5 The argument of this whole section is developed in greater detail in my paper "The Foundations of Foundationalism," *Nous* 14 (1980): 547–65.

6 For some examples of the influence of doxastic ascent arguments, see Wilfrid Sellars's writing in epistemology: e.g., "Empiricism and the Philosophy of Mind," in *Science, Perception and Reality* (London: Routledge & Kegan Paul, 1963), ch. 5, especially section VIII, and particularly p. 168. Also I. T. Oakley, "An Argument for Skepticism Concerning Justified Beliefs," *American Philosophical Quarterly* 13 (1976): 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. D. Armstrong, *Belief, Truth and Knowledge* (Cambridge, 1973); and that of F. Dretske, A. Goldman, and M. Swain, whose relevant already published work is included in *Essays on Knowledge and Justification*, ed. G. Pappas and M. Swain (Ithaca and London, 1978). But the theory is still under development by Goldman and 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 The main ideas in this essay were first presented in a seminar of 1976–7 at the University of Texas. I am grateful to those who made that seminar a valuable stimulus.

13 The Elements of Coherentism

Laurence BonJour

1 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 [in an earlier section]. 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.[1] 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.

Second, it is also worth emphasizing at the outset that I am concerned here only with coherence theories which purport to provide a response to skepticism. My view thus differs from those of several recent coherence theorists, most notably Michael Williams but also, to a lesser extent, Gilbert Harman and Keith Lehrer, who depart from foundationalism not only in their account of the struc-

ture of empirical justification but also with regard to the goals or purposes of an epistemological theory, by holding that such a theory need not attempt to provide a "global" account of justification or to answer "global" varieties of skepticism.

Third, the dialectical motive for coherentism depends heavily on the unacceptability of the externalist position. It is thus crucially important that a coherentist view itself avoid tacitly slipping into a nonfoundationalist version of externalism. If coherentism is to be even a dialectically interesting alternative, the coherentist justfication must, in principle at least, be accessible to the believer himself.

The aim of this chapter is to begin the task of formulating a coherence theory which satisfied the foregoing structures by, first, considering in detail some of the main ingredients of such a view, including the idea of nonlinear or holistic justification, the concept of coherence itself, and the presumption concerning one's grasp of one's own system of beliefs; and, second, elaborating the leading objections which such a position must face. The upshot of the chapter will be the hardly surprising conclusion that a central, very likely decisive, issue with respect to coherence theories is whether they can somehow make room for a viable concept of *observation*.

2 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? Briefly, 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, nonskeptical 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.[2] 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 justification 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 the particular situation. And hence, a coherence theory wll 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.

How tenable is such a nonlinear conception of empirical justification? Of the three crucial transitions represented in this obviously quite schematic account,

only the third, from step (3) to step (4), is reasonably unproblematic, depending as it does on the inferential relations that obtain between the justificandum belief and the other beliefs of the system; in effect it is this transition which is made when an inferential justification is offered in an ordinary context of local justification, with the other steps being taken for granted. But the other two transitions are highly problematic, and the issues that they raise are crucial for understanding and assessing the very conception of a coherence theory.

The transition from step (1) to step (2), from the relations obtaining between particular beliefs to the attribution of the holistic property of coherence to the empirical system as a whole, is rendered problematic by the obscurity of the central concept of coherence itself. A fully adequate explication of coherence is unfortunately not possible within the scope of this book (nor, one may well suspect, within the scope of any work of manageable length). But I will attempt to render the concept manageably clear in the next section, where I will also suggest that the clarity of the concept of coherence is not, surprisingly enough, a very crucial issue in assessing the plausibility of coherence theories *vis-à-vis* their nonskeptical opponents.

The problems relating to the other problematic transition in the schematic account, that from step (2) to step (3), are, in contrast, more serious, indeed critical. What is at issue here is the question of the connection between coherence and epistemic justification: why, if a system of empirical beliefs is coherent (and more coherent than any rival system), is it thereby justified *in the epistemic sense*, that is, why is it thereby likely to be true? I will address this question in section 5, where the standard set of objections to coherence theories will be developed in further detail.

3 The Concept of Coherence

What, then, is coherence? Intuitively, coherence is a matter of how well a body of beliefs "hangs together": how well its component beliefs fit together, agree or dovetail with each other, so as to produce an organized, tightly structured system of beliefs, rather than either a helter-skelter collection or a set of conflicting subsystems. It is reasonably clear that this "hanging together" depends on the various sorts of inferential, evidential, and explanatory relations which obtain among the various members of a system of beliefs, and especially on the more holistic and systematic of these. Thus various detailed investigations by philosophers and logicians of such topics as explanation, confirmation, probability, and so on, may be reasonably taken to provide some of the ingredients for a general account of coherence. But the main work of giving such an account, and in particular one which will provide some relatively clear basis for *comparative* assessments of coherence, has scarcely been begun, despite the long history of the concept.

My response to this problem, for the moment at least, is a deliberate – though, I think, justified – evasion. It consists in pointing out that the task of giving an adequate explication of the concept of coherence is not uniquely or even prima-

rily the job of coherence theories. This is so because coherence – or something resembling it so closely as to be subject to the same sort of problem – is, and seemingly must be, a basic ingredient of virtually all rival epistemological theories as well. We have already seen that weak foundationalism essentially involves an appeal to coherence. And it seems clear that even moderate and strong foundationalisms cannot avoid an appeal to something like coherence in giving an account of knowledge of the past, theoretical knowledge, and other types of knowledge which (on any view) go beyond direct experience. Thus it is not surprising that virtually all of the leading proponents of comprehensive foundationalist views, whether weak, moderate, or strong, employ the notion of coherence in their total epistemological accounts – though sometimes under other names, such as "congruence" (Lewis) or "concurrence" (Chisholm).[3] Even "contextualist" views, which attempt to repudiate the whole issue of global justification, make a similar appeal. The conclusion strongly suggested is that something like coherence is indispensable to any nonskeptical epistemological position which is even *prima facie* adequate. And if this is so, the absence of an adequate explication of coherence does not count against coherence theories any more than against their rivals.

