13

Evidence and Entitlement

Diagnostic approaches to scepticism

Foundationalism and the coherence theory represent *direct* responses to scepticism. Taking the sceptic's arguments at face value, they attack one or more of his acknowledged premises. Diagnostic responses take a more roundabout approach.

We can make a rough-and-ready distinction between two diagnostic strategies. *Therapeutic* diagnosis treats sceptical problems as *pseudo-problems* generated by misuses or misunderstandings of language. On this approach, sceptical claims and arguments do not really make sense. The problem is, however, that we seem to understand such things well enough to appreciate how they generate a space of epistemological theories, structured by the possible forms of direct response. *Theoretical* diagnosis takes a different approach, questioning the naturalness or intuitiveness of the sceptic's arguments rather than their intelligibility. It aims at showing sceptical arguments to be much more complex and theory-laden than their proponents want to admit, deepening our understanding of such arguments by making available for critical scrutiny their *unacknowledged* theoretical preconceptions.[1]

I do not want to ride the distinction between therapeutic and theoretical diagnosis too hard. It may turn out that the sceptic's presuppositions extend into the theory of meaning. If they do, a theoretical diagnosis of the sceptic's presuppositions may encourage second thoughts about how well we understand everything he says. I introduce the distinction for two reasons. First, I want to head off a tendency to be too quick to accuse the sceptic of not making sense: for example by charging that his idea of a reality that transcends experience is obviously unintelligible. Second, I want to distance myself from the ideal of a *purely* therapeutic approach to philosophical problems. By this I mean the ideal of exposing philosophical problems as illusory while avoiding any theoretical commitments of one's own. We achieve a lot if we get the sceptic to share the burden of theory. We do not have to imagine that we ourselves escape it entirely.

A diagnostic approach is promising because of the conditions scepticism must meet to be a problem and not just a puzzle. First and foremost, scepticism must be intuitive, in the sense that it must exploit only our most everyday, lowest-common-denominator ideas about knowledge and justification. If sceptical arguments turn out to depend on contentious theoretical ideas, we can see scepticism as reflecting on those ideas, rather than on our everyday epistemic practices. Scepticism will prove to be rooted, not in the human condition, but in a particular genre of epistemological theorizing

Finally, a diagnostic approach is needed because sceptical arguments appear (or can be made to appear) intuitive. (That is one reason why they cannot simply be dismissed.) But perhaps appearances are misleading. The task of theoretical diagnosis is to show that this is so.

Let us now turn to scepticism itself, beginning with Agrippan scepticism.

Sceptical commitments

When I introduced the distinction between personal and evidential justification, I noted that many philosophers have been attracted to a strongly 'intellectualist' account of the relation between them, according to which personal justification is subject to what I called the *Prior Grounding Requirement*. My claim is that the Prior Grounding conception of justification must be presupposed by the Agrippan argument, *if it is to amount to an argument for radical and general scepticism*. In other words, scepticism is a serious problem only if it is *not* natural or intuitive.

To see why this is so, we need to look more closely at what the Prior Grounding conception of justification involves. I suggest that we analyse it into four sub-principles:

(PG1) *No Free Lunch Principle.* Epistemic entitlement—personal justification—does not just accrue to us: it must be earned by epistemically responsible behaviour.

(PG2) *Priority Principle.* It is never epistemically responsible to believe a proposition true when one's grounds for believing it true are less than adequate.

(PG3) *Evidentialism.* Grounds are evidence: propositions that count in favour of the truth of the proposition believed.

(PG4) *Possession Principle.* For a person's belief to be adequately grounded, it is not sufficient for there merely to be appropriate evidence for it. Rather, the believer himself or herself must possess (and make proper use of) evidence that makes the proposition believed (very) likely to be true.

These four principles accord well with the thought that knowledge requires both personal justification and adequate grounding, while also explaining why this distinction is not one of deep theoretical significance. (PG1) and (PG2) imply the uniform subordination of personal justification to grounding. By (PG2), believing on less than adequate grounds is always irresponsible and hence, by (PG1), never justified. Call this 'the Dependence Thesis'. (PG3) and

(PG4) add to this a strongly *internalist* account of what it is for someone's belief to be adequately grounded. On this view, a person's grounds must be evidence in the strong sense: further beliefs—or if not beliefs, some other personal cognitive state—in virtue of which he holds the belief in question and to which he has immediate cognitive access. A person's belief may have been formed by a method that is *in fact* reliable. It may even have been the result of a process that, in the circumstances, ensures that it is true. But this externalist form of grounding, where a person is not necessarily aware of the factors that make his belief truth-reliable, is just what (PG3) and (PG4) exclude. This exclusion is strongly motivated by the dependence thesis. Externalist 'grounding' is something that 'just happens'. It therefore fails to provide the kind of earned entitlement that epistemic justification requires. The Prior Grounding conception of justification involves ideas about the relation between justification, responsibility, and grounding that are made for each other.

