

# Why Coherence Is not Enough: A Defense of Moderate Foundationalism

*James van Cleve*

## I

Foundationalism has been characterized as the view that “the knowledge which a person has at any time is a structure or edifice, many parts and stages of which help to support each other, but which as a whole is supported by its own foundation” (Chisholm, 1964). To unpack the metaphor, we may say that the foundation consists of *basic beliefs* – beliefs that a subject is justified in holding even in the absence of any justifying reason for them – and all other justified beliefs derive their justification at least in part from such basic beliefs.

The classical argument for foundationalism is an infinite regress argument going back to Aristotle: unless there were basic beliefs, every justified belief would rest on further justified beliefs, and so on without end. An infinite regress is not in fact the only alternative to basic beliefs, but the other alternatives are equally problematic, or so foundationalists maintain. To canvass all the options, let us set down four propositions that jointly imply the existence of an infinite regress of justified beliefs:

- 1 Some beliefs are justified.
- 2 No belief is justified unless some other belief serves as a reason for it.
- 3 One belief cannot serve as a reason justifying another unless the first is itself justified.
- 4 If A serves as a reason justifying B, then B cannot serve (directly or indirectly) as a reason justifying B.

These four propositions jointly entail the existence of an infinite regress of justified propositions. We are therefore faced with five alternatives: accept the infinite regress or reject one of the four assumptions.

*Skeptics* (of the universal ilk) deny 1, maintaining that no beliefs whatever are justified. *Foundationalists* deny 2, maintaining that some beliefs are justified in the absence of reasons. *Positists* (not to be confused with positivists) deny 3, maintaining that chains of justifying reasons can terminate in reasons that are not justified themselves, but are simply individual or societal posits.<sup>1</sup> *Coherentists* deny 4, maintaining that beliefs can be justified in virtue of relations of mutual support. *Infini-*  
*tists* accept all four assumptions and the resulting infinite regress.

All five options have their takers. Skepticism and infinitism are both defended elsewhere in this volume (see the essays by Fumerton and Klein). Positism finds advocacy in Wittgenstein’s *On Certainty* and in assorted postmodern thinkers (Wittgenstein, 1969). But I think it is fair to say that the leading contenders among the five options are foundationalism and coherentism.

Coherentism is sometimes characterized as a view that sanctions circular reasoning, but that is an oversimplified construal of it. Coherentists do not typically endorse simple loops in which A justifies B, B justifies C, and C justifies A; rather, they envision vast webs of belief in which everything is supported by some significant portion of the remaining beliefs: A by B and C, B by A, D, J, and K, and so on.

For their part, foundationalists do not typically deny the power of coherence to contribute to the overall epistemic status of a body of belief. They simply insist that coherence cannot do all the work on its own – there must be at least a modicum of intrinsic credibility or non-inferential warrant possessed by basic beliefs before coherence can have its amplifying effect.

Laurence Bonjour distinguishes three grades of foundationalism (Bonjour, 1985, pp. 26–30). According to *strong* foundationalism, basic beliefs are “not just adequately justified, but also *infallible, certain, indubitable, or incorrigible*” (Bonjour, 1985, pp. 26–7). According to *moderate* foundationalism, the non-inferential warrant possessed by basic beliefs need not amount to absolute certainty or any of the other privileged statuses just mentioned, but it must be “sufficient by itself to satisfy the adequate-justification condition for knowledge” (Bonjour, 1985, p. 26). Finally, according to *weak* foundationalism,

basic beliefs possess only a very low degree of epistemic justification on their own, a degree of justification insufficient by itself either to satisfy the adequate-justification condition for knowledge or to qualify them as acceptable justifying premises for further beliefs. Such beliefs are only “initially credible,” rather than fully justified. (Bonjour, 1985, p. 28)

We must rely on coherence among such initially credible beliefs to amplify their level of warrant up to the point where it is adequate for knowledge.

