

Laurence Bonjour, *The Structure of Empirical Knowledge*

A basic problem for foundationalism

The fundamental concept of moderate foundationalism, as of empirical foundationalism generally, is the concept of a basic empirical belief. It is by appeal to basic beliefs that the threat of an infinite regress is to be avoided and empirical knowledge given a secure foundation. But a new problem now arises: how can there be any empirical beliefs which are thus basic? For although this has often been overlooked, the very idea of an epistemically basic empirical belief is more than a little paradoxical. On what basis is such a belief supposed to be justified, once any appeal to further empirical premises is ruled out? Chisholm's theological analogy, cited earlier, is most appropriate: a basic empirical belief is in effect an epistemological unmoved (or self-moved) mover. It is able to confer justification on other beliefs, but, in spite of being empirical and thus contingent, apparently has no need to have justification conferred on it. But is such a status any easier to understand in epistemology than it is in theology? How can a contingent, empirical belief impart epistemic "motion" to other empirical beliefs unless it is itself in "motion"? (Or, even more paradoxically, how can such a belief epistemically "move" itself?) Where does the noninferential justification for basic empirical beliefs come from?

This difficulty may be developed a bit by appealing to the account of the general concept of epistemic justification which was presented [earlier]. I argued there that the fundamental role which the requirement of epistemic justification serves in the overall rationale of the concept of knowledge is that of a *means* to truth; and accordingly that a basic constraint on any account of the standards of justification for empirical knowledge is that there be good reasons for thinking that following those standards is at least likely to lead to truth. Thus if basic beliefs are to provide a secure foundation for empirical knowledge, if inference from them is to be the sole basis upon which other empirical beliefs are justified, then that feature, whatever it may be, by virtue of which a particular belief qualifies as basic must also constitute a good reason for thinking that the belief is true. If this were not so, moderate foundationalism would be unacceptable as an account of epistemic justification.

This crucial point may be formulated a bit more precisely, as follows. If we let

Laurence Bonjour, *The Structure of Empirical Knowledge* (Cambridge: Harvard University Press, 1985).

ϕ represent the feature or characteristic, whatever it may be, which distinguishes basic empirical beliefs from other empirical beliefs, then in an acceptable foundationalist account a particular empirical belief B could qualify as basic only if the premises of the following justificatory argument were adequately justified:

- (1) B has feature ϕ .
 - (2) Beliefs having feature ϕ are highly likely to be true.
- Therefore, B is highly likely to be true.

If B is to actually *be* basic, then presumably premise (1) would have to be true as well, but I am concerned here only with what would have to be so for it to be reasonable to *accept* B as basic and use it to justify other beliefs.

Clearly it is possible that at least one of the two premises of the argument might be justifiable on a purely *a priori* basis, depending on the particular choice of ϕ . It does not seem possible, however, that *both* premises might be thus justifiable. B is after all, *ex hypothesi*, an empirical belief, and it is hard to see how a particular empirical belief could be justified on a purely *a priori* basis. Thus we may conclude, at least provisionally, that for any acceptable moderate foundationalist account, at least one of the two premises of the appropriate justifying argument will itself be empirical.

The other issue to be considered is whether, in order for B to be justified for a particular person A (at a particular time), it is necessary, not merely that a justification along the above lines exist in the abstract, but also that A himself be in cognitive possession of that justification, that is, that he believe the appropriate premises of forms (1) and (2) and that these beliefs be justified *for him*. [Earlier,] I argued tentatively that such cognitive possession by the person in question is indeed necessary, on the grounds that *he* cannot be epistemically responsible in accepting the belief unless *he himself* has access to the justification; for otherwise, *he* has no reason for thinking that the belief is at all likely to be true. No reason for questioning this claim has so far emerged.

But if all this is correct, we get the disturbing result that B is not basic after all, since its justification depends on that of at least one other empirical belief. It would follow that moderate foundationalism is untenable as a solution to the regress problem—and an analogous argument would show weak foundationalism to be similarly untenable.

It will be helpful in the subsequent discussion to have available a slightly more explicit statement of this basic antifoundationalist argument:

- (1) Suppose that there are *basic empirical beliefs*, that is, empirical beliefs (a) which are epistemically justified, and (b) whose justification does not depend on that of any further empirical beliefs.
- (2) For a belief to be epistemically justified requires that there be a reason why it is likely to be true.
- (3) For a belief to be epistemically justified for a particular person requires that this person be himself in cognitive possession of such a reason.
- (4) The only way to be in cognitive possession of such a reason is to believe *with justification* the premises from which it follows that the belief is likely to be true.

- (5) The premises of such a justifying argument for an empirical belief cannot be entirely *a priori*; at least one such premise must be empirical.

