CHAPTER 13

The Raft and the Pyramid

Ernest Sosa

Contemporary epistemology must choose between the solid security of the ancient foundationalist pyramid and the risky adventure of the new coherentist raft. Our main objective will be to understand, as deeply as we can, the nature of the controversy and the reasons for and against each of the two options. But first of all we take note of two underlying assumptions.

1 Two Assumptions

(A1) Not everything believed is known, but nothing can be known without being at least believed (or accepted, presumed, taken for granted, or the like) in some broad sense. What additional requirements must a belief fill in order to be knowledge? There are surely at least the following two: (a) it must be true, and (b) it must be justified (or warranted, reasonable, correct, or the like).

(A2) Let us assume, moreover, with respect to the second condition A1(b): first, that it involves a normative or evaluative property; and, second, that the relevant sort of justification is that which pertains to knowledge: epistemic (or theoretical) justification. Someone seriously ill may have two sorts of justification for believing he will recover: the practical justification that derives from the contribution such belief will make to his recovery and the theoretical justification provided by the lab results, the doctor's diagnosis and prognosis, and so on. Only the latter is relevant to the question whether he knows.

2 Knowledge and Criteria

a. There are two key questions of the theory of knowledge:
   (i) What do we know?
   (ii) How do we know?
   The answer to the first would be a list of bits of knowledge or at least of types of knowledge: of the self, of the external world, of other minds, and so on. An answer to the second would give criteria (or canons, methods, principles, or the like) that would explain how we know whatever it is that we do know.

b. In developing a theory of knowledge, we can begin either with a(i) or with a(ii). Particularism would have us begin with an answer to a(i) and only then take up a(ii) on the basis of that answer. Quite to the contrary, methodism

Originally published in *Midwest Studies in Philosophy, Vol. 5: Studies in Epistemology* (Minneapolis: University of Minnesota Press, 1980), pp. 3–25; an appendix to this paper is drawn from Ernest Sosa, "How Do You Know?" *American Philosophical Quarterly* 11 (1974), pp. 113–22.

would reverse that order. The particularist thus tends to be antiskeptical on principle. But the methodist is as such equally receptive to skepticism and to the contrary. Hume, for example, was no less a methodist than Descartes. Each accepted, in effect, that only the obvious and what is proved deductively on its basis can possibly be known.

c. What, then, is the obvious? For Descartes it is what we know by intuition, what is clear and distinct, what is indubitable and credible with no fear of error. Thus for Descartes basic knowledge is always an infallible belief in an indubitable truth. All other knowledge must stand on that basis through deductive proof. Starting from such criteria (canons, methods, etc.), Descartes concluded that knowledge extended about as far as his contemporaries believed.[1] Starting from similar criteria, however, Hume concluded that both science and common sense made claims far beyond their rightful limits.

d. Philosophical posterity has rejected Descartes's theory for one main reason: that it admits too easily as obvious what is nothing of the sort. Descartes's reasoning is beautifully simple: God exists; no omnipotent perfectly good being would descend to deceit; but if our common sense beliefs were radically false, that would represent deceit on His part. Therefore, our common sense beliefs must be true or at least cannot be radically false. But in order to buttress this line of reasoning and fill in details, Descartes appeals to various principles that appear something less than indubitable.

e. For his part, Hume rejects all but a minuscule portion of our supposed common sense knowledge. He establishes first that there is no way to prove such supposed knowledge on the basis of what is obvious at any given moment through reason or experience. And he concludes, in keeping with this methodism, that in point of fact there really is no such knowledge.

3 Two Metaphors: The Raft and the Pyramid

Both metaphors concern the body or system of knowledge in a given mind. But the mind is of course a more complex marvel than is sometimes supposed. Here I do not allude to the depths plumbed by Freud, nor even to Chomsky's. Nor need we recall the labyrinths inhabited by statesmen and diplomats, nor the rich patterns of some novels or theories. We need look no further than the most common, everyday beliefs. Take, for instance, the belief that driving tonight will be dangerous. Brief reflection should reveal that any of us with that belief will join to it several other closely related beliefs on which the given belief depends for its existence or (at least) its justification. Among such beliefs we could presumably find some or all of the following: that the road will be icy or snowy; that driving on ice or snow is dangerous; that it will rain or snow tonight; that the temperature will be below freezing; appropriate beliefs about the forecast and its reliability; and so on.

How must such beliefs be interrelated in order to help justify my belief about the danger of driving tonight? Here foundationalism and coherentism disagree, each offering its own metaphor. Let us have a closer look at this dispute, starting with foundationalism.

Both Descartes and Hume attribute to human knowledge an architectonic structure. There is a nonsymmetric relation of physical support such that any two floors of a building are tied by that relation: one of the two supports (or at least helps support) the other. And there is, moreover, a part with a special status: the foundation, which is supported by none of the floors while supporting them all.

With respect to a body of knowledge K (in someone's possession), foundationalism implies that K can be divided into parts K_1, K_2...such that there is some nonsymmetric relation R (analogous to the relation of physical support) which orders those parts in such a way that there is one – call it F – that bears R to every other part while none of them bears R in turn to F.

According to foundationalism, each piece of knowledge lies on a pyramid such as that shown in Figure 13.1. The nodes of such a pyramid (for a proposition P relative to a subject S and a time *t*) must obey the following requirements:

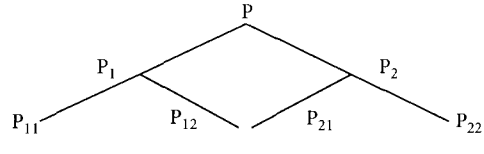

Figure 13.1

a. The set of all nodes that succeed (directly) any given node must serve jointly as a base that properly supports that node (for S at *t*).
b. Each node must be a proposition that S is justified in believing at *t*.
c. If a node is not self-evident (for S at *t*), it must have successors (that serve jointly as a base that properly supports that node).
d. Each branch of an epistemic pyramid must terminate.

For the foundationalist Descartes, for instance, each terminating node must be an indubitable proposition that S believes at *t* with no possibility of error. As for the nonterminal nodes, each of them represents inferential knowledge, derived by deduction from more basic beliefs.

Such radical foundationalism suffers from a fatal weakness that is twofold: (a) there are not so many perfectly obvious truths as Descartes thought; and (b) once we restrict ourselves to what is truly obvious in any given context, very little of one's supposed common sense knowledge can be proved on that basis. If we adhere to such radical foundationalism, therefore, we are just wrong in thinking we know so much.

Note that in citing such a "fatal weakness" of radical foundationalism, we favor particularism as against the methodism of Descartes and Hume. For we reject the methods or criteria of Descartes and Hume when we realize that they plunge us in a deep skepticism. If such criteria are incompatible with our enjoyment of the rich body of knowledge that we commonly take for granted, then as good particularists we hold on to the knowledge and reject the criteria.

If we reject radical foundationalism, however, what are we to put in its place? Here epistemology faces a dilemma that different epistemologists resolve differently. Some reject radical foundationalism but retain some more moderate form of foundationalism in favor of a radically different coherentism. Coherentism is associated with idealism – of both the German and the British variety – and has recently acquired new vigor and interest.

The coherentists reject the metaphor of the pyramid in favor of one that they owe to the positivist Neurath, according to whom our body of knowledge is a raft that floats free of any anchor or tie. Repairs must be made afloat, and though no part is untouchable, we must stand on some in order to replace or repair others. Not every part can go at once.

According to the new metaphor, what justifies a belief is not that it be an infallible belief with an indubitable object, nor that it have been proved deductively on such a basis, but that it cohere with a comprehensive system of beliefs.

