7 The architecture of knowledge

On the mountain in the distance before me, I see the huge oak and tulip trees swaying, with their leaves turned upward revealing the lighter green of their undersides. Waves of green, light and then dark, seem to cross the surface of the upper region from west to east as the leaves show the colors of their inner and outer sides. Parts of the hillside seem almost to breathe in and out as the trees bend away from me and back. It is a familiar sight and I immediately realize that there is a wind. My belief that there is a wind is based on my belief that the trees are swaying. It is also justified on the basis of that belief. And if I know that there is a wind, I know it on the basis of my belief that they are swaying. In each case, one belief is inferentially based on another.

To what extent does this relation in which one belief is based on another represent the structure of our belief systems as a whole? The question is especially pertinent to epistemology as applied to the common cases in which our beliefs constitute knowledge, as they so often seem to. Might perceptual beliefs, for instance, form a foundation on which others are inferentially built? Or are the former just a stopping place on the way to something yet more basic, or perhaps merely a place where we usually stop pursuing further premises, though we might go on seeking them and find deeper grounds that support perceptual beliefs?

These questions represent perennial issues, and we shall see many versions of the *foundationalist view* – the classical position on them – and various opposing theories. The questions also take us, as often happens in epistemology, into questions about the nature of mind as well as questions directly about justification and knowledge. This is certainly to be expected where the central topic is the structure of knowledge and justification; for knowledge is apparently constituted by belief, and, in epistemology, justification is important chiefly in connection with belief. It is appropriate, then, to begin an exploration of the structure of knowledge and justification with some major points about the structure of a person's body of beliefs.

Inferential chains and the structure of belief

As in discussing inference, it is useful to call the kind of inferential belief, justification, and knowledge just illustrated *indirect*. For one has such beliefs, justification, and knowledge only on the basis of, and thereby *through*, other beliefs, justification, or knowledge. By contrast, my belief that the trees are swaying is direct. I believe this simply because I see it, not on the basis of something else I believe.

Infinite inferential chains

The natural picture just sketched can be challenged. Perhaps all our beliefs could be *indirect*. If they could be, could not all justification of belief, and all our knowledge, be indirect? An adequate epistemology requires answers to these questions about the structure of a body of belief, justification, or knowledge. In exploring them, I will talk above all about knowledge and justification, and especially about knowledge. But what we know (propositionally) we believe; and the kind of justification epistemology is chiefly concerned with is that of belief. The structure of my knowledge and justification, then, is chiefly that of a certain body of my beliefs.

I am not talking about knowledge in the abstract, as we sometimes do. We speak, for instance about the extent of "human knowledge." Some of this knowledge is solely in books, and not remembered by anyone. Thus, some scientific knowledge might be of propositions no one actually believes, propositions available to us should we need them, but not objects of actual belief. We *can* talk about the structure of such knowledge in the abstract, say about whether all the propositions of scientific knowledge can be systematized by certain basic laws of physics and chemistry. Then these basic laws would be geometrical axioms, and the other laws, like its theorems, would be derivable from the basic laws. But that is not my topic. I am exploring how people's beliefs may actually be structured.

I want to start with a simple example. When I am being very cautious, my belief that the trees are swaying could be based on my belief that I have a visual impression of swaying. Could the latter belief also be based on another one? What might that be? Might I now believe that it *seems* to me that I have a visual impression of swaying, and base my belief that I have that impression on this new belief? This is doubtful. I cannot base one belief on another simply because I want to.

This example shows that the view that what we believe, and certain relations between our beliefs, are entirely under the direct control of our wills – a strong version of *doxastic voluntarism* (voluntarism about belief) – is a mistake. Suppose, for instance, that I want to believe someone's testimony. If it seems false, I cannot make myself believe it just by willing to believe it. I also lack direct voluntary control over what my beliefs are based on; if I already know first-hand that I am gravely ill, I cannot, simply by willing it, *base* my belief of this on someone's testimony that it is so.[1]

Even if one cannot base one belief on another at will, it might still seem that a sequence of beliefs, each based on the next, could go on without limit. But could I, for instance, believe what seems the next proposition in the evidential series, the involuted proposition that it appears to me that it seems to me that I have a visual impression of swaying? I suppose I could (though not simply at will). Still, I do not see that I would now come to hold anything on the *basis* of believing this strange proposition.

Suppose, however, that I did come to hold, on the basis of this involuted proposition, that it seems to me that I have an impression of swaying. I cannot in this way manufacture an *inferential chain* of beliefs – a chain in which each belief is based on the next – running to infinity. Nor do I already have an infinite set of appropriate beliefs as raw material waiting to be brought to consciousness – if indeed I *can* have an *infinite* number of beliefs (particularly outside mathematics, where it may seem that I can have an infinite number corresponding to the series 2 is even, 4 is even, 6 is even, etc.).[2]

Circular inferential chains

So far, however, I have ignored another way in which it might be thought to be possible that every belief is indirect: by virtue of lying not at the origin of an infinite chain, but instead in a circular chain. Imagine that I could hold one belief on the basis of a second and a second on the basis of a third, and so on, until we come full circle and get to a belief I hold on the basis of the first. Then all my beliefs would be indirect, yet I need not have infinitely many. To assess this, recall my belief that there is a swaying. Might there be a circular chain of beliefs here? For instance, could my belief that it appears to me that it seems to me that I have a visual impression of swaying be based on my belief that there is a swaying? This is far from clearly possible.

Suppose for the sake of argument that I do have a circular chain of beliefs, each based on the next. This raises a problem. First, there is good reason to think that (a) one belief is based on a second only if the second is at least in part causally responsible for (one's holding) the first. For instance, if I believe there is a wind, on the basis of my believing that the trees are swaying, then I believe that there is a wind, at least in part *because* I believe that the trees are swaying. Second, there is good reason to think that (b) if one thing is in part causally responsible for a second and the second is in part causally responsible for a third, then the first is in part causally responsible for the third. But together these two points imply that (c) in a circular chain of beliefs, each based on the next, every belief is in part causally responsible for, and thus a partial cause of, itself. That seems impossible. To see why, let us explore how such a circle might go in a simple case.

Imagine a circle of three beliefs, each based on the next. (1) I believe there is a wind. I believe this on the basis of (2) my believing there is a swaying of the trees; I believe that there is this swaying, on the basis of (3) my believing I have an impression of such swaying; and I believe that I have this impression, on the basis of believing there is a wind. This case would be a *circular causal chain*, one whose last link is connected to its first in the same way that each is connected to its successor. For, given point (a), belief (1) is in part causally responsible for belief (3), and, given point (b), (3) is in part causally responsible for (1). This implies, however, given (b), that (1) is in part causally responsible for itself. That is apparently impossible. The belief would be holding itself up by its bootstraps.

If the bootstraps problem shows that circular causal chains of this kind are not possible, then there cannot be a circular chain of beliefs each based on the next; for on the highly plausible assumptions, (a) through (c), this would have to be a circular causal chain. (We have not assumed that the imagined chain implies that some belief must be *based* on itself, only that such chains imply a belief's being in part causally responsible for itself; this basis relation implies more than a causal connection.)

It may seem that a wheel is a model of a circular causal chain of the relevant kind and that something must therefore be wrong with the reasoning just noted. Consider a wheel standing on the ground in a line running east and west, and imagine the wheel having eight equal sections and an axle, each section consisting of a pie-slice segment with its apex at the axle. Does each section not support the next, so that each is "based on" or rests on the others and ultimately (in the eighth link) on itself?

If we distinguish between the relation of *being connected with* and that of *supporting*, the answer no longer seems clear. Granted that if one section is connected to a second, it *will* support the second *if* a force is applied to the second in the direction of the first. But a wheel with eight such connected sections can exist in empty space with no such forces acting on it. Mere connectedness between segments does not imply any actual support relations, only a readiness to enter them.

Consider, then, the realistic case in which the wheel is on the ground. Gravity exerts a downward force on the entire wheel. Here, however, the ground supports the entire wheel, and each segment of the wheel that has a segment above it supports that segment, with the two top sections (whose common seam, we may assume, runs straight up from the center of the wheel to its highest point) being the only ones plausibly said to support each other directly. But notice that each of the top sections supports the other *with respect to* a different force. There is a westward force in the case of the western section's support of the eastern one (which would fall backwards to the east if disconnected from its western counterpart because all the seams become unfastened); and there is an eastward force in the case of the eastern section's support of the western one (which would fall backwards to the west if disconnected from its eastern counterpart because all the seams become unfastened).

Each top section of the wheel, then, pulls on the other in the opposite direction, with the result being a balance. In no case do we get a force in *one* direction that goes fully around the circle with the result that any section supports itself in that same direction. The forces on the two top sections are, as described in physics, equal and opposite.

Returning to the case of belief, there the support in question – the kind of cognitive support given by one belief to a second that is based on it – is

also in one direction. It is, in good part, support with respect to three dimensions: conviction, explanation, and memory. Consider this cognitive force in relation to a common case, that of a conclusion belief being based on a belief of a premise for it, such as a point made by a respected friend in favor of the conclusion. My premise belief tends to increase or buttress my conviction in my conclusion belief, to explain (in part and in some broadly causal way) why I hold that belief, and to help me remember my conclusion. This is not the kind of support relation that a belief may be plausibly thought to bear *to itself*.

One might think that a belief of a self-evident proposition can be in part causally responsible for itself and in that way support itself. But this seems at best an inaccurate way of saying that such a proposition is not believed because one believes *something else*. That is normally so; normally, one believes it because one grasps the appropriate conceptual relation(s) it expresses. In any case, our concern is beliefs in general, not just beliefs of self-evident propositions.