The foregoing response is dialectically cogent in defending coherence theories against other, nonskeptical epistemologies, but it must be admitted that it is of little use *vis-à-vis* the skeptic, who may well argue that what it shows is that all nonskeptical epistemologies are fundamentally flawed by virtue of their dependence on this inadequately explicated concept. But although this challenge must be taken seriously, it is far from obvious that it is even close to being decisive. A better account of coherence is beyond any doubt something devoutly to be sought; but it is, I think, quite plausible to say, as Ewing does, that what proponents of coherence "are doing is to describe an ideal that has never yet been completely clarified but is none the less immanent in all our thinking,"[4] and to hold on this basis that our intuitive grasp of this notion, though surely not ideally satisfactory, will suffice so long as the only alternative is skepticism – which itself carries, after all, a significant burden of implausibility.

In any case, however, there is little point in talking at length about coherence without a somewhat clearer idea of what is involved. Thus I will attempt to provide in this section a reasonable outline of the concept of coherence, while recognizing that it falls far short of what would be ideal. The main points are: first, coherence is not to be equated with mere consistency; second, coherence, as already suggested, has to do with the mutual inferability of the beliefs in the system; third, relations of explanation are one central ingredient in coherence, though not the only one; and, fourth, coherence may be enhanced through conceptual change.

First. A serious and perennial mistake in discussing coherence, usually committed by critics but occasionally also by would-be proponents of coherence theories, is to assume that coherence means nothing more than logical consistency, the absence of explicit contradiction.[5] It is true that consistency is one requirement for coherence, that inconsistency is obviously a very serious sort of

incoherence. But it is abundantly clear, as many coherentists have pointed out, that a system of beliefs might be perfectly consistent and yet have no appreciable degree of coherence.

There are at least two ways in which this might be so. The more obvious is what might be called *probabilistic inconsistency.* Suppose that my system of beliefs contains both the belief that P and also the belief that it is extremely improbable that P. Clearly such a system of beliefs may perfectly well be logically consistent. But it is equally clear from an intuitive standpoint that a system which contains two such beliefs is significantly less coherent than it would be without them and thus that probabilistic consistency is a second factor determining coherence.

Probabilistic consistency differs from straightforward logical consistency in two important respects. First, it is extremely doubtful that probabilistic inconsistency can be entirely avoided. Improbable things do, after all, sometimes happen, and sometimes one can avoid admitting them only by creating an even greater probabilistic inconsistency at another point.[6] Second, probabilistic consistency, unlike logical consistency, is plainly a matter of degree, depending on (a) just how many such conflicts the system contains and (b) the degree of improbability involved in each case. Thus we have two initial conditions for coherence, which we may formulate as follows:

(1) A system of beliefs is coherent only if it is logically consistent.[7]
(2) A system of beliefs is coherent in proportion to its degree of probabilistic consistency.

But these two requirements are still not enough. Imagine a set of beliefs, each member of which has simply no bearing at all on the subject matter of any of the others, so that they make no effective contact with each other. This lack of contact will of course assure that the set is both logically and probabilistically consistent by ruling out any possibility of conflict; but it will also assure that the members of the set fail to hang together in any very significant way. Thus consider the following two sets of propositions, A and B. A contains "this chair is brown," "electrons are negatively charged," and "today is Thursday." B contains "all ravens are black," "this bird is a raven," and "this bird is black." Clearly both sets of propositions are free of contradiction and are also probabilistically consistent. But in the case of A, this consistency results from the fact that its component propositions are almost entirely irrelevant to each other; though not in conflict, they also fail to be positively related in any significant way. And for this reason, set A possesses only a very low degree of coherence. In the case of set B, in contrast, consistency results from the fact that the component propositions, rather than being irrelevant to each other, fit together and reinforce each other in a significant way; from an epistemic standpoint, any two of them would lend a degree of positive support to the third (though only very weak support in two out of the three cases). Thus set B, though obviously much too small to have a really significant degree of coherence, is much more coherent than set A. As the classical proponents of coherence have always insisted, coher-

ence must involve some sort of positive connection among the beliefs in question, not merely the absence of conflict.

Second. But what sort of positive connection is required and how strong must it be? The obvious answer to the first question is that the connections in question are inference relations namely, any sort of relation of content which would allow one belief or set of beliefs, if justified, to serve as the premise(s) of a cogent epistemic-justificatory argument for a further belief. The basic requirement for such an inference relation, as suggested in the earlier discussion of epistemic justification, is that it be to some degree truth-preserving; any sort of relation which meets this requirement will serve as an appropriate positive connection between beliefs, and no other sort of connection seems relevant here.

This much would be accepted by most, if not all, proponents of coherence theories. The main thing that divides them is the issue of how close and pervasive such inferential connections are required to be. One pole with regard to this issue is represented by the classical absolute idealists. Blanshard's formulation is typical:

> Fully coherent knowledge would be knowledge in which every judgment entailed, and was entailed by, the rest of the system.[8]

(In interpreting this formulation it is important to remember that Blanshard, like many others in this tradition, believes in synthetic entailments and indeed holds the admittedly dubious view that causal connections are one species of entailment.) The main problem with this view is that it is quite impossible even to imagine a system of beliefs which would satisfy such a requirement; as Blanshard himself admits, even such a system as Euclidean geometry, often appealed to as a paradigm of coherence, falls far short.[9] Thus it is plausible to weaken the requirement for coherence at least to the degree advocated by Ewing, who requires only that each proposition in a coherent system be entailed by the rest taken together, not that the reciprocal relation hold.[10] (We will see shortly that weakening the requirement in this way creates a problem which forces Ewing to add a further, related requirement.)