Therefore, the justification of a supposed basic empirical belief must depend on the justification of at least one other empirical belief, contradicting (1); it follows that there can be no basic empirical beliefs. . . .

The elements of coherentism

The very idea of a coherence theory

In light of the failure of foundationalism, it is time to look again at the apparent alternatives with regard to the structure of empirical justification which were distinguished in the discussion of the epistemic regress problem. If the regress of empirical justification does not terminate in basic empirical beliefs, then it must either (1) terminate in unjustified beliefs, (2) go on infinitely (without circularity), or (3) circle back upon itself in some way. As discussed earlier, alternative (1) is clearly a version of skepticism and as such may reasonably be set aside until all other alternatives have been seen to fail. Alternative (2) may also be a version of skepticism, though this is less clear. But the more basic problem with alternative (2) is that no one has ever succeeded in amplifying it into a developed position (indeed, it is not clear that anyone has even attempted to do so); nor do I see any plausible way in which this might be done. Failing any such elaboration which meets the objections tentatively developed earlier, alternative (2) may also reasonably be set aside. This then leaves alternative (3) as apparently the only remaining possibility for a nonskeptical account of empirical knowledge.

We are thus led to a reconsideration of the possibility of a coherence theory of empirical knowledge. If there is no way to justify empirical beliefs apart from an appeal to other justified empirical beliefs, and if an infinite sequence of distinct justified beliefs is ruled out, then the presumably finite system of justified empirical beliefs can only be justified from within, by virtue of the relations of its component beliefs to each other—if, that is, it is justified at all. And the idea of *coherence* should for the moment be taken merely to indicate whatever property (or complex set of properties) is requisite for the justification of such a system of beliefs.

Obviously this rather flimsy argument by elimination carries very little weight by itself. The analogous argument in the case of foundationalism lead to an untenable result; and that failure, when added to the already substantial problems with coherence theories which were briefly noted above, makes the present version even less compelling. At best it may motivate a more open-minded consideration of coherence theories than they have usually been accorded, such theories having usually been treated merely as dialectical bogeymen and only rarely as serious epistemological alternatives.

It will be useful to begin by specifying more precisely just what sort of coherence theory is at issue here. In the first place, our concern is with coherence theories of *empirical justification* and not coherence theories of *truth*; the latter hold that truth is to be simply *identified* with coherence (presumably coherence

with some specified sort of system). The classical idealist proponents of coherence theories in fact generally held views of both these sorts and unfortunately failed for the most part to distinguish clearly between them. And this sort of confusion is abetted by views which use the phrase “theory of truth” to mean a theory of the *criteria* of truth, that is, a theory of the standards or rules which should be appealed to in deciding or judging whether or not something is true; if, as is virtually always the case, such a theory is meant to be an account of the criteria which can be used to arrive at a rational or warranted judgment of truth or falsity, then a coherence theory of truth in that sense would seem to be indiscernible from what is here called a coherence theory of justification, and quite distinct from a coherence theory of the very nature or meaning of truth. But if such confusions are avoided, it is clear that coherence theories of empirical justification are both distinct from and initially a good deal more plausible than coherence theories of empirical truth and moreover that there is no manifest absurdity in combining a coherence theory of justification with a *correspondence* theory of truth. Whether such a combination is in the end dialectically defensible is of course a further issue and one to which I will return in the final chapter of this book. . . .

Linear versus nonlinear justification

The initial problem is whether and how a coherence theory constitutes even a *prima facie* solution to the epistemic regress problem. Having rejected both foundationalism and the actual-infinite-regress position, a coherentist must hold, as we have seen, that the regress of empirical justification moves in a circle—or, more plausibly, some more complicated and multidimensional variety of closed curve. But this response to the regress will seem obviously and utterly inadequate to one who approaches the issue with foundationalist preconceptions. Surely, his argument will go, such a resort to circularity fails to solve or even adequately confront the problem. Each step in the regress is a justificatory argument whose premises must be justified *before* they can confer justification on the conclusion. To say that the regress moves in a circle is to say that at some point one (or more) of the beliefs which figured earlier as a conclusion is now appealed to as a justifying premise. And this response, far from solving the problem, seems to yield the patently absurd result that the justification of such a belief depends, indirectly but still quite inescapably, on *its own* logically prior justification: it cannot be justified unless it is already justified. And thus, assuming that it is not justified in some independent way, neither it nor anything which depends upon it can be genuinely justified. Since empirical justification is always ultimately circular in this way according to coherence theories, there can on such a view be in the end no empirical justification and no empirical knowledge.