As Bonjour notes, weak foundationalism could be regarded as a hybrid view, mixing together foundational and coherentist elements. In fact, Susan Haack prefers to call it “foundherentism,” which she illustrates with the example of a crossword puzzle. Experience corresponds to the clues, which give an initial presumption in favor of certain beliefs or entries in the puzzle; the initial beliefs are then confirmed by the way in which they interlock with other entries (or sometimes discarded because they do not fit in). Thus does coherence amplify (or its absence erode) the initial warrant possessed by basic beliefs (Haack, 1993).<sup>2</sup>

But if coherence can elevate the epistemic status of a set of beliefs in this way, what prevents it from generating warrant entirely on its own, without any need for basic beliefs? This is a question that has been asked by several authors, including Bonjour:

The basic idea is that an initially low degree of justification can somehow be magnified or amplified by coherence, to a degree adequate for knowledge. But how is this magnification or amplification supposed to work? How can coherence, not itself an independent source of justification on a foundationalist view, justify the rejection of some initially credible beliefs and enhance the justification of others? (Bonjour, 1985, p. 29)

The implied suggestion is that if coherence can do what the weak foundationalist allows, it can also do what the thoroughgoing coherentist says it can do.

In the next section, I address this challenge. I argue that there is indeed a good rationale for the weak foundationalist's insistence that before coherence can do its work, there must be initially credible inputs for it to work upon. But I also argue that the level of initial credibility cannot be as low as that envisioned by weak foundationalists, who should therefore upgrade their position to moderate foundationalism. Finally, in section III, I offer some critical reflections on the coherentist alternative defended by Catherine Elgin.

## II

An excellent example of a weak foundationalist theory is provided by C. I. Lewis's theory of memory knowledge. It has two elements:

First; whatever is remembered, whether as explicit recollection or merely in the form of our sense of the past, is *prima facie* credible because so remembered. And second; when the whole range of empirical beliefs is taken into account, all of them more or less dependent upon memorial knowledge, we find that those which are most credible can be assured by their mutual support, or as we shall put it, by their *congruence*. (Lewis, 1946, p. 334)<sup>3</sup>

Lewis defines a congruent set as one in which any member is more probable given the rest than it is on its own:

A set of statements, or a set of supposed facts asserted, will be said to be congruent if and only if they are so related that the antecedent probability of any one of them will be increased if the remainder of the set can be assumed as premises. (AKV, p. 338)

A point on which Lewis repeatedly insists is that congruence alone cannot generate probability or warrant. Rather, some of the statements must have initial credibility, which congruence can then amplify:

The feature of such corroboration through congruence that should impress us, is the requirement that the items exhibiting these congruent relationships must – some of them at least – be *independently given facts* or have a probability which is antecedent. (AKV, p. 352)

How *much* probability must the congruent items have? Only a “slight” amount, Lewis tells us, illustrating his point with the example of individually unreliable witnesses who tell the same story:

Our previous example [AKV, p. 239] of the relatively unreliable witnesses who independently tell the same circumstantial story, is another illustration of the logic of congruence; and one which is more closely typical of the importance of relations of congruence for determination of empirical truth in general. For any of these reports,

taken singly, the extent to which it confirms what is reported may be slight. And antecedently, the probability of what is reported may also be small. But congruence of the reports establishes a high probability of what they agree upon, by principles of probability determination which are familiar: on any other hypothesis than that of truth-telling, this agreement is highly unlikely. (AKV, p. 346)

But how much probability is a “slight” amount? In one place Lewis says, “Anything sensed as past is just a little more probable than that which is incompatible with what is remembered and that with respect to which memory is blank” (AKV, p. 358). He thereby implies that a remembered piece of information has a probability greater than 0.5, given that it is remembered. (This is because it is an axiom of the probability calculus that  $P(\sim h, e) = 1 - P(h, e)$ ; thus a proposition has greater probability than its negation if and only if it has probability greater than 0.5.) Thus according to Lewis’s version of weak foundationalism, the congruence of a set of remembered items can raise their level of justification arbitrarily high, but only if the items have initial credibility amounting to a probability (given that we seem to remember them) greater than 0.5.