At the opposite extreme is Lewis's account of "congruence," a concept which plays a crucial role in his account of memory knowledge:

> A set of statements . . . will be said to be congruent if and only if they are so related that the antecedent probability of any one of them will be increased if the remainder of the set can be assumed as given premises.[11]

This is obviously an extremely weak requirement. A system of beliefs which satisfied it at only the most minimal level would possess a vastly lower degree of systematic interconnection than that envisaged by the idealists, in two significantly different respects. First, reducing the requirement from entailment to merely some increase in probability obviously allows a weakening of the inferential connections which constitute coherence. But this is no objection to Lewis's account, so long as it is understood that coherence is a matter of degree, and

that a lower degree of inferential interconnection carries with it only a lower degree of coherence. Second, however, Lewis's account, and indeed Ewing's as well, by making the inferential connection between the individual belief in question and the rest of the system one-way rather than reciprocal, creates the possibility that a system of beliefs could count as coherent to as high a degree as one likes by being composed of two or more subsystems of beliefs, each internally connected by strong inference relations but none having any significant connection with the others. From an intuitive standpoint, however, it is clear that such a system, though coherent to some degree, would fall very far short of ideal coherence. Ideal coherence requires also that the entire system of beliefs form a unified structure, that there be laws and principles which underlie the various subsystems of beliefs and provide a significant degree of inferential connection between them. We are obviously very close here to the ideal of a "unified science," in which the laws and terms of various disparate disciplines are reduced to those of some single master discipline, perhaps physics; while such a specific result is not essential for coherence, it would represent one way in which a high degree of coherence could be achieved, and something in this general direction seems to be required.

Ewing attempts to meet this difficulty by adding as a separate requirement for coherence the condition that no set of beliefs smaller than the whole system be logically independent of the rest of the system,[12] and a similar requirement could be added to Lewis's account as well. It would be better, however, to make this further aspect of coherence also a matter of degree, since there are obviously many intermediate cases between a completely unified system and a system with completely isolated subsystems. Putting all of this together results in the following two additional conditions for coherence:

(3) The coherence of a system of beliefs is increased by the presence of inferential connections between its component beliefs and increased in proportion to the number and strength of such connections.

(4) The coherence of a system of beliefs is diminished to the extent to which it is divided into subsystems of beliefs which are relatively unconnected to each other by inferential connections.

It should be noted that condition (3), in addition to summarizing the preceding discussion, includes one important idea which did not emerge explicitly there: each individual belief can be involved in many different inferential relations, and the degree to which this is so is also a determinant of coherence.

Third. The foregoing account, though it seems to me to be on the right track, is obviously still extremely sketchy. One way to reduce this sketchiness somewhat is to consider the major role which the idea of *explanation* plays in the overall concept of coherence. As I have already suggested by mentioning the ideal of unified science, the coherence of a system of beliefs is enhanced by the presence of explanatory relations among its members.

Indeed, if we accept something like the familiar Hempelian account of explanation, this claim is to some extent a corollary of what has already been said.

According to that account, particular facts are explained by appeal to other facts and general laws from which a statement of the explanandum fact may be deductively or probabilistically inferred; and lower-level laws and theories are explained in an analogous fashion by showing them to be deducible from more general laws and theories.[13] Thus the presence of relations of explanation within a system of beliefs enhances the inferential interconnectedness of the system simply because explanatory relations *are* one species of inference relations.

Explanatory connections are not just additional inferential connections among the beliefs of a system, however; they are inferential connections of a particularly pervasive kind. This is so because the basic goal of scientific explanation is to exhibit events of widely differing kinds as manifestations of a relatively small number of basic explanatory principles. As Hempel remarks: "What scientific explanation, especially theoretical explanation, aims at is . . . an objective kind of insight that is achieved by a systematic unification, by exhibiting the phenomena as manifestations of common underlying structures and processes that conform to specific, testable, basic principles."[14] What Hempel calls "systematic unification" is extremely close to the concept of coherence.

One helpful way to elaborate this point is to focus on the concept of *anomaly*. For my purposes, an anomaly is a fact or event, especially one involving some sort of recurring pattern, which is claimed to obtain by one or more of the beliefs in the system of beliefs, but which is incapable of being explained (or would have been incapable of being predicted) by appeal to the other beliefs in the system.[15] (Obviously such a status is a matter of degree.) The presence of such anomalies detracts from the coherence of the system to an extent which cannot be accounted for merely by appeal to the fact that the belief in an anomalous fact or event has fewer inferential connections to the rest of the system than would be the case if an explanation were available. In the context of a coherentist position, such beliefs will have to be inferentially connected to the rest of the system in other, nonexplanatory ways if there is to be any justification for accepting them (see the discussion of observation in Chapter 6), and such connections may be very extensive. The distinctive significance of anomalies lies rather in the fact that they undermine the claim of the allegedly basic explanatory principles to be genuinely basic, and thus threaten the overall coherence of the system in a much more serious way. For this reason, it seems advisable to add one more condition to our list of conditions for coherence:

(5) The coherence of a system of beliefs is decreased in proportion to the presence of unexplained anomalies in the believed content of the system.[16]

Having insisted on the close connection between coherence and explanation, we must nonetheless resist the idea that explanatory connections are all there is to coherence. Certain proponents of coherentist views, notably Sellars and Harman, have used phrases like "explanatory coherence" in speaking of coherence, seeming to suggest (though I doubt whether any of those using it really intend such a suggestion) that coherence depends *entirely* on explanatory con-

nections.[17] One could of course adopt a conception of coherence which is restricted in this way, but there is no reason at all – from an epistemological standpoint – to do so. The epistemologically significant concept of coherence is bound up with the idea of *justification*, and thus any sort of inference relation which could yield some degree of justification also enhances coherence, whether or not such a relation has any explanatory force.