The crucial, though tacit, assumption which underlies this seemingly devastating line of argument is the idea that inferential justification is essentially *linear* in character, that it involves a one-dimensional sequence of beliefs, ordered by the relation of epistemic priority, along which epistemic justification is passed from the earlier to the later beliefs in the sequence via connections of inference. It is just this linear conception of justification which generates the regress problem in the

first place. So long as it remains unchallenged, the idea that justification moves in a circle will seem obviously untenable, and only moderate or strong foundationalism will be left as an alternative: even weak foundationalism cannot accept a purely linear view of justification, since its initially credible beliefs are not sufficiently justified on that basis alone to serve as linear first premises for everything else. Thus the primary coherentist response to the regress problem cannot be merely the idea that justification moves in a circle, for this would be quite futile by itself; rather such a position must repudiate the linear conception of justification in its entirety.

But what is the alternative? What might a nonlinear conception of justification amount to? As suggested briefly [earlier], the main idea is that inferential justification, despite its linear appearance, is essentially systematic or holistic in character: beliefs are justified by being inferentially related to other beliefs in the overall context of a coherent system.

The best way to clarify this view is to distinguish two importantly different levels at which issues of empirical justification can be raised. The epistemic issue on a particular occasion will usually be merely the justification of a single empirical belief, or small set of such beliefs, within the context of a cognitive system whose overall justification is (more or less) taken for granted; we may call this the *local* level of justification. But it is also possible, at least in principle, to raise the issue of the overall justification of the entire system of empirical beliefs; we may call this the *global* level of justification. For the sort of coherence theory which will be developed here—and indeed, I would argue, for any comprehensive, non-skeptical epistemology—it is the issue of justification as it arises at the latter, global, level which is in the final analysis decisive for the determination of empirical justification in general. This tends to be obscured in practice, I suggest, because it is only issues of the former, local, sort which tend to be explicitly raised in actual cases. (Indeed, it may well be that completely global issues are never in fact raised outside the context of explicitly epistemological discussion; but I cannot see that this in any way shows that there is something illegitimate about them.)

It is at the local level of justification that inferential justification *appears* linear. A given justificandum belief is shown to be justified by citing other premise-beliefs from which it correctly follows via some acceptable pattern of inference. Such premise-beliefs may themselves be challenged, of course, with justification being offered for them in the same fashion. But there is no serious danger of an infinite regress at this level, since the justification of the overall system of empirical beliefs, and thus of most of its constituent beliefs, is *ex hypothesi* not at issue. One quickly reaches premise-beliefs which are dialectically acceptable in that particular context and which can thus function there rather like the foundationalist's basic beliefs. (But these *contextually basic beliefs*, as they might be called, are unlikely to be only or even primarily beliefs which would be classified as basic by any plausible version of foundationalism.)

If, on the other hand, no dialectically acceptable stopping point were reached, if the new premise-beliefs offered as justification continued to be challenged in turn, then (according to the sort of coherence theory with which I am concerned) the epistemic dialogue would if ideally continued eventually circle back upon itself, giving the appearance of a linear regress and in effect challenging the entire system of empirical beliefs. At this global level, however, the previously harmless

illusion of linearity becomes a serious mistake. According to the envisaged coherence theory, the relation between the various particular beliefs is correctly to be conceived, not as one of linear dependence, but rather as one of mutual or reciprocal support. There is no ultimate relation of epistemic priority among the members of such a system and consequently no basis for a true regress. Rather the component beliefs of such a coherent system will ideally be so related that each can be justified in terms of the others, with the direction of argument on a particular occasion of local justification depending on which belief (or set of beliefs) has actually been challenged in that particular situation. And hence, a coherence theory will claim, the apparent circle of justification is not in fact vicious *because it is not genuinely a circle*: the justification of a particular empirical belief finally depends, not on other particular beliefs as the linear conception of justification would have it, but instead on the overall system and its coherence.

According to this conception, the fully explicit justification of a particular empirical belief would involve four distinct main steps or stages of argument, as follows:

- (1) The inferability of that particular belief from other particular beliefs and further relations among particular empirical beliefs.
- (2) The coherence of the overall system of empirical beliefs.
- (3) The justification of the overall system of empirical beliefs.
- (4) The justification of the particular belief in question, by virtue of its membership in the system.

The claim of a coherence theory of empirical justification is that each of these steps depends on the ones which precede it. It is the neglecting of steps (2) and (3), the ones pertaining explicitly to the overall cognitive system, that lends plausibility to the linear conception of justification and thus generates the regress problem. And this is a very seductive mistake: since the very same inferential connections between particular empirical beliefs are involved in both step (1) and step (4), and since the issues involved in the intervening steps are very rarely (if ever) raised in practical contexts, it becomes much too easy to conflate steps (1) and (4), thus leaving out any explicit reference to the cognitive system and its coherence. The picture which results from such an omission is vastly more simple; but the price of this simplicity, according to coherence theories, is a radical distortion of the very concept of epistemic justification—and also, in the end, skepticism or something tantamount to it. . . .