At the heart of the Agrippan argument is the apparently fatal trilemma: any attempt to justify a belief must open a vicious regress, end with a brute assumption, or go in a circle. The sceptic concludes that no one is ever justified in believing one thing rather than another. Given the distinctions we have just drawn, we see that this conclusion concerns epistemic entitlement: personal justification. However, all the sceptic's argument shows is that there are limits to our capacity to *give reasons* or *cite evidence*. This is a point about grounding. To get from what he argues to what he concludes, the sceptic must take it for granted that no belief is responsibly held unless it rests on adequate and citable evidence. He needs the Prior Grounding Requirement. More precisely, he needs the Dependence Thesis (to link responsibility with grounding) and Strong Internalism (to identify grounding with the possession of evidence). Nothing less will do the job.

Another model

If this diagnosis is right, there are various ways to meet Agrippan scepticism, depending on which of the sceptic's presuppositions we decide to reject and how we reject them. Radical externalists reject them all. They have no interest in (PG1) and (PG2), since they deny that knowledge requires any kind of personal justification. This allows them to dispense with (PG3) and (PG4) as well since, detached from all questions of responsibility, grounding need have no links with the possession of citable evidence. However, since I have argued against radical externalism, this approach is not open to me. In my view, (PG1) to (PG4) are not wholly without point. But they are unacceptable as the sceptic interprets them.

We can preserve the link between knowledge and justification without accepting the Prior Grounding requirement. As we have noted, though so far without much elaboration, we can see justification as exhibiting what Robert Brandom calls a 'default and challenge structure'.[2] The difference between the 'Prior Grounding' and 'Default and Challenge' conceptions of justification is like that between legal systems that treat the accused as guilty unless proved innocent and those that do the opposite, granting presumptive innocence and throwing the burden of proof onto the accuser. Adopting the second model, epistemic entitlement is the default status of a person's beliefs and assertions. One is entitled to a belief or assertion (which, remember, is an implicit knowledge-claim, unless clearly qualified) in the absence of appropriate 'defeaters': that is, reasons to think that one is *not* so entitled.

Appropriate defeaters cite reasonable and relevant error-possibilities. There are two main types. Non-epistemic defeaters cite evidence that one's assertion is false: this evidence might be purely negative, or it might be positive evidence for the truth of some incompatible claim. Epistemic defeaters give grounds for suspecting that one's belief was acquired in an unreliable or irresponsible way. Here the objector concedes that his interlocutor's claim or belief might be true but denies that it is well grounded. The types are not exclusive. Sceptical scenarios are meant to work both ways.

There is something right about (PG1). The status of epistemic subject does not come with mere sentience: it has to be earned through training and education. But the sceptic (and the traditional epistemologist) give the No Free Lunch principle a much stronger reading: they take it to require that entitlement must be earned by taking *specific positive steps* in *each* situation in which entitlement is claimed. This is what allowing for default entitlements lets us deny. However, this is not to say that personal justification is completely independent of the ability to give grounds for what one believes, so there is some point to (PG2) as well. What we should reject is only the idea that a responsible believer's commitment to providing grounds is *unrestricted*. A claim to knowledge involves a commitment to respond to whatever appropriate challenges emerge, or to withdraw the claim should no effective defence be available. In claiming knowledge, I commit myself to my belief's *being* adequately grounded—formed by a reliable method—but not to my having *already established* its well-groundedness. This sort of defence is necessary only given an appropriate challenge: a positive reason to think that I reached my belief in some unreliable manner.

Notice that, on this view, the grounding required by knowledge can be understood, in appropriate cases, in the way that externalists recommend: I would only have to defend such externalist grounding if some appropriate doubt were raised. At the same time, if a belief of mine, no matter how responsibly held, were not well grounded, it would not amount to knowledge. Even though we do not have an unrestricted commitment to give grounds, our commitment to adequate grounding is always a potential entry point for

criticism. The evidentialist conception of grounding, captured by (PG3) and (PG4), is the appropriate conception for the situation where, in order to maintain entitlement, I have to meet a challenge. Meeting challenges means citing evidence (to defeat defeaters). But being able to cite evidence is not the *sine qua non* of being justified.

It is easy to miss the fact that the practice of justifying is only activated by finding oneself in the context of a properly motivated challenge. Since we do not go around stating the obvious, we do not always have to wait for challenges to *emerge*. Interesting claims are typically *not* justified by default: that is why they are interesting. However, the fact that we enter claims in the face of standing objections, automatically triggering the Defence Commitment, should not mislead us into overlooking the connection between the existence of motivated challenges and the obligation to produce positive evidence. Overlooking this connection will lead us to transform the ever-present possibility of contextually appropriate demands for evidence into a unrestricted insistence on grounds, encouraging us to move from fallibilism to radical scepticism. We can and should resist the invitation.

Is scepticism intelligible?

Let us look more closely at how rejecting the Prior Grounding Requirement in favour of a default and challenge model of justification gets us out of Agrippan scepticism.