4 A Coherentist Critique of Foundationalism

What reasons do coherentists offer for their total rejection of foundationalism? The argument that follows below summarizes much of what is alleged against foundationalism. But first we must distinguish between subjective states that incorporate a propositional attitude and those that do not. A propositional attitude is a mental state of someone with a proposition for its object: beliefs, hopes, and fears provide examples. By way of contrast, a headache does not incorporate any such attitude. One can of course be conscious of a headache, but the headache itself does not constitute or incorporate any attitude with a proposition for its object. With this distinction in the background, here is the antifoundationalist argument, which has two lemmas – a(iv) and b(iii) – and a principal conclusion.

a. (i) If a mental state incorporates a propositional attitude, then it does not give us direct contact with reality, e.g., with pure experience, unfiltered by concepts or beliefs.
   (ii) If a mental state does not give us direct contact with reality, then it provides no guarantee against error.
   (iii) If a mental state provides no guarantee against error, then it cannot serve as a foundation for knowledge.
   (iv) Therefore, if a mental state incorporates a propositional attitude, then it cannot serve as a foundation for knowledge.
b. (i) If a mental state does not incorporate a propositional attitude, then it is an enigma how such a state can provide support for any hypothesis, raising its credibility selectively by contrast with

its alternatives. (If the mental state has no conceptual or propositional content, then what logical relation can it possibly bear to any hypothesis? Belief in a hypothesis would be a propositional attitude with the hypothesis itself as object. How can one depend logically for such a belief on an experience with no propositional content?)

(ii) If a mental state has no propositional content and cannot provide logical support for any hypothesis, then it cannot serve as a foundation for knowledge.

(iii) Therefore, if a mental state does not incorporate a propositional attitude, then it cannot serve as a foundation for knowledge.

c. Every mental state either does or does not incorporate a propositional attitude.

d. Therefore, no mental state can serve as a foundation for knowledge. (From a(iv), b(iii), and c.)

According to the coherentist critic, foundationalism is run through by this dilemma. Let us take a closer look.[2]

In the first place, what reason is there to think, in accordance with premise b(i), that only propositional attitudes can give support to their own kind? Consider practices – e.g., broad policies or customs. Could not some person or group be justified in a practice because of its consequences: that is, could not the consequences of a practice make it a good practice? But among the consequences of a practice may surely be found, for example, a more just distribution of goods and less suffering than there would be under its alternatives. And neither the more just distribution nor the lower degree of suffering is a propositional attitude. This provides an example in which propositional attitudes (the intentions that sustain the practice) are justified by consequences that are not propositional attitudes. That being so, is it not conceivable that the justification of belief that matters for knowledge be analogous to the objective justification by consequences that we find in ethics?

Is it not possible, for instance, that a belief that there is something red before one be justified in part because it has its origins in one's visual experience of red when one looks at an apple in daylight? If we accept such examples, they show us a source of justification that serves as such without incorporating a propositional attitude.

As for premise a(iii), it is already under suspicion from our earlier exploration of premise b(i). A mental state M can be nonpropositional and hence not a candidate for so much as truth, much less infallibility, while it serves, in spite of that, as a foundation of knowledge. Leaving that aside, let us suppose that the relevant mental state is indeed propositional. Must it then be infallible in order to serve as a foundation of justification and knowledge? That is so far from being obvious that it seems more likely false when compared with an analogue in ethics. With respect to beliefs, we may distinguish between their being true and their being justified. Analogously, with respect to actions, we may distinguish between their being optimal (best of all alternatives, all things considered) and their being (subjectively) justified. In practical deliberation on alternatives for action, is it inconceivable that the most *eligible* alternative *not* be objectively the best, all things considered? Can there not be another alternative – perhaps a most repugnant one worth little if any consideration – that in point of fact would have a much better total set of consequences and would thus be better, all things considered? Take the physician attending to Frau Hitler at the birth of little Adolf. Is it not possible that if he had acted less morally, that would have proved better in the fullness of time? And if that is so in ethics, may not its likeness hold good in epistemology? Might there not be justified (reasonable, warranted) beliefs that are not even true, much less infallible? That seems to me not just a conceivable possibility, but indeed a familiar fact of everyday life, where observational beliefs too often prove illusory but no less reasonable for being false.

If the foregoing is on the right track, then the antifoundationalist is far astray. What has led him there?

As a diagnosis of the antifoundationalist argument before us, and more particularly of its second lemma, I would suggest that it rests on an Intellectualist Model of Justification.

According to such a model, the justification of belief (and psychological states generally) is parasitical on certain logical relations among propositions. For example, my belief (i) that the streets are wet, is justified by my pair of beliefs (ii) that it

is raining, and (iii) that if it is raining, the streets are wet. Thus we have a structure such as this:

B(Q) is justified by the fact that B(Q) is grounded on (B(P), B(P⊃Q)).

And according to an Intellectualist Model, this is parasitical on the fact that

P and (P⊃Q) together logically imply Q.

Concerning this attack on foundationalism I will argue (a) that it is useless to the coherentist, since if the antifoundationalist dilemma impales the foundationalist, a form of it can be turned against the coherentist to the same effect; (b) that the dilemma would be lethal not only to foundationalism and coherentism but also to the very possibility of substantive epistemology; and (c) that a form of it would have the same effect on normative ethics.

a. According to coherentism, what justifies a belief is its membership in a coherent and comprehensive set of beliefs. But whereas being grounded on B(P) and B(P⊂Q) is a property of a belief B(Q) that yields immediately the logical implication of Q and P and (P⊂Q) as the logical source of that property's justificatory power, the property of being a member of a coherent set is not one that immediately yields any such implication.

   It may be argued, nevertheless, (i) that the property of being a member of a coherent set would supervene in any actual instance on the property of being a member of a particular set *a* that is in fact coherent, and (ii) that this would enable us to preserve our Intellectualist Model, since (iii) the justification of the member belief B(Q) by its membership in *a* would then be parasitical on the logical relations among the beliefs in *a* which constitute the coherence of that set of beliefs, and (iv) the justification of B(Q) by the fact that it is part of a coherent set would then be *indirectly* parasitical on logical relations among propositions after all.

   But if such an indirect form of parasitism is allowed, then the experience of pain may perhaps be said to justify belief in its existence parasitically on the fact that P logically implies P! The Intellectualist Model seems either so trivial as to be dull, or else sharp enough to cut equally against both foundationalism and coherentism.

b. If (i) only propositional attitudes can justify such propositional attitudes as belief, and if (ii) to do so they must in turn be justified by yet other propositional attitudes, it seems clear that (iii) there is no hope of constructing a complete epistemology, one which would give us, in theory, an account of what the justification of any justified belief would supervene on. For (i) and (ii) would rule out the possibility of a finite regress of justification.

c. If only propositional attitudes can justify propositional attitudes, and if to do so they must in turn be justified by yet other propositional attitudes, it seems clear that there is no hope of constructing a complete normative ethics, one which would give us, in theory, an account of what the justification of any possible justified action would supervene upon. For the justification of an action presumably depends on the intentions it embodies and the justification of these, and here we are already within the net of propositional attitudes from which, for the Intellectualist, there is no escape.

It seems fair to conclude that our coherentist takes his antifoundationalist zeal too far. His antifoundationalist argument helps expose some valuable insights but falls short of its malicious intent. The foundationalist emerges showing no serious damage. Indeed, he now demands equal time for a positive brief in defense of his position.

5 The Regress Argument

a. The regress argument in epistemology concludes that we must countenance beliefs that are justified in the absence of justification by other beliefs. But it reaches that conclusion only by rejecting the possibility in principle of an infinite regress of justification. It thus opts

for foundational beliefs justified in some noninferential way by ruling out a chain or pyramid of justification that has justifiers, and justifiers of justifiers, and so on *without end.* One may well find this too short a route to foundationalism, however, and demand more compelling reasons for thus rejecting an infinite regress as vicious. We shall find indeed that it is not easy to meet this demand.

b. We have seen how even the most ordinary of everyday beliefs is the tip of an iceberg. A closer look below the surface reveals a complex structure that ramifies with no end in sight. Take again my belief that driving will be dangerous tonight, at the tip of an iceberg, (I), as presented in Figure 13.2. The immediate cause of my belief that driving will be hazardous tonight is the sound of raindrops on the windowpane. All but one or two members of the underlying iceberg are as far as they can be from my thoughts at the time. In what sense, then, do they form an iceberg whose tip breaks the calm surface of my consciousness?