To this latter aspect of Lewis’s theory BonJour has raised an objection. There is no need, he says, for Lewis’s requirement that memory reports or other cognitive deliverances have initial credibility:

What Lewis does not see, however, is that his own example shows quite convincingly that no antecedent degree of warrant or credibility is required. For as long as we are confident that the reports of the various witnesses are genuinely independent of each other, a high enough degree of coherence among them will eventually dictate the hypothesis of truth telling as the only available explanation of their agreement – even, indeed, if those individually reports initially have a high degree of *negative* credibility, that is, are much more likely to be false than true (for example, in the case where all of the witnesses are known to be habitual liars). (BonJour, 1985, pp. 147–8)

We are now presented with a clear-cut issue to investigate: in order for the congruence of a set of items to raise their credibility to near 1, what level of antecedent credibility is required? Must it be greater than that of their negations and thus greater than 0.5, as Lewis maintains? Or may it be less than 0.5, as BonJour implies when he says that the reports of the witnesses (or of our memory) may be more likely false than true?

For light on this question, we may look to discussion of a traditional topic in probability theory: how to assess the probability that independent witnesses who agree in their testimony are telling the truth. This is a problem to which a number of classical authors have proposed answers. One of the standard answers is due to George Boole (1952, p. 364):<sup>4</sup>

Let  $p$  be the general probability that  $A$  speaks the truth,  $q$  the general probability that  $B$  speaks the truth; it is required to find the probability that, if they agree in a statement, they both speak the truth. Now agreement in the same statement implies that they either both speak truth, the probability of which beforehand is  $pq$ , or that they both speak falsehood, the probability of which beforehand is  $(1 - p)(1 - q)$ . Hence the probability

beforehand that they will agree is  $pq + (1 - p)(1 - q)$  and the probability that if they agree, they will agree in speaking the truth is accordingly expressed by the formula

$$w[= P(\text{A and B speak truly, they agree})] = \frac{pq}{pq + (1 - p)(1 - q)}$$

For an explanation of the rationale behind Boole's formula, I must refer the reader to what I have said elsewhere.<sup>5</sup> Here there is space only to note the bearing of his formula on the issue separating Lewis and BonJour. Suppose that A and B each tell the truth 60 percent of the time; that is, suppose that  $p$  and  $q$  are each equal to 0.6. The reader may verify that in this case,  $w = 0.69$ . That is, setting A's credibility and B's each equal to 0.6, the probability that X is true given that they each testify to it is 0.69. More generally, if we plug in any numbers greater than 0.5 as  $p$  and  $q$ ,  $w$  will be greater than the mean of  $p$  and  $q$ . So far we have an illustration of the point, common ground for Lewis and BonJour, that congruence can boost credibility.

But what if we plug in values of  $p$  and  $q$  equal to or less than 0.5? If 0.5 goes in, 0.5 comes out; and if  $p$  and  $q$  are each less than 0.5, the output value will be less than their mean. For example, if  $p$  and  $q$  are each 0.1,  $w$  is approximately 0.01. In other words, if the witnesses have what BonJour calls "negative credibility" (credibility less than 0.5), the probability of a statement given that they both testify to it is not enhanced but diminished! So if Boole's formula is correct, Lewis is vindicated and BonJour refuted: with initial credibilities less than 0.5, coherence makes things worse rather than better.

There is reason to be suspicious of Boole's formula, however. In deriving it, he tacitly assumes that there are only two possible answers to the questions put to the witnesses – true or false. That is why he can equate agreement with "both speak truly or both speak falsely." In a more realistic scenario, the witnesses would be asked multiple-choice questions, and if they agreed in giving the same answer out of ten possible choices (let us say), their agreement would be much more impressive.

A formula allowing for an arbitrary number of possible answers to questions has been devised by Michael Huemer (1997).<sup>6</sup> For simplicity's sake, Huemer assumes that the witnesses have the same level of credibility, so  $p = q$ . If  $n$  is the number of possible answers and X is the answer on which independent witnesses A and B agree, then Huemer's formula may be written as follows (omitting details of the derivation):

$$P(X, \text{A says X} \ \& \ \text{B says X})[= W] = \frac{np^2 - p^2}{np^2 - p^2 + 1}$$

This formula agrees with Boole's in the special case where  $n = 2$ , but gives dramatically different results in the cases not covered by Boole's. Specifically, it enables coherence to have its amplifying affect even when the credibility level of the witnesses is below 0.5, just so long as it is greater than the chance or random guessing level of  $1/n$ . For example, if  $p$  is only 0.3, but there are ten possible answers (say, ten possible last digits in a glimpsed license plate number), then  $w = 0.62$ . If  $p = 0.3$  and  $n = 100$ , then  $w = 0.86$ . For any value of  $p$ , just so long as it is greater than zero, we can bring the final probability of X (i.e. its probability given that the witnesses

agree on it) as close as we like to certainty by making  $n$  high enough. With a more complex version of the formula, we can also make  $w$  higher by increasing the number of witnesses.