On balance, then, it is reasonable to conclude not only that we *have* direct beliefs, such as beliefs about colors before us and beliefs of self-evident propositions, but also that we could not have *only* indirect beliefs. Neither infinite nor circular chains of indirect beliefs are possible for us.

The epistemic regress problem

Is knowledge like belief in this, so that some of it is direct, or could all our knowledge be indirect, that is, based on other knowledge we have? It may seem that this is possible, and that there can be an infinite *epistemic regress* – roughly, an infinite series of knowings each based on the next.

It is especially likely to appear that indirect knowledge need not always be based on direct knowledge, if one stresses that, very commonly, 'How do you know?' can be repeatedly answered, and one then supposes that we stop answering only for practical reasons having to do with our patience or ingenuity. Let us explore this issue by assuming for the sake of argument that there is indirect knowledge and seeing what this implies.

Assume that a belief constituting indirect knowledge is based on knowledge of something else, or at least on a further belief. The further knowledge or belief might be based on knowledge of, or belief about, something still further, and so on. Call this sequence an *epistemic chain*. It is simply a chain of beliefs with at least the first constituting knowledge, and each belief linked to the previous one by being based on it.

It is often held that there are just four possible kinds of epistemic chain. Two kinds are unanchored and do not end; two kinds are anchored and do end. First, an epistemic chain might be infinite, hence entirely unanchored. Second, it might be circular, hence also unanchored. Third, it might end with a belief that is not knowledge, and thus (figuratively speaking) be anchored in sand. Fourth, it might end with a belief that constitutes direct knowledge, and thus be anchored in bedrock. Our task is to assess these chains as possible sources of knowledge or justification. This is a version of the epistemic regress problem.

Infinite epistemic chains

The first possibility is difficult to appreciate. Even if I could have an infinite number of beliefs, how would I ever know anything if knowledge required an infinite epistemic chain? To know, and thus to *learn*, the simplest kind of thing, such as that there is a green field before me, I would apparently have to know an infinite number of things.

It is doubtful that, given our psychological make-up, we *can* know, or even believe, infinitely many things. It might seem that we can have an infinite set of arithmetical beliefs, say that 2 is larger than 1, that 3 is larger than 2, and so forth. But surely for a finite mind there will be some point or other at which the relevant proposition cannot be grasped (the point might be different for different people or even the same person at different times). Imagine the "largest" proposition a supercomputer could formulate after years of work. It could easily be too long to understand and so cumbersome that one could not even take in a formulation of it. One would be unable to remember enough about the first part of it when one gets to the end; one could thus never understand the whole thing. What we cannot understand we cannot believe; and what we cannot believe we cannot know.[3]

Even if we could have infinite sets of beliefs, however, infinite epistemic chains apparently could not account for all, and probably not for any, of our knowledge. In the case of some beliefs, such as the belief that if some dogs are pets, some pets are dogs, I cannot even find *any* belief I hold that yields another link (a belief this one seems to be based on). The proposition is luminously self-evident, and it is difficult even to imagine a further proposition I would consider a good premise on the basis of which I would believe it if I thought I needed a premise for it. Thus, I find it unclear how this belief could be grounded, as knowledge, by any epistemic chain, much less by an infinite one.

In any event, how might infinite epistemic chains help us account for any other knowledge (or justified belief)? Notice that many kinds of infinite chain are possible. No one has provided a plausible account of what kind might generate justification or knowledge. But some restrictions are badly needed. For any proposition, an infinite chain can be imagined (in outline) that may be claimed to provide support for the proposition. Thus, even for a proposition one believes to be obviously false, one would find it easy to imagine beliefs to back it up; and though one could not continue doing this to infinity, one could nonetheless claim that one *has* the infinite set required to support the original belief.

Take the obviously false proposition that I weigh at least 500 pounds. I could back up a belief of this by claiming that if I weigh at least 500.1

pounds, then I weigh at least 500 (which is self-evident), and that I *do* weigh at least 500.1 pounds. I could "defend" *this* by appeal to the propositions that I weigh at least 500.2 pounds, and that if I do, then I weigh at least 500.1. And so forth, until the challenger is exhausted. A chain like this can be infinite; hence, no matter how ridiculous a proposition I claim to know, there is no way to catch me with a claim I cannot back up in the same way. Given such resources, anything goes. But nothing is accomplished.

Circular epistemic chains

The possibility of a circular epistemic chain as a basis of knowledge has been taken much more seriously. It might seem that if there cannot be a circular causal chain of indirect beliefs, each based on the next, then there cannot be a circular epistemic chain either. But perhaps knowledge can be based on premises in a way that differs from the way belief is based on them; perhaps, for instance, my *knowledge* that there is a wind could be somehow based on my belief that the leaves are swaying, even though my *belief* that there is a wind is not based on any further belief. We would then have a circle of knowledge, but not of belief, and no causal bootstraps problem. If this is possible, it may turn out to be important. But how realistic is it?

Does any of our knowledge really emerge from circular epistemic chains? Let us try to go full circle. I know there is a wind. I know this on the basis of the swaying of the trees. Now I think I know they are swaying because I *see* them sway. But it might be argued that my seeing this is only the causal basis of my *belief* that they are swaying, and I just do not notice that it is only on the basis of, say, my knowledge that I have a visual impression of swaying that I *know* they are swaying. Perhaps. But how far can this go?

I do not see how to go full circle, unless I *think up* propositions I do not originally believe, hence do not originally know. If I do not originally have any belief of them, then I (originally) have no *justified* belief or knowledge of the premise they express, and thus no belief appropriate to serve as a link in the epistemic chain or play any supporting role toward my original knowledge.

Suppose, however, that I do think up a suitable set of evidential propositions, come to know them, and make my way full circle. Suppose, for instance, that I get as far as knowledge that it seems to me that I have a visual impression of swaying. Might I know this on the basis of knowing that there is a wind (the first link)? How would knowledge that there is a wind justify my belief that it seems to me that I have a visual impression of tree swaying? I apparently know introspectively, not perceptually or inferentially, that I have the impression of swaying. Other difficulties also beset the circular approach. But these problems alone cast sufficient doubt on it to suggest that we consider the remaining options.

Epistemic chains terminating in belief not constituting knowledge

The third possibility for the structure of epistemic chains, that an epistemic chain terminates in a belief which is not knowledge, can be best understood if we recall that in discussing the transmission of knowledge, we noted both source conditions and transmission conditions. If the third possibility can be realized, then knowledge can originate through a belief of a premise that is not known. On the basis of believing that there is a swaying, for example, I might know that there is a wind, even though I do not know that there is a swaying. The regress is thus stopped by grounding knowledge on something else, but not in the *way* it is normally grounded in experience or reason.

Is this possible? In one kind of case it is not. Suppose that (in foggy conditions) I simply guess that what I see is a swaying of trees, but happen to be right. Might I then know there is a wind anyway, provided there is? Surely not; knowledge cannot be grounded in such guesswork, even when the guess is correct.

Imagine, however, that although I do not know there is a swaying, I do hear some sounds that might indicate swaying, and I make an *educated* guess and am thereby justified, to *some* extent, in believing that there is. If, on the basis of this somewhat justified belief that there is a swaying, I now believe that there is a wind, and there is, do I know this?

The answer is not clear. But that would be no help to proponents of the third possibility, who claim that knowledge can arise from belief which does not constitute knowledge. For it is equally unclear, and for the same sort of reason, whether my guess that there is a swaying is *sufficiently* educated – say, in terms of how good my evidence is – to give me (a weak kind of) knowledge that there is a swaying. If it is clear that my guess is *not* sufficiently educated to yield this knowledge, then I also do not know there is a wind. If it is clear that the guess is educated enough, I apparently do know that there is a wind, but my knowledge would be based on other knowledge, hence would not realize the third possibility.

Notice something else. In the only cases in which the third kind of chain is at all likely to ground knowledge, there is a degree – perhaps a substantial degree – of justification. If there can be an epistemic chain which ends with belief that is not knowledge only because the chain ends, in this way, with justification, then it appears that we are at least in the general vicinity of knowledge. We are at most a few degrees of justification away. The sand has turned out to be rather firm; it is at least close to being firm enough to support knowledge.[4]

Epistemic chains terminating in knowledge

The fourth possibility is the one apparently favored by common sense: epistemic chains end in direct knowledge – in the sense that they have direct knowledge as their last link. That knowledge, in turn, is apparently grounded

(anchored, if you like) in experience or in reason, and this non-inferential grounding explains how it is (epistemically) direct: it arises, directly, from perception, memory, introspection, or reason (or indeed from testimony, provided this has an appropriate grounding in at least one of the first four).

The ground-level knowledge just described could not be inferential; otherwise the chain would not end without a further link. To illustrate, normally I know that there is a swaying just because I see that there is. Hence, the chain grounding my knowledge that there is a wind is anchored in my perception.

Such experientially or rationally grounded epistemic chains may differ in many ways. Here are four. They differ in *composition*, in the sorts of beliefs constituting them. They differ in the kind of *transmission* they exhibit; it may be deductive, inductive, or combine both deductive and inductive links. Epistemic chains also differ in their ultimate grounds, the *anchors* of the chains, which may be experiential or rational; and epistemic chains may vary in justificational *strength*, the degree of justification they give to the initial belief.

Different proponents of the fourth possibility have held various views about the character of the *foundational knowledge*, that is, of the beliefs constituting the knowledge that makes up the final link of the epistemic chain that is anchored in experience or reason. Some philosophers, for instance, have thought that the appropriate beliefs must be infallible, or at least indefeasibly justified. But this is not implied by anything said here. All that the fourth possibility requires is *direct* knowledge, knowledge not based on other knowledge (or on justified belief).