Coherentist observation: an example

Consider then the following example of (putative) observational knowledge: As I sit at my desk (or so I believe), I come to have the belief, among very many others, that there is a red book on the desk. In fact, of course, the content of the belief is a good deal more precise and specific than the formulation just given would suggest: I do not believe simply that there is a red book on the desk, but rather that there is a book of a certain approximate size, of an approximately rectangular shape, which is a certain fairly specific shade of red, and so on. But what matters

for the moment is that I do not *infer* that there is a red book on the desk, nor does the belief result from any other sort of deliberative or ratiocinative process, whether explicit or implicit. Rather it simply occurs to me, “strikes me,” in a manner which is both involuntary and quite coercive; such a belief is, I will say, *cognitively spontaneous*. It is cognitive spontaneity which marks the belief as putatively observational, as what Sellars calls a “language-entry transition,” in a way which can be recognized from within the system of beliefs.

At first glance, such a belief represents as clear a paradigm of an observational belief, indeed of observational knowledge, as one could want. How then is it justified? It is reasonably obvious what the various foundationalist views which were examined earlier and found wanting would say, but what might our envisaged coherentist account of observation offer as an alternative? How might the justification of such a belief, considered as an observation, depend on coherence with or inferability from other beliefs in my overall system of beliefs, on the availability of something like a justificatory argument?

There are several obvious but crucial facts (or at least things which I believe to be facts) concerning the belief and its context which can plausibly serve as the premises of a justificatory argument. Presumably these are things that I know, but what matters for the moment is that I believe them—and that these further beliefs are themselves justified in some manner or other.

First, the belief in question is a cognitively spontaneous belief of a certain, reasonably definite kind K_1 , which we may specify, somewhat misleadingly, by saying that it is a visual belief about the color and general classification of a “medium-sized physical object.” The reason that this is apt to be misleading is that the term “visual” suggests a classification in terms of causal etiology, whereas what is intended here is a classification concerned only with the intrinsic character and content of the belief, however it may in fact have been caused. Thus hallucinatory or dream beliefs of the right sort could qualify as visual in this sense, despite having been caused in some way having no connection at all with the physiological machinery of vision. We might better describe such beliefs as “putatively visual” or “apparently visual,” but I will not bother with this terminological refinement here.

Second, the conditions of observation are of a specifiable sort C_1 : the lighting is good, I am reasonably close to the apparent location of the object, my eyes are functioning normally, and so on. It is common to speak of “standard conditions,” but these may vary substantially for different sorts of cases; it will thus be less confusing to assume an actual listing of the conditions, though I will not attempt to give a complete one here.

Third, it is a true law of nature concerning me and a large, though indefinite class of relevantly similar observers (where a rough specification of an appropriate sort of observers can be taken to be part of the specified conditions) that our cognitively spontaneous beliefs of that kind in conditions of the sort specified are highly reliable, that is, very likely to be true.

Since I believe all of these things, I am in a position to offer the following justificatory argument for the original belief:

- (1) I have a cognitively spontaneous belief of kind K_1 that there is a red book on the desk.

- (2) Conditions C_1 obtain.
 - (3) Cognitively spontaneous visual beliefs of kind K_1 in conditions C_1 are very likely to be true.
- Therefore, my belief that there is a red book on the desk is very likely to be true.
- Therefore, (probably) there is a red book on the desk.

Obviously this is very far from the end of the matter: if my belief is to be genuinely justified by appeal to this argument, the premises of the argument must themselves be justified; and if the resulting account of observation is to be genuinely coherentist, these further justifications must also make no appeal to basic beliefs. . . . For the moment, the point is that the justification of my original belief is, on this account, not somehow intrinsic or primitive, as would be the case for versions of foundationalism like Quinton's, but is rather dependent on the background and context provided by my other beliefs. This is the basic claim which a coherentist account of observation must make for *all* varieties of observation.

Answers to objections

Answers to standard objections (I) and (II)

The coherentist account of observation and introspection offered in the previous chapter provides the last of the main ingredients needed for the formulation (as opposed to the metajustification) of a coherentist account of empirical justification. In the present chapter I will explore the shape of such a theory in more detail, by considering whether and how it can meet various objections. I begin in this section with a reconsideration of the first two of the standard objections to coherence theories which were formulated [earlier]; this will also yield a significant modification in the theory itself. The second section will then formulate and attempt to answer a number of additional objections which arise in connection with the view in question, following which the final section will summarize the overall position which results and touch briefly on the justification of memory knowledge.