A simple example (borrowed from Lehrer who in turn borrowed it from Bromberger) may help to illustrate this point.[18] Suppose that I am standing three feet from a pole which is four feet high. Next to my foot is a mouse, and on top of the pole is perched an owl. From these conditions I may obviously infer, using the Pythagorean theorem, that the mouse is five feet from the owl. This inference is surely adequate to justify my believing that the mouse is five feet from the owl, assuming that I am justified in believing these other propositions. And intuitively speaking, this inferential connection means that the belief that the mouse is five feet from the owl coheres with the rest of my beliefs to quite a significant extent. But none of this has any apparent connection with explanation. In particular, as Lehrer points out, this inference does not in any way help to *explain* why the mouse is so close to the owl. Thus it is a mistake to tie coherence too closely to the idea of explanation. Of course, it is still true that the coherence of the system in question would be enhanced by adding an explanation for the presence of the mouse in such close proximity to the owl: given the usual behavior of mice around owls, the presence of the mouse at that distance is an explanatory and predictive anomaly. The point is simply that coherence is also enhanced by inferential connections of a nonexplanatory sort.

Fourth. The final point is really just a corollary of the one just made. To the extent that coherence is closely bound up with explanation and systematic unification, achieving a high degree of coherence may well involve significant conceptual change. This point is most clear in the area of theoretical science, though it has much broader application. A typical situation of theoretical explanation involves one or more anomalies at the "observational" level: apparently well-established facts formulated in the available system of concepts for which no adequate explanation seems to be available in those terms. By devising a new system of the oretical concepts the theoretician makes an explanation available and thus enhances the coherence of the system. In this way the progress of theoretical science may be plausibly viewed as a result of the search for greater coherence.[19]

The foregoing account of coherence is a long way from being as definitive as desirable. I submit, however, that it does indeed identify a concept which, in Ewing's phrase, is "immanent in all our thinking," including all our most advanced scientific thinking; and also that the concept thus identified, though vague and sketchy in many ways, is nonetheless clear enough to make it reasonable to use it, albeit with caution, in dealing with the sorts of epistemological issues under discussion here. In particular, it seems clear that the concept is not so vague as to be at all easy to satisfy.

4 The Doxastic Presumption

I have so far considered two of the elements which are arguably essential to a viable coherence theory: the idea of nonlinear justification and the concept of coherence itself. A third essential element is the presumption regarding one's grasp of one's own system of beliefs which I mentioned briefly at the end of the previous chapter; this is required, I will suggest, if our coherence theory is to avoid a relapse into externalism. (A fourth ingredient is the coherentist conception of observation; and a fifth, on a somewhat different level, is the metajustificatory argument for such a theory.)

It will be useful, before attempting to say in detail what the presumption in question amounts to and what it is supposed to do, to see more clearly why it is needed in the first place. According to a coherence theory of empirical justification, as so far characterized, the epistemic justification of an empirical belief derives entirely from its coherence with the believer's overall system of empirical beliefs and not at all from any sort of factor outside that system. What we must now ask is whether and how the fact that a belief coheres in this way is cognitively accessible to the believer himself, so that it can give *him* a reason for accepting the belief.

It would be possible, of course, to adopt an externalist version of coherentism. Such a view would hold that the person whose belief is justified need himself have no cognitive access to the fact of coherence, that his belief is justified if it in fact coheres with his system of beliefs, whether or not such coherence is cognitively accessible to him (or, presumably, to anyone). But such a view is unacceptable for essentially the same reasons which were offered against foundationalist versions of externalism and, as discussed earlier, seems to run counter to the whole rationale for coherence theories. (If externalism were acceptable in general, the foundationalist versions would obviously be far simpler and more plausible.) But if the fact of coherence is to be accessible to the believer, it follows that he must somehow have an adequate grasp of his total system of beliefs, since it is coherence with this system which is at issue. One problem which we will eventually have to confront is that it seems abundantly clear that no actual believer possesses an *explicit* grasp of his overall belief system; if such a grasp exists at all, it must be construed as tacit or implicit, which creates obvious problems for the claim that he is actually, as opposed to potentially, justified.

The problem at issue in this section is, however, more immediate and more serious. For whether the believer's grasp of his own system of beliefs is construed as explicit or implicit, of what can that grasp possibly consist except a set of empirical metabeliefs, *themselves in need of justification*, to the effect that he has such and such specific beliefs? How then are these metabeliefs themselves to be justified? If a return to foundationalism is to be avoided, the answer must apparently be that these metabeliefs too are justified by virtue of their coherence with the rest of my system of beliefs. And the problem is that it is absolutely clear that such an answer is unacceptable: it is beyond any doubt viciously

circular to claim that the metabeliefs which constitute the believer's grasp of his system of beliefs are themselves justified by virtue of their coherence with that system – even if the nonlinear view of justification articulated earlier is accepted in its entirety. How can my metabelief B_2 that I have a certain other belief B_1 be justified for me by appeal to the fact that B_2 coheres with my total system of beliefs if my very grasp of that system depends on the justification of B_2 and other similar beliefs? How, that is, can my reason for accepting B_2 be its coherence with my total system of beliefs when I have no justification apart from the appeal to B_2 and similar beliefs for thinking that I even have that system of beliefs? The shift to holism is of no help here, since the very possibility of a nonexternalist holism depends on my having a cognitive grasp of my total system of beliefs and its coherence which is prior to the justification of the particular beliefs in the system. It is quite clear, therefore, that this grasp, upon which any nonexternalist appeal to coherence must depend, cannot itself be justified by appeal to coherence.[20] And thus the very idea of a coherence theory of empirical justification threatens to collapse.

