davidson's much discussed remarks concerning interpretation. Davidson (1981) argues, perhaps not too implausibly, that in interpreting another's language I must (methodologically) assume that most of what that person asserts and believes is true. He then argues that even an *omniscient* being would need to employ such a principle of charity when interpreting ordinary claims about the world. But if an *omniscient* being would need to believe that most of what we believe is true, then most of what we believe is true! The argument seems too good to be true and, of course, it is. Foley and I have pointed out that Davidson's argument actually needs a premise asserting the existence of an omniscient being to secure its conclusion.¹² To fully understand how so many people could have taken seriously the argument, we would need to examine more fully the presuppositions of the counterfactual conditional about the omniscient being. As this would take us too far afield, I simply refer the reader to our article on Davidson's "theism."

It was obviously going to be an uphill battle to make good the claim that there is something formally self-refuting in the skeptic's position. It would be more than a little odd to be able to reach the conclusion that a position with which people have been fascinated for thousands of years was in some deep sense literally unintelligible. The charge of epistemic self-refutation, however, is likely to be more significant. A skeptic who is in the business of undermining presuppositions about the rationality of our beliefs must be careful that the foundations on which the skepticism gets built are not undermined in the process. In what ways might an argument be epistemically self-refuting?

Consider again classical skeptical arguments for strong local skepticism concerning propositions describing the physical world. You will recall that a crucial step in the skeptical argument is to establish the possibility of sensations occurring without the physical objects we take those sensations to represent. In arguing that this is a genuine possibility, the skeptic appeals to the possibility of dreams, hallucinations, or surreptitious manipulation of that part of the brain directly responsible for producing sensations. But if the skeptic appeals to facts about the causal dependency of sensations on brain events, or facts about the hallucinogenic character of certain drugs, or previous occasions on which dream experiences seemed indistinguishable from the experiences of waking life, then surely the skeptic must be justified in believing these

truths about the physical world in order to generate the conclusion that one cannot be justified in believing any proposition about the physical world. The skeptical conclusion, however, entails that the skeptic has no justification for believing these commonplace truisms about the causal conditions surrounding sensation.

Is there any way for the skeptic to avoid the charge of epistemic self-refutation? I think there is, and that it involves careful use of modal operators. To be sure, many of the classic skeptics (or philosophers seriously advancing for consideration skeptical arguments) appeared to appeal to facts about the physical world in advancing their argument. Thus, in his refutation of epistemological direct realism, Hume certainly did start talking about what happens to your visual sensation when you press against the side of your eye.¹³ And Descartes does sometimes describe actual dreams he remembers having.¹⁴ But given the purpose that appeal to these skeptical hypotheses is designed to serve, these philosophers need not have made any claims about what actually happens. One need only appeal to the *intelligibility* of vivid dreams, hallucinations, or the causal dependency of sensations on a physical world to establish that there is no logical connection between the occurrence of sensations and the existence of physical objects. Claims about possibilities are not in any obvious sense contingent, and would not fall within the scope of the propositions under attack by the skeptic advancing local skepticism with respect to the physical world. It must be conceded that the most convincing proof of a possibility is an actuality. Descartes's dream argument would hardly have fascinated us so much were we never to have vivid dreams. Indeed, Descartes probably would not have thought of the argument in the first place (nor would he even have been understood) were people not convinced that dreams do occur. And if people did not believe the scientific data on the brain, one would not get very far appealing to the possibility of producing sensation by manipulating the brain. But a clever enough philosopher could have formulated these hypotheses as possibilities even without the empirical evidence that the possibilities are actual and could have reached the appropriate conclusions about whether or not our access to the physical world was direct. Since the appeal to possibilities presupposes no knowledge or rational belief about the physical world, a skeptical conclusion based in part on such appeals involves no epistemic self-refutation.

