

A Defense of Strong Foundationalism

Simple questions often have the most awkward consequences. One of the simplest questions that may be legitimately asked of our everyday judgments can be put in four short words: "How do you know?" And the remarkable thing about this question is that it can be iterated: in any normal situation, an answer to the question will be some claim or set of claims that can reasonably be subjected to the same question all over again.

Any conversation in which one side sincerely asks this question and the other sincerely attempts to answer it is going to take one of a very limited number of forms. Either the conversation will go on forever, or it will stop. If it goes on forever, either the answerer comes back around in a circle, eventually repeating some of the claims that were initially called into question, or he goes on forever giving endless new reasons for believing the claims put forward earlier. If the conversation stops, either the answerer has just given up (an understandable reaction in everyday life!) or he has reached some claim so basic that he can fairly be said to be justified in believing it without having reasoned his way to it via other beliefs.

Foundationalism is the position that all justified empirical beliefs are either basic, in something like the final sense given above, or else have supporting lines of reasons that can be traced one way or another back to beliefs that are basic. It doesn't matter that no two sane people would have a conversation like this: the important point for foundationalists is that for justified beliefs the underlying reasons are there and *could* be produced, under ideal conditions, if necessary. ⁽¹⁾ If we drew a diagram of such a conversation, it would have the form of an upside-down tree, branching out when two or more reasons support a belief at a given point. But ultimately every branch could be traced down to a belief that is basic, a point where that branch of the evidence tree comes to an end.

Foundationalists do not all agree among themselves as to just what sort of justification the basic beliefs have to possess, or as to what sort of justifying relations enable us to reason our way from the foundations "up" to our everyday beliefs. Descartes, whose principal work *Meditations on First Philosophy* is a milestone of foundationalist thought, took a strong position on both points: the basic beliefs must be certainties, and the inference relations leading from those beliefs to higher-level beliefs must be deductively valid, allowing for no possibility of error at any step along the way. But the trend in recent decades has been to weaken both requirements, demanding only that the basic beliefs possess *some* degree of intrinsic plausibility rather than absolute certainty, and allowing the use of various inference relations that are not air-tight instead of insisting on deductive inference at every step. Robert Audi, a noted contemporary foundationalist, adopts this doubly modest position in reaction to the rather obvious deficiencies of the Cartesian position. Unlike the Cartesian, Audi can make room for justified belief regarding things of which we are not absolutely certain. And the modest requirements he places on basic beliefs allow him to start with a wider set of grounds than the classical strong foundationalist has at his disposal, thereby (apparently) enhancing his chances of providing a really good justification for beliefs about the external world. In both ways, moderate foundationalism seems to be a more desirable position than its austere cousin.

What seems to have been overlooked, or at least left unexplored, is a position that separates the two parts of the strong Cartesian position. Audi is certainly right that non-deductive inference relations are necessary for us to arrive at interesting conclusions and that such inferences can have justificatory force. But in contrast to Audi's doubly moderate position I think that Descartes is right about the need for certainties as the bases of empirical knowledge. Though much of our empirical knowledge is of course less than certain, there are compelling reasons to believe that our everyday empirical beliefs rest on a foundation of certainties.

Any attempt to revive strong foundationalism will have to answer three obvious questions: Are these strong foundations necessary? Are they available? And are they sufficient to ground empirical knowledge of the everyday sort?

I. Are Strong Foundations Necessary?

A Modest Proposal

If we can give a satisfactory answer to skeptical challenges without insisting on a foundation of empirical certainties, then it certainly makes sense to do so: we take on fewer responsibilities in this sort of argument if we start out making only more modest claims. But if fallible foundations turn out to be inadequate to stop the skeptical regress then we will have to turn to strong foundationalism for a way out.

Moderate foundationalism, in its non-skeptical form, maintains that there is a knower S of whom the following claims are all true:

1. S has some basic (i.e. non-inferentially justified) empirical beliefs
2. S has some justified nonbasic empirical beliefs
3. Every branch of an evidence tree supporting any of S's nonbasic empirical beliefs terminates in a basic empirical belief
4. Some of S's basic empirical beliefs are less than certain for S.

The crucial point where moderate foundationalism differs from strong foundationalism is, of course, point 4; strong foundationalists will deny this claim while affirming the other three. In order to understand the motivation for strong foundationalism, consider the conversation we started with, picking up at the point where the imaginary conversationalists have gotten down to a basic belief.

Knower: ... and I know that because I know that Z.

Skeptic: Okay, I can see that Z would be a good reason there -- provided that you really know it. So now, of course, I have to ask: How --

Knower: Stop! You've asked that question dozens of times so far, but this time it won't work. Z is a *basic* belief: I am justified in believing it, but I don't have or need any argument for it.

Skeptic: What kind of crazy position is that? Can you really mean that you just stop somewhere and dig in your heels and refuse to give any further reasons?