Here I will assume that the members of (I) are beliefs of the subject, even if unconscious or subconscious, that causally buttress and thus justify his prediction about the driving conditions.

Can the iceberg extend without end? It may appear obvious that it cannot do so, and one may jump to the conclusion that any piece of knowledge must be ultimately founded on beliefs that are *not* (inferentially) justified or warranted by other beliefs. This is a doctrine of *epistemic foundationalism.*

Let us focus not so much on the *giving* of justification as on the *having* of it. *Can* there be a belief that is justified in part by other beliefs, some of which are in turn justified by yet other beliefs, and so on without end? Can there be an endless regress of justification?

c. There are several familiar objections to such a regress:

(i) *Objection*: "It is incompatible with human limitations. No human subject could harbor the required infinity of beliefs."
Reply: It is mere presumption to fathom with such assurance the depths of the mind, and especially its unconscious and dispositional depths. Besides, our object here is the nature of epistemic justification in itself and not only that of such justification as is accessible to humans. Our question is not whether humans could harbor an infinite iceberg of justification. Our question is rather whether *any* mind, no matter how deep, could do so. Or is it ruled out *in principle* by the very nature of justification?

(ii) *Objection*: "An infinite regress is indeed ruled out in principle, for if justification were thus infinite how could it possibly end?"
Reply: (i) If the end mentioned is *temporal*, then why must there be such an end? In the first place, the subject may be eternal. Even if he is not eternal, moreover, why must belief acquisition and justification occur seriatim? What precludes an infinite body of beliefs

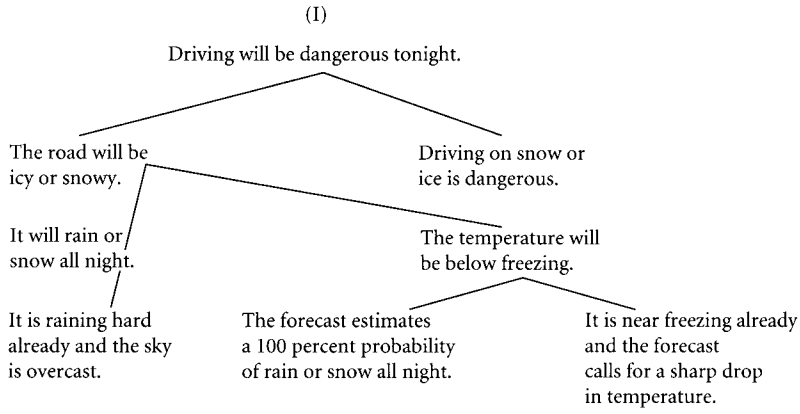

Figure 13.2

acquired at a single stroke? Human limitations may rule this out for humans, but we have yet to be shown that it is precluded in principle, by the very nature of justification. (ii) If the end mentioned is justificatory, on the other hand, then to ask how justification could possibly end is just to beg the question.

(iii) *Objection*: "Let us make two assumptions: first, that S's belief of q justifies his belief of p only if it works together with a justified belief on his part that q provides good evidence for p; and, second, that if S is to be justified in believing p on the basis of his belief of q and is to be justified in believing q on the basis of his belief of r, then S must be justified in believing that r provides good evidence for p via q. These assumptions imply that an actual regress of justification requires belief in an infinite proposition. Since no one (or at least no human) can believe an infinite proposition, no one (no human) can be a subject of such an actual regress."[3]

Reply: Neither of the two assumptions is beyond question, but even granting them both, it may still be doubted that the conclusion follows. It is true that each finitely complex belief of form "r provides good evidence for p via $q_1 \ldots q_n$" will *omit* how some members of the full infinite regress are epistemically tied to belief of p. But that seems irrelevant given the fact that for each member r of the regress, such that r is tied epistemically to belief of p, there *is* a finite belief of the required sort ("r provides good evidence for p via $q_1 \ldots q_n$") that ties the two together. Consequently there is no apparent reason to suppose – even granted the two assumptions – that an infinite regress will require a single belief in an infinite proposition, and not just an infinity of beliefs in increasingly complex finite propositions.

(iv) *Objection*: "But if it is allowed that justification extend infinitely, then it is too easy to justify any belief at all or too many beliefs altogether. Take, for instance, the belief that there are perfect numbers greater than 100. And suppose a mind powerful enough to believe every member of the following sequence:

(σ1) There is at least one perfect number > 100
There are at least two perfect numbers > 100
There are at least three perfect numbers > 100

If such a believer has no other belief about perfect numbers save the belief that a perfect number is a whole number equal to the sum of its whole factors, then surely he is *not* justified in believing that there are perfect numbers greater than 100. He is quite unjustified in believing any of the members of sequence (σ1), in spite of the fact that a challenge to any can be met easily by appeal to its successor. Thus it cannot be allowed after all that justification extend infinitely, and an infinite regress is ruled out."

Reply: We must distinguish between regresses of justification that are actual and those that are merely potential. The difference is *not* simply that an actual regress is composed of actual beliefs. For even if all members of the regress are actual beliefs, the regress may still be *merely potential* in the following sense: while it is true that *if* any member *were* justified then its predecessors *would* be, still none is in fact justified. Anyone with our series of beliefs about perfect numbers in the absence of any further relevant information on such numbers would presumably be the subject of such a merely potential justificatory regress.

(v) *Objection*: "But defenders of infinite justificatory regresses cannot distinguish thus between actual regresses and those that are merely potential. There is no real distinction to be drawn between the two. For if any regress ever justifies the belief at its head, then every regress must always do so. But obviously not every regress does so (as we have seen by examples), and hence no regress can do so."[4]

Reply: One can in fact distinguish between actual justificatory regresses and merely potential ones, and one can do so both abstractly and by examples.

What an actual regress has that a merely potential regress lacks is the property of containing only justified beliefs as members. What they both share is the property of containing no member without successors that would jointly justify it.

Recall our regress about perfect numbers greater than 100; i.e., there is at least one; there are at least two; there are at least three; and so on. Each member has a successor that would justify it, but no member is justified (in the absence of further information external to the regress). That is therefore a merely potential infinite regress. As for an actual regress, I see no compelling reason why someone (if not a human, then some more powerful mind) could not hold an infinite series of actually justified beliefs as follows:

($\sigma 2$) There is at least one even number
There are at least two even numbers
There are at least three even numbers

It may be that no one could be the subject of such a series of justified beliefs unless he had a proof that there is a denumerable infinity of even numbers. But even if that should be so, it would not take away the fact of the infinite regress of potential justifiers, each of which is actually justified, and hence it would not take away the fact of the actual endless regress of justification.

The objection under discussion is confused, moreover, on the nature of the issue before us. Our question is *not* whether there can be an infinite potential regress, each member of which would be justified by its successors, such that the belief at its head is justified in virtue of its position there, at the head of such a regress. The existence and even the possibility of a single such regress with a belief at its head that was *not* justified in virtue of its position there would of course settle that question in the negative. Our question is, rather, whether there can be an actual infinite regress of justification, and the fact that a belief at the head of a potential regress might still fail to be justified despite its position does *not* settle this question. For even if there can be a merely potential regress with an unjustified belief at its head, that leaves open the possibility of an infinite regress, each member of which is justified by its immediate successors working jointly, where every member of the regress is in addition actually justified.

6 The Relation of Justification and Foundationalist Strategy

The foregoing discussion is predicated on a simple conception of justification such that a set of beliefs β conditionally justifies (*would* justify) a belief X iff, necessarily, if all members of β are justified then X is also justified (if it exists). The fact that on such a conception of justification actual endless regresses – such as ($\sigma 2$) – seem quite possible blocks a straightforward regress argument in favor of foundations. For it shows that an actual infinite regress cannot be dismissed out of hand.

Perhaps the foundationalist could introduce some relation of justification – presumably more complex and yet to be explicated – with respect to which it could be argued more plausibly that an actual endless regress is out of the question.