If Huemer's formula is correct, then, BonJour is vindicated and Lewis refuted. Initial credibilities need not be greater than 0.5; they need only be greater than the chance level of  $1/n$  (where  $n$  = the number of possible answers). If enough witnesses agree without collusion in giving the same answer from among a large enough number of choices, it becomes overwhelmingly likely that they are correct, even if their initial level of credibility was scarcely above zero.

Is BonJour right, then? Can coherence alone be a source of warrant without need of inputs with initial credibility? An answer of yes would be too hasty, for there is another requirement of initial credibility we have yet to consider. The requirement at which we have so far demurred is the requirement that what is reported must be more probable than not (and thus have probability greater than 0.5) given that a witness (or an ostensible memory) attests to it. If several individually unreliable reporters agree without collusion, then the fact to which they bear common witness may have high probability in the end. But in attaching a high final probability to the fact attested, we are of course taking for granted that the various witnesses *do* testify to it. If we had reason to think that the courtroom and all its proceedings were happening only in a dream or a novel, the fact that the ostensible reports hang together would count for little. And so it is with the reports of memory, the senses, and cognitive systems more generally: coherence among them lends high final credibility only on the assumption that the reports genuinely occur.

How, then, do we know *these* things: that witness *A* does say *X*, that I do ostensibly remember *Y*, that I do seem to see *Z*? Many foundationalists would say that these are the grounds on which the rest of our knowledge rests, and that they must themselves be matters of basic knowledge.<sup>7</sup> Lewis himself famously maintained that nothing can be probable unless something is certain, and among the certainties he placed the facts that I do have this or that presentation of sense or memory. His insistence on certainty is controversial,<sup>8</sup> but it seems to me that a good case can be made that there must at least be high intrinsic credibility – perhaps high enough to constitute knowledge – attaching to the facts that such-and-such cognitive states (be they experiences, ostensible memories, or beliefs at large) are actually taking place. If this is right, we must not only abjure pure coherentism: we must also adopt a moderate rather than a weak foundationalism.

I see only one plausible alternative to an assumption of high initial credibility or knowledge-sufficient warrant at the foundational level, and that is the view that the promptings of sense or memory function as *external* conditions of knowledge. An external condition of knowledge is a condition that makes knowledge possible regardless of whether it is itself known to obtain. For example, in Goldman's reliability theory, if a subject comes to believe *p* as the result of a reliable process, his belief is knowledge regardless of whether the subject knows anything about the reliability of the process (Goldman, 1979). Perhaps the facts that I have such-and-such ostensible perceptions or memories could function in this external way, contributing to my knowledge even if not themselves known. The idea would be that my ostensible perceiving and remembering are not pieces of evidence on which I conditionalize when

their epistemic status is high enough;<sup>9</sup> instead, they are facts whose mere obtaining confers credibility on their contents.

I turn now to another point at which I think a theory of knowledge that invokes coherence must make a concession to either foundationalism or externalism.

The example of the witnesses who agree is in one respect a drastically oversimplified case of coherence. The agreement of the witnesses is literal identity, or at least logical equivalence, of content: witness 1 says X and so does witness 2. But the coherence that figures in epistemology is typically a much looser sort of hanging together. The coherence of ostensible memories is not their all being memories that p, for the same p or something logically equivalent. Nor is the coherence of beliefs or cognitions generally like that. Rather, it is a type of coherence that is exemplified by the following items:

I seem to remember seeing a skunk last night;  
I seem to remember smelling a skunk last night;  
I seem to remember that the lid was on the garbage can when I went to bed;  
I now see that the can has been knocked over and trash strewn about;  
There was a skunk here last night;

and so on. In other words, it is not identity or even equivalence of content, but rather something like the relation Lewis calls congruence: a matter of each item being more probable given the rest than it is on its own.