It will prove convenient to consider first objection (II), which alleges that empirical justification, as understood by a coherentist, involves no *input* from the extratheoretic world. In light of the discussion of observation, we are already in a position to see that at least part of this objection is mistaken. It need not be true, as the objection alleges, that coherentist justification is purely a matter of the internal relations within the system of beliefs. For if the system in question contains beliefs to the effect that recognizable kinds of cognitively spontaneous beliefs are likely to be true, and if beliefs of these kinds indeed occur, then such beliefs will be at least provisionally justified in a way which does not depend at all on the relation between their assertive content and the rest of the system. They can thus constitute input in at least the minimal sense of being new elements of the system which are not merely derived inferentially from the earlier elements. And such beliefs need not merely augment the system but may also force the alteration or abandonment of parts of it: either because the (putative) observational belief is directly inconsistent with one or more of the previous beliefs in

the system or because such alteration will, in light of the new beliefs, enhance the overall coherence of the system. Of course the observational beliefs could themselves be rejected as a result of such conflict, though if this is done very often, the law which specifies the degree of reliability of that particular sort of observational belief will also have to be revised.

Thus any new observational belief which conflicts with other parts of the system forces a choice between at least two alternative ways of revising the system. The primary basis for making this choice is the relative coherence of the alternatives, though there is another important constraint, of a rather different sort, which will be mentioned momentarily. In this way a coherence theory can allow for a system of beliefs to be tested against the results of (putative) observation and revised accordingly.

There are, however, two important issues with respect to the foregoing suggestion which need to be discussed. First, though such beliefs may constitute input in the minimal sense just specified, is there any reason to think that they genuinely constitute input in the full sense involved in the objection, that is, input from the extratheoretic world? This question can, indeed must, be discussed on two different levels. On an empirical level, operating within the cognitive system, the standard explanation given for the occurrence of such beliefs is that they are *caused* in regular ways by the world; and moreover, it is very hard to think of any alternative explanation which could be offered at this level for the existence of significant numbers of cognitively spontaneous beliefs which are at least largely in agreement with each other. Thus such beliefs will normally be at least *claimed* within the system to constitute extratheoretic input. Of course it can still be asked whether there is any reason to think that such a claim is true; but this is merely a specific case of the general issue of whether coherentist justification is truth-conducive. Thus a complete answer to the input problem will, not surprisingly, depend on the outcome of that later discussion.

The second issue, of more immediate concern, is whether a coherence theory of empirical justification, while perhaps allowing in the way just indicated for the *possibility* of input into the system of beliefs, does not also permit there to be a system of justified empirical beliefs which lacks such input. For suppose that a particular system of beliefs simply fails to attribute a sufficient degree of reliability to enough kinds of cognitively spontaneous beliefs to yield a significant degree of input (or alternatively fails to attribute reliability to those introspective beliefs which are essential for the reliable recognition of other kinds of reliable spontaneous beliefs). One might arbitrarily construct a system of beliefs with this feature; or alternatively, it might be produced gradually (and perhaps unintentionally) if conflicts between putative observations and other beliefs in the system are always settled by rejecting the observations. Such a system would fail to have any effective input from outside the system. But there seems to be no reason why it might not still possess the highest possible degree of coherence and hence be epistemically justified according to the coherentist account offered so far. And this is surely a mistaken, even absurd result.

This point is, I believe, essentially sound. What it shows is that any adequate account of empirical knowledge must *require* putative input into the cognitive system, not merely allow for the possibility of such input. For, as was already argued in the initial statement of objection (II), without input of some sort any

agreement which happened to exist between the cognitive system and the world could only be accidental and hence not something which one could have any good reason to expect. Thus, as a straightforward consequence of the idea that epistemic justification must be truth-conducive, a coherence theory of empirical justification must require that in order for the beliefs of a cognitive system to be even candidates for empirical justification, that system must contain laws attributing a high degree of reliability to a reasonable variety of cognitively spontaneous beliefs (including in particular those kinds of introspective beliefs which are required for the recognition of other cognitively spontaneous beliefs).

This requirement, which I will refer to as the *Observation Requirement*, is obviously quite vague, and I can see no way to make it very much more precise without going into vastly more detail than is possible here. The underlying idea is that any claim in the system which is not justified *a priori* should in principle be capable of being observationally checked, either directly or indirectly, and thereby either confirmed or refuted. But whether or not this is so in a given system depends not only on the modes of observation available in that system, but also on the inferential interconnectedness of the system. In a fairly tight-knit system, the Observation Requirement could thus be interpreted less stringently than would be necessary in a looser system.