Knower: No, not at all. What I'm saying is that Z is a belief that I'm *justified* in holding, not just that it's something I'd like to believe or that my peers let me get away with believing.

Skeptic: I guess I see where you're coming from, but now I have to ask: are you *really* justified in holding Z? Anybody can say, "I'm justified in this belief," and sometimes it may actually be true. But is it true here?

The skeptic has a good point; this is a question all foundationalists have to answer. And this is just the point where things begin to get difficult for the moderate foundationalist.

Making It All Up

If I claim that I have met the Queen of England and you are in a skeptical mood, you may ask me how I know. The typical response for me would be to give you my reasons, but nothing absolutely prevents me from saying, "Oh, I don't have reasons -- I just believe that I've met her." But if I gave this response you would quite reasonably conclude that I am some sort of nut; the belief that I've met a person of such prominence is not the sort of thing one can be justified in believing without having at least some reasons, even if they amount only to faded memories. Anyone who is allowed to get away with this sort of thing can end up claiming to be "justified" in literally any belief. And none of us will take him seriously, for a very good reason: in the absence of any evidence that he really is justified, he may as well be making it all up.

There is a tricky point here about what counts as "evidence" that a person is justified in holding a certain belief. Foundationalists all agree that in order to be justified in believing that P, where P is some belief that S has *inferred* from other beliefs, S needs to be justified in believing the premises from which P has been inferred. To know that S is justified in believing that P, then, I would need to know that S justifiably believes the premises on which it is based, and also that the method of inference being used is a legitimate one -- that the premises really do support S's belief that P. The second part of this can be put more crisply. Epistemologists speak of the connection between premises and conclusions as the sort of thing that can be expressed by an *epistemic principle* with a conditional form; roughly, if anyone makes such-and-such an inference from justified beliefs, then the person is justified in believing this proposition.

For basic beliefs, however, there are no such premises. In this case, the evidence will have to take the form of an epistemic principle regarding the way in which the belief is formed, a principle that states (roughly) that whenever anyone forms a belief in manner X, it is a justified belief. In neither case are we solely concerned with S's evidence for the truth of P; these epistemic principles come in only when we are looking for evidence for the truth of the claim that *S is justified in believing that P*. The individual in question doesn't necessarily have to know what the principle is; no one has to become a philosopher in order to know such mundane facts

as "I have a headache." But the principle does have to be *true*. And someone who wants to know whether S's belief that he has a headache is really justified may reasonably hold out for an epistemic principle (whether known to S or not) that does *not* sound arbitrary and that indicates why that belief really is justified without any inferential support.

In the case at hand, this means that the skeptical question amounts to doubt that the supposedly basic beliefs really are underwritten by a true epistemic principle. It isn't enough to point out that they *might* be underwritten by a true principle; this would leave open the alternative that they are not, and the skeptic cannot lose this argument unless it can be shown that the alleged basic beliefs really do arise in such a way that one cannot fail to be justified in believing them.

Modesty On the Rocks

Moderate Foundationalism holds that there are basic beliefs that are less than certain. For a wide range of concepts of justification this amounts to saying that they are more or less probable, but that this probability falls short of a one hundred percent guarantee. But this raises an interesting question: is there such a thing as probability, in the sense that counts for justification, that is *not* based on an inference from other beliefs?

Suppose you meet a friend on the sidewalk and in the course of conversation he says, "It's probably going to rain today." What role does the word 'probably' have in this sentence? That will depend on your friend; it could indicate a guarded utterance, meaning roughly "I think it will rain, but don't call me a liar if it turns out not to," or it could be an expression of confidence, meaning "I have virtually no doubt that it will rain." These paraphrases, however, do not bring out the epistemic significance of your friend's remark; to say that it will probably rain is (normally) to indicate that there are good, though not conclusive, reasons for believing that it will rain.⁽²⁾ This is perfectly compatible with the implication that one does not wish to be held accountable for actions undertaken on one's testimony, or at the other extreme that one has great confidence; having good reasons is a paradigm case of well-placed confidence. But it goes beyond these mere psychological factors and focuses our attention on the feature that is critical for the discussion of modest foundationalism: probability arises from a relation between the probable proposition and a body of evidence.

This simple fact about probability creates a fatal dilemma for moderate foundationalism. If there are basic beliefs that are merely probable, then they are not basic at all; they are inferred, probable in relation to some other beliefs that support them. The focus of our inquiry shifts back to the supporting beliefs, and the dilemma gets started there once again: either they are basic or they are not. If they are not, we have to go back still further. If they are basic, then they cannot be merely probable.

Slipping Out the Back Door

A natural response to this argument is disbelief. I know many things by memory, for example, but I'm far from certain of the things I remember. Surely memory is a fallible source of justified beliefs? Or what about faint perception? At the very edge of my hearing range, I seem to hear a low throb and the grinding of gravel -- I think, but am not sure, that a car has turned into my

driveway. Obviously I am not certain of this. But if perception is a fallible source of justified basic beliefs, something must be wrong with this attack on moderate foundationalism.