What are these coherence-constituting relations of probability founded upon, and how do we know that they obtain?

One answer has been given by Russell: "It is only by assuming laws that one fact can make another probable or improbable" (Russell, 1948, p. 188). Perhaps Russell goes too far in requiring strict laws in order for one fact to make another probable, but it is plausible that we at least require rough empirical generalizations. Where do these generalizations come from? Presumably, they are inferred inductively from particular facts gathered by memory. And now the following difficulty emerges: ostensible memories give rise to knowledge only with the help of coherence; coherence depends on laws or empirical generalizations; and such generalizations can be known only with the help of memory. In short, we cannot get coherence without the help of laws, and if memory does not suffice on its own to give knowledge of particular facts from which the laws are inferred, we cannot get laws without the help of coherence. It appears to follow that we cannot have any knowledge from memory unless the occurrence of ostensible memories is *prima facie* sufficient for knowledge. Such was Russell's own conclusion:

memory is a premise of knowledge. . . . When I say memory is a premise, I mean that among the facts upon which scientific laws are based, some are admitted solely because they are remembered. (Russell, 1948, pp. 188–9).

Note the word "solely." Russell is saying that individual memories must be capable of giving rise to knowledge on their own, without benefit of coherence. This is compatible, of course, with allowing that the warrant provided by memory is defeasible,

as Russell did allow. But the resulting view is nonetheless a foundationalism of memory knowledge stronger than that of Lewis, who required only an initial “slight” presumption in favor of the truth of any ostensible memory. Russell’s view accords to memory greater epistemic powers than that: ostensibly remembering that *p* is a source of prima facie warrant that, if undefeated (and if *p* is true) is strong enough for knowing that *p*. In BonJour’s terms, we have again advanced from weak to moderate foundationalism, this time as regards the contents of ostensible memories rather than the occurrences of them as mental events.

Russell’s argument assumes that coherence and the laws that underlie it contribute to our knowledge only if they are themselves known. As in the case of the occurrences of our ostensible memories, one could challenge this assumption by going external, holding that coherence does its work regardless of whether the subject knows it obtains. This is the second point at which I believe coherentism can avoid a concession to foundationalism only by making a concession to externalism. The fact that *p*, *q*, and *r* do cohere with one another, as well as the facts that they are deliverances of our cognitive systems to begin with, are facts that must either function externally or be known foundationally.

### III

As the debate between foundationalists and coherentists has progressed, each side has moved in the direction of the other. Contemporary foundationalists are seldom foundationalists of the strong Cartesian variety: they do not insist that basic beliefs be absolutely certain. They also typically allow that the elements in a system of belief can acquire enhanced justification through their coherence. On the other side, many coherentists admit that coherence alone is not the sole source of justification – there must be some initially credible inputs before coherence can work its wonders. Is there anything more to disagree about, or do foundationalists and coherentists now meet in the middle?

There are indeed still points of difference. To highlight several of them, I shall discuss the broadly coherentist views of Catherine Elgin as developed in her book *Considered Judgment* (1996; cited as CJ hereafter) and in her contribution to this volume.

Elgin characterizes herself as a proponent of reflective equilibrium.<sup>10</sup> As she conceives of it, reflective equilibrium has two chief requirements: “The components of a system in reflective equilibrium must be reasonable in light of one another, and the system as a whole reasonable in light of our initially tenable commitments.” (CJ, p. 107; see also pp. ix, 13, and 127–8). It is by the second requirement that Elgin distinguishes her view from a pure coherentism: the components of a system in equilibrium must be answerable not just to one another, but also to our initially tenable commitments.<sup>11</sup>

The second requirement puts a “tether” on permissible systems (CJ, pp. 10, 107, 128), thereby enabling Elgin to avoid some of the objections to pure coherentism. For example, one of the standard objections to coherentism is that the contents of a consistent fairy tale would be a body of warranted propositions (see Schlick, 1973,

p. 419). Not so for Elgin, since the propositions in the story may be reasonable in light of each other without being reasonable in light of our initially tenable commitments.