Notice that the Observation Requirement does *not* stipulate that the cognitively spontaneous beliefs to which reliability is attributed must actually *be* reliable, even as judged from within the system. Nor does it place any restriction on the sort of taxonomy which can be employed in specifying particular classes of such beliefs. Obviously it is part of the background concept of observation that observational beliefs are reliable and also at least implicitly that observational beliefs will fall into something like natural kinds with each kind having a distinctive causal etiology. But these conditions need not be built into the Observation Requirement, since failure to satisfy them will virtually guarantee that the system will not both remain coherent and continue to satisfy the Observation Requirement as stated, at least not in the long run. To attribute reliability to beliefs which are not in fact reliable or to lump together beliefs of very different sorts (which will be affected by different sorts of conditions) is almost certain to lead to eventual incoherence. The Observation Requirement should, however, be understood to include the requirement, common to all adequate theories of knowledge, that a user of the system must make a reasonable effort to seek out relevant, possibly conflicting observations, if his beliefs are to be justified.

Thus understood, the Observation Requirement effectively guarantees that a cognitive system which satisfies it will receive at least apparent input from the world and hence that empirical justification will not depend merely on the internal relations of a static belief system; it thus provides the basic answer to objection (II).

It is important to understand clearly the status of the Observation Requirement within a coherentist position. The need for the requirement is *a priori*: it is, for reasons already indicated in the original discussion of objection (II), an *a priori* truth that empirical knowledge of an independent world is not possible without input from that world; and it also seems to be true *a priori*, in light of my earlier discussion of foundationalism, that such input can only be understood in terms of something very close to Sellars's idea of token credibility which does not

derive from type credibility and hence in terms of cognitively spontaneous beliefs which are justified, at least in part, in virtue of that status. Hence, according to a coherence theory, it is an *a priori* truth that a cognitive system must attribute reliability to some members of the general class of cognitively spontaneous beliefs, to the extent indicated, *if* it is to contain empirical knowledge. But for a given system, it is *not an a priori* truth that the antecedent of this conditional is satisfied and hence also not an *a priori* truth that its consequent is satisfied—or even that it epistemically ought to be satisfied. Whether any varieties of cognitively spontaneous beliefs are in fact reliable and hence should be recognized as such is an empirical issue to be decided, purely on the basis of coherence, within the cognitive system. It is logically conceivable, relative to a particular system, that no variety of cognitively spontaneous belief is in fact sufficiently reliable and hence that this system will be unable to satisfy the Observation Requirement in the long run while remaining coherent. The Observation Requirement says not that such a situation could occur, but only that if it did occur, there would in consequence be no empirical justification and no empirical knowledge.

Thus the Observation Requirement, as it functions within a coherentist position, might be described, perhaps a bit ponderously, as a regulative metaprinciple, as opposed to a first-level epistemic principle. It does not impinge directly on issues of empirical justification; these are decided entirely by appeal to coherence. Rather the Observation Requirement provides a partial basis for *categorizing* or *classifying* the results yielded by such a system. This is one difference between a coherence theory of the present sort and that version of weak foundationalism which attributes some degree of initial credibility to all cognitively spontaneous beliefs. According to such a foundationalist view, it is true *prior* to any appeal to coherence that cognitively spontaneous beliefs have this minimal degree of credibility—for which no adequate justification is or ever can be offered. Whereas for a coherence theory, *all* epistemic justification of empirical beliefs depends on coherence.

What then would be the status of contingent and superficially empirical beliefs belonging to a coherent system of beliefs which violates the Observation Requirement? I suggest that they be thought of as analogous to beliefs—or at least belief-like states—which are a product of sheer imagination or which are the mental correlate of literary fiction. It is a consequence of the holism which is part and parcel of a coherence theory that the distinction between genuine empirical description and these other categories of thought or discourse is not to be drawn at the level of particular beliefs or statements but only at the level of systems. And the empirical thrust of a cognitive system is precisely the implicit claim that its component beliefs will agree, in general at least, with those classes of cognitively spontaneous beliefs which it holds to be reliable; while one who presents or regards a given body of propositions as purely imaginative or fictional commits himself to no such claim. (Thus the Observation Requirement might be viewed as a kind of rough analogue of the old positivist verifiability criterion of empirical meaningfulness, transmuted so as to apply to systems of beliefs rather than to isolated beliefs or statements.)