Earlier, I argued that the problem of skepticism with respect to the physical world is not the most fundamental of the local skeptical challenges. Without rational beliefs about the past, one could not even learn

anything about connections between sensations that one might hope to use in justifying belief about the physical world. But strong skeptical conclusions about the rationality of belief in past events is much more likely to end up involving epistemic self-refutation. The problem is more severe, of course, the stricter one's conception of what the past is. In the most extreme view, "now" refers to an instant in time, and the epistemological problem of access to the past is the problem of how one can reasonably infer the occurrence of *any* past events from data that is available to consciousness *now*, this instant. Reasoning itself, including skeptical reasoning, takes place over time. The skeptic moves step by step to a skeptical conclusion. If the skeptic implicitly accepts the principle of inferential justification, then to be justified in reaching a skeptical conclusion each step in the reasoning process must itself be justified. But if one has no justification for believing anything about the past, if one is an epistemic prisoner of present consciousness, how could the skeptic be justified in believing the premises of the skeptical argument?

One attempt to get around the problem is to expand "now" to include what is sometimes called a specious present. Consciousness is capable of grasping directly and immediately an expanse of time and a sequence of events that are included in that limited time. If the specious present is "large" enough and the skeptic's reasoning is quick enough, one might try to avoid the charge of epistemic self-refutation this way. Needless to say, it would not be easy to decide on the duration of a specious present. Unless it grows to an implausible size, the skeptic must think a lot faster than I can to escape epistemic self-refutation in this way.

Curley (1978) has argued that a skeptic can avoid charges of epistemic self-refutation by holding the premises of the skeptical argument to much lower epistemic standards than those challenged by the conclusion. This escape certainly would be in principle available to the philosopher arguing only for weak skepticism with respect to the past. One could easily reconcile the conclusion that we cannot know with absolute certainty anything about the past with the claims that one has some reason to believe some propositions about the past, and that those somewhat rational beliefs are sufficient to entitle one to accept rationally the skeptical conclusion. But this reply is clearly unavailable to the philosopher arguing for strong skepticism.

Ultimately, it may be that the charge of epistemic self-refutation will stick in the case of the most extreme skepticism about the past. The reason again is straightforward. As we already have had occasion to

note, all strong global skepticism is epistemically self-refuting. If one concludes that one has no epistemic reason for believing anything at all, then it follows that one has no epistemic reason for believing that one has no epistemic reason for believing anything at all. Further, one has no epistemic reason for believing anything upon which one bases one's conclusion that one has no epistemic reason for believing anything at all. That is precisely why so few skeptics have been strong global skeptics. In particular, skeptics have almost always presupposed a kind of unproblematic access to some foundational empirical data and to the legitimacy of the reasoning on which their skeptical conclusions depend. But an extreme, strong skepticism with respect to the past challenges the unproblematic access to reasoning and seems to run the danger of spilling over into areas not under skeptical attack.

If the skeptic has difficulty denying the charge of epistemic self-refutation to this very fundamental sort of strong local skepticism, perhaps the next best move is to absorb it in a way that leaves the force of the skeptical argument intact. Is one in a position, after all, to dismiss an argument on the grounds that it is epistemically self-refuting? Well, if strong skepticism with respect to the past is epistemically self-refuting, then by definition, the skeptic is not epistemically rational in believing at least some of the premises of the argument. But if the argument is a valid argument and if the anti-skeptic *believes* the premises of the argument, the anti-skeptic is hardly out of the woods.

When I was a child I owned a "magic" eight ball whose function was to predict the future. You asked the eight ball a question that could be answered "Yes" or "No," shook the ball, and a "Yes" or "No" floated to a transparent opening in the ball. Now suppose we lived in a culture in which people took the eight ball to be a reliable guide to the future. If you asked the eight ball "Will it rain tomorrow?" and the eight ball answered "Yes," then, according to the members of our culture, you would be epistemically justified in believing that it will rain. Let us also suppose that our culture contains a few annoying skeptics who do not see what possible grounds one could have for supposing that the eight ball's answers are reliable predictors of the future, and who advocate strong skepticism toward conclusions reached via eight-ball reasoning. Finally, let us suppose that one day a skeptic gets the bright idea of asking the eight ball whether or not conclusions reached via eight-ball reasoning are rational, and the eight ball answers "No." To the chagrin of the supporters of commonsense eight-ball reasoning, the results of the experiment are duplicated again and again. In what position are eight-ball reasoners left?