Is there any solution to this problem? Most proponents of coherence theories seem, surprisingly enough, either to take the believer's grasp of his own system of beliefs entirely for granted, or simply to ignore the issue of whether their envisaged coherentist justification is accessible to the believer himself. And the obvious conclusion, suggested by some foundationalists in passing, is that this problem shows that even an intended coherence theory must involve an irreducibly foundationalist element, that one's grasp of one's own system of beliefs must be justified in a foundationalist manner, even if everything else depends on coherence. But if the antifoundationalist arguments offered in an earlier chapter are genuinely cogent, no such retreat to foundationalism is available here, and skepticism looms as the only conclusion unless a further alternative can be found.

It was suggested earlier that an a priorist version of foundationalism (or quasi-foundationalism) might attempt to solve the problem of how the empirical claim that I have a certain belief is to be justified by maintaining that the existence of the justificandum belief is *presupposed* by the very raising of the issue of justification, so that the metabelief in question is not in need of justification, while still being available as a justifying premise. The normal justificatory issue, on this view, is whether the believer is justified in holding a certain belief *which he does in fact hold*, not whether such a belief would be somehow justified in the abstract independently of whether he holds it, nor even the hypothetical issue of whether it would be justified *if* he held it (though these other questions can, of course, also be asked). But since the basic unit of justification for a coherence theory is an entire system of beliefs, the analogous claim within the context of such a position is that the raising of an issue of empirical justification *presupposes* the existence of some specifiable *system* of empirical beliefs – or rather, as I will explain below, of *approximately* that system; the primary justificatory issue is whether or not, under the presumption that I do indeed hold approximately the system of beliefs which I believe myself to hold, those beliefs are justified. And thus the suggested solution to the problem raised in this section is that the

grasp of my system of beliefs which is required if I am to have cognitive access to the fact of coherence is dependent, in a sense yet to be adequately clarified, on this *Doxastic Presumption*, as I will call it, rather than requiring further justification.

But how exactly is this presumption to be understood? Three issues need to be considered: First, what is the significance of the qualifier "approximately" as it occurs in the above formulations of the presumption? Second, how exactly is this presumption supposed to function within the overall system of empirical knowledge? How exactly is it supposed to certify or secure (even the choice of word here is uncertain) one's grasp of one's system of beliefs? And third, what is the bearing of the Doxastic Presumption on issues pertaining to skepticism? – does it not amount to begging the question against a certain perhaps unusual, but nonetheless quite possible, version of skepticism? I will consider each of these questions in turn.

First. I have noted that the Doxastic Presumption is only that my representation of my overall system of beliefs is *approximately* correct. The point of the qualifier is that although assessments of coherence can be made only relative to a system of beliefs of which one has some prior grasp or representation, this does not mean that no aspect of that representation can be questioned. On the contrary, it is perfectly possible to raise the issue of whether I have a certain particular belief or reasonably small set of beliefs which I believe myself to have, and then to answer this question by appeal to the coherence or lack of coherence between the metabelief that I have the specific belief(s) in question and the rest of the system as I represent it – the existence of the rest of the system, but not of those particular beliefs, being presupposed. What is *not* possible is to question whether my grasp of my system of beliefs might be wholly or largely mistaken and then resolve *this* question by appeal to coherence: the raising of this issue would leave me with no sufficiently ample grasp of my system of beliefs which would not beg the question and relative to which coherence might be judged.

Second. It might seem plausible, at first glance, to construe the Doxastic Presumption as constituting a further *premise* to be employed in the justificatory arguments or at least as functioning like such a premise. But only a little reflection will show that such an interpretation is quite untenable. For what might such a premise say? The only apparent possibility is that it would say that my metabeliefs to the effect that I have certain beliefs may be presumed to be true, without requiring justification. And it is immediately obvious that such a premise would do me no good relative to the problem under discussion here. For to apply it in any useful fashion, I would need further premises to the effect that I do in fact believe myself to have such and such specific beliefs, and the justification of these further premises would obviously be just as problematic as before.

Thus the Doxastic Presumption, if it is to solve the problem, cannot function like a premise. It is rather a characterization of something which is, from the standpoint of a coherence theory, a basic and unavoidable feature of cognitive *practice*. Epistemic reflection, according to such a theory, *begins* from a (per-

haps tacit) representation of myself as having (approximately) such and such a specific system of beliefs: only relative to such a representation can questions of justification be meaningfully raised and answered. This representation is presumably a product of something like ordinary introspection (as understood from within the system), but whereas most introspective beliefs can be justified by appeal to coherence, the metabeliefs which constitute this representation cannot be thus justified in general for the reasons already considered. The issue of their justification can be raised and answered in particular, relatively confined cases which are for some reason especially problematic. But apart from such cases, such metabeliefs must be presumed to be correct in order for the process of justification to even get started. And this is what the Doxastic Presumption says.

Thus the Doxastic Presumption does not, strictly speaking, function at all in the normal workings of the cognitive system. Rather it simply describes or formulates, from the outside, something that I unavoidably *do*: I assume that the beliefs constituting my overall grasp of my system of beliefs are, by and large, correct.