Direct knowledge need not be of self-evident propositions, or constituted by indefeasibly justified belief. Introspective beliefs illustrate this. The proposition that I am now thinking about knowledge is not self-evident. It is not even *self*-evident to *me*. First, it is evident to me, not in itself, as is the proposition that if some dogs are pets then some pets are dogs, but on the basis of my conscious experience. Second, since I realize that my reflections can sometimes merge into daydreaming, I do not even consider it rock-solidly true in the way I do self-evident propositions. But surely I do have direct knowledge of the proposition.

The epistemic regress argument

What we have just seen suggests a version of the *epistemic regress argument*. It starts with the assumption that

(1) if one has any knowledge, it occurs in an epistemic chain.

Epistemic chains are understood to include the special case of a single link, such as a perceptual or a priori belief, which constitutes knowledge by virtue of being anchored directly (non-inferentially) in one's experience or reason.[5]

The argument then states that

(2) the only possible kinds of epistemic chain are the four mutually exclusive kinds just discussed: the infinite, the circular, those terminating in beliefs that are not knowledge, and those terminating in direct knowledge.

Its third, also restrictive premise is that

(3) knowledge can occur only in the fourth kind of chain.

And the argument concludes that

(4) if one has any knowledge, one has some direct knowledge.[6]

A similar argument was advanced by Aristotle, and versions of this regress argument have been defended ever since.[7]

As proponents of the argument normally understand (1), it implies that any given instance of indirect knowledge depends on at least one epistemic chain for its status *as* knowledge. So understood, the argument clearly implies the further conclusion that any indirect knowledge a person has *epistemically depends on*, in the sense that it cannot be knowledge apart from, an appropriate inferential connection, via some epistemic chain, to some direct knowledge that the person has.

Given this dependence assumption, the regress argument would show not only that if there is *in*direct knowledge, there is direct knowledge, but also that if there is indirect knowledge, that very knowledge is *traceable* to some direct knowledge as its foundation. One could trace an item of indirect knowledge to some premise for it, and, if there is a premise for that, to the next premise, and so on until the chain is anchored in a basic source of knowledge.

A similar argument applies to justification. We simply speak of *justificatory chains* and proceed in a parallel way, substituting justification for knowledge; and we arrive at the conclusion that if one has any justified beliefs, one has some directly justified beliefs. Similarly, if one has any indirectly justified belief, it exhibits *justificational dependence* on an epistemic chain appropriately linking it to some directly justified belief one has, that is, to a foundational belief.

Foundationalism and coherentism

These two sets of conclusions constitute the heart of the position called *epistemological foundationalism*. The first set, concerning knowledge, may be interpreted as the thesis that the structure of a body of knowledge, such as yours or mine, is foundational, where this is taken to imply that any indirect

(hence non-foundational) knowledge there is depends on direct (and thus in a sense foundational) knowledge. The superstructure, one might say, rests on the foundations. The second set of conclusions, regarding justification, may be interpreted as the thesis that the structure of a body of justified beliefs is foundational, where this is taken to imply that any indirectly (hence non-foundationally) justified beliefs there are depend on directly (thus in a sense foundationally) justified beliefs.

In both cases, different foundationalist theories may diverge in the kind and degree of dependence they assert. A strong foundationalist theory of justification might hold that indirectly justified beliefs derive *all* their justification from foundational beliefs; a moderate theory might maintain only that the former would not be justified apart from the latter, and the theory might grant that other factors, such as coherence of a belief with others one holds that are *not* in the chain can add to its justification.

None of the foundationalist theses I have stated says anything about the *content* of a body of knowledge or of justified belief, though proponents of foundationalism usually specify, as René Descartes does in his *Meditations on First Philosophy* (first published in 1641), what sorts of content they think appropriate. Foundationalism, as such, thus leaves open what, in particular, is believed by a given person who has knowledge or justified belief and what *sorts* of propositions are suitable material for the foundational beliefs. I want to talk mainly about foundationalism regarding knowledge, but much of what I say can be readily applied to justified belief.

Foundationalism has been criticized on a number of points. Let us focus on the most important objections that stem from the best alternative theory of the structure of knowledge, *coherentism*. There are many versions of coherentism, including some that seem to be based mainly on the idea that if an epistemic circle is large enough and sufficiently rich in content, it can generate justification and account for knowledge. But we have seen serious difficulties besetting circular chains. I therefore want to formulate a more plausible version of coherentism.

The central idea underlying coherentism is that the justification (justifiedness) of a belief depends on its coherence with other beliefs one holds. The *unit of coherence* – roughly, the range of the beliefs that must cohere in order for a belief among them to derive justification from their coherence – may be as large as one's entire set of beliefs (though of course some may figure more significantly in producing the coherence than others, say because of differing degrees of closeness to one another in their subject matter).

The variability of the unit of coherence would be accepted by a proponent of the circular view, but the thesis I want to explore differs from that view in not being *linear*: it does not construe justification or knowledge as emerging from an inferential line going from premises to that conclusion, and from other premises to the first set of premises, and so on, until we return to the original proposition as a premise.

In the circular coherentist view, no matter how wide the circle, there is a *line* from any one belief in a circular epistemic chain to any other. In practice we may never *trace* the entire line, as by inferring one thing we know from a second, the second from a third, and so on, until we re-infer the first. Still, on this view, there is such a line for every belief that constitutes knowledge. Thus, the kinds of problems we encountered earlier regarding circular epistemic chains must be resolved (as I doubt they can be) if the view is to be sustained.

Holistic coherentism

Coherentism need not be linear. It may be holistic. To see how a holistic theory of knowledge (and justification) works, consider a question that evokes a justification. John wonders how I know, as I sit reading, that the wind is blowing. I say that the leaves are rustling. He then asks how I know that Sally is not just making this noise by walking in the high grass. I reply that the high grass is too far away. He now wonders whether I can distinguish rustling leaves from the sound of a quiet car on the pebbled driveway. I reply that what I hear is too much like a whisper to be the crunchy sound of pebbles under tires.

Patterns of justification

In giving this kind of justification, I apparently go only one step along the inferential line: just to my belief that the leaves are rustling. For my belief that there is a wind is based on this belief about the leaves. After that, I do not even mention anything that this belief, in turn, is based on. Rather, I defend my beliefs as appropriate in terms of an entire pattern of mutually cohering beliefs I hold. And I may cite many different parts of the pattern. For instance, I might have said that walking through high grass sounds different from windblown leaves. On the coherentist view, then, beliefs representing knowledge do not have to lie in a grounded chain; they fit a coherent pattern, and their justification emerges from their fitting that pattern in an appropriate way.

Consider a different sort of example. A gift is delivered to you with its card apparently missing. The only people you can think of who send you gifts at this time of year live in Washington and virtually never leave, but this is from Omaha. That origin does not cohere well with your hypothesis that it was sent by your Washington benefactors, the Smiths. Then you open it and discover that it is frozen steak. You realize that this can be ordered from anywhere. But it is not the sort of gift you would expect from the Smiths. A moment later you recall that you recently sent them cheese. You suppose that they are probably sending something in response. Suddenly you remember that they once asked if you had ever tried frozen gourmet

steaks, and when you said you had not they replied that they would have to give you some one of these days.

You now have a quite coherent pattern of beliefs and might be justified in believing that it was they who sent the package. If you come to believe this on the basis of the pattern, you presumably have a justified belief. When you at last find their card at the bottom of the box, then (normally) you would *know* that they sent the package.

The crucial things to notice in this example are how, initially, a kind of *incoherence* with your standing beliefs prevents your justifiedly believing your first hypothesis (that the box came from the Smiths) and how, as relevant pieces of the pattern developed, you became justified in believing, and (presumably) came to know, that the Smiths sent it. Arriving at a justified belief, on this view, is more like answering a question in the light of a whole battery of relevant information than like deducing a theorem by successive inferential steps from a set of luminous axioms.

A coherentist response to the regress argument

It is important to see how, using examples like those just given, holistic coherentism can respond to the regress argument. It need *not* embrace the possibility of an epistemic circle (though its proponents need not reject that either). Instead, it can deny the premise that there are only the four kinds of possible epistemic chains so far specified. There is a fifth: a chain terminating with belief that is *psychologically direct*, yet *epistemically indirect* (or, if we are talking of coherentism about justification, *justificationally indirect*). This is in effect to grant foundationalists that they are right about our psychology, while insisting that they are wrong about epistemology. Let me explain.

The idea is that although a terminal, direct belief is not *psychologically* based on any other, as where it is inferentially grounded on another, its justification nonetheless *is* based on other beliefs. Hence, the last link is, as belief, direct, yet, as knowledge, indirect, not in the usual sense that it is inferential but in the broad sense that the belief constitutes knowledge only by virtue of receiving support from other knowledge or belief. This belief is *psychologically foundational but epistemically dependent*. Its justification depends on a pattern of supporting beliefs.

To illustrate all this, consider again my belief that there is a swaying of the trees. It is psychologically direct because it is simply grounded, causally, in my vision and is not inferentially based on any other belief. Yet (the coherentist might argue) my *knowledge* that there is such a movement is not epistemically direct. It is epistemically, but not inferentially, based on the coherence of my belief that there is a rustling with my other beliefs, presumably including many that represent knowledge themselves. It is thus knowledge *through*, but not by inference from, other knowledge – or at least not through justified beliefs. The knowledge is therefore epistemically indirect. Hence, it is at best misleading to call the *knowledge*, as opposed to the belief expressing it, direct at all.