Plausible as this response sounds, it is mistaken. In the memory case, there is something that I am absolutely certain about, namely, that I *seem* to have these memories. Doubt only arises at the point where I ask whether I am remembering *correctly*, but these doubts would not even make sense unless I could independently determine what I *seem* to remember. Something similar goes for the case of faint perception. What is absolutely certain in the perceptual case is that I really do have a certain sort of experience (which I am inclined to describe in the indicated fashion). This is the foundation: even if I'm dreaming, it remains unshaken. What is uncertain is whether my instinctive conjecture about the cause of this experience is correct. Perception gives rise to many spontaneous beliefs -- beliefs that I do not *try* to have, and for which I am not normally conscious of any explicit argument. And these beliefs are undoubtedly fallible. But they are not basic.

It is important to stress this point: strong foundationalists, no less than their moderate cousins, are free to use beliefs that arise from faint perception, fallible memory and uncertain testimony in the process of justification. What they insist is that everyone who makes legitimate use of such beliefs, even if he is a moderate foundationalist (and therefore committed to denying the need to rely on certainties), is able to do so only because there is a deeper level at which there is something of which he is certain and on which the higher-level belief depends for its justification.

There is one other way to try to slip out the back door and evade the force of the argument against moderate foundationalism. We might try looking at perceptual beliefs simply as a class of beliefs that arise spontaneously, and count up the number of times that they turn out to be correct. If a high proportion of the perceptual beliefs that I form spontaneously turn out to be true, isn't this evidence that my perceptual beliefs are highly likely to be true? And can't we then say that they are foundational, though merely probable, after all?

This response turns on an equivocation in the use of the word 'probable.' In the sense in which this term (or something functionally equivalent) is needed for justification, it indicates something about the rationality of our beliefs. But as it is used in the preceding paragraph, 'probable' hasn't got anything to do with rationality: it has to do with *success*. The frequency with which my beliefs turn out true doesn't say anything about my rationality: they may turn out to be true in virtue of factors that have nothing to do with the quality of my reasons. (Perhaps some brilliant and slightly twisted scientists have decided to prime my brain with spontaneous beliefs about high-energy physics.) It may strain our credulity to think that I could spontaneously produce a whole string of esoteric but true pronouncements about physics without knowing anything about the field, but this does not prove that I am really being rational: we are simply strongly inclined to think that I must have some evidence that I'm not letting on about. The mere fact that some beliefs arise spontaneously and are true does not mean that they are justified, much less basic in the sense that foundationalism requires.

So there is no way out of the dilemma: strong foundations are necessary if we are to have justified inferred beliefs.

II. Are Strong Foundations Available?

Is there any empirical knowledge -- knowledge of facts, rather than of mathematical and logical truths -- that we literally cannot be wrong about it? Strong foundationalism asserts that there is and that all the rest of our empirical knowledge is based on these secure beliefs. But the argument of the preceding section does not give us any assurance that these foundations exist. It could turn out that they are necessary but unavailable; and in that case, the skeptic would turn out to be right after all.

Headaches and Hamburgers

Imagine a situation in which you believe, for the normal sort of reasons, that you have a hamburger in your hands. Is there even the slightest possibility that you could be wrong about this? Although from a practical point of view it is hard to take doubts about this seriously, we have to admit that it is just barely possible that you could be mistaken. It is possible, for example, that you are having an exceptionally vivid dream; it seems there's a hamburger there, but you're really asleep in your bed with no hamburgers around. Scenarios like this show that however secure our normal empirical beliefs seem in practice, they are not quite certain. It is possible for us to hold these beliefs and be wrong about them.

By contrast, your belief that you have a headache isn't open to this sort of challenge. You may be awake or asleep, drunk or sober -- it makes no difference. If you believe that you have a headache, you are *right*. What is more, your belief is *justified*. This is not at all like the bizarre case where you have (much to your surprise) a series of spontaneous true beliefs about linear accelerators while remaining clueless of physics. When it comes to your headache, you are aware without any inference or possibility of slippage of the very factors that make your belief a true one.

This same point can be extended to any sort of experience -- not just headaches, but visual sensations, tactile perception, apparent memories, and even rushes of emotion. In each case, there is a level at which you cannot be wrong about your experience. It has just those qualities that you believe it to have. Someone who tries to argue that you do not have a headache may succeed in making your headache go *away* (though the reverse process is more likely to occur), but no one can make you not have a headache at the very time that you are attending to the experience. Headaches, unlike hamburgers, do not have any hidden features. There is a difference between seeming to have a hamburger and really having one, and we can imagine circumstances (however bizarre) under which the two could come apart. But if your doctor tells you that you merely *seem* to yourself to be in desperate pain and that you are in fact feeling fine, then it is time to find a new doctor.