There is, however, a more straightforward strategy open to the foundationalist. For he *need not* object to the possibility of an endless regress of justification. His essential creed is the more positive belief that every justified belief must be at the head of a terminating regress. Fortunately, to affirm the universal necessity of a terminating regress is *not* to deny the bare possibility of a nonterminating regress. For a single belief can trail at once regresses of both sorts: one terminating and one not. Thus the proof of the denumerably infinite cardinality of the set of evens may provide for a powerful enough intellect a *terminating* regress for each member of the *endless* series of justified beliefs:

($\sigma 2$) There is at least one even number
There are at least two even numbers
There are at least three even numbers

At the same time, it is obvious that each member of ($\sigma 2$) lies at the head of an actual endless regress of justification, on the assumption that each member is conditionally justified by its successor, which is in turn actually justified.

"Thank you so much," the foundationalist may sneer, "but I really do not need that kind of help. Nor do I need to be reminded of my essential creed, which I know as well as anyone. Indeed my rejection of endless regresses of justification is only a means of supporting my view that every justified belief must rest ultimately on foundations, on a terminating regress. You reject that strategy much too casually, in my view, but I will not object here. So we put that strategy aside. And now, my helpful friend, just what do we put in its place?"

Fair enough. How then could one show the need for foundations if an endless regress is not ruled out?

7 Two Levels of Foundationalism

a. We need to distinguish, first, between two forms of foundationalism: one *formal*, the other *substantive*. A type of *formal foundationalism* with respect to a normative or evaluative property ϕ is the view that the conditions (actual and possible) within which ϕ would apply can be specified in general, perhaps recursively. *Substantive foundationalism* is only a particular way of doing so, and coherentism is another.

Simpleminded hedonism is the view that:

(i) every instance of pleasure is good,
(ii) everything that causes something good is itself good, and
(iii) everything that is good is so in virtue of (i) or (ii) above.

Simpleminded hedonism is a type of formal foundationalism with respect to the good.

Classical foundationalism in epistemology is the view that:

(i) every infallible, indubitable belief is justified,
(ii) every belief deductively inferred from justified beliefs is itself justified, and
(iii) every belief that is justified is so in virtue of (i) or (ii) above.

Classical foundationalism is a type of formal foundationalism with respect to epistemic justification.

Both of the foregoing theories – simpleminded hedonism in ethics, and classical foundationalism in epistemology – are of course flawed. But they both remain examples of formal foundationalist theories.

b. One way of arguing in favor of formal foundationalism in epistemology is to formulate a convincing formal foundationalist theory of justification. But classical foundationalism in epistemology no longer has for many the attraction that it had for Descartes, nor has any other form of epistemic foundationalism won general acceptance. Indeed epistemic foundationalism has been generally abandoned, and its advocates have been put on the defensive by the writings of Wittgenstein, Quine, Sellars, Rescher, Aune, Harman, Lehrer, and others. It is lamentable that in our headlong rush away from foundationalism we have lost sight of the different types of foundationalism (formal vs. substantive) and of the different grades of each type. Too many of us now see it as a blur to be decried and avoided. Thus our present attempt to bring it all into better focus.

c. If we cannot argue from a generally accepted foundationalist theory, what reason is there to accept formal foundationalism? There is no reason to think that the conditions (actual and possible) within which an object is spherical are generally specifiable in nongeometric terms. Why should we think that the conditions (actual and possible) within which a belief is epistemically justified are generally specifiable in nonepistemic terms?

So far as I can see, the main reason for accepting formal foundationalism in the absence of an actual, convincing formal foundationalist theory is the very plausible idea that epistemic justification is subject to the supervenience that characterizes normative and evaluative properties generally. Thus, if a car is a good car, then any physical replica of that car must be just as good. If it is a good car in virtue of such properties as being economical, little prone to break down,

etc., then surely any exact replica would share all such properties and would thus be equally good. Similarly, if a belief is epistemically justified, it is presumably so in virtue of its character and its basis in perception, memory, or inference (if any). Thus any belief exactly like it in its character and its basis must be equally well justified. Epistemic justification is supervenient. The justification of a belief supervenes on such properties of it as its content and its basis (if any) in perception, memory, or inference. Such a doctrine of supervenience may itself be considered, with considerable justice, a grade of foundationalism. For it entails that every instance of justified belief is founded on a number of its nonepistemic properties, such as its having a certain basis in perception, memory, and inference, or the like.

But there are higher grades of foundationalism as well. There is, for instance, the doctrine that the conditions (actual and possible) within which a belief would be epistemically justified *can be specified* in general, perhaps recursively (and by reference to such notions as perception, memory, and inference).

A higher grade yet of formal foundationalism requires not only that the conditions for justified belief be specifiable, in general, but that they be specifiable by a simple, comprehensive theory.

d. Simpleminded hedonism is a formal foundationalist theory of the highest grade. If it is true, then in every possible world goodness supervenes on pleasure and causation in a way that is recursively specifiable by means of a very simple theory.

Classical foundationalism in epistemology is also a formal foundationalist theory of the highest grade. If it is true, then in every possible world epistemic justification supervenes on infallibility cum indubitability and deductive inference in a way that is recursively specifiable by means of a very simple theory.

Surprisingly enough, coherentism may also turn out to be formal foundationalism of the highest grade, provided only that the concept of coherence is itself both simple enough and free of any normative or evaluative admixture. Given these provisos, coherentism explains how epistemic justification supervenes on the nonepistemic in a theory of remarkable simplicity: a belief is justified if it has a place within a system of beliefs that is coherent and comprehensive.

It is a goal of ethics to explain how the ethical rightness of an action supervenes on what is not ethically evaluative or normative. Similarly, it is a goal of epistemology to explain how the epistemic justification of a belief supervenes on what is not epistemically evaluative or normative. If coherentism aims at this goal, that imposes restrictions on the notion of coherence, which must now be conceived innocent of epistemically evaluative or normative admixture. Its substance must therefore consist of such concepts as explanation, probability, and logical implication – with these conceived, in turn, innocent of normative or evaluative content.

e. We have found a surprising kinship between coherentism and substantive foundationalism, both of which turn out to be varieties of a deeper foundationalism. This deeper foundationalism is applicable to any normative or evaluative property ϕ, and it comes in three grades. The *first* or lowest is simply the supervenience of ϕ: the idea that whenever something has ϕ its having it is founded on certain others of its properties which fall into certain restricted sorts. The *second* is the explicable supervenience of ϕ: the idea that there are formulable principles that explain in quite general terms the conditions (actual and possible) within which ϕ applies. The *third* and highest is the easily explicable supervenience of ϕ: the idea that there is a *simple* theory that explains the conditions within which ϕ applies. We have found the coherentist and the substantive foundationalist sharing a primary goal: the development of a formal foundationalist theory of the highest grade. For they both want a simple theory that explains precisely how epistemic justification supervenes, in general, on the nonepistemic. This insight gives us an unusual viewpoint on some recent attacks against foundationalism. Let us now consider as an example a certain simple form of argument distilled from the recent antifoundationalist literature.[5]

8 Doxastic Ascent Arguments

Several attacks on foundationalism turn on a sort of "doxastic ascent" argument that calls for closer scrutiny.[6] Here are two examples:

A. A belief B is foundationally justified for S in virtue of having property F only if S is justified in believing (1) that most at least of his beliefs with property F are true, and (2) that B has property F. But this means that belief B is not foundational after all, and indeed that the very notion of (empirical) foundational belief is incoherent.

It is sometimes held, for example, that perceptual or observational beliefs are often justified through their origin in the exercise of one or more of our five senses in standard conditions of perception. The advocate of doxastic ascent would raise a vigorous protest, however, for in his view the mere fact of such sensory prompting is impotent to justify the belief prompted. Such prompting must be coupled with the further belief that one's senses work well in the circumstances, or the like. For we are dealing here with *knowledge*, which requires not blind faith but *reasoned* trust. But now surely the further belief about the reliability of one's senses itself cannot rest on blind faith but requires its own backing of reasons, and we are off on the regress.