I suppose die-hard proponents of eight-ball reasoning can argue that our skeptic who concludes that it is irrational to believe the dictates of the eight ball, is certainly in no position to use the eight ball's "answers" to reach that conclusion. This seems right, of course. The skeptical argument that proceeds from observations about what the eight ball indicates and reaches a conclusion about the illegitimacy of eight-ball reasoning is epistemically self-refuting. But should the eight-ball reasoners be celebrating? Can they go on as before, trusting the predictions of the eight ball? If it is obvious that the skeptic's eight-ball argument is epistemically self-refuting, it seems equally obvious that the anti-skeptic cannot continue to embrace the unproblematic legitimacy of eight-ball reasoning. The skeptical argument has revealed an internal problem for the anti-skeptic despite the fact that the argument is epistemically self-refuting. In the same way, I would suggest, a skeptical argument with the conclusion that beliefs about the past are epistemically irrational is not something one can dismiss just because one concludes that even the skeptic would have to rely on some beliefs about the past in order to reach the skeptical conclusion. As long as the anti-skeptic shares belief in the premises that the skeptic acknowledges one has no epistemic reason to accept, and as long as the anti-skeptic has no reason to reject the legitimacy of the skeptic's reasoning, the anti-skeptic cannot simply dismiss the import of the argument on the grounds that it is epistemically self-refuting.

It is perhaps worth emphasizing again that as I presented them, most skeptical arguments for strong local skepticism are not epistemically self-refuting. The premises that the skeptic would need to believe rationally in order to infer the skeptical conclusion fall outside the class of propositions under skeptical attack. A fundamental and extreme sort of local skepticism concerning belief about the past may encounter difficulties with epistemic self-refutation, but the argument will still be a thorn in the side of the anti-skeptic until the anti-skeptic can figure out which of the skeptic's premises should be abandoned.

In the preceding discussion we distinguished charges of formal and epistemic self-refutation that might be leveled against the skeptic's arguments. In addition to these relatively precise criticisms of skepticism there is a somewhat more nebulous charge that is related to criticisms associated both with epistemic commonsensism and charges of self-refutation. That criticism attempts to denigrate the importance of the skeptical challenge by observing that it is in some way impossible to take skepticism seriously. In fact, most philosophers have never treated skepticism as a viable option. Even those who are interested in the

skeptical challenge are primarily interested in finding the correct way of refuting what is assumed to be an illegitimate conclusion. Everything the skeptic *does* belies the seriousness of the skeptical position. The famous philosophical skeptics, after all, have gone to great lengths to publish their skeptical treatises and in doing so have made clear that they assume that there are other minds (and so implicitly other bodies). Skeptics who managed to survive made the same inductive inferences rejected as illegitimate in their skeptical "mode." They also argued with their fellow philosophers, past and present, and in doing so placed complete faith in the existence of a past revealed to them through memory. The very activity of philosophy, Butchvarov argues, is one that presupposes the existence of an external world, a past, and other minds, and the philosophers who realize this can take philosophy seriously only if they reject strong skepticism.¹⁵ Reaching a skeptical conclusion is incompatible with taking oneself and one's work seriously. Philosophical skepticism is not then a serious philosophical position.

The charge invoked here is *not* the charge of formal epistemic self-refutation. The claim is not that the skeptic's conclusion entails that the skeptic has no reason to believe the conclusion. The alleged self-refutation is more subtle. It amounts to the claim that in embracing a skeptical conclusion as a serious philosophical position the skeptic is implicitly engaging in behavior that makes sense only against the backdrop of a set of beliefs that are incompatible with radical skepticism.

The first step in a skeptical response to this sort of criticism involves the obvious distinction between believing something and rationally believing something. Of course, there are no skeptics who withhold belief with respect to the questions of whether there is a past, a physical world (in some sense of "physical"), other minds, and regularities that can be safely projected into the future. But it is by no means obvious that to believe *P* one must believe that it is epistemically rational to believe *P*. I know all sorts of very religious people who seem to quite happily concede that their theism is epistemically irrational. Earlier, when discussing the question of whether the concept of justified belief is a normative concept, we very briefly raised the question of the extent to which belief is something we *control*. I sidestepped the issue of whether it was in principle possible to produce belief through an act of will, but it is difficult to deny that as a matter of empirical fact we simply find ourselves believing all sorts of things. It is not until our first philosophy class that many of us even raise the question of whether the things we take for granted are epistemically rational for us to accept. Indeed, it seems to me that philosophers too often forget that the questions raised