Reference and Incorrigibility

Strong foundations, on the account we are developing, have a special property known as *incorrigibility*. Literally the term means "uncorrectability," but in epistemology it has a more technical meaning: for my belief in some contingent proposition p to be incorrigible entails that,

necessarily, if I believe that p, then p is true. How can any contingent belief have this property? What kind of relation will guarantee the correctness of the belief?

The answer lies in reference. When I have a particular experience and express this fact to you, I am referring to something that I am directly aware of. Given the constraints of natural language, I am likely to try to use shared terms in order to convey to you the quality of the experience as I have it, but I do not have to describe it to myself in order to have the experience. Perhaps the closest linguistic construction to the belief I form is "I am experiencing *this*," where the italicized term picks out the experience in question by denoting it. A belief formed in this fashion cannot go wrong, for a very simple reason: if there were nothing for the term *this* to refer to, it would not be possible to form the belief at all.

It is important, because of the use we have in mind for these foundations, that incorrigible beliefs have one other property. In addition to being immune to error, they must be immune to being unjustified. Since they are supposed to function as basic beliefs, this means that there must be some intrinsic feature they possess that satisfies the minimal internal constraints on justified belief: the truth of the proposition must be not only guaranteed by our act of believing, it must be bound up in the proper way with our mode of access to the subject matter of these propositions. But from an internalist point of view, these constraints are satisfied by referentially formed beliefs. The relevant content is something that lies within the individual's field of awareness, and the very existence of that content is what makes the belief true. This is internal justification in the fullest sense that a basic belief can possibly have.

Objections and Replies

Every epistemological position has its critics, and strong foundationalism has more than most. Since these criticisms are largely responsible for the recent flight to moderate foundationalism, it is worth our while to address four of them here.

A fair number of critics of strong foundationalism confuse the classification of experiences with the content of those experiences. The first criticism we need to consider falls into this category. Bruce Aune advances an argument against strong foundations along the following lines:

1. All cognitively significant judgments involve predication.
2. All predication involves comparison.
3. All comparison relies on memory.
4. Memory is fallible.

-
5. All cognitively significant judgments are fallible.

The conclusion of this argument is, of course, incompatible with strong foundationalism. The basic beliefs of strong foundationalism are supposed to be cognitively significant (otherwise they can do no epistemic work), but they are also supposed to be certainties. If the argument works, it proves that strong foundations are entirely devoid of content.

In fact, it proves too much. If Aune's argument were sound, it would establish the impossibility of any cognitively significant judgments whatsoever. In order to make comparisons we must have something to compare our present experiences to; we must already have had and named a similar experience. But how, then, does anyone's mental life get off the ground? There must be a *first* cognitively significant experience, and if later comparisons are possible at all then it must be possible to assign to this first one some sort of description, even something as vague as "that unpleasant smell," in a noncomparitive way.

Classifying our experiences into groups that have relevant similarities is a fallible business, if only because we have to remember what the earlier experiences were like in order to make the comparison and memory is a fallible thing. But this doesn't mean that our experiences have no distinctive characteristics or that we create their characteristics by an act of mental sorting. On the contrary: sorting and classification make no sense unless the things sorted have some definite characteristics on the basis of which the sorting can be done.

A second objection turns on the question of empirical evidence for or against basic beliefs. It is possible (for all we know) that a device could be constructed that would track neurological activity in your brain and give scientists excellent reason to believe that you are in a certain mental state -- say, the state of feeling no pain. Suppose that such a cerebroscope were tested thousands of times without failing: when the device indicated that people were in pain, they confirmed it; and when it indicated that they were not, the test subjects also agreed. If this device were put on your head and indicated that you felt fine, could this count as evidence against your own belief that you have a splitting headache?

If the issue were a matter of piling up external evidence, it might. But in fact, you do not reason your way to the conclusion that you have a headache by noting your bloodshot eyes, the thermometer on the counter and the bottle of aspirin clutched in your trembling hand. Your belief is formed referentially and cannot go wrong. To insist that a cerebroscope could give us evidence to the contrary is to miss the point about these beliefs. They are not up for grabs in a contest of evidence; because of the special way that they are formed, they are not sensitive to inductive evidence. Insisting that they are open to refutation is insisting that they are not, after all, incorrigible. And in this context that is question-begging -- even if the person urging for the test is wearing a lab coat.

A third objection is that our typical empirical beliefs simply don't seem to be based on first-person foundations. In fact, it is difficult for most people to look out over a garden and "see it as" a patchwork of colors, a pure visual experience without any level of interpretation imposed on it. In her recent book *Evidence and Inquiry*, Susan Haack takes this objection to be seriously damaging to any epistemology built on phenomenal data.⁽³⁾

At one level this observation is perfectly legitimate, but to use it as an objection to strong foundationalism confuses psychological priority with epistemic priority. It is certainly true that the first explicit thought we have when looking at a garden is of the real, three dimensional physical garden, and it may be very difficult for us, unless we are professional artists, to create a mindset in which we can look out over the garden and be aware only of a collage of colors. But it is likewise difficult for a reader to focus on a line of clear prose and see it as a pattern of dark and light regions on a page. Yet when we look at a child learning to read it is obvious that this is precisely what is going on; and the development from a child to a mature reader is a continuous process. The moral we draw in the reading case is not that awareness of marks on the page is irrelevant to reading, but that increasing competence (mercifully) makes much of the process subconscious. Similarly, the answer to this challenge to strong foundationalism is that an adult's awareness of visual, tactile, and auditory stimuli is often subconscious but not therefore irrelevant to justification of empirical beliefs.