This coherentist view grants, then, that the belief element *in* my knowledge is non-inferentially grounded in perception and is in that sense direct. But this is just a kind of psychological directness: there is no belief through which I hold the one in question in the way that I hold a conclusion belief on the basis of premise beliefs. But there are beliefs through which the belief constitutes knowledge: those with which it coheres even though it is not based on them. The basis relation between beliefs and the counterpart premise–conclusion relation between propositions are simply not the only producers of coherence.

One could insist that if a non-inferential, thus psychologically direct, belief constitutes knowledge, this *must* be direct knowledge. But the coherentist would reply that in that case there will be two kinds of direct knowledge: the kind the foundationalist posits, which derives from grounding in a basic experiential or rational source, say perception or reflection, and the kind the coherentist posits, which derives from coherence with other beliefs and not from being based on those sources. Why not classify the directness of knowledge in terms of what it *evidentially* depends on and the directness of belief in terms of what it *psychologically* depends on? This is surely a plausible response.

Is the holistic coherentist trying to have it both ways? Not necessarily. Holistic coherentism can grant that a variant of the regress argument holds for belief, since the only kind of belief chain that it is psychologically realistic to attribute to us is the kind terminating in direct (non-inferential) belief. But even on the assumption that knowledge is constituted by (certain kinds of) beliefs, it does not follow that direct belief which is knowledge is also direct *knowledge*.

Thus, the coherentist is granting *psychological foundationalism*, which says (in part) that if we have any beliefs at all, we have some direct ones, yet denying epistemological foundationalism, which says that, assuming there is any knowledge at all, there is knowledge which is epistemically (and normally also psychologically) direct. Holistic coherentism may grant experience and reason the status of psychological foundations of our entire structure of beliefs. But it gives them no place, independently of coherence, in generating justification or knowledge.[8]

The nature of coherence

As I have described holistic coherentism, it avoids some of the major problems for linear coherentism. But there remain serious difficulties for it. First, what is coherence? Second, what reason is there to think that coherence *alone* counts toward the justification of a belief, or toward its truth, as it must in some way if it is to give us the basis of a good account of knowledge?

It turns out to be very difficult to explain what coherence is. It is not mere mutual consistency, though inconsistency is the clearest case of *in*coherence. Two propositions having nothing to do with each other, say that 7 + 5 = 12 and that carrots are nourishing, are mutually consistent but do not exhibit coherence.

Coherence and explanation

Coherence is sometimes connected with *explanation*. Certainly, if the Smiths' sending the package explains why the card bears their names, then my belief of the first proposition coheres with my belief of the second (other things being equal). What explains something makes it understandable; and *making understandable* is a coherence-generating relation between propositions (as well as between other kinds of things).

Probability is also relevant to coherence. If the probability of the proposition that the Smiths sent the steaks is raised in the light of the proposition that I sent them cheese, this at least counts in favor of my belief of the first cohering with my belief of the second. But how are we to understand the notions of explanation and of probability? Let us consider these questions in turn.

Does one proposition (genuinely) explain another so long as, if the first is (or at least is assumed to be) true, then it is clear why the second is true? Apparently not; for if that were so, then the proposition that a benevolent genie delivered the box explains why it arrived. In any event, if that proposition did explain why the box arrived, would I be justified in believing it because my believing it coheres with my believing that I know not what other source the box might have come from? Surely not.

Even if we can say what notion of explanation is relevant to understanding coherence, it will remain very difficult to specify when an explanatory relation generates *enough* coherence to create justification. For one thing, consider cases in which a proposition, say that Jill hurt Jack's feelings, would, if true, very adequately explain something we believe, such as that Jack is upset. Believing Jill did this might cohere well with his being upset, but that would not, by itself, justify our believing it. There are too many possible competing explanations we might just as well accept.

Similar points hold for probability. Not just any proposition I believe which, if true, would raise the probability of my hypothesis that the gift is from the Smiths will strengthen my justification for believing that it is. Consider, for example, the proposition that the Smiths send such gifts to all their friends. Suppose I have no justification for believing this, say because I have accepted it only on the basis of testimony which I should see to be unreliable. Then, although the proposition, if true, raises the probability of my hypothesis (since I am among their friends) and (let us assume) coheres with what I already believe, I am not entitled to believe it, and my believing it will not add to my justification for believing that the Smiths sent the box.

It might be replied that this belief about the Smiths' habits does not cohere well with *other* things I believe, such as that people do not generally behave like that. But suppose I believed nothing about the Smiths' or other people's habits of gift-giving that conflicts with the Smiths' being so generous, and I happened, without grounds, to believe the Smiths to be both generous and rich. Then there might be a significant degree of coherence between my belief that the Smiths send gifts to all their friends and my other beliefs; yet my forming the belief that they give gifts to all their friends still would not strengthen my justification for my hypothesis that the steak is from them.

Coherence as an internal relation among cognitions

These examples suggest the second problem. So far as we do understand coherence, what reason is there to think that by itself it generates any justification or counts toward truth at all? Whatever coherence among beliefs is, it is an *internal* relation, in the sense that whether it holds among beliefs is a matter of how those beliefs (including their propositional content, which is intrinsic to them) are related *to one another*. It is not a matter of anything outside one's system of beliefs, such as one's perceptual experience. Now why could there not be a vast number of equally coherent systems of beliefs that are mutually incompatible, so that no two of them can be without at least some falsehood? If there can be, why should my having one of these coherent systems provide any reason to think *my* beliefs, rather than those of someone with one of the "opposing" systems, are justified or represent knowledge?

This is part of what might be called the *isolation problem*: the problem of explaining why coherent systems of beliefs are not readily isolated from truth, and thus do not contain knowledge, which implies truth. There is also a problem of explaining why there is not a similar isolation from justification, which seems in some way to point toward truth, roughly in the sense that what justifies a belief "indicates" its truth, and indicates it in proportion to the degree of justification. Why should coherence by itself imply that any of the cohering beliefs is justified or constitutes knowledge, when both justification and knowledge point toward truth as something external to the belief system? It is not as though coherentists could count on the implication's being guaranteed by God; and nothing else seems to assure us of it.

Consider a schizophrenic who thinks he is Napoleon. If he has a completely consistent story with enough interlocking details, his belief system may be superbly coherent. He may even be able to explain quite coherently why there are coherent belief systems that conflict with his, such as those of his psychiatrists. If coherence alone generates justification, however, we must say that each system is equally well justified – assuming their belief systems are exactly as coherent as his. We need not attribute knowledge to any of the systems, since any of them might contain falsehood.

But is it plausible to say that a system of beliefs is highly justified even when there is no limit to the number of radically different yet equally justified belief systems – even on the part of other people with experience of many of the same things the beliefs are about – that are incompatible with it in this thoroughgoing way? The question is especially striking when we realize that two equally coherent systems, even on the part of the same person at different times, might differ not just on one point but on *every* point: each belief in one might be opposed by an incompatible belief in the other.

To appreciate the significance of the possibility of multiple coherent systems of belief that are mutually incompatible, recall the plausible assumption that a well-justified belief may reasonably be considered *true*. If, however, the degree of justification of a belief is entirely a matter of its support by considerations of coherence, no degree of justification by itself can carry any greater presumption of truth than is created by the same degree of support from coherence on the part of a belief of the contradictory proposition. Thus, if "Napoleon" (unlike his historical namesake) has a sufficiently coherent set of beliefs yielding justification of his belief that he fought in and won the Battle of Waterloo, this belief may be as well-justified as his psychiatrists' belief that he was not even born at the time.

If this coherentist picture of justification is correct, is there any reason to think that a belief supported solely by considerations of coherence is true or even justified? And if Napoleon's and the psychiatrists' belief systems are equally coherent, how can we justify our apparently quite reasonable tendency to regard their belief systems as more likely to represent truths, and on that count more likely to contain knowledge, than his?

Granted, the psychiatrists' belief that he was born long after the battle coheres with our beliefs. But why should our own beliefs be privileged over equally coherent conflicting sets? And why should agreement even with nearly everyone's beliefs, say about Napoleon's being dead, be a factor, unless we assume that some element other than coherence, such as perception or memory, confers justification without drawing on coherence? If coherence is the only source of justification, it is not clear how perception or memory or introspection contributes to justification. Moreover, even what seems the highest degree of justification, such as we have for simple introspective beliefs and beliefs of self-evident truths, provides us no presumption of truth or knowledge.

Coherence, reason, and experience

This brings us to a third major problem for coherentism: how can it explain the role of experience and reason as sources of justification and knowledge? Certainly experience and reason *seem* to be basic sources of justification and knowledge. Coherentists themselves commonly *use* beliefs from these sources to illustrate coherent bodies of beliefs that are good candidates for knowledge. How can holistic coherentism explain the role of these sources in relation to justification and knowledge?

Why is it, for instance, that when I have a vivid experience of the kind characteristic of seeing a green field, I am apparently justified (though prima facie rather than indefeasibly justified), simply by that experience, in believing that there is a green field before me? And why do I seem so very strongly justified, simply on the basis of my rational grasp of the proposition that if some dogs are pets then some pets are dogs, in believing this?

One thing a coherentist might say here is that in fact many of our beliefs are *causally* and non-inferentially based on perception or on the use of reason; and given these similarities of origin, it is to be expected that they often cohere with one another. Hence, although we do not, and do not need to, infer propositions like those just cited from any others that might provide justifying evidence for them, they *do* cohere with many other things we believe, and this coherence is what justifies them.