B. A belief B of proposition P is foundationally justified for S only if S is justified in believing that there are no factors present that would cause him to make mistakes on the matter of the proposition P. But, again, this means that belief B is not foundational after all and indeed that the notion of (empirical) foundational belief is incoherent.

From the vantage point of formal foundationalism, neither of these arguments seems persuasive. In the first place, as we have seen, what makes a belief foundational (formally) is its having a property that is nonepistemic (not evaluative in the epistemic or cognitive mode), and does not involve inference from other beliefs, but guarantees, via a necessary principle, that the belief in question is justified. A belief B is made foundational by having some such nonepistemic property that yields its justification. Take my belief that I am in pain in a context where it is caused by my being in pain. The property that my belief then has, of being a self-attribution of pain caused by one's own pain is, let us suppose, a nonepistemic property that yields the justification of any belief that has it. So my belief that I am in pain is in that context foundationally justified. Along with my belief that I am in pain, however, there come other beliefs that are equally well justified, such as my belief that someone is in pain. Thus I am foundationally justified in believing that I am in pain only if I am justified in believing that someone is in pain. Those who object to foundationalism as in A or B above are hence mistaken in thinking that their premises would refute foundationalism. The fact is that they would not touch it. For a belief is no less foundationally justified for having its justification yoked to that of another closely related belief.

The advocate of arguments like A and B must apparently strengthen his premises. He must apparently claim that the beliefs whose justification is entailed by the foundationally justified status of belief B must in some sense function as a *necessary source* of the justification of B. And this would of course preclude giving B foundationally justified status. For if the *being justified* of those beliefs is an *essential* part of the source of the justification of B, then it is ruled out that there be a wholly *non-epistemic* source of B's justification.

That brings us to a second point about A and B, for it should now be clear that these cannot be selectively aimed at foundationalism. In particular, they seem neither more nor less valid objections to coherentism than to foundationalism, or so I will now argue about each of them in turn.

A′. A belief X is justified for S in virtue of membership in a coherent set only if S is justified in believing (1) that most at least of his beliefs with the property of thus cohering are true, and (2) that X has that property.

Any coherentist who accepts A seems bound to accept A′. For what could he possibly appeal to as a relevant difference? But A′ is a quicksand of endless depth. (How is he justified in believing A′ (1)?

Partly through justified belief that *it* coheres? And what would justify *this*? And so on...)

B′. A belief X is justified for S only if S is justified in believing that there are no factors present that would cause him to make mistakes on the subject matter of that belief.

Again, any coherentist who accepts B seems bound to accept B′. But this is just another road to the quicksand. (For S is justified in believing that there are no such factors only if... and so on.)

Why are such regresses vicious? The key is again, to my mind, the doctrine of supervenience. Such regresses are vicious because they would be logically incompatible with the supervenience of epistemic justification on such nonepistemic facts as the totality of a subject's beliefs, his cognitive and experiential history, and as many other nonepistemic facts as may seem at all relevant. The idea is that there is a set of such nonepistemic facts surrounding a justified belief such that no belief could possibly have been surrounded by those very facts without being justified. Advocates of A or B run afoul of such supervenience, since they are surely committed to the more general views derivable from either A or B by deleting "foundationally" from its first sentence. In each case the more general view would then preclude the possibility of supervenience, since it would entail that the source of justification *always* includes an *epistemic* component.

9 Coherentism and Substantive Foundationalism

a. The notions of coherentism and substantive foundationalism remain unexplicated. We have relied so far on our intuitive grasp of them. In this section we shall consider reasons for the view that substantive foundationalism is superior to coherentism. To assess these reasons, we need some more explicit account of the difference between the two.

By coherentism we shall mean any view according to which the ultimate sources of justification for any belief lie in relations among that belief and other beliefs of the subject: explanatory relations, perhaps, or relations of probability or logic.

According to substantive foundationalism, as it is to be understood here, there are ultimate sources of justification other than relations among beliefs. Traditionally these additional sources have pertained to the special content of the belief or its special relations to the subjective experience of the believer.

b. The view that justification is a matter of relations among beliefs is open to an objection from alternative coherent systems or detachment from reality, depending on one's perspective. From the latter perspective the body of beliefs is held constant and the surrounding world is allowed to vary, whereas from the former perspective it is the surrounding world that is held constant while the body of beliefs is allowed to vary. In either case, according to the coherentist, there could be no effect on the justification for any belief.

Let us sharpen the question before us as follows. Is there reason to think that there is at least one system B′, alternative to our actual system of beliefs B, such that B′ contains a belief X with the following properties:

(i) in our present nonbelief circumstances we would not be justified in having belief X even if we accepted along with that belief (as our total system of beliefs) the entire belief system B′ in which it is embedded (no matter how acceptance of B′ were brought about); and

(ii) that is so despite the fact that belief X coheres within B′ at least as fully as does some actual justified belief of ours within our actual belief system B (where the justification of that actual justified belief is alleged by the coherentist to derive solely from its coherence within our actual body of beliefs B).

The coherentist is vulnerable to counterexamples of this sort right at the surface of his body of beliefs, where we find beliefs with minimal coherence, whose detachment and replacement with contrary beliefs would have little effect on the coherence of the body. Thus take my belief that I have a headache when I do have a splitting headache, and let us suppose that this *does* cohere within my present body of beliefs. (Thus I have

no reason to doubt my present introspective beliefs, and so on. And if my belief does *not* cohere, so much the worse for coherentism, since my belief is surely justified.) Here then we have a perfectly justified or warranted belief. And yet such a belief may well have relevant relations of explanation, logic, or probability with at most a small set of other beliefs of mine at the time: say, that I am not free of headache, that I am in pain, that someone is in pain, and the like. If so, then an equally coherent alternative is not far to seek. Let everything remain constant, *including* the splitting headache, except for the following: replace the belief that I have a headache with the belief that I do *not* have a headache, the belief that I am in pain with the belief that I am *not* in pain, the belief that someone is in pain with the belief that someone is *not* in pain, and so on. I contend that my resulting hypothetical system of beliefs would cohere as fully as does my actual system of beliefs, and yet my hypothetical belief that I do *not* have a headache would not therefore be justified. What makes this difference concerning justification between my actual belief that I have a headache and the hypothetical belief that I am free of headache, each as coherent as the other within its own system, if not the actual splitting headache? But the headache is *not* itself a belief nor a relation among beliefs and is thus in no way constitutive of the internal coherence of my body of beliefs.

Some might be tempted to respond by alleging that one's belief about whether or not one has a headache is always *infallible*. But since we could devise similar examples for the various sensory modalities and propositional attitudes, the response given for the case of headache would have to be generalized. In effect, it would have to cover "peripheral" beliefs generally – beliefs at the periphery of one's body of beliefs, minimally coherent with the rest. These peripheral beliefs would all be said to be infallible. That is, again, a possible response, but it leads to a capitulation by the coherentist to the radical foundationalist on a crucial issue that has traditionally divided them: the infallibility of beliefs about one's own subjective states.

What is more, not all peripheral beliefs are about one's own subjective states. The direct realist is probably right that some beliefs about our surroundings are uninferred and yet justified. Consider my present belief that the table before me is oblong. This presumably coheres with such other beliefs of mine as that the table has the same shape as the piece of paper before me, which is oblong, and a different shape than the window frame here, which is square, and so on. So far as I can see, however, there is no insurmountable obstacle to replacing that whole set of coherent beliefs with an equally coherent set as follows: that the table before me is square, that the table has the same shape as the square window frame, and a different shape than the piece of paper, which is oblong, and so on. The important points are (a) that this replacement may be made without changing the rest of one's body of beliefs or any aspect of the world beyond, including one's present visual experience of something oblong, not square, as one looks at the table before one; and (b) that it is so, in part, because of the fact (c) that the subject need not have any beliefs about his present sensory experience.

Some might be tempted to respond by alleging that one's present experience is *self-intimating*, i.e., always necessarily taken note of and reflected in one's beliefs. Thus if anyone has visual experience of something oblong, then he believes that he has such experience. But this would involve a further important concession by the coherentist to the radical foundationalist, who would have been granted two of his most cherished doctrines: the infallibility of introspective belief and the self-intimation of experience.