The fourth and most serious objection comes from a distinction between having a belief and having an experience. This objection is widely discussed but often presented in a confusing fashion, so it will help to lay it out in standard form:

1. If believing that I am having an experience on the one hand, and actually having the experience on the other, are not the same thing, then it is possible in principle for one of them to occur without the other occurring.

2. They are not the same thing.

3. It is possible for me to have a belief that I am having a particular experience without actually having the experience, and hence for my belief to be wrong.

There are multiple issues lurking just beneath the premises here. First, it is important to remember that the strong foundationalist is only urging the incorrigibility of beliefs that are referentially formed. Since it is possible to believe the same *proposition* in multiple ways, we must distinguish between the content of beliefs and the way in which belief is formed and remember that the latter may be epistemically significant. One obvious case where this comes up is in cases of the same proposition believed at two different times in two different ways. Yesterday, when you believed that you had a headache, you formed the belief referentially. Today you believe *that you had a headache yesterday*. The same factors make both the earlier and the later belief true; but today you believe this truth not referentially but rather by inferring it from memory, written records, photographs of your flushed face, etc. Strong foundationalists will not underwrite this latter belief as incorrigible.

Given this distinction, it could just turn out that the antecedent of premise 1 is false when applied to the sorts of beliefs that strong foundationalists have in mind. Since strong foundationalists are not demanding that all beliefs be explicit, they are free to construe belief in a modest way that is separable from and does not entail the existence of explicit verbalized judgment. Having an

experience at all, on this view, might come out to be equivalent to having a tacit referential belief. There is some room for flexibility in the strong foundationalist position here.

But regardless of whether we take this line, there is a much greater problem with the argument: it is invalid. The error can be seen clearly if we look at the following parallel argument:

1. If being a bear and being an animal are not the same thing, then it is possible in principle for one of them to be instantiated without the other being instantiated.
2. They aren't the same thing.

-
3. It is possible for there to be a bear that is not an animal.

The point is that the truth of premises 1 and 2 only guarantees that being a bear and being an animal can come apart *in one direction*. From the fact that we may have A without B, it does not follow that we may have B without A. Carried up to the objection above, this criticism amounts to pointing out that even if we were to allow that 1 is true because we may have experiences without forming tacit referential beliefs, it does not follow that it is possible for us to form referential beliefs without having the relevant experiences. Indeed, the referential connection shows why this is not possible. And with that, the fourth objection to the availability of strong foundations collapses.

III. Are Strong Foundations Sufficient?

Merely showing that we do all have basic beliefs about which we cannot be wrong does not solve all of our problems. A skeptic might well respond to the arguments presented so far with a shrug of his shoulders. "You've shown, I'll grant, that there are some empirical certainties -- beliefs about which you cannot be wrong," he might say, "and that may be interesting for its own sake. But how are you going to reason your way from these first-person experiential statements to claims about a world of mind-independent objects? It's a long way from 'I seem to see something red' to 'Here's an apple!'"

This gap between appearance and reality is the final challenge that strong foundationalists confront. How can we move from the way things *seem* to justified beliefs about the way things *are*? There are really two parts to this difficulty. First, is there any way to form beliefs about mind-independent reality? And second, supposing that there is, can we have good reasons to believe in the existence of that reality? Can we really support realism about our surroundings on these first-person foundations? Both questions, according to strong foundationalists, can be given an affirmative answer.

The Means of Ascent

The form of reasoning we need is known as *explanatory inference*, or inference to the best explanation. In a wide range of situations, from criminology to medicine, from particle physics

to automotive repair, we attempt to find explanations that will unify and render comprehensible a welter of facts. Strong foundationalism can make use of this form of inference in order to address the skeptical challenge of ascending from first-person basic beliefs to a justified superstructure.

Explanatory inference enables us to have justified beliefs because of a simple but powerful feature of confirmation theory. Given a theory T, which is neither guaranteed nor ruled out by the known facts, and a belief e which is also neither guaranteed nor ruled out by anything else we know, the following relationship holds:

IC If T raises the probability of e, then e is evidence for (raises the probability of) T.