Of the various things we believe, which have initial tenability? In Elgin's view, they *all* do, if we genuinely believe them: anything actually held has some initial tenability or presumption in its favor (CJ, pp. 101–2). The presumption may only be slight and it may be lost in the end, but it is there in the beginning.

What about the various principles of logic, evidence, and method whereby some things are reasonable in light of others? In Elgin's view, these have the same status as everything else: they are initially tenable if held, and they may gain or lose in tenability depending on how they fit in with everything else (CJ, p. 104).

Why is Elgin's view as so far set forth not simply a form of weak foundationalism, in which initially tenable claims function as basic beliefs? She cites two differences: unlike the justification that attaches to foundationalism's basic claims, initial tenability can be lost; it can also be augmented through coherence (CJ, p. 110).

It is not clear to me, however, that either of these features should be regarded as a prerogative of coherentists alone. In a typology of possible foundationalisms, Roderick Firth has suggested that the minimal tenet of foundationalism is simply this: basic beliefs have some measure of initial warrant that is not derived from coherence. This level of warrant may be increased by coherence with other statements or diminished, even to the vanishing point, by lack of coherence with other statements.<sup>12</sup> Firth's minimal view thus incorporates both of the features Elgin sees as antithetical to foundationalism.<sup>13</sup>

Nonetheless, I see two other questions on which foundationalists are apt to disagree with Elgin. First, are all commitments initially tenable, or only those in some specially marked out class? Second, are all commitments likewise revisable, or are some immune from subsequent rejection? On each question, Elgin takes the more egalitarian stand (CJ, pp. 101–2 and 121).

Using terminology from Michael Huemer, we may say that the first issue is the issue of phenomenal conservatism versus a more general doxastic conservatism (Huemer, 2001, pp. 99–115). Phenomenal conservatism is the view that if anything seems to be the case, one is *prima facie* justified in accepting it. The seemings can include perceptual seemings (it looks to me as if there is a red object over there), memorial seemings (I seem to remember being chased by a dog one day on my way to kindergarten), and intellectual seemings (it strikes me as self-evident that the relation of equality is transitive). Doxastic conservatism is the more sweepingly democratic view that anything the subject believes has some presumption in its favor – the products of wishful thinking and superstition no less than the deliverances of perception and memory. Foundationalists are typically phenomenal conservatives, while Elgin (in her book) is a doxastic conservative.

In her contribution to this volume, Elgin no longer espouses doxastic conservatism, or at any rate holds that coherentists need not be committed to it. She says that perceptual deliverances may be assigned special weight, provided they do not have it *a priori*. They have it only in virtue of coherence considerations: by accepting the deliverances of perception, we get systems that remain coherent over time (that is, they continue to jibe with new deliverances of perception). But might it not likewise be true that by accepting fantasies, we get systems that remain coherent with future con-

tents of fantasy – especially if the fantasizer has a one-track mind? So it seems to me that there is a privileging of perception presupposed and not explained by this coherentist underpinning for it.

I turn now to the second point: are all commitments on a par as regards the possibility of revision? Elgin holds, with Quine, that in the quest for reflective equilibrium, anything may be revised. There is no commitment that may not be sacrificed in order to maximize the tenability of the entire system. Here I would like to protest that there are certain principles of logic, at least, that cannot be given up, because they are framework principles without which coherence could scarcely be defined. What would happen if we gave up the law of non-contradiction? It is not clear that there could any longer be such a thing as what Quine calls a recalcitrant experience, forcing changes elsewhere in the system. If a new deliverance stood in contradiction to things we already accepted, we could simply accept it with a “What – me worry?” shrug.