10 The Foundationalist's Dilemma

The antifoundationalist zeal of recent years has left several forms of foundationalism standing. These all share the conviction that a belief can be justified not only by its coherence within a comprehensive system but also by an appropriate combination of observational content and origin in the use of the senses in standard conditions. What follows presents a dilemma for any foundationalism based on any such idea.

a. We may surely suppose that beings with observational mechanisms radically unlike ours might also have knowledge of their environment. (That seems possible even if the radical difference in observational mechanisms precludes overlap in substantive concepts and beliefs.)

Table 13.1

Human	*Extraterrestrial being*
Visual experience	ϕ experience
Experience of something red	Experience of something F
Belief that there is something red before one	Belief that there is something F before one

b. Let us suppose that there is such a being, for whom experience of type ϕ (of which we have no notion) has a role with respect to his beliefs of type ϕ analogous to the role that our visual experience has with respect to our visual beliefs. Thus we might have a schema such as that in Table 13.1.

c. It is often recognized that our visual experience intervenes in two ways with respect to our visual beliefs: as cause and as justification. But these are not wholly independent. Presumably, the justification of the belief that something here is red derives at least in part from the fact that it originates in a visual experience of something red that takes place in normal circumstances.

d. Analogously, the extraterrestrial belief that something here has the property of being F might be justified partly by the fact that it originates in a ϕ experience of something F that takes place in normal circumstances.

e. A simple question presents the foundationalist's dilemma: regarding the epistemic principle that underlies our justification for believing that something here is red on the basis of our visual experience of something red, is it proposed as a fundamental principle or as a derived generalization? Let us compare the famous Principle of Utility of value theory, according to which it is best for that to happen which, of all the possible alternatives in the circumstances, would bring with it into the world the greatest balance of pleasure over pain, joy over sorrow, happiness over unhappiness, content over discontent, or the like. Upon this fundamental principle one may then base various generalizations, rules of thumb, and maxims of public health, nutrition, legislation, etiquette, hygiene, and so on. But these are all then derived generalizations which rest for their validity on the fundamental principle. Similarly, one may also ask, with respect to the generalizations advanced by our foundationalist, whether these are proposed as fundamental principles or as derived maxims or the like. This sets him face to face with a dilemma, each of whose alternatives is problematic. If his proposals are meant to have the status of secondary or derived maxims, for instance, then it would be quite unphilosophical to stop there. Let us turn, therefore, to the other alternative.

f. On reflection it seems rather unlikely that epistemic principles for the justification of observational beliefs by their origin in sensory experience could have a status more fundamental than that of derived generalizations. For by granting such principles fundamental status we would open the door to a multitude of equally basic principles with no unifying factor. There would be some for vision, some for hearing, etc., without even mentioning the corresponding extraterrestrial principles.

g. It may appear that there is after all an idea, however, that unifies our multitude of principles. For they all involve sensory experience and sensible characteristics. But what is a sensible characteristic? Aristotle's answer appeals to examples: colors, shapes, sounds, and so on. Such a notion might enable us to unify perceptual epistemic principles under some more fundamental principle such as the following.

> If σ is a sensible characteristic, then the belief that there is something with σ before one is (prima facie) justified if it is based on a visual experience of something with σ in conditions that are normal with respect to σ.

h. There are at least two difficulties with such a suggestion, however, and neither one

can be brushed aside easily. First, it is not clear that we can have a viable notion of sensible characteristics on the basis of examples so diverse as colors, shapes, tones, odors, and so on. Second, the authority of such a principle apparently derives from contingent circumstances concerning the reliability of beliefs prompted by sensory experiences of certain sorts. According to the foundationalist, our visual beliefs are justified by their origin in our visual experience or the like. Would such beliefs be equally well justified in a world where beliefs with such an origin were nearly always false?

i. In addition, finally, even if we had a viable notion of such characteristics, it is not obvious that fundamental knowledge of reality would have to derive causally or otherwise from sensory experience of such characteristics. How could one impose reasonable limits on extraterrestrial mechanisms for noninferential acquisition of beliefs? Is it not possible that such mechanisms need not always function through sensory experience of any sort? Would such beings necessarily be denied any knowledge of the surroundings and indeed of any contingent spatiotemporal fact? Let us suppose them to possess a complex system of true beliefs concerning their surroundings, the structures below the surface of things, exact details of history and geography, all constituted by concepts none of which corresponds to any of our sensible characteristics. What then? Is it not possible that their basic beliefs should all concern fields of force, waves, mathematical structures, and numerical assignments to variables in several dimensions? This is no doubt an exotic notion, but even so it still seems conceivable. And if it is in fact possible, what then shall we say of the noninferential beliefs of such beings? Would we have to concede the existence of special epistemic principles that can validate their noninferential beliefs? Would it not be preferable to formulate more abstract principles that can cover both human and extraterrestrial foundations? If such more abstract principles are in fact accessible, then the less general principles that define the human foundations and those that define the extraterrestrial foundations are both derived principles whose validity depends on that of the more abstract principles. In this the human and extraterrestrial epistemic principles would resemble rules of good nutrition for an infant and an adult. The infant's rules would of course be quite unlike those valid for the adult. But both would still be based on a more fundamental principle that postulates the ends of well-being and good health. What more fundamental principles might support both human and extraterrestrial knowledge in the way that those concerning good health and well-being support rules of nutrition for both the infant and adult?

11 Reliabilism: An Ethics of Moral Virtues and an Epistemology of Intellectual Virtues

In what sense is the doctor attending Frau Hitler justified in performing an action that brings with it far less value than one of its accessible alternatives? According to one promising idea, the key is to be found in the rules that he embodies through stable dispositions. His action is the result of certain stable virtues, and there are no equally virtuous alternative *dispositions* that, given his cognitive limitations, he might have embodied with equal or better total consequences, and that would have led him to infanticide in the circumstances. The important move for our purpose is the stratification of justification. Primary justification attaches to virtues and other dispositions, to stable dispositions to act, through their greater contribution of value when compared with alternatives. Secondary justification attaches to particular acts in virtue of their source in virtues or other such justified dispositions.

The same strategy may also prove fruitful in epistemology. Here primary justification would apply to *intellectual* virtues, to stable dispositions for belief acquisition, through their greater contribution toward getting us to the truth. Secondary

justification would then attach to particular beliefs in virtue of their source in intellectual virtues or other such justified dispositions.[7]

That raises parallel questions for ethics and epistemology. We need to consider more carefully the concept of a virtue and the distinction between moral and intellectual virtues. In epistemology, there is reason to think that the most useful and illuminating notion of intellectual virtue will prove broader than our tradition would suggest and must give due weight not only to the subject and his intrinsic nature but also to his environment and to his epistemic community. This is a large topic, however, to which I hope some of us will turn with more space, and insight, than I can now command.