This principle of Incremental Confirmation is a deductive consequence of the axioms of probability theory, and although it is not necessary for a given individual to know IC in order to be justified by reasoning in accordance with it, the principle itself can be known *a priori*. Most people working in the field today acknowledge it as a workable definition of evidence.⁽⁴⁾ When a scientist takes a new theory and derives an unexpected result from it, a result which turns out to be correct, there is strong reason to take the new theory seriously. This is, in outline, what happened when Einstein used the theory of relativity to predict that the apparent positions of certain stars during the full solar eclipse in 1919 would deviate from the apparent positions predicted by Newtonian physics.

Strong foundationalists can put the principle to work in the following fashion:

1. I have the surprising experience E.
2. But if there is a lawn in front of me, then of course I would have such experiences.

Therefore it is more probable than before that:

3. There is a lawn in front of me.

This argument is not deductive: the premises may be true while the conclusion in 3 is false. But it is not presented as a deductive argument; on the contrary, the argument is intended to be non-deductive. Strong foundationalists are not committed to the position that *all* justified beliefs are certainties! This is not a concession to moderate foundationalism; the merely probable beliefs at the higher levels do not run afoul of the dilemma presented in the first section, because unlike basic beliefs they are *inferred*. It is precisely in our inferred beliefs that we would expect probability rather than certainty.

The general form of the simplified argument given here is this: whenever we successfully anticipate and control our experiences on the basis of our theories, the theories gain in credibility. Even if we are not using exact numbers to quantify the degree of probability that a particular theory has, it remains true that under the conditions described the credibility of T given B is greater than the antecedent credibility of T. It is relatively simple to show that multiple lines of evidence supporting the same belief raise its credibility with dramatic speed. But this is precisely what we confront in daily life. The smell, sight, touch and taste I have from a

cheeseburger are all independent sensory clues that are well-explained by the actual existence of a cheeseburger in my hands but not well-explained by any rival hypothesis: they provide an overlapping and mutually reinforcing set of sensory evidence in favor of the belief that there really is a cheeseburger. Since we are constantly engaged in this sort of anticipation and control of our experience, the total confirmatory effect is overwhelming.

An important feature of this method of ascending from foundations -- which is popular with moderate as well as with strong foundationalists -- is that it leaves room for the introduction of new concepts. By contrast, any attempt to use pure rule-guided inductive extrapolation would not allow us to break outside of the circle of first-person experiences. From premises about falling apples, rule-guided extrapolation can at best give us a conclusion about more falling apples, or more vaguely about falling things. It permits us to omit concepts but not to introduce (for example) the notion of gravity as part of an explanation for why apples fall. Similarly, from premises of the form "I am experiencing like *this*" inductive extrapolation will, at most, permit us to infer conclusions of the form "In the future I will have similar experiences." Rule-guided extrapolation can only take the concepts involved in the premises and give us more of the same. But explanatory inference places no restrictions on the genesis of the concepts involved. The important epistemic feature is the relation between the hypothesis of realism on the one hand and our experience on the other. How the hypothesis arises is of no consequence.

Concept Empiricism and Concept Rationalism

This freedom accorded to us by explanatory inference would not get us very far if concept empiricism were true. According to concept empiricism, a position advocated somewhat inconsistently in Locke and with greater rigor by Hume, we cannot frame any concepts that are not given to us in experience. If all that we have is first-person experience, then we will never get around to a serious and independent third-person discourse: we will not be able to talk about objects except as a sort of shorthand for clumps of our individual experience.

This is a point on which phenomenologists and direct realists are in agreement. Where they disagree is just in the direction they take this principle. Phenomenologists, insisting that we are restricted to first-person data, maintain that we never get out and that realist talk about hamburgers is just a disguised way of talking about actual and possible experiences (visual, tactile, taste). Direct realists, insisting that we inhabit a world of real, mind-independent objects, demand that we must get our third-person concepts directly in experience: we must be *directly* aware of hamburgers. Many contemporary moderate foundationalists allow at least spontaneous beliefs about physical objects to stand as foundational; this position naturally lends itself to a form of direct realism.

Only indirect realists are blocked out by this principle. And this brings up an important point of contrast: unlike its moderate cousin, strong foundationalism of the sort we have been investigating does not have any third-person physical-object statements among the basic empirical beliefs and is a version of indirect realism. Strong foundationalists are therefore committed to concept rationalism -- the position that it is possible for us to create and employ concepts that go beyond experience.

There are good reasons to believe that concept empiricism is false. Empiricists have a notoriously hard time accounting for the genesis even of sensory concepts. Hume admitted in a moment of weakness that there were some sensory concepts that could be formed without having been given in experience, as when we look at a spectrum from which one color has been omitted and our minds "fill in" the intermediate shade. So much the more with concepts like "electron" or "gravity" does it strain our credulity to be asked to believe that these are either given directly in experience (as direct realists insist) or else just shorthands for aspects of our experience (as phenomenologists maintain).