by the skeptic are *philosophical* questions. They are questions raised by people who have a certain kind of philosophical curiosity that arises naturally from a very unnatural kind of activity. I argued earlier that there may well be different concepts of epistemic rationality and that some of these concepts might have particular relevance to *philosophy*. It may be that the philosopher is interested in and *wants* a kind of justification that ordinary people do not even think about in their day-to-day lives. The philosophical skeptic may best be construed as telling the philosopher that this kind of justification is unavailable. In every other walk of life people must get used to the idea that they cannot have everything they want, and the skeptic might maintain that it is a kind of perverted optimism to suppose that the *kind* of justification that would satisfy the *kind* of curiosity that afflicts the epistemologist is there to be found. This is a theme to which I return in the final chapter.

Notes

1. Klein (1981) calls this *iterative* skepticism.
2. Of course, when the necessary truths and falsehoods and the logical entailments get complex, the story may be different. In particular, if time is required to "see" the truth or entailment, then memory may be involved in the discovery of such truth or entailment. If memory is required to discover complex necessary truth or entailment, then any skepticism with respect to memory will infect all knowledge that explicitly or implicitly relies on it.
3. Although not neurophysiologists, the Martians in Bradbury's *The Martian Chronicles* caused the initial visitors from Earth no end of trouble through their telepathic ability to induce hallucinatory experiences. At the risk of exposing the shallowness of my aesthetic sensibilities, I would also mention the film *Total Recall*, whose plot centers on the question of whether the protagonist is really on a trip to Mars or is enjoying the experiences of an artificially induced fantasy.
4. Sartre (1966, part 4) would certainly reject the claim that our beliefs about others always involve inference if this is intended to be a phenomenological description of how our minds work. On the other hand, with philosophers in this tradition it is notoriously difficult to get an unambiguous statement of their views about the epistemological significance of their phenomenological observations.
5. In contrast, for example, to the more pivotal role they play in, say, Stroud 1984.
6. Indeed, some—Gilbert Harman, for example—have argued that inductive reasoning is itself a species of reasoning to the best explanation. See Harman 1965, 1970.

7. For a detailed discussion of this move by the skeptic, see Alston 1993, chapt. 4.

8. I argue in chapter 6 that it depends on the plausibility of externalist metaepistemologies.

9. Plantinga (1993b, chap. 12) appears to argue that it is difficult to see how natural selection would favor even true beliefs. The right false belief coupled with the right odd desires will do wonders for increasing the probability of my survival.

10. Hume 1888, p. 187.

11. Putnam 1981, chap. 1.

12. Fumerton and Foley 1985.

13. See Hume 1888, p. 210.

14. In his famous opening remarks in Meditation I.

15. See Butchvarov 1992.

Chapter Three

Internalist and Externalist Foundationalism

The traditional debate over skepticism has largely presupposed the framework of foundationalism. With the rise of the internalism/externalism debate in epistemology, however, it is apparent that there are radically different ways to understand foundational justification. In this chapter we begin by examining the traditional epistemic regress argument for foundationalism. Before presenting what I take to be the most important traditional conception of foundational justification, I examine the internalism/externalism controversy, or more appropriately, *controversies*, so that we may better understand the sense in which traditional foundationalism is (and is not) committed to internalism. This paves the way for a detailed examination in the next chapter of paradigm externalist versions of foundationalism. Having distinguished importantly different senses in which views about justification may be internalist or externalist, I examine in some detail what I take to be the most plausible traditional account of foundationalism, an account that is standardly regarded as internalist. As we shall see, however, one must be very careful to distinguish the senses in which this view is and is not committed to internalism. I conclude by distinguishing a *conceptual* from an *epistemic* regress argument for foundationalism.

The Principle of Inferential Justification and the Epistemic Regress Argument for Foundationalism

We saw in the last chapter that the skeptic relies heavily on the principle of inferential justification to support strong local skepticism with respect to various sorts of propositions. Ironically, perhaps, that same