Coherence and the a priori

This response by way of associating the coherence of beliefs with their causal basis is more plausible for perceptual beliefs than for beliefs of simple self-evident a priori truths, at least if coherence is construed as more than consistency and as related to explanation, probability, and justification. For notice that, unlike the proposition that there is a green field before me, the proposition that if some dogs are pets, then some pets are dogs apparently need not explain, render probable, or justify anything else I believe. Nor is it obvious that anything else I believe need explain, render probable, or justify my believing this proposition. Where is the need for coherence as a requirement for my justification? I may *have* other beliefs that cohere with this one, but my justification for it does not seem to *derive* from such coherence. Yet my belief of this proposition is justified to about as high a degree as is any belief I have.

By contrast, the proposition that there is a green field before me perhaps does cohere, in a way that might serve coherentism, with other things I believe: that there is grass there, that I am on my front porch, and so on; and there appear to be some explanatory and probability relations among these propositions. For instance, the proposition that there is a green field before me adds to the probability that I am on my porch; and that I am on that porch partly explains why I see a green field.

A coherentist might respond to the difference just indicated by qualifying the coherence view, applying it only to beliefs of empirical, rather than a priori, propositions.[9] This move could be defended on the assumption that propositions known a priori are necessarily true and hence are not appropriately said to be made *probable* by other propositions, or to be explained by them in the same way empirical propositions are explained. In support of this it might be argued that although we can explain the *basis* of a necessary

truth and thereby show *that* it holds, still, since it cannot fail to hold, there is no explaining *why* it, *as opposed to something else*, holds.

This is plausible but inconclusive reasoning. We may just as reasonably say that we can sometimes explain why a necessary truth holds and in doing so explain why a contrasting proposition is false. Imagine that someone mistakenly takes a certain false proposition to be a theorem of logic and cannot see why a closely similar, true proposition is a theorem. If we now prove the correct one step by step, with accompanying examples, we might thereby explain why this theorem, as opposed to the other proposition, is true.

So far as explanation is central to coherence, then, coherentism apparently owes us an account of knowledge of at least some necessary truths. But suppose that it can account for knowledge of *some* necessary truths. There remain others, such as simple, luminously self-evident ones, for which it cannot offer anything plausibly said to explain why they hold, or any other way of accounting for knowledge of them as grounded in coherence.

Consider how one might explain why, if it is true that Jane Austen wrote *Persuasion*, then it is not false that she did. If someone did not see this, it would probably not help to point out that no proposition is both true and false. For if one needs to have the truth of such a clear and simple instance of this general truth explained, one presumably cannot understand the general truth either. But suppose this is not so, and that one's grasp of the general truth is somehow the basis of one's seeing the particular truth that instantiates it. Then the same point would apply to the general truth: there would apparently be nothing plausibly said to explain to one why *it* is true.

Coherence and the mutually explanatory

It might now be objected that the general truth that no proposition is both true and false, and the instances of it, are *mutually explanatory*: its truth explains why they hold, and their truth explains why it holds; and this is the chief basis of their coherence with one another. But is it really possible for one proposition to explain another *and* the other to explain it? If what explains why the grass is wet is that there is dew on it, then the same proposition – that there is dew on it – is not explained by the proposition that the grass is wet (instead, condensation explains why it is wet).

Reflection on other purported examples of mutual explanation also suggests that two propositions cannot explain each other. It might seem that a man could say something because his wife did, and that she could say it because he did. But notice how this has to go to make good sense. One of them would have to say it first to cause the other to. But then we would have a case in which something like this occurs: her saying it explains why he says it, later (this could be so even if her saying it is explained by her believing he *thinks* it). His saying it earlier than she does might still explain her saying it. But then the fact that he says it at a given time does not both explain and get explained by her saying it at some particular time.

When we carefully specify what explains something, we seem to find that the latter, carefully specified, does not explain the former. In the case where she says something because he did, earlier, and he says it because she did, earlier than he did, we would have a kind of *reciprocal* explanation, wherein a *kind* of thing, here spousal affirmation, explains and is explained by another thing of the same kind. But this is not a *mutual* explanation, wherein the very *same* thing explains and is explained by a second thing. The first may look like the second, but it is quite different.[10]

Perhaps mutual explanation of the kind the coherentist apparently needs – as opposed to reciprocal explanation and other sorts involving two-way relations – is somehow possible. But until a good argument for it is given, we should conclude that even if an explanatory relation between propositions is sufficient for a belief of one of the propositions to cohere with a belief of the other, coherentism does not in general provide a good account of knowledge of self-evident truths.

If coherentism applies only to empirical beliefs, however, and not to beliefs of a priori propositions, then it is not a *general* theory of justification or knowledge and leaves us in need of a non-coherentist account of a priori justification (and knowledge). In any case, it would be premature to conclude that coherentism does account for empirical justification. Let us return to the perceptual case.

Epistemological versus conceptual coherentism

It might seem that we could decisively refute the coherence theory of justification by noting that one might have only a single belief, say that there is a green field before one, and that this lone belief might still be justified. For there would be a justified belief that coheres with no other beliefs one has. But could one have just a single belief? Could one, for instance, believe that there is a green field before one, yet not believe, say, that it has any vegetation? It is not clear that one could; and foundationalism does not assume this possibility, though the theory may easily be wrongly criticized for implying it.

Foundationalism is in fact consistent with *one* kind of coherentism, namely, a coherence theory of the acquisition and function of concepts – for short, the *coherence theory of concepts*. According to this theory, concepts are what they are partly in relation to one another, and a person acquires concepts, say of (physical) objects and shapes, and of music and sounds, only in relation to one another and must acquire an entire set of related concepts in order to acquire any concept. The concept of an object in some way includes that of shape (if only the notion of something bounded), as that of music includes the concept of sound. This may be why any object must have some shape or other, and why anything that makes music produces some sound. One cannot (fully) acquire object concepts without acquiring some shape concepts, or (fully) acquire the concept of music without acquiring that of sound.

If the coherence theory of concepts is sound, foundationalists must explain how it squares with their epistemology. The central point they may appeal to is a distinction between *grounding conditions* for belief and *possession conditions* for it. What grounds a belief in such a way as to justify it or render it an item of knowledge is largely independent of *what* other beliefs one must have, and what concepts one must have, to be able to *hold* the first belief. Perhaps I cannot believe that music is playing if I do not have a concept of sound; I may even have to believe sounds with a certain structure to be occurring. And perhaps I could not have acquired these and other relevant concepts one at a time. Indeed, it may be (as suggested in Chapter 5) that at least normally we cannot acquire concepts without acquiring some knowledge or justified belief. Still, what it is that justifies a belief can be a matter of how the belief is grounded; it need not be a matter of the coherence conditions required for *having* the belief.

If, however, coherence relations are essential for holding a belief at all, they are on that ground necessary for, and – in ways that will soon be apparent – important in understanding, the belief's being justified. The point here is simply that we cannot treat conditions for having a belief at all as doing the more specific job of grounding its justification. By and large beliefs can be possessed without being justified, and there is commonly a good distance between meeting the conditions for simply having beliefs and meeting the standards for justification in holding them.

Coherence, incoherence, and defeasibility

We must directly ask, then, whether my justification for believing that there is a green field out there *derives* from the coherence of the belief with others. Let us first grant an important point by focusing on a line of reasoning that seems to lead many philosophers to think it does derive from coherence. Suppose this visual belief turns out to be *in*coherent with a second belief, such as that one is standing where one seems to see the field around one yet *feels* no grass on the smooth ground beneath one's feet and can walk right across the area without feeling any. Then the first belief may *cease* to be justified. Incoherence, then, defeats my justification.

This defeating role of incoherence is important, but it shows only that our justification is *defeasible* – liable to being outweighed (overridden) or undermined – should sufficiently serious incoherence arise. It does not show that justification is produced by coherence in the first place, any more than a wooden cabin's being destroyed by fire shows that it was produced by the absence of fire. In the case in which I feel no grass beneath my feet, the justification of my visually grounded belief is outweighed: my better justified beliefs, including the conviction that a field must have a certain texture, make it more reasonable for me to believe that there is *not* a field here.

A major lesson that emerges here is that we cannot tell what the basis of something is just from the range of things that outweigh it, much less conclude that this basis *is* the absence of the things that destroy it. Incoherence is absent where there are mutually irrelevant beliefs as well as where there are mutually coherent ones. Mutual irrelevance between two sets of beliefs certainly does not make one of them a justificational or epistemic basis for the other.

Two important questions arise here. First, could incoherence outweigh justification of a belief in the first place if we were not *independently* justified in believing something to the effect that a proposition incoherent with certain other ones is, or probably is, false? Second, are the other relevant propositions not precisely the kind for which, directly or inferentially, we have some degree of justification through the basic experiential and rational sources? Foundationalists are likely to answer the first negatively and the second affirmatively.

There is also a different kind of defeat of justification: our justification can be simply *undermined*. We can cease to be justified in believing a proposition, though we do not become justified in believing it false, as one does where counter-evidence demands a belief contrary to the initial one. Suppose I cease to see a bird on a branch when, without obscuring my line of sight to the bird, I move six feet to my left. This could justify my believing that I might be hallucinating. This belief is incoherent with, and thereby undermines the justification of, my visual belief that the bird is there, though it does not by itself justify my believing that there is *no* bird there.

Again, I am apparently justified, independently of coherence, in believing that my seeing the bird there is incoherent with my merely hallucinating it there. It seems that coherence has the role it does in justification largely because *some* beliefs are justified independently of it.