A possible misunderstanding needs to be blocked here. Concept rationalists are not committed to a strong belief-voluntarism. It is highly implausible to say that, unaided by experience, I could believe *that there are objects* by a sheer effort of will. Just forming the concept of a mind-independent object, however, isn't the same as believing that there really are such objects, any more than forming the concept of a unicorn puts me at odds with established zoological consensus. Concept rationalism is simply the position that, contrary to concept empiricists, we can sometimes have concepts that are not constructions from our sensory or introspective data. And the idea of a persisting material object is one such concept.

Two Kinds of Priority

In at least one respect, the metaphor of "foundations" can be misleading. When we think of the foundations of a house, it is obvious that the supporting bricks must be laid down earlier than those that rest on them; and once the house is built, the foundations cannot change without tearing the whole edifice down. If the analogy were perfect, we would expect that basic beliefs, in order to give any epistemic support to a higher-level empirical belief, must be formed before the higher-order belief is formed and must thereafter remain as its sole support. But as a picture of human knowledge this is paradoxical, for at least two reasons. First, the basic beliefs allowed us by strong foundationalism are momentary. They involve our present experience (including, of course, present memory experiences, present beliefs, etc.), but in a moment those experiences will be gone, replaced at best with experiences that are qualitatively similar. So the foundations seem to have vanished! Second, we often attempt to support something we already believe by acquiring *more* evidence for its truth. I believe that I have left my wallet in the car because of a memory of setting it on the seat beside me; then I go out to the car and actually have the experience that I describe as seeing the wallet on the seat. But if strong foundationalism is right, the only lines of reasoning that count as evidence are those that can be traced to basic beliefs. So we seem to be acquiring basic beliefs *after* we already believe the hypothesis in question. Can strong foundationalism account for this?

Failure to make a distinction between epistemic and chronological priority can make foundationalism, and particularly strong foundationalism, look hopelessly rigid. But when the distinction is made, the "paradoxical" consequences turn out to be tame. Strong foundationalism, like any theory of knowledge, is a position about justification *at a given point in time*. It is, if you like, a philosophical claim about "snapshots" of justified belief; it lays out conditions that must be met if one's belief that P, at time t, is to count as justified. Strong foundationalists are not committed to any further claims about the evolution of our belief systems *across time*. In particular, they are free to acknowledge that a purely experiential, first-person basis for

knowledge will be constantly changing: old memories will fade, new experiences will offer fresh support for our beliefs, and so forth in a sort of cognitive "movie." Since this sort of development is not ruled out by foundationalism, we are free to acknowledge that many of our beliefs will actually predate, chronologically, our acquisition of the evidence that we *now* have to support them. What strong foundationalism requires is simply that our beliefs, to be justified at any given time, must *at that time* rest on a structure of supporting reasons that can be traced back *epistemically* to empirical certainties at their basis.

Restrictions

Even if we abandon concept empiricism, there is an objection to all forms of indirect realism that remains to be addressed. According to this criticism, indirect realism locks us behind a veil of perceptions. We may *guess* that we inhabit a world of mind-independent objects, but how can we ever hope to check up on such a guess? If I doubt whether a newspaper has reported a politician's speech accurately, I can hunt down an audio clip and check it out for myself. But according to indirect realism we can never get at the world directly; we can never manage to get a firm grip on the objects that are supposedly causing our experiences. "To know that there was [such a causal connection]," Jonathan Bennett writes, "we should need independent access to empirical facts about the objective realm."⁽⁵⁾ As a consequence, the objection runs, we can never tell whether our conjectures are correct. Indirect realism dooms us to skepticism.

In his article in defense of phenomenalism W. T. Stace parses out the intuition behind this objection by placing a restriction on our legitimate inferences that runs, roughly, as follows:

R We can come to know that A causes B only if A is the sort of thing with which, at least some of the time, we can be directly acquainted.

This restriction seems reasonable so long as the only mode of inference available to us is the method of cross-checking our perceptions against the facts themselves. If we can never, so to speak, grab the real world by the scruff of the neck and hold it up against our sensory experiences to see if they match, then it seems we can never have good reasons to believe facts about the world beyond our senses. But on a closer look, it turns out that R is in conflict with IC, the principle of incremental confirmation that we looked at earlier. According to that principle, *any* successful theory that meets the conditions spelled out can receive enhanced credibility for its successes.

And in point of fact, we constantly use IC in circumstances where R is violated, and it seems obvious that such reasoning can be justificatory. Consider black holes. By definition, their mass is so great that light cannot escape from them, so we cannot see them directly. Nevertheless, the behavior of nearby visible objects may give us excellent reason to believe that there is a black hole in a given region. We infer the existence of the black hole because it enables us to give a good causal explanation of the visible phenomena we do in fact observe. Or consider contemporary research on fundamental particles. Some of these particles are so small that it is not even theoretically possible to see them. To get around this difficulty, physicists use cloud chambers for studying nuclear transformations: they observe trails of water droplets that condense in the wake of a charged particle as it ionizes the water vapor through which it passes.

If a theory predicts that a particular sort of particle will be produced in a given reaction, scientists look for its telltale vapor trail -- a trail that is caused by the particle -- rather than for the particle itself.