Third. But does not the Doxastic Presumption, or rather the aspect of cognitive practice which it reflects, amount to begging the question against a certain from of skepticism, namely, that form which would question whether my representation of my own system of beliefs is in fact accurate? The answer is that it would be begging the question if it purported to be an answer to such a skeptical challenge but that as proposed here no such answer is intended. It would be possible, of course, to argue that if it is correct that empirical justification is only possible relative to a specific system of beliefs whose existence is presumed, then it follows that skepticism of the sort in question simply makes no sense; the underlying idea would be that a question is meaningful only if there is some way, at least in principle, in which it can be answered. But I can see no reason to accept such a view, amounting as it does to a version of verificationism. What the discussion leading up to the Doxastic *Presumption* shows is precisely that a coherence theory of empirical justification cannot, in principle, answer this form of skepticism; and this seems to me to count in favor of the skeptic, not against him.

Thus the position advocated here holds that such a version of skepticism, though certainly unusual, is perfectly coherent (and thus that it would be desirable to be able to answer it) but also concedes that such an answer is unfortunately in principle not available for a coherence theory. However, the failure to answer one version of skepticism does not in any way mean that there is no point in attempting to answer others. The effect of the Doxastic Presumption is precisely to distinguish a version of skepticism which cannot be successfully answered from others which perhaps can. Even if it is not possible in general to justify my representation of my own system of beliefs, it may yet be possible to argue successfully relative to the presumption that this representation is (approximately) correct that the beliefs which I hold are justified in a sense which makes them genuinely likely to be true; and this would be a significant epistemological result, even if not quite the one which would be ideally desirable.

There is one more important point about the Doxastic Presumption to be noted here. Obviously a person's system of beliefs changes and develops over time as new beliefs are added and old ones abandoned or forgotten. And it is clear on reflection that one's grasp of these changes is just as incapable of being justified in general by appeal to coherence as is one's grasp of the system at a moment. Thus the Doxastic Presumption must be understood to include the presumption that one's grasp of this temporal dimension of one's system of beliefs is also approximately correct.

The foregoing will suffice for an initial discussion of the Doxastic Presumption. For the moment the central point is that something like this presumption seems to be unavoidable if a coherentist position is to even get started. Nothing like a justification for the presumption has been offered for the simple reason that if it is properly understood, none is required: there can obviously be no objection to asking what follows about the justification of the rest of my beliefs from the presumption that my representation of my own system of beliefs is approximately correct. The only questions needling to be asked are: first, whether it is possible to justify my representation of my own system of beliefs, rather than having to presume that it is correct (I have argued that it is not); and, second, whether the epistemological issue which results from this presumption is still worth bothering with (I have suggested that it is).

5 The Standard Objections

There is obviously much which is problematic in the very tentative and fragmentary picture of a coherence theory of empirical justification which has so far emerged in this chapter, and many important questions and problems remain to be considered. But even if the conception were otherwise acceptable, there would still remain the three standard and extremely forceful objections to coherence theories – objections which have usually been thought to destroy any plausibility which such a view might possess. As will become clear, these objections are not entirely independent of one another and indeed might be plausibly regarded as merely different facets of one basic point. But each of them possesses enough independent plausibility and intuitive force to warrant separate consideration.

(I) The alternative coherent systems objection. According to a coherence theory of empirical justification, at least as so far characterized, the system of beliefs which constitutes empirical knowledge is epistemically justified *solely* by virtue of its internal coherence. But such an appeal to coherence will never even begin to pick out one uniquely justified system of beliefs, since on any plausible conception of coherence, there will always be many, probably infinitely many, different and incompatible systems of belief which are equally coherent. No nonarbitrary choice between such systems can be made solely on the basis of coherence, and thus all such systems, and the beliefs they contain, will be equally justified. And this will mean in turn, since all or virtually all consistent beliefs will belong to some such system, that we have no more reason to think that the

beliefs we actually hold are true than we have for thinking that any arbitrarily chosen alternative belief is true – a result which is surely tantamount to skepticism and which obviously vitiates entirely the concept of epistemic justification by destroying its capacity to discriminate between different empirical beliefs.

A clear conception of this objection requires that it not be exaggerated, as it frequently is. Sometimes it is said that if one has an appropriately coherent system, an alternative coherent system can be produced simply by negative all of the components of the first system. This would be so if coherence amounted simply to consistency; but once it is seen that such a conception of coherence is much too limited, there is no reason to accept such a claim. Nor is it even minimally plausible that, as is sometimes suggested, a "well written novel," or indeed anything remotely resembling an actual novel, would have the degree of coherence required to be a serious alternative to anyone's actual system of beliefs. What would be missing in both cases is the pervasive inferential and especially explanatory connections needed for a high degree of coherence.

But even without these exaggerations, the objection is obviously very forceful. One suggestive way to elaborate it is by appeal to the idea of alternative possible worlds. Without worrying about whether there are infinitely many possible worlds or whether all possible worlds are capable of being given equally coherent descriptions, it seems enormously obvious that there are at least very many possible worlds, differing in major ways from the actual world, which are capable of being described in equally coherent ways. But then a standard of justification which appeals only to internal coherence has no way of choosing among the various systems of beliefs which would correctly describe these various possible worlds; such a standard is apparently impotent to justify believing in one of these worlds as opposed to any of the others. The skeptic need ask for nothing more.