We are now also in a position to offer an answer to objection (I), the alternative coherent systems objection. But once it is clear that a coherence theory can allow for, indeed insist upon, the possibility that a cognitive system which is

coherent at one time may be rendered incoherent, and thereby in need of revision, by subsequent observational input, this objection needs some major reformulation. If it is to be interesting, the objection cannot be merely that *at a given time* there may be many equally coherent but incompatible systems between which a coherence theory provides no basis for decision. This claim is surely correct but does not constitute an objection to coherence theories, since an analogous claim would also hold for virtually any imaginable theory of knowledge, including all of the standard foundationalist views: on *any* account of the standards of epistemic justification, it is quite possible, even likely, that there will be competing sets of empirical claims which at a particular time are tied for the status of most justified and between which those standards offer *at that time* no basis for decision. This is neither alarming nor particularly surprising. The most that it seems reasonable to expect of an epistemological account is that it make it possible for such ties to be broken *in the long run*.

Thus if it is to constitute a serious objection to a coherence theory of the sort in question here, objection (I) must be interpreted to mean that even in the long run and with the continued impact of (putative) observational beliefs, there will always be multiple, equally coherent empirical systems between which a coherence theory will be unable to decide. But once the possibility of observational input is appreciated, it is no longer clear why this claim should be accepted, or at least why it is thought to be any more plausible in relation to a coherence theory than it is in relation to other theories of knowledge. The basic rationale for the original version of the objection was that alternative coherent systems could, at least in principle, be constructed arbitrarily. But such an arbitrarily constructed system will not in general satisfy the Observation Requirement; and if one should be so constructed as to initially satisfy that requirement, there is no reason to think that it would remain coherent as (putative) observations accumulate, even if it were coherent in the beginning. Thus the possibility of arbitrary invention seems to provide no real support for the envisaged objection.

One useful way to put this point is to say that a coherence theory which incorporates the indicated conception of observation bases justification not on the *static* coherence of a system of beliefs considered in the abstract but rather on the *dynamic* coherence of an ongoing system of beliefs which someone actually accepts. Only such an actually functioning system can contain cognitively spontaneous beliefs and thereby satisfy the Observation Requirement. For this reason, the possibility of arbitrarily constructing a coherent system in the abstract has no bearing on such a theory.

Once the possibility of arbitrary invention is set aside, is there any other reason for thinking that the possibility of alternative coherent systems is a serious problem for this sort of coherence theory? I can think of only one further way of pressing such an objection. According to a coherence theory of the sort in question, the classification of a given sort of cognitively spontaneous belief as reliable and hence as a species of observation is not in any way an *a priori* matter but rather depends entirely on the extent to which such a classification yields a maximally coherent system. But suppose that relative to a given person's cognitive system there are two disjoint classes of cognitively spontaneous beliefs, such that: if the beliefs in one class are classified as observational, one system results and remains coherent in the long run; while if the beliefs in the second class are

classified as observational, a different, incompatible system results which is equally coherent and remains equally coherent in the long run; whereas if the beliefs in both classes are classified as observational, a system with a much lower degree of coherence, too low to meet the requirement for justification, results. (There could be more than two such classes, but I will neglect this possibility for the sake of simplicity.) A coherence theory seems to provide no basis for choosing between these two coherent systems. And this might not be so for some versions of foundationalism, depending on just what kinds of cognitively spontaneous beliefs are involved.

Is such a situation a genuine possibility? Could it perhaps be produced by a Cartesian demon, if not in some more ordinary way? The issue is extremely difficult, and I have been unable to devise any really compelling argument in either direction. But there is at least one consideration to be noted. For the situation to work as described, it must be the case that the cognitively spontaneous beliefs in each of the two classes are, when taken separately, strongly in agreement with each other and quite coherent. But then the internal agreement and coherence of these two classes of beliefs are facts which must be explained by *any* total view which such a person might adopt, on pain of serious anomaly and hence greatly reduced coherence. As already briefly suggested, the obvious explanation of the internal agreement and coherence of such a class of beliefs is that it is caused in such a way as to genuinely reflect an objective reality. But if such an explanation is ruled out, as it must be for one of these two classes of beliefs by either of the two cognitive systems in question, then some alternative explanation must be found. And thus for the choice between the two cognitive systems to be genuinely symmetrical in the way supposed by the objection, each would have to have such an alternative, reasonably satisfactory explanation of this sort for the agreement and coherence of the observation beliefs of the other, and the two explanations would have to be equally good (other things being equal). And while such a situation *may* still be a possibility, I can see no reason to think that it is likely enough to constitute a serious objection to our proposed coherence theory.

We have, in any case, obviously come very far from the original version of objection (I). Instead of the claim that there will *always* be indefinitely many equally coherent and incompatible cognitive systems, between which a coherence theory provides no basis for decision, we have now the claim that there *might possibly* be two (or, an even more questionable possibility, more than two) such systems between which a coherence theory could not decide (but for which some foundationalist views *might* provide a basis for decision). This is a very weak objection, if indeed it is still an objection at all.