Here is another difficulty for a coherentism that holds everything revisable, at least if we understand this as “anything could be justifiably rejected,” symbolized as “(p)possibly  $J \sim p$ .” Suppose there is some proposition  $q$  (the law of non-contradiction, perhaps?) whose truth is necessary for anything to belong to a coherent system and therefore necessary for anything to be justified. Since  $\sim q$  entails that nothing is justified, we now have *possibly  $J$  (nothing is justified)*, which is absurd.<sup>14</sup>

A related question about the status of the rules and principles of logic and evidence is whether they have force only because a subject is committed to them. I gather Elgin would say yes, but I say no. Consider a system of beliefs containing the elements  $p$ ,  $q$ , and  $p \rightarrow \sim q$ . Suppose it is not a sheer fact of logic, independent of anything the subject believes, that the system is inconsistent and in need of revision. Could we make the system intolerably inconsistent by adding the principle: if any two of  $\{p, q, \text{ and } p \rightarrow \sim q\}$  are true, then the third must be false? No, for anyone who could live comfortably with the original system could live comfortably with the expanded system. Such is the lesson the tortoise taught Achilles (Carroll, 1895).

I wish to raise one more question about the role of logical and evidential relations in Elgin’s coherentism. Elgin defines both coherence and equilibrium in terms of the relation “ $p$  is reasonable in light of  $q$ ,  $r$ , and  $s$ .” What is the required epistemic status of this relation (or of the logical and other relations on which it supervenes)? Must such relations be known to hold among the propositions in a system before the propositions are warranted for the subject? And if so, how does such knowledge arise?

I see three possibilities to consider in answer to these questions. First, the holding of coherence-constituting relations might be regarded as an external condition of knowledge, making knowledge possible regardless of whether the subject knows that such relations hold. Although this seems to me a good way for a coherentist to go, I gather that Elgin would not find it congenial. On more than one occasion, she expresses her dissatisfaction with externalist stratagems in epistemology (CJ, pp. 22, 46, 51; this volume, p. 163). Second, it might be held that logical relations and other relations of support are known to hold because they are necessary relations, apprehensible *a priori*. But this, too, is an option Elgin would reject. In the first place, it would be a concession to foundationalism; in the second place, she has Quinean qualms about there being any propositions at all that are true necessarily and known

*a priori* (CJ, pp. 40–6 and 53–7). Third, it might be held that coherence-constituting relations are known to hold because the propositions saying that they hold are part of the best overall coherent system that is reasonable in light of our antecedent commitments.<sup>15</sup> But that way lies an infinite regress. A proposition *p* affirming a relation of coherence would be justified only because the subject is justified in believing that *p* belongs to a coherent system. That belief in turn would be justified only because the subject is justified in believing that *p belongs to a coherent system* belongs to a coherent system, and so on. Even if it is the same system every step of the way, we still get a regress in which ever more complicated propositions must be believed with justification to belong to the system. The untenability of such a regress suggests to me that we should go instead with one of the first two options, agreeing with the externalist that coherence propositions need not be known at all or with the foundationalist that they are known because they are either basic propositions or propositions inferrable from basic propositions.

## IV

I have not lived up to the title of this essay, for I have not offered a complete defense of moderate foundationalism. I hope nonetheless to have shown that an internalist coherentism cannot be a satisfactory theory of justification. We must be either externalists or foundationalists, and if we are foundationalists, our foundationalism must be of the moderate rather than the weak variety.

## Notes

- 1 The distinction between positivism and foundationalism is lost on those who cannot hear the word “justified” as anything but a past participle, implying that some act or relation of justifying has occurred whereby a belief is justified by something else that serves as a reason for it. For foundationalists, “justified” simply connotes a favorable epistemic status, which a belief may have even though the subject has no reason for it. In this connection, another term, such as “evident” or “credible,” might be less misleading than “justified.”
- 2 Haack summarizes her theory in Haack (2000). She does not classify foundherentism as weak foundationalism, believing it essential to foundationalism that the foundations not receive support from other elements in the structure.
- 3 Hereafter I shall cite this work in the text as AKV.
- 4 This paper was first published in 1857.
- 5 James Van Cleve, “Can coherence generate warrant?” (in preparation).
- 6 I discuss Huemer’s work at some length in “Can coherence generate warrant?”
- 7 For foundationalists who hold that propositions about the physical world are always derived, not basic, that *A* says *X* would not be basic after all. It would rest, however, on deliverances that are basic.
- 8 For further discussion of Lewis’s view, see Van Cleve (1977). For a method of assigning probabilities in relation to evidence without assuming that the evidence is certain (in the sense of having probability 1), see Jeffrey (1983).
- 9 By “conditionalizing” I mean drawing the inference “*h* has probability *n* given evidence *e*, and *e* is certainly true; therefore, *h* has probability *n*.”