(II) The input objection. The second objection is somewhat more elusive, but also perhaps more fundamental. Coherence is purely a matter of the *internal* relations between the components of the belief system; it depends in no way on any sort of relation between the system of beliefs and anything external to that system. Hence if, as a coherence theory claims, coherence is the sole basis for empirical justification, it follows that a system of empirical beliefs might be adequately justified, indeed might constitute empirical knowledge, in spite of being utterly out of contact with the world that it purports to describe. Nothing about any requirement of coherence dictates that a coherent system of beliefs need receive any sort of *input* from the world or be in any way causally influenced by the world. But this is surely an absurd result. Such a self-enclosed system of beliefs, entirely immune from any external influence, cannot constitute empirical knowledge of an independent world, because the achievement of even minimal descriptive success in such a situation would have to be either an accident or a miracle, not something which anyone could possibly have any reason to expect – which would mean that the beliefs involved would not be epistemically justified, even if they should somehow happen to be true. This objection is most obviously forceful against a coherentist position, like my own, which adopts a realist conception of independent reality. But in fact it is cogent

vis-à-vis any position, including at least most versions of idealism, which does not simply identify the individual believer's limited cognitive system with its object: how can a system of beliefs be justified in a sense which carries with it likelihood of truth, while at the same time being entirely isolated from the reality, however that be understood, which it purports to describe?

Though intuitively forceful, this objection is also rather vague – mainly because of the vagueness of the crucial notion of "input." It would, however, be a mistake to attempt too precise a specification here, prior to the development of a more specific theory. The rough idea is that some of the elements in the cognitive system must be somehow shaped or influenced by the world outside the system;[21] and that this must be not just something which might or might not happen to occur, but rather in some way an essential requirement for the justification of the system. But just what precise form such input might take is a matter to be specified by a particular theory.[22]

(III) The problem of truth. The final objection of the three is the most fundamental of all. Recall that one crucial part of the task of an adequate epistemological theory is to show that there is an appropriate connection between its proposed account of epistemic justification and the cognitive goal of *truth*. That is, it must be somehow shown that justification as conceived by the theory is *truth-conducive*, that one who seeks justified beliefs is at least likely to find true ones. All this is by now quite familiar. The objection is simply that a coherence theory will be unable to accomplish this part of the epistemological task unless it also adopts a coherence theory of truth and the idealistic metaphysics which goes along with it – an expedient which is both commonsensically absurd and also dialectically unsatisfactory.

Historically, the appeal to a coherence theory of truth was made by the absolute idealists and, in a slightly different but basically parallel way, by Peirce. These philosophers attempted to solve the problem of the relation between justification and truth by in effect construing truth as simply *identical* with justification-in-the-long-run. Thus an idealist, having adopted a coherence theory of epistemic justification, might argue that only by adopting a coherence theory of truth could the essential link between justification and truth be secured: obviously if truth is long-run, ideal coherence, it is plausible to suppose that it will be truth-conducive to seek a system of beliefs which is as coherent as one can manage to make it at the moment.[23] Something like this seems also to be the essential motivation behind Peirce's version of the pragmatic conception of truth in which truth is identified with the ideal, long-run outcome of scientific inquiry; whether this amounts to precisely a coherence theory of truth depends on just how Peirce's rather obscure account of justification is properly to be understood, but it is at least similar. The same underlying motivation also seems present, albeit less clearly, in other versions of pragmatism.

Obviously, given such a construal of truth, there will be no difficulty of principle in arguing successfully that one who accepts justified beliefs will in the long run be likely to find true ones. But such a gambit is nonetheless quite unsatisfactory in relation to the basic problem at issue, even if the intuitive and commonsensical objections to such accounts of truth are discounted. The whole

point, after all, of seeking an argument connecting justification and truth is to provide a rationale or metajustification for the proposed standard of epistemic justification by showing that adopting it leads or is likely to lead to the attainment of truth. But the force of such a metajustification depends on the independent claim to acceptance of the concept of truth which is invoked. If – as seems to be the case both historically and dialectically with respect to the specific concepts of truth under discussion here – the only rationale for the chosen concept of truth is an appeal to the related standard of justification, then the proposed metajustification loses its force entirely. It is clearly circular to argue both (1) that a certain standard of epistemic justification is correct because it is conducive to finding truth, conceived in a certain way, and (2) that the conception of truth in question is correct because only such a conception can connect up in this way with the original standard of justification. Such a defense would obviously be available to the proponent of *any* proposed standard of epistemic justification, no matter how silly or counterintuitive or arbitrary it might be: all he has to do is adopt his own nonstandard conception of truth as justification-in-the-long-run (in his idiosyncratic sense of justification). The moral of the story is that although any adequate epistemological theory must confront the task of bridging the gap between justification and truth, the adoption of a nonstandard conception of truth, such as a coherence theory of truth, will do no good unless that conception is independently motivated.[24] Therefore, it seems that a coherence theory of justification has no acceptable way of establishing the essential connection with truth. A coherentist standard of justification, it is claimed, can be a good test only for a coherentist conception of truth, so that to reject the coherence theory of truth commits one also to the rejection of any such account of justification.[25]

Of these three objections, (III) is the most basic and (I) is the most familiar. It is (II), however, which must be dealt with first, since the answer to it turns out, not surprisingly, to be essential for answering the other two objections. My view is that the point advanced in (II) must in the end simply be accepted: a cognitive system which is to contain empirical knowledge must somehow receive input of some sort from the world. And this means that the purest sort of coherence theory turns out, as the objections claim, to be indeed unacceptable. I will argue, however, that this need not mean a return to foundationalism (which has already been shown to be hopeless), that a theory which is recognizably coherentist – and more important, which is free of any significant foundationalist ingredients – can allow for such input.

Notes

1 Rescher's "coherence theory of truth" in Rescher (1973) is a coherence theory of the criteria of truth, not of the nature of truth. Though leaving the connection with justification somewhat obscure, this book contains an excellent discussion of the distinction between the two sorts of theories of truth.