Positive and negative epistemic dependence

Examples like these show that it is essential to distinguish *negative epistemic dependence* – which is a form of defeasibility – from *positive epistemic dependence* – the kind beliefs bear to the sources from which they *derive* any justification they have or, if they represent knowledge, derive their status as knowledge. The defeasibility of a belief's justification by incoherence does not imply, as coherentists have commonly thought, that this justification positively depends on coherence. If my well is my source of water, I (positively) depend on it. The possibility that people could poison it does not make their nonmalevolence part of my source of water, or imply a (positive) dependence on them, such as I have on the rainfall. Moreover, it is the rainfall that *explains* both my having the water and its level.

So it is with perceptual experience as a source of justification. Foundationalists need not claim that justification does not depend negatively on anything else, for as we have seen they need not claim that justification must be indefeasible. Its vulnerability to defeat can be construed as a kind of dependence. A belief's justification is, then, not *completely* independent of the

justification of other beliefs, actual or hypothetical. But negative dependence does not imply positive dependence. Justification can be defeasible by incoherence, and thus outweighed or undermined should incoherence arise, without owing its existence to coherence in the first place.

Coherence and second-order justification

There is something further that may be considered supportive of coherentism, and in assessing it we can learn more about both coherentism and justification. If one sets out to *show* that one's belief is justified, one has to cite propositions that cohere with the one in question, say the proposition that there is a green field before me. In some cases, these are not even propositions one already believes. Often, in defending the original belief, one forms new beliefs, such as the belief one acquires, in moving one's head, that one can vividly see the changes in perspective that go with seeing a bat in flight.

The process versus the property of justification

More importantly, these new, back-up beliefs are especially appropriate to the *process of justifying* one's belief; and the result of that process is (a kind of) *showing* that the original belief is justified, together (in typical cases) with one's forming a certain *second-order belief* – so called because it is a belief about a belief (such as a perceptual one) which is not itself about any other belief. In this case the second-order belief is to the effect that the first-order belief is justified. Thus, coherence is important in showing that a belief is justified and is in *that* sense an element in a typical kind of process of justification.

The moment we reflect on this point, however, we may wonder why the beliefs appropriate to *showing* that a belief is justified are required for its *being* justified in the first place. There is no good reason to think they need be. Indeed, why should one's simply having a justified belief imply even that one is (situationally) justified in holding beliefs appropriate to showing that it is justified? It would seem that just as one can be virtuous even if one does not know how to defend one's good character against attack or even show that one has good character at all, one can have a justified belief even if, in response to someone who doubts that one has it, one could not show that one does.

Justifying a second-order belief is a sophisticated process. The process is particularly sophisticated if the second-order belief concerns a special property like the justification of the original belief. Simply being justified in a belief about the color of an object is a much simpler matter.

Confusion is easy here because of the way we often speak of justification. Consider the question of how a simple perceptual belief "is justified." The very phrase is ambiguous. For all it tells us, the question could be 'By what process, say of reasoning, has the belief been (or might it be) justified?' or, on the other hand, 'In virtue of what is the belief justified (possessed of the property of justifiedness)?' These are two very different questions. But much of our talk about justification makes it easy to run them together. A belief said to be "justified" could be one that *has* justification or one that *has been* justified; and asking for someone's justification could be either a request for justifying factors or an invitation to recount the process by which the person has in fact justified the belief.

Does coherentism have any plausible argument, not grounded in the mistakes just pointed out, for the (positive) dependence of perceptual justification on coherence? I do not see that it does, though given how hard it is to discern precisely what coherence is, we cannot be confident that no direct argument is forthcoming.

Granted, one could point to the oddity of saying things like, 'I am justified in believing that there is a green field before me, but I cannot justify the belief'. Coherentists might think this is odd because they tend to hold that if one has a justified belief, one can surely give a justification for it by appeal to other beliefs one holds that cohere with it. But look closely. Granted that, commonly, in asserting something I suggest that I *can* justify it in some way or other (particularly if the belief I express is not grounded in a basic source), yet here I deny that very suggestion. Still, it could well be that it is my *asserting* that my belief is justified, rather than its being so, that gives the appearance that I must be able to give a justification to the belief.

In asserting that I am justified, after all, I have not, or not merely, *expressed* a first-order justified belief, something a normal child of three can do; I have *ascribed* first-order justification to my belief. That requires some sophistication. More important, even foundationalists who hold that we are typically directly justified in, say, perceptual beliefs may deny that normally we are directly justified in these sophisticated ascriptions of justification. To hold that there are non-inferentially justified beliefs does not in the least commit one to holding that ascriptions of justification itself are thus justified.

Beliefs, dispositions to believe, and grounds of belief

To be sure, when I say that there is a green field before me, I can give a justification: for instance, that I see it. But first, *giving* a justification is not equivalent to claiming that one has it. The first cites a justifier and need not employ the concept of justification; the latter employs that sophisticated concept and need not cite a justifier. Second, note that *before* the question of justification arises I need not even *believe* that I see the field. That question leads me to focus on my circumstances, in which I first had a belief solely about the *field*, not about my own perceptual relation to it.

To be sure, when I said there is a green field before me, I did have a *disposition*, based on my visual experience, to form the belief that I *see* the field, and this is largely why, in the course of justifying that belief, I then form the further belief that I *do* see it. But a *disposition to believe* something does not

imply one's already having a *dispositional belief* of it: here I tend to form the belief that I see the field if, as I view it, the question whether I see it arises; yet I need not have subliminally believed this already.

Thus, the justification I offer for my belief that there is a green field before me is not by appeal to coherence with other beliefs I already hold – such as that I saw the field and heard the swishing grass beneath my feet – but by reference to a basic source, sensory experience. It is thus precisely the kind of justification that foundationalists are likely to consider appropriate for a non-inferential belief. Indeed, one consideration favoring foundationalism about both justification and knowledge, at least as an account of our justificational practices in everyday life (including much scientific practice) is that typically we cease offering justification or defending a knowledge claim precisely when we reach one or more of the basic sources.[11]

Suppose, however, that I would be dumbfounded if asked, in clear daylight, what justifies me in believing there is a green field before me. Would it follow that I am not justified? No, for I might be simply unable to marshal my quite ample justificatory resources. Coherentism offers no good argument to show that being justified requires being able to show that one is, any more than having good character entails being able to show that one has it.

Justification, knowledge, and artificially created coherence

There is one further point here. If coherentism regards justification as deriving from coherence alone, then it accords *no* justificatory weight to experiential or rational grounding except insofar as they contribute to coherence. Our examples cast much doubt on this view.

Consider a related implication of coherentism. If I seek the best justified body of beliefs possible – which is surely a rational goal – then I am free to consider adopting, or to manipulate my brain to cause myself to form, an entirely new system of beliefs. Would its coherence alone guarantee that it contains justified beliefs? It might contain none of the experiential and a priori beliefs I now have; and for all coherence requires it may contain no beliefs based on experience or reason at all.

A superbly coherent system of beliefs I might acquire could even run counter to my experience. Even if I *see* a square field of green grass before me, I might coherently believe that there is an oval field of brown shrubbery there, since my other beliefs might support this. I could, for instance, coherently believe that when I seem to see green grass I am having a hallucination caused by brown shrubbery. There is no limit to the number of beliefs for which one might be able thus to rationalize away the states and events that it is natural to call the *evidence of the senses*.

We are apparently incapable of changing our belief systems in this way. But suppose that we could do so by properly setting a neurological machine to instill an optimally coherent set of beliefs and remove the rest. Would

that be rational from the point of view of maximizing the justification of one's beliefs? I doubt this, particularly if, in seeking justification, we aim, as we normally do, at discovering or retaining *truths*.

A coherentist might reply that if we are talking not only about justification but also about *knowledge*, then we must give some special role to beliefs (and perhaps dispositions to believe) grounded in experience and reason, for if we ignore these sources we cannot expect our justified beliefs to be true, hence cannot expect them to constitute knowledge.[12]

Now, however, we face what seems an artificial separation between what justifies a belief and what is plausibly taken to count towards its truth. If, because it implies truth, knowledge must in some way reflect experience or reason, should not justification, which seems in some way to count toward truth, also reflect them? Is it reasonable to suppose that what justifies a belief may in no way count towards its truth?

It is not reasonable to separate justification and knowledge in this way (even though in some ways they are very different); nor have coherentists generally thought that it is (though some have held a justification-based coherence theory of truth of a kind to be discussed in Chapter 8). Often, what motivates asking for a justification of a belief is doubt that it is true; and if so, then the view that what justifies a belief has no tendency whatever to count toward its truth seems plainly mistaken. Moreover, if we can know a priori, as I believe may be possible (and will explore in Chapter 10), that perceptual and rational grounding of beliefs count, in some way, toward their truth, why may we not know equally well that they count toward justifying beliefs?

Moderate foundationalism

There is far more to say about both foundationalism and coherentism. But if what has emerged here is on the right track, then the problems confronting coherentism are more serious than those confronting foundationalism. The most serious problems for foundationalism are widely taken to be the difficulties of specifying source conditions for justification and knowledge and, second, of accounting, on the basis of those sources and plausible transmission principles, for all that we seem to know. The first of these problems is addressed in Part One, which describes the basic sources and illustrates how they generate direct – though not indefeasible – knowledge, and direct (though again not generally indefeasible) justification. The second problem is treated in Chapter 6, which indicates many ways in which, even without actual inferences, knowledge and justification can be transmitted from beliefs which are justified, or represent knowledge, by virtue of being grounded in the basic sources, to other beliefs. Both problems are difficult, and they have not been completely solved here. But enough has been said to make clear along what lines they can be dealt with in a foundationalist framework.