- 10 The term was coined by John Rawls (1971, pp. 20f). It refers to a state of affairs in which specific judgments and general principles have been brought into agreement with each other through a process of mutual adjustment.
- 11 Elsewhere Elgin adds other requirements. One is that there must not be a competing system that is more tenable overall (Elgin, 1996, p. 107). Another is that the best explanation of the system's coherence must be that it is at least roughly true ("Non-foundationalist epistemology: holism, coherence, and tenability," this volume).
- 12 See Firth (1964, pp. 466–7) for his progressively weaker foundationalisms.
- 13 Even so staunch a foundationalist as Roderick Chisholm incorporates Elgin's two allegedly distinctive features to some extent. Although the warrant belonging to Chisholmian "self-presenting" propositions can be neither increased nor diminished by their relations to other propositions, the same is not true of the warrant possessed by "indirectly evident" propositions. Their epistemic status is defeasible, and it may be raised by concurrence (Chisholm's term for coherence). See Chisholm (1977, chapters 2 and 4).
- 14 In fairness to Elgin, I note that she may give special status to the law of non-contradiction. She says that tenable systems must at least be logically consistent, or else be all-entailing (CJ, p. 136). Against this rationale, however, I note that in revisionary relevance logic, contradictions do not entail everything.
- 15 Perhaps this is what coherentists should say in response to Russell's question about how the laws underlying coherence are known.

## References

- BonJour, L. (1985) *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.
- Boole, G. (1952) On the application of the theory of probabilities to the question of the combination of testimonies or judgments. In *Studies in Logic and Probability*. London: C. A. Watts.
- Carroll, L. (1895) What the Tortoise said to Achilles. *Mind*, 4, 275–80.
- Chisholm, R. M. (1964) The myth of the given. In R. M. Chisholm et al., *Philosophy*. Englewood Cliffs, NJ: Prentice Hall.
- Chisholm, R. M. (1977) *Theory of Knowledge*, 2nd edn. Englewood Cliffs, NJ: Prentice Hall.
- Elgin, C. (1996) *Considered Judgment*. Princeton, NJ: Princeton University Press.
- Firth, R. (1964) Coherence, certainty, and epistemic priority. *Journal of Philosophy*, 61, 545–57. Reprinted in R. Chisholm and R. J. Swartz (eds, 1973) *Empirical Knowledge*. Englewood Cliffs, NJ: Prentice Hall.
- Goldman, A. (1979) What is justified belief? In G. Pappas (ed.), *Justification and Knowledge*. Dordrecht: D. Reidel.
- Haack, S. (1993) *Evidence and Inquiry*. Oxford: Blackwell.
- Haack, S. (2000) A foundherentist theory of empirical justification. In E. Sosa and J. Kim (eds), *Epistemology*. Oxford: Blackwell.
- Huemer, M. (1997) Probability and coherence justification. *Southern Journal of Philosophy*, 35, 463–72.
- Huemer, M. (2001) *Skepticism and the Veil of Perception*. Lanham, MD: Rowman & Littlefield.
- Jeffrey, R. (1983) *The Logic of Decision*, 2nd edn. Chicago: University of Chicago Press.
- Lewis, C. I. (1946) *An Analysis of Knowledge and Valuation*. LaSalle, IL: Open Court.
- Rawls, J. (1971) *A Theory of Justice*. Cambridge, MA: Harvard University Press.
- Russell, B. (1948) *Human Knowledge: Its Scope and Limits*. New York: Simon & Schuster.

- Schlick, M. (1973) The foundation of knowledge. In R. Chisholm and R. J. Swartz (eds), *Empirical Knowledge*. Englewood Cliffs, NJ: Prentice Hall.
- van Cleve, J. (1977) Probability and certainty: a re-examination of the Lewis–Reichenbach debate. *Philosophical Studies*, 32, 323–34.
- Wittgenstein, L. (1969) *On Certainty*. Oxford: Blackwell.