2 As already noted, some theories combine an appeal to coherence with a rejection of the global issue of justification. See, for example, Sklar (1975) and Williams (1980).

Because I see no warrant for dismissing the global issue in this way, such views seem to me to constitute merely complicated versions of skepticism.

3 See Lewis (1946, chap. II); and Chisholm (1977, chap. 4).

4 Ewing (1934, p. 231).

5 Perhaps the clearest example of this rather pervasive mistake is Scheffler (1967, chap. 5). Another interesting case is Rescher (1973): the "coherence criterion of truth" advocated there uses consistency to segregate propositions into maximally consistent subsets, and then chooses among those subsets on a variety of different bases, none of which have much to do with the standard idea of coherence. (Rescher's later development of the same view does, however, employ a more traditional notion of coherence, though in a different place.) See Rescher (1977), and also BonJour (1976).

6 One pervasive case of this sort is worth explicit notice: it often happens that my system of beliefs makes it extremely probable that an event of a certain general description will occur, while providing no guidance as to which of a very large number of alternative, more specific possibilities will realize this description. If there is nothing more to be said, each of the specific possibilities will be very improbable, simply because there are so many, while at the same time it will be highly probable that one of them will occur. In such a case, adding a new belief (arrived at through observation or in some other way) that one of the possible specific events has actually occurred will bring with it a measure of probabilistic inconsistency – but less than would result from excluding all such specific beliefs, thereby coming into conflict with the more general one.

7 It might be questioned whether it is not an oversimplification to make logical consistency in this way an absolutely necessary condition for coherence. In particular, some proponents of relevance logics may want to argue that in some cases a system of beliefs which was sufficiently rich and complex but which contained some trivial inconsistency might be preferable to a much less rich system which was totally consistent. And there are also worries such as the Preface Paradox. But while I think there may be something to be said for such views, the issues they raise are too complicated and remote to be entered into here.

8 Blanshard (1939, p. 264).

9 Ibid., p. 265.

10 Ewing (1934, p. 229).

11 Lewis (1946, p. 338), Chisholm's definition of "concurrence" in Chisholm (1977) is very similar.

12 Ewing (1934, pp. 229–30).

13 For refinements, see, e.g., the title essay in Hempel (1965).

14 Hempel (1967, p. 83).

15 On the Hempelian view, explanation and prediction involve the same sorts of inferential relations within the system of beliefs, differing only as to whether the fact in question is known prior to drawing the appropriate inference. Although this view is not uncontroversial, I will assume that it is at least approximately correct.

16 Such a situation of anomaly may of course involve probabilistic inconsistency in the sense explained above, but it need not do so in any very straightforward way; the set of basic explanatory principles does not normally include an explicit rider to the effect that anything which it cannot subsume is thereby rendered improbable.

17 For a version of such a position, see Lehrer (1974, chap. 7). (The position in question is one which Lehrer is criticizing, not one he wishes to advocate.)

18 Lehrer (1974, pp. 166–67).

19 For an elaboration of this view of scientific theories, see Wilfrid Sellars, "The Language of Theories," reprinted in Sellars (1963, pp. 106–26).

20 A coherence theory, at least as construed here, does not somehow reject any notion of epistemic priority. Its claim is rather that what look superficially like relations of epistemic priority and posteriority among individual beliefs turn out to be relations of reciprocal support in relation to a system of beliefs which is genuinely prior. Thus no appeal to the nonlinear conception of justification will help if it is the very existence of such a system which is in question.

21 I will ignore here the alternative possibility that a person's system of beliefs might be likely to be true because those beliefs shape reality rather than the other way around.

22 Many foundationalist views also fail to address this issue in any clear way, either by offering no real account of the status of the foundational beliefs or by merely appealing to the commonsensical belief that certain beliefs are somehow justified.

23 For the clearest version of this approach, see Blanshard (1939, chaps. 25 and 26).

24 A related objection also afflicts currently fashionable verificationist accounts of truth.

25 It is worth noting, however, that foundationalist views seem to face at least a somewhat analogous problem if it is true that they must appeal to coherence or something like it in their accounts of knowledge of the past, theoretical knowledge, and so on.

Works Cited

Blanshard, Brand (1939). *The Nature of Thought.* London: Allen and Unwin.

BonJour, Laurence (1976). "Rescher's Idealistic Pragmatism." *Review of Metaphysics* 29: 702–26.

Chisholm, Roderick (1977). *Theory of Knowledge*, 2nd edn. Englewood Cliffs, NJ: Prentice-Hall.

Ewing, A. C. (1934). *Idealism.* London: Methuen.

Hempel, Carl G. (1965). *Aspects of Scientific Explanation.* New York: Free Press.

—— (1967). *Philosophy of Natural Science.* Englewood Cliffs, NJ: Prentice-Hall.

Lehrer, Keith (1974). *Knowledge.* Oxford: Oxford University Press.

Lewis, C. I. (1946). *An Analysis of Knowledge and Valuation.* La Salle, IL,: Open Court.

Rescher, Nicholas (1973). *The Coherence Theory of Truth.* Oxford: Oxford University Press.

—— (1977). *Methodological Pragmatism.* New York: New York University Press.

Scheffler, Israel (1967). *Science and Subjectivity.* New York: Bobbs-Merrill.

Sellars, Wilfred (1963). *Science, Perception and Reality.* London: Routledge and Kegan Paul.

Sklar, Lawrence (1975). "Methodological Conservatism." *Philosophical Review* 84: 374–400.

Williams, Michael (1980). "Coherence, Justification, and Truth." *Review of Metaphysics* 34: 243–72.