Summary

1. *Two assumptions*: (A1) that for a belief to constitute knowledge it must be (a) true and (b) justified; and (A2) that the justification relevant to whether or not one knows is a sort of epistemic or theoretical justification to be distinguished from its practical counterpart.
2. *Knowledge and criteria.* Particularism is distinguished from methodism: the first gives priority to particular examples of knowledge over general methods or criteria, whereas the second reverses that order. The methodism of Descartes leads him to an elaborate dogmatism whereas that of Hume leads him to a very simple skepticism. The particularist is, of course, antiskeptical on principle.
3. *Two metaphors: the raft and the pyramid.* For the foundationalist every piece of knowledge stands at the apex of a pyramid that rests on stable and secure foundations whose stability and security do not derive from the upper stories or sections. For the coherentist a body of knowledge is a free-floating raft every plank of which helps directly or indirectly to keep all the others in place, and no plank of which would retain its status with no help from the others.
4. *A coherentist critique of foundationalism.* No mental state can provide a foundation for empirical knowledge. For if such a state is propositional, then it is fallible and hence no secure foundation. But if it is *not* propositional, then how can it possibly serve as a foundation for belief? How can one infer or justify anything on the basis of a state that, having no propositional content, must be logically dumb? An analogy with ethics suggests a reason to reject this dilemma. Other reasons are also advanced and discussed.
5. *The regress argument.* In defending his position, the foundationalist often attempts to rule out the very possibility of an infinite regress of justification (which leads him to the necessity for a foundation). Some of his arguments to that end are examined.
6. *The relation of justification and foundationalist strategy.* An alternative foundationalist strategy is exposed, one that does not require ruling out the possibility of an infinite regress of justification.
7. *Two levels of foundationalism.* Substantive foundationalism is distinguished from formal foundationalism, three grades of which are exposed: first, the supervenience of epistemic justification; second, its explicable supervenience; and, third, its supervenience explicable by means of a simple theory. There turns out to be a surprising kinship between coherentism and substantive foundationalism, both of which aim at a formal foundationalism of the highest grade, at a theory of the greatest simplicity that explains how epistemic justification supervenes on nonepistemic factors.
8. *Doxastic ascent arguments.* The distinction between formal and substantive foundationalism provides an unusual viewpoint on some recent attacks against foundationalism. We consider doxastic ascent arguments as an example.
9. *Coherentism and substantive foundationalism.* It is argued that substantive foundationalism is superior, since coherentism is unable to account adequately for the epistemic status of beliefs at the "periphery" of a body of beliefs.
10. *The foundationalist's dilemma.* All foundationalism based on sense experience is subject to a fatal dilemma.
11. *Reliabilism.* An alternative to foundationalism of sense experience is sketched.

Appendix[8]

What one is rationally justified in believing obviously depends on the data in one's possession. But what data one has can depend on how much and how well one investigates. Consider, therefore, the following possibility. What if A is rationally justified in believing x given his body of data D_1 whereas B is not rationally justified in believing x given his body of data D_2, where D_2 includes D_1 but is much more extensive as a result of A's irresponsible negligence and B's commendable thoroughness? The present account might unfortunately grant A knowledge while denying it to B, for A's neglect so far has no bearing on any epistemic pyramid.

We have considered a situation where someone lacks knowledge owing to his misuse of his cognitive equipment, either by letting it idle when it should be functioning or by busily employing it dysfunctionally. Another situation where someone lacks knowledge despite having rationally justified correct belief might be called the Magoo situation – where S lacks adequate equipment to begin with (relative to the question in hand: whether p).[9] It is because of this type of lack that despite his extensive experience with cable cars, Mr Magoo does not know that his cable car will arrive safely when, unknown to him, bombs are raining all around it. Of course, even if you have less than 20–20 vision you can still know that there is an elephant in front of you when you see one there. So not just any defect will make your equipment inadequate for a judgment on the question whether p. I would venture that it must be a defect that prevents you from acquiring information that (i) a normal inquirer in the epistemic community would acquire in that situation *and* (ii) makes a difference to what you can reasonably conclude on the question whether p (or at least to how reasonably you can draw the conclusion).

The possibility of inadequate cognitive equipment requires a further and more striking departure from the traditional conception of knowledge. Despite having warranted correct belief, someone may lack knowledge owing to his neglectful data-collection. There lack of knowledge could be traced back to epistemic irresponsibility, to substandard performance blamed on the investigator. In the present example, blame is out of place. By hypothesis, Magoo conducts impeccable "inquiry" both in arriving at his data and on the basis of his data. But he still falls short of knowledge, despite his warranted, correct belief. His shortcoming is substandard equipment, for which we may suppose him to be blameless. Hence something other than epistemic justification or correct belief can help determine what one knows or does not know. Even if one correctly believes that p with full rational justification and free of irrational or neglectful unbelief, one may still be in no position to know, because of faulty cognitive equipment.

In all of the foregoing cases, someone misses or is liable to miss available information which may be highly relevant and important and may make a difference to what he can conclude on the question in hand. In each case, moreover, he seems culpable or discredited in some sense: he would seem less reliable than otherwise for his role in any such case. But there appear to be situations where again someone misses available information with no culpability *or* discredit. Harman gives an example where S reads in a newspaper that some famous person has been assassinated, but does not read the next edition, where all reports of the assassination are denied by highly authoritative and trustworthy people. If practically the whole country reads the next edition and people don't know what to believe, does S alone know of the assassination, provided the next edition is in fact a pack of lies?[10] I suppose we would be inclined to say that he does not know (especially if had he read the next edition, *he* would not have known what to believe). But what if only two or three people get a chance to read the next edition before it is recalled by the newspaper? Should we now say that out of millions who read the first story and mourn the loved leader not one knows of his death? I suppose we would be inclined to say that the fake edition and the few deceived by it make no difference concerning what everybody else knows. It seems plausible to conclude that knowledge has a further "social aspect," that it cannot depend on one's missing or blinking what is generally known.

Our departures from the traditional conception of knowledge put in relief the relativity of

knowledge to an epistemic community. This is brought out most prominently by the requirement that inquirers have at least *normal* cognitive equipment (e.g., normal perceptual apparatus, where that is relevant). But our new requirement – that inquirers not lack or blink generally known relevant information – also brings out the relativity. A vacationer in the woods may know that *p* well enough for an average vacationer, but he won't have the kind of knowledge his guide has. A guide would scornfully deny that the tenderfoot really knows that *p*. Relative to the epistemic community of guides (for that area) the tenderfoot lacks relevant generally known information, and misses relevant data that the average guide would grasp in the circumstances.

These departures from the traditional account may make better sense if we reflect that the honorific term "knowledgeable" is to be applied only to those who are reliable sources of information, surely an important category for a language-using, social species.

We have now taken note of two types of situation where correct, fully warranted belief falls short of knowledge owing to no neglect or faulty reasoning or false belief. Despite commendable thoroughness and impeccable reasoning unspoiled by falsehood, one may still fail to be "in a position to know," owing either to faulty cognitive equipment or to missed generally known information. I am not suggesting that these are the only ways to be out of position to know. I have no complete list of epistemic principles describing ways of arriving at a position to know or of being blocked from such a position. My suggestion is only that there are such principles, and that in any case we must go beyond the traditional emphasis by epistemologists on warrant and reasoning as determinants of knowledge. Despite the importance of warranted correct belief in determining what we know, the Gettier examples show that it is not alone enough to guarantee knowledge. What is more, warranted correct belief supported by reasoning *unspoiled by falsehood* seems immune to Gettier examples, but it still falls short of knowledge, as we have seen.

My conclusion is that to understand knowledge we must enrich our traditional repertoire of epistemic concepts with the notion of *being in a position to know (from the point of view of a K, e.g., a human being)*. Thus a proposition is evident (from the point of view of a *K*) to a subject only if *both* he is rationally justified in believing it *and* he is in a position to know (from the *K* point of view) whether it is true. It may *be* (and not just appear) evident to Magoo from *his* point of view that he will reach the other side safely, but it seems wrong to say of Magoo as he steps into the cable car with bombs raining all around that it *is* quite evident to him that he will arrive safely. It seems wrong *for whom* to say this? For one of us, naturally; that is, for a normal human from *his* point of view. And since a normal human could not help seeing and hearing the bombs, from the human point of view Magoo is not in a position to know that he will arrive safely, inasmuch as he is missing relevant information that a normal human would gather in the circumstances. Hence Magoo does not have *human* knowledge that he will arrive safely, for it is not evident to him from the human point of view that he will so arrive.

Consider this account:

(A) S knows that *p* iff
- (a) it is true that *p*;
- (b) S believes that *p*; and
- (c) there is a non-defective epistemic pyramid for S and the proposition that *p*.