Of course, skeptics about empirical knowledge are likely to be skeptics about science. But the point is not that current scientific theories are true: it is that this mode of non-deductive inference is reasonable regardless of whether the particular conclusions we draw from it turn out to be correct. If skeptics want to attack realism, they had better find a means that does not rely on the dubious restriction R.

Rivals

A final obstacle to the use of strong foundations in defense of realism is one suggested by Descartes himself. All of our first-person experiences could, in principle, be brought about by an extremely powerful being (Descartes calls him a "deceiver of the utmost power and cunning") whose goal was to deceive us into thinking that we inhabit a physical world. Granting that the hypothesis of realism predicts our experience and is therefore confirmed by it, it seems that the hypothesis of an evil deceiver does so as well. The two theories are empirically equivalent: they offer us precisely the same expectations. Hence they will apparently be confirmed in tandem, with the unhappy result that neither one can rise to more than a 50% credibility level no matter how much evidence we get. This isn't very encouraging.

The objection, if it goes through at all, demolishes all forms of foundationalism, weak as well as strong. Some philosophers have tried to block it by claiming that the idea of such a deceiver makes no sense or is conceptually incoherent, but this seems to violate our very strong intuition that it is at least *possible*, however unlikely, that we should be thus deceived. Foundationalists will need to tackle this one directly.

The best approach to this difficulty starts with a look at confirmation theory. Even when two theories both offer successful predictions and both rise in credibility as a result, they are not necessarily confirmed by the same amounts. In particular, if theory T_1 has a higher initial credibility than theory T_2 , and they both raise the likelihood of an empirical truth e to the same extent, then it can be shown from the axioms of probability theory that T_1 will get a bigger "boost" from the successful prediction of e than T_2 does. The gap between their credibilities will widen in favor of the initially more credible theory with each successful confirmation.

And it would seem that the initial credibility of the existence of an evil deceiver is at least somewhat lower than that of ordinary realism. One reason for this is that in the deceiver scenario there is at least one thing that is never represented to us by any sensation: the deceiver himself, who must stay "hidden" through this process if the Cartesian scenario is to remain empirically equivalent to realism. For this reason, there is no explanatory advantage to invoking the existence of the deceiver: we can (and do) manage to get around without hypothesizing his existence, and there is no empirical or scientific problem for which his existence could, in principle, be of the slightest interest to us. In the Cartesian scenario we have not only the deceiver but also all of his individual mental states which are the causes of our particular experiences, not to mention the structure of whatever sort of (non-deceptive) world the deceiver

actually inhabits. The existence of the deceiver over and above these particular causes of our experiences is a "fifth wheel." Though we may grant that his existence is logically possible, it is perfectly reasonable for us not to take it seriously since we have at hand a simpler, more plausible and therefore better-confirmed explanation for our experience: that we inhabit a real world of mind-independent objects.

The Wider Picture

Why should anyone care what theory of knowledge is correct? Epistemology is pretty obviously *hard*, and at first blush it doesn't seem that any particular scientific achievement or practical course of action is likely to be affected by our choice of strong foundationalism over weak foundationalism or even some non-foundational theory. Why should we rack our brains over theories that don't have any practical application?

The simple but sobering answer is that human beings have an overwhelming desire to know. Many of the scientific achievements we most admire have been brought about because people had a burning desire to figure out how things really work, to get it *right*. Any theory of knowledge that writes off these achievements as an illusion or an accident will have a catastrophic effect on our world view. Closer to home, we would like to think that our beliefs about cars, houses, trees and sidewalks are in some important sense better off than the beliefs of people who take their cues from astrologers and crystal balls. If a theory of knowledge offers no defense against the argument that we are all ultimately required to retreat to a point of unargued, unjustified commitment, then we are wrong: there is no ultimate epistemic difference between our beliefs and the beliefs of devout horoscope-readers, and it is merely a historical accident that horoscope-readers are currently a minority of the population.

A sound and defensible theory of knowledge is our only *rational* line of defense against conceptual anarchy. This is the best and perhaps the only reason to take epistemology seriously; but it is also a sufficient one.

ENDNOTES

1. Or at least, this is the form that foundationalism takes when it is advanced by an internalist. But since I find the case for epistemic internalism wholly convincing, I will take it for granted that we are speaking among internalists here.
2. See J. L. Mackie, *Truth, Probability and Paradox* (Oxford: Clarendon Press, 1973), 158ff for a discussion of this point.
3. Susan Haack, *Evidence and Inquiry* (Cambridge, MA: Blackwell, 1993), pp. 107-8.
4. Most, but not quite all. Peter Achinstein maintains in *The Nature of Explanation* (New York: Oxford University Press, 1983) that the condition is neither necessary nor sufficient. I find his arguments unpersuasive, but the dispute would take us far afield here.
5. *Locke, Berkeley, Hume: Central Themes* (New York: Oxford University Press, 1971), p. 70.