Thus the first two of the standard and supposedly fatal objections to coherence theories have little real force against a version of coherentism which incorporates the proposed account of observation. This does not mean, of course, that such a position is finally defensible. There remains the third of the standard objections, the problem of truth. But before considering that objection, there are several other serious objections which need to be formulated and assessed.

Some further objections

The first three objections concern the proper classification of the view presented: whether it is genuinely a version of coherentism, and its relation to foundationalism.

Objection 1. The view presented is not genuinely a version of coherentism. Rather it is a version of weak foundationalism in which the foundational beliefs are the person's metabeliefs about the composition of his own system of beliefs, that is, those beliefs specified by the Doxastic Presumption.

Reply. It must be conceded that there is *something* to this objection. Indeed, I have insisted at several points that one's reflective grasp of the composition of one's own system of beliefs provides an essential starting point for this version, or indeed for any plausible version, of coherentism. No nonexternalist appeal to coherence is possible without a grasp of the system of beliefs relative to which coherence is to be judged. This grasp may be, as I have suggested, in part defeasible, but it is not dispensable. And there can be no real objection to characterizing the central role which the metabeliefs that make up this grasp of one's own system of beliefs play for a coherentist position by saying that they constitute the *foundation* of empirical knowledge for such a view, so long as it is clearly understood that "foundation" here does not carry with it the implications which it would possess within a standard foundationalist view. For no claim is being made that these metabeliefs possess any sort of intrinsic or independent justification or warrant of any kind (nor would such a claim be defensible in light of the earlier antifoundationalist arguments). Rather the approximate correctness of these beliefs is an essential presupposition for coherentist justification, and both such justification itself and any resulting claim of likelihood of truth must be understood as relativized to this presupposition. In this respect, then, the present view is fundamentally different from weak foundationalism in a way which makes it only confusing to assimilate the two, in spite of the admitted parallels between them.

Objection 2. The view presented is not genuinely a version of coherentism. It is a version of weak foundationalism in which the initially credible foundational beliefs are just the cognitively spontaneous beliefs. Such beliefs must be regarded as having some small, defeasible degree of justification if coherence with them is to confer justification on anything else; and the effect of the Observation Requirement is to confer on them just such a status.

Reply. Though not entirely without merit, this objection has rather less to be said for it than the preceding one. Obviously the status of cognitively spontaneous beliefs is very special for a view of the sort in question, and obviously that status is conferred at least in part by the Observation Requirement. Moreover, it would be possible to formulate a version of weak foundationalism, or something very close to weak foundationalism, in which cognitively spontaneous beliefs were accorded some degree of initial or independent warrant, and such a view would have fairly close structural similarities to the version of coherentism suggested here.

But in spite of this, the main claims made in the objection are mistaken. First, it is simply not necessary in order for such a view to yield justification to suppose that cognitively spontaneous beliefs have some degree of initial or independent

credibility. One way to see this is to consider a parallel example taken, surprisingly enough, from C.I. Lewis. Lewis's account of memory knowledge is a version of weak foundationalism: memory beliefs are claimed to have some antecedent degree of warrant simply by virtue of being memory beliefs, and this is then amplified by appeal to coherence (which Lewis, as we have seen, calls "congruence"). In arguing for his account, Lewis considers the example of "relatively unreliable witnesses who independently tell the same circumstantial story." The point of the example is that: "For any one of these reports, taken singly, the extent to which it confirms what is reported may be slight . . . 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; the story that any one false witness might tell being one out of so very large a number of equally possible choices." And he adds that this result would still follow even if one of the witnesses were to tell a different story. 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 individual 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). And by the same token, so long as apparently cognitively spontaneous beliefs are genuinely independent of each other, their agreement will eventually generate credibility, without the need for any initial degree of warrant.

Secondly, there is no reason why the Observation Requirement should be regarded as in fact conferring such an initial degree of warrant on cognitively spontaneous beliefs. The main point is that it is quite consistent with the Observation Requirement, as explained above, that no cognitively spontaneous belief of any kind might turn out to be warranted: this would be so, for example, if no class of such beliefs turned out to be in internal agreement to any significant degree. But such a result would not seem to be possible for a weak foundationalist view, according to which the largest consistent (or coherent?) class of basic beliefs will seemingly have to be justified to some degree, even if perhaps not enough to satisfy the requirement for knowledge. There is thus a quite fundamental distinction between the two views.

QUESTIONS

- 1 What is meant by a "basic belief"?
- 2 According to BonJour, what is required in order for a belief to be "justified"?
- 3 According to BonJour, what four steps are involved in the justification of an empirical belief?
- 4 What is a cognitively spontaneous belief?
- 5 Are cognitively spontaneous beliefs also *basic* beliefs, according to BonJour? (Explain.)