Every node of such a pyramid must be true *and* evident. And for every node *n* that has successors, the successors must serve as grounds that give the subject S rational warrant for believing *n*. What now seems too narrow about this account emerges with the explanation of what a pyramid of knowledge is, and of what the evident is. For in this explanation what is evident to S is identified with what S is rationally justified in believing. But it now seems plain that for *x* to be evident to S, *two* conditions must be satisfied: (i) that S be rationally justified in believing *x*, and (ii) that S be in a position to know whether *x* is true. And we must also take note of the relativity of knowledge to an epistemic community. Let us therefore replace (A) with the following:

(B) S knows (from the *K* point of view) that *p* iff
- (a) it is true that *p*;
- (b) S believes that *p*; and
- (c) there is a non-defective epistemic pyramid (from the *K* point of view) for S and the proposition that *p*.

Every node of such a pyramid must now be true and evident from the *K* point of view.

Normally when epistemologists discuss knowledge (of the colors and shapes of surrounding objects, of one's own or one's neighbor's mental states, and so on), they plainly do so from the *human* point of view. But other points of view are possible even in ordinary conversation. The expert/layman distinction is replicable in many different contexts, and with each replication we have a new epistemically relevant distinction in points of view, with expert knowledge on one side and layman knowledge on the other.

Neither Magoo nor the newspaper reader who alone has not seen the new edition is in a position to know (from the human point of view) about the relevant subject matter. Thus we can understand their ignorance and, by parity of reasoning, the ignorance of all those who are out of position to know that *p* because they lack either adequate cognitive equipment or relevant information that is generally known to those who have taken an epistemic stand on the question whether *p* (where to suspend judgment *is* to take an epistemic stand, whereas to be totally oblivious to the matter is not).

What it is for S's belief that *p* to be fully grounded has been explained by means of our epistemic pyramids. That answer points in the right direction, but it can be made more precise: e.g., by clarifying the grounding relation. Moreover, we have found that a fully grounded correct belief is not necessarily knowledge, and this for at least two reasons: (i) it may rest directly or indirectly on some false ground, and (ii) the believer may not be in a position to know.

We have tried to allow for these possibilities by broadening epistemic pyramids, by making room for our new epistemic notion of being-in-a-position-to-know, and by noting that to support knowledge epistemic pyramids must be non-defective, i.e., must contain no false nodes. But pyramids are objectionable for other reasons as well: (i) they may mislead by suggesting that terminal nodes provide a "foundation" in one or another undesirable sense, or by suggesting that terminal nodes must come first in time, so that one may later build on them; (ii) more seriously, there is an unacceptable vagueness in the very idea of such a pyramid, which derives mainly from the vagueness of the "grounding" relation in terms of which pyramids were defined. What follows is an attempt to solve these problems by switching pyramids upside down into trees.

Let us emphasize, however, that this will not commit one to a picture of knowledge according to which there is a bedrock of self-evident propositions. It is perfectly consistent with the present theory that part of what makes *any* proposition evident is its coherence with a network of mutually supporting propositions. Since there is bound to be a multitude of such coherent networks, however, a non-arbitrary narrowing of the field must be supported by something other than coherence.

We turn finally to an account (C) according to which S knows that *p* provided that both (a) S correctly believes that *p*,[11] and (b) there is a set of propositions that fully and non-defectively renders it evident to S that *p* (where a set "non-defectively renders it evident to S that *p*" if and only if it does so without attributing to S any false belief).[12]

Supposing this account correct, every bit of knowledge has a tree like that shown in Figure 13.3. Note that each node of such a tree is a proposition. Thus the "root" node is the-proposition-that-p_1, and the first terminal node (from the left) the- proposition-that-p_{111}.[13]

There is an important difference between these trees and our earlier pyramids. Except for terminal nodes, every node of a tree is an epistemic proposition, whereas not a single node of a pyramid need be epistemic at all. Pyramids display propositions that are evident to A (*not* propositions that such and such other propositions are evident to S), and they also show which propositions ground (for S) any proposition for which S has grounds. Trees display true epistemic propositions concerning S and they also show what "makes these propositions true" *via* epistemic principles. A tree must do this for *every* epistemic proposition that constitutes one of its nodes. That is to say, trees contain no epistemic terminal nodes. It is in this sense that trees provide *complete* epistemic explanations of the truth of their root nodes.

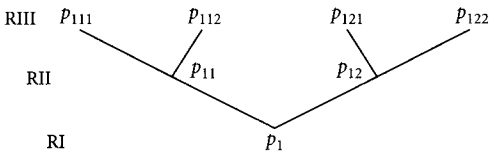

Figure 13.3

Notes

1 But Descartes's methodism was at most partial. James Van Cleve has supplied the materials for a convincing argument that the way out of the Cartesian circle is through a particularism of basic knowledge. See James Van Cleve, "Foundationalism, Epistemic Principles, and the Cartesian Circle," *The Philosophical Review* 88 (1979), pp. 55–91 and E. Sosa and J. Kim (eds), *Epistemology: An Anthology* (Malden, MA; Oxford: Blackwell, 2000), pp. 242–60. But this is, of course, compatible with methodism on inferred knowledge. Whether Descartes subscribed to such methodism is hard (perhaps impossible) to determine, since in the end he makes room for all the kinds of knowledge required by particularism. But his language when he introduces the method of hyperbolic doubt, and the order in which he proceeds, suggest that he did subscribe to such methodism.

2 Cf. Laurence BonJour, "The Coherence Theory of Empirical Knowledge," *Philosophical Studies* 30, pp. 281–312, and, especially, Michael Williams, *Groundless Belief* (Oxford: Blackwell, 1977); and Bonjour, "Can Empirical Knowledge Have a Foundation?," this vol., ch. 10.

3 Cf. Richard Foley, "Inferential Justification and the Infinite Regress," *American Philosophical Quarterly* 15 (1978), pp. 311–16.

4 Cf. John Post, "Infinite Regresses of Justification and of Explanation," *Philosophical Studies* 34 (1980).

5 The argument of this whole section is developed in greater detail in my paper "The Foundations of Foundationalism," *Nous* 14, pp. 547–65.

6 For some examples of the influence of doxastic ascent arguments, see Wilfrid Sellars's writing in epistemology: "Empiricism and the Philosophy of Mind," in *Science, Perception and Reality* (London: Routledge and Kegan Paul, 1963), esp. section VIII, and particularly p. 168. Also I. T. Oakley, "An Argument for Skepticism Concerning Justified Belief," *American Philosophical Quarterly* 13 (1976), pp. 221–8; and BonJour, "Can Empirical Knowledge Have a Foundation?."

7 This puts in a more traditional perspective the contemporary effort to develop a "causal theory of knowing." From our viewpoint, this effort is better understood not as an attempt to *define* propositional knowledge, but as an attempt to formulate fundamental principles of justification.

Cf. the work of D. Armstrong, *Belief, Truth and Knowledge* (Cambridge: Cambridge University Press, 1973), and that of F. Dretske, A. Goldman, and M. Swain, whose relevant already published work is included in G. Pappas and M. Swain (eds), *Essays on Knowledge and Justification* (Ithaca and London, 1978). But the theory is still under development by Goldman and by Swain, who have reached general conclusions about it similar to those suggested here, though not necessarily – so far as I know – for the same reasons or in the same overall context.

8 From "How Do You Know?" *American Philosophical Quarterly* 11, 2 (1974), pp. 113–22.

9 The Magoo situation is the situation of that unfortunate nearsighted and hearing-impaired cartoon character who fortunately escapes disaster at every turn.

10 Gilbert Harman, "Induction," in Marshall Swain (ed.), *Induction, Acceptance and Rational Belief* (Dordrecht: Reidel, 1970), esp. Sect. IV, pp. 95–7.

11 Whether knowledge entails belief at all is of course a vexed question of long standing, but there is no room for it here. A helpful and interesting discussion is found in Keith Lehrer's "Belief and Knowledge," *Philosophical Review* 77 (1968), pp. 491–9.

12 In what follows, the relativity of knowledge to an epistemic community is left implicit, as it normally is in ordinary thought and speech.

13 Strictly speaking, what we have here is obviously a *partial tree schema*. For convenience, however, I speak of trees even when I mean partial tree schemata. Also, it should not be thought that every tree must have exactly three ranks (RI, RII, and RIII). On the contrary, a tree may have any number of ranks, so long as it has more than one.