The role of coherence in moderate foundationalism

Still another problem for foundationalism is the difficulty of accounting for the place of coherence in justification. But this is not a crippling difficulty for the kind of foundationalism I have described, which need not restrict the role of coherence any more than is required by the regress argument. Indeed, while (pure) coherentism grants nothing to foundationalism beyond perhaps its underlying psychological picture of how our belief systems are structured, foundationalism can account for some of the insights of coherentism, for instance the point that we need a coherence theory of the acquisition and function of concepts.

More positively, foundationalism can acknowledge a significant role for coherence in relation to justification and can thereby answer one traditional coherentist objection. I have in mind a kind of *moderate foundationalism*: a foundationalist view of knowledge or justification which (1) takes the justification of foundational beliefs to be at least typically defeasible; (2) is not *deductivist*, that is, does not demand that principles governing the inferential transmission of knowledge or justification be deductive (i.e., require entailment as opposed to probability as a condition for transmission); and (3) allows a significant role for coherence by requiring, not that inferentially justified beliefs derive *all* their justification from foundational ones, but only that they derive enough of it from the latter so that they would remain justified if any other justification they have were eliminated.[13] Some versions are more moderate than others, but the most plausible ones give coherence at least two roles.

The first role moderate foundationalism may give to coherence, or strictly speaking to incoherence, is negative: incoherence may defeat justification or knowledge, even of a directly justified (foundational) belief, as where my justification for believing I may be hallucinating prevents me from knowing, or even remaining justified in believing, that the green field is before me. (If this is not ultimately a role for coherence itself, it is a role crucial for explaining points stressed by coherentism.)

Second, moderate foundationalism can employ a principle commonly emphasized by coherentists, though foundationalists need not grant that the justification or truth of the principle is based on coherence and will tend to treat it as a transmission principle accounting for generation of inferential justification or as a combinatorial principle applying to the simultaneous testimony of sources of non-inferential justification. I refer to an *independence principle*: that the larger the number of independent mutually consistent factors one believes to support (or to constitute evidence for) the truth of a proposition, the better one's justification for believing it (other things being equal). This principle can explain, for instance, why my justification for believing that the box of steaks is from the Smiths increases as I acquire new beliefs, each of which I believe independently supports that conclusion.[14] In part, the idea is that evidential relations generate coherence; hence by giving the former a justificatory role, foundationalism can account for a good many

of the cases in which coherence appears to yield justification. That appearance may be due not to coherence itself, but to its basis in the kinds of elements foundationalism takes to ground justification.

Similar principles consistent with foundationalism can accommodate other cases in which coherence enhances justification, say those in which a proposition's explaining, and thereby cohering with, something one justifiably believes tends to confer some degree of justification on that proposition.

Moderate foundationalism and the charge of dogmatism

Moderate foundationalism contrasts with *strong foundationalism*, which, in one form, is deductivist, takes foundational beliefs as indefeasibly justified, and allows coherence only a minimal role. To meet these conditions, strong foundationalists may reduce the basic sources of justification to reason and consciousness. The easiest way to do this is to take the skeptical view (considered in Chapter 10) that our only justified beliefs are either a priori or introspective.

Moreover, since strong foundationalists are committed to the indefeasibility of foundational justification, they would not grant that incoherence can defeat the justification of foundational beliefs. They would also refuse to concede to coherentism, and hence to any independence principle they recognize, any more than a minimal positive role, say by insisting that if a belief is supported by two or more independent cohering sources, its justification is increased at most additively, that is, at most by bringing together the justification transmitted separately from each relevant basic source.[15]

By contrast, what moderate foundationalism denies regarding coherence is only that it is a basic source of justification. Coherence by itself is not sufficient for justification. Thus, the independence principle does not apply to sources that have *no* justification. At most, it allows coherence to raise the level of justification originally drawn from other sources to a level *higher* than it would have if those sources were not mutually coherent.

Similarly, if inference is a basic source of coherence (as some coherentists seem to have believed), it is not a basic source of justification. It may lead to justification, as where one strengthens one's justification for believing someone's testimony by inferring the same point from someone else's. But inference *alone* does not generate justification: I might infer any number of propositions from several I already believe merely through wishful thinking; yet even if I thus arrive at a highly coherent set of beliefs, I have not thereby increased my justification for believing any of them. My premises, based in the way they are on desire, are ill-grounded.

At this point it might occur to one that the main problems faced by coherentism could be solved by taking *coherence with experience* to be required by coherentism as a condition for the coherence of a body of beliefs of the kind we normally have. This is, to be sure, *not* how coherence is characteristically understood by coherentists; they typically take it to be a relation

among beliefs or their propositional contents or other items that may be said to be true or false, or some combination of these.[16] Might it be, however, that leading coherentists misrepresent the resources of their own theory? Could they claim, for instance, that if my visual experience contains an appearance of a printed page, then my believing there is one before me coheres with my experience and is thus justified?

If we think this, we must ask how a coherentist view that gives a crucial epistemological role to coherence of beliefs with experience differs from a moderate foundationalism. One would, after all, be insisting that in order to contain justified beliefs about the world, a person's belief system would in some sense depend on experience. This gives an essential role to foundations of justification (or knowledge) – grounds of belief that are not true or false and do not themselves admit of justification. It is true that the view would also require coherence among beliefs as an essential element; but a moderate foundationalist could agree that coherence is necessary within a body of justified beliefs such as normal people have, yet insist that this coherence is not a basic source of justification rather than, chiefly, a *product* of the elements, such as grounding in experiential and rational sources, that *are* basic.

If coherentists cannot show that coherence among beliefs is a basic source of justification – as it is far from clear they can – then requiring coherence with experience to make their theory plausible yields a view that is apparently at least compatible with a moderate foundationalism and may well be a version of that view. This may be a welcome conclusion for epistemologists uncommitted on the foundationalism–coherentism issue, but it would be unwelcome to philosophers in the coherentist tradition.[17]

Suppose, however, that moderate foundationalism is correct. We must not suppose that this theory leads easily to an adequate, detailed picture of a typical body of knowledge or justified belief. Moderate foundationalism as so far described – mainly structurally – tells us only what sort of structure a body of knowledge or of justified belief has. It says that if one has any knowledge or justified belief, then one has some direct knowledge or directly justified belief, and any other knowledge or justified belief one has is traceable to those foundations. A belief direct and foundational at one time may be indirect and non-foundational at another; it may gain or lose justification; it may have any kind of content; and some foundational beliefs may be false or unjustified or both.

By leaving this much open, however, moderate foundationalism avoids a narrow account of what is needed for knowledge and justification and allows many routes to their acquisition. For similar reasons, it avoids *dogmatism*, in the sense of an attitude of self-assured certainty, especially concerning claims that are neither self-evident nor obvious. In addition to avoiding this *attitudinal dogmatism*, it rejects, for the same sorts of reasons, at least one version of *epistemological dogmatism* – the one ascribing to us indefeasible justification, epistemic certainty, or the like, where these attributions are unwarranted by

our evidence. For moderate foundationalism allows alternative kinds of foundational beliefs for different people and under different circumstances; and, by acknowledging the imperfect reliability of the experiential sources and of many inferences from the beliefs they generate, it also explains why it is so difficult to know that one has knowledge or justified belief, and hence important to be open to the possibility of mistakes.

Moderate foundationalism even allows that a person may not always be able to see the truth of a self-evident proposition; one might, for instance, lack the conceptual resources for adequately understanding it. This should induce humility about how extensive our knowledge is even regarding what is in principle readily known. Ignorance can occur where one would least expect it. The position also treats reason as a fallible source of belief: we can easily take a false proposition to be true on the basis of a specious sense of its being a priori. This should induce humility about how confident we are entitled to be. Error can occur where it might seem impossible. Foundationalism is committed to unmoved movers; it is not committed to unmovable movers. It leaves open, moreover, just what knowledge is, and even whether there actually is any. These questions must still be faced.

Notes

1 Clearly, there could be devices or strategies by which one can manipulate one's beliefs; what I deny is that one can control belief "at will" (simply by willing it) the way one can normally raise an arm at will. The point is not that the will has *no* power over belief. For wide-ranging critical discussion of doxastic voluntarism see William P. Alston, 'The Deontological Conception of Epistemic Justification,' *Philosophical Perspectives* 2 (1983), 257–99; my 'Doxastic Voluntarism and the Ethics of Belief,' *Facta Philosophica* I, 1 (1999), 87–109 and reprinted in Matthias Steup (ed.), *Knowledge, Truth, and Duty* (Oxford: Oxford University Press, 2001); and, for critical discussion of Alston's position, Steup's contribution to that collection.

2 There is dispute about whether people can have infinite sets of beliefs. I have offered some reasons for doubting this (and cited some of the relevant literature) in 'Dispositional Beliefs and Dispositions to Believe,' *Nous* 28 (1994), 419–34.

3 Granted, one could look at the formulation, say by tracing it along a mile-long print-out, and believe *that* it expresses a truth; but the point is that one could not grasp, and so could not believe, the *truth* that it expresses. Of course, if we are talking about infinity, the relevant formulations would approach an infinite number of miles in length. For an extensive discussion of the prospects for epistemological infinitism, see Peter D. Klein, 'Human Knowledge and the Infinite Regress of Reasons,' *Philosophical Perspectives* 13 (1999).

4 In a well-argued and highly instructive paper on the question whether inferential knowledge must be grounded in knowledge of an appropriate premise, Claudio de Almeida (taking off from an example of Peter Klein's) argues that there are important exceptions. Here is a representative case. The trustworthy department secretary told me, last Thursday (and knew), that I have an appointment this Friday. Now, asked whether I am free at the relevant time Friday, I say 'No, the secretary told me on Wednesday that I have an appointment then'. Plainly, I can know I have the appointment, though the belief I express now as a basis is false, since I have the wrong day. There is much to say, but three points must serve. (1), on my account of testimony-based knowledge (ch. 5), I would know that I have the appointment *non-inferentially*. The false belief that she told

me on Wednesday is *offered* as a ground of my belief that I have the appointment on Friday, but the ground on which I *know* that is her *attesting* to it. (2) Suppose my *statement* 'She told me on Wednesday that I have an appointment then' does in a way express my ground. The ground may be *that she told me* that I have an appointment then, with 'on Wednesday' functioning like a parenthesis, such as 'and by the way it was on Wednesday'. Then I *do* know my ground. (3) What if I have forgotten her testimony, am told that she gave it on Wednesday, and infer, apparently from this falsehood, that I have the appointment. Now we need a theory. One move is to make a distinction useful in any case, between the ground's being, as in (2), *that she told me* that I have an appointment then, and its being something like: *It was on Wednesday* that she told me that I have an appointment then (where the time is important in my thinking). In the latter case I would not know; in the former I presumably would. Still, do I, in the former case, *believe* the true proposition that she told me that I have an appointment then, or am I only disposed to believe it (actually believing only the larger proposition "containing" it)? This is just one of the important questions de Almeida forces us to explore. See 'Knowledge and Benign Falsehoods,' forthcoming.

5 An item of knowledge can occur in more than one epistemic chain, as where you have two entirely independent sets of premises showing the same conclusion. The regress argument requires one chain, but it allows more than one.

6 We may also draw the more general conclusion that if there is any knowledge, there is some direct knowledge. This more general conclusion follows only on the assumption that if there is any knowledge, then there is at least one knower who has it. This is self-evident for the main sense of 'knowledge'; but if we think of certain books as containing knowledge and then imagine the possibility that all knowers cease to exist while the books live on, it may then seem that there would be (residual) knowledge without there being any knowers (though even here there would *have been* knowers). Such *unpossessed knowledge* is discussed in some detail in Chapter 9.

7 See Aristotle's *Posterior Analytics*, Books I and II. His argument is importantly different in at least one respect: he spoke of the foundational items as "indemonstrable," which implies that there *cannot* be any deeper foundations. The regress argument as stated here implies only that one's foundational knowledge is of something that (at the time) one has not demonstrated. This leaves open that one might later demonstrate it by appeal to something "deeper."

8 The possibility of combining psychological foundationalism with epistemological coherentism seems quite open to Wilfrid Sellars, a leading coherentist. See, for example, his 'The Structure of Knowledge,' in Hector-Neri Castañeda (ed.), *Action, Knowledge, and Reality: Essays in Honor of Wilfrid Sellars* (Indianapolis: Bobbs-Merrill, 1975).

9 This is the position taken by Laurence BonJour in *The Structure of Empirical Knowledge* (Cambridge, Mass.: Harvard University Press, 1985). It should perhaps be noted that he has rejected the coherentist epistemology of this book in, e.g., 'The Dialectic of Foundationalism and Coherentism,' in John Greco and Ernest Sosa (eds), *The Blackwell Guide to Epistemology* (Oxford: Blackwell, 1999).

10 Does the fact that the topmost eastern section of the wheel is in place not explain why the topmost western section, which is contiguous with it, is in place, and isn't the converse also true? Only, I suspect, if this comes to saying that given these facts we can infer that each *is* in place. Why each is in place is explained by the same thing: the overall pattern of forces including the support provided by the ground. Each is in place because the gravitational force pulling it backward and downward is matched by a gravitational force pulling it forward and holding it up: both phenomena are indeed explained by the "same thing" – the qualitatively identical forces – but not by the same thing in the sense of the other, qualitatively identical phenomenon. Explanation by two phenomena that are "exactly alike" exhibits a kind of mutuality, but it is not the same as explanation of each of two exactly similar phenomena in terms of the other.

11 On the topic of practices of justification, Ludwig Wittgenstein's *On Certainty* (Oxford, 1969) is a valuable source. He is often cited as stressing that there comes a point at which one says "My spade is turned" (a foundationalist metaphor).

12 This line of thought is suggested by what Laurence BonJour, in *The Structure of Empirical Knowledge*, calls "the observation requirement." For extensive discussion of the theory he puts forth there and of coherentism in general, especially that of Keith Lehrer, see John W. Bender (ed.), *The Current State of the Coherence Theory* (Dordrecht and Boston: Kluwer, 1989).

13 A slightly different formulation may be required if, for the sorts of reasons to be given in Chapter 8, knowledge does not entail justification; but the formulation given will serve here. Here and elsewhere the reference to foundational beliefs is to those that are justified; I also omit an other-things-equal clause appropriate after the 'if' in clause (3). For a highly detailed statement of a moderate foundationalism, see Paul K. Moser, *Knowledge and Evidence* (Cambridge and New York: Cambridge University Press, 1989).

14 The independence principle cited here is not the only one that seems sound. For instance, it is plausible to hold that one's justification for a proposition also rises (other things being equal) the larger the number of factors one is appropriately aware of that *do* support it, whether or not one believes them to do so. I should add that (as I suggest below) the independence principle should probably specify factors one *justifiedly* takes to support the belief in question. Certainly *more* justification is conferred (other things equal) by factors justifiedly taken to support the belief than by those unjustifiedly taken to do so.

15 It is a strong foundationalism, especially the kind found in Descartes' *Meditations*, that is influentially criticized by Richard Rorty in *Philosophy and the Mirror of Nature* (Princeton, NJ: Princeton University Press, 1979). Many of Rorty's criticisms do not hold for the moderate foundationalism developed in this chapter. His doubts about the very idea that the mind is a "mirror of nature," however, may cut against at least the majority of plausible epistemological theories, depending on how much in built into the metaphor of a mirror. This book as a whole can be seen as a case for *some* kind of realist epistemology, and some aspects of Rorty's challenge are treated at least implicitly in Chapter 10 and in parts of other chapters, such as the sections on phenomenalism and truth.

16 Keith Lehrer provided an influential statement of this view in *Knowledge* (Oxford: The Clarendon Press, 1974): having said that "complete justification is a matter of coherence *within* a system of beliefs" (p. 17, emphasis added), he added, "There is no exit from the circle of ones [*sic*] own beliefs from which one can sally forth to find some exquisite tool to measure the merits of what lies within the circle of subjectivity" (pp. 17–18). Such sensory states as an impression of green grass are among the excluded tools. Further indications of why a coherentist view disallows appeal to experiential and other non-truth-valued states as justificatory are given by Wilfrid Sellars, 'The Structure of Knowledge,' in Hector-Neri Castañeda (ed.) *Action, Knowledge, and Reality* (Indianapolis: Bobbs-Merrill, 1975); and Donald Davidson, 'A Coherence Theory of Truth and Knowledge,' in Dieter Hendrich (ed.), *Kant oder Hegel* (Stuttgart: Klett-Cotta, 1983), 432–8. It should be noted that Davidson has written an afterword to this paper. Here he says, regarding "The main thrust of 'A Coherence Theory,'" that "the important thesis for which I argue is that belief is intrinsically veridical." See 'Afterthoughts, 1987,' in Sven Bernecker and Fred Dretske (eds), *Knowledge: Readings in Contemporary Epistemology* (Oxford: Oxford University Press, 2000), p. 427. This view bears a foundationalist interpretation: beliefs, being intrinsically veridical, are as such prima facie justified, even if weakly; hence they are defeasibly foundational. Incoherence would defeat them; but particularly if, as Davidson says here, "Coherence is nothing but consistency" (p. 427), it would not be plausible to take coherence to be a basic source of justification.

17 The idea of enriching coherentism by making coherence with experience an essential element in coherentist justification is proposed and defended by Jonathan L. Kvanvig and Wayne D. Rigg, 'Can a Coherence Theory Appeal to Appearance States?,' *Philosophical Studies* 67 (1992), 197–217. This paper deserves study. Here I raise just one

difficulty. Although they grant that "coherentism arises historically because of dissatisfaction with the foundationalists' picture" (p. 199), they characterize a foundationalist warranting relation in a way that does not distinguish it from the relation coherentists take to confer justification.

> One such account could claim that a belief is foundationally warranting just in case the evidence for it is an appearance state involving the same content as that of the belief. For example ... perhaps my belief that something is red is intrinsically warranting because it appears to me that something is red.
>
> (p. 199)

A foundationalist will not take an appearance state, such as a sensory impression of red, to have the same content as a belief: a propositional content in virtue of which the belief is true or false, e.g. "that something is red." Rather, the experiential content is qualitative; it may be *appropriate to* certain propositions but is not itself truth-valued. Such a content might be *an appearance of red* but not the proposition that "it appears to me that something is red." The latter *is* a candidate to enter into a coherence relation with beliefs or their contents. Perhaps Kvanvig and Rigg are thinking of experiential justification of belief as possibly working through beliefs or other states which have propositional content and truth value; this could explain why they find such justification available to coherentism. If, however, experiential justification could work that way, then one could *still* have a coherent system of beliefs that goes against experience. Beliefs *about* one's states – such as the (appearance) "belief that something is red" – would have to play a role, but those states would not be any kind of bedrock grounding these beliefs, even if the beliefs happened to be based on them. The problem, then, is that either the coherence-with-experience approach assimilates coherentism to a kind of foundationalism or it fails to capture the role of experience, which seems essential for a body of justified beliefs about the world.

Part Three

The nature and scope of justification and knowledge