The Agrippan sceptic is committed to the Prior Grounding Requirement because he *must* assume that the question 'How do you know that?' or 'Why do you believe that?' can *always* reasonably be entered. Implicitly, therefore, he must deny that, to be reasonable, a challenge to a knowledge-claim itself needs to be motivated by reasons. His position is that, simply in virtue of entering a claim or holding a belief, anyone with pretensions to epistemic responsibility accepts an unrestricted commitment to demonstrate entitlement to opinion. Given the Prior Grounding Requirement, this position is entirely reasonable. If all responsible believing is believing-on-evidence, the sceptic is entitled to ask for the evidence to be produced. Absent this requirement, however, this question is not reasonable at all.

The crucial feature of the Default and Challenge conception is that it saddles challengers, as well as claimants, with justificational obligations. Assuming the Prior Grounding Requirement, a request to back up a belief or assertion needs no justification: in conceding an unrestricted commitment to produce grounds, the claimant grants the sceptic's entitlement to request them. The sceptic acquires the right to issue naked challenges.

If we reject the Prior Grounding Requirement, however, the sceptic loses this right. Entitlement to enter a challenge must itself be earned by finding specific reasons for questioning either the truth of the target belief or the

claimant's entitlement to hold it, which means that naked challenges are out of order. The question 'How do you know?' or 'Mightn't you be making a mistake?' can always reasonably be met with 'What do you have in mind?' or 'What mistake do you think I might be making?' If no answer is forthcoming, no challenge has been entered and no response is required.

Rejecting the Prior Grounding Requirement thus defangs Agrippa's trilemma. There is no presumption that requests for further justification can be repeated indefinitely. At some point, they are brought to an end by default entitlements. Since these are genuine entitlements, and also subject to the Defence Commitment, they are not mere assumptions. But since they are *default* entitlements, they do not depend on any kind of citable evidence. In particular, they do not have to be self-evidencing or intrinsically credible. To be sure, default entitlement is always provisional. But this is fallibilism, not radical scepticism.

Summing up, the Prior Grounding Requirement generates a vicious regress of justification by enforcing a gross asymmetry in the justificational responsibilities of claimants and challengers. Because claimants are saddled with a standing obligation to cite evidence, challengers are accorded a standing license to request that it be cited. Since the entitlement to issue a challenge need not be earned, naked challenges are always in order. Thus, whatever a claimant offers to back up a claim, a new challenge is automatically authorized. Rejecting the Prior Grounding Requirement cancels the standing obligation to cite evidence, removes the license to issue naked challenges, and stops the regress in its tracks.

Blocking the regress this way raises deep questions about whether the sceptic can question the legitimacy of our beliefs in the general way he aspires to. On the Default and Challenge conception, which insists that claimants and challengers share justificational responsibilities, no move in the game of giving and asking for reasons is presuppositionless. Quite the reverse: all moves depend for their legitimacy—perhaps even for their full intelligibility—on commitments currently not under scrutiny, at least some of which have the status of default entitlements. This applies to challenges, as much as to claims. A motivated, thus concrete, challenge will presuppose a large background of default entitlements. All questioning, hence all positive justifying, takes place *in some definite justificational context*, constituted by a complex and often largely tacit array of current entitlements. In abstraction from all such contexts, epistemic questions simply get no purchase. It follows that although (perhaps) *any* belief may be challenged given appropriate stage-setting, there is no possibility of questioning the legitimacy of our beliefs in the *collective* way that the philosophical sceptic aspires to. On a Default and Challenge conception of justification, there is no room for either the sceptic's global doubts or the traditional epistemologist's global reassurances. Both foundationalism and the coherence theory, which try to provide such reassurance, are solutions

in search of a problem. To reinstate the legitimacy of his highly generalized doubts, the sceptic will need some further theoretical resources.

I am greatly in sympathy with this line of thought. But we need to go carefully. What, exactly, have we accomplished? Have we sidelined the sceptic (and the traditional epistemologist) by delegitimating the idea of global justification? Or have we argued for the stronger conclusion that the sceptic's doubts *do not really make sense*?

Even the issue of delegitimation is not straightforward. In a sense, the Default and Challenge conception could be seen as responding to the demand for global reassurance, assuring us that beliefs not justified by evidence can nevertheless be justified by default. Of course, this is not at all what the sceptic and the traditional epistemologist have in mind. What they want is a kind of explicit general grounding, which the Default and Challenge conception expressly excludes. Nevertheless, the fact that the Default and Challenge conception can be seen as giving an account of how, in general, our beliefs might be justified shows that we should be hesitant to rule the sceptic's questions unintelligible merely on the basis of their unusual generality.

The issue of the intelligibility of scepticism requires careful handling. Sceptical reflections are inferentially articulated: for example, in the way they constrain what counts as a response to sceptical questioning. How, then, can the sceptic be talking utter nonsense? On the other hand, pushed far enough, theoretical diagnosis engenders second thoughts about whether the sceptic is altogether intelligible. A genuinely naked challenge—one that cannot be further explained—is no more a challenge than 'Because' is an explanation. If I want to make an objection, I must be able to say *what it is* that I am objecting to. If I can't, no intelligible objection has been made. When the sceptic asks 'Why do you say that?' and we reply 'What do you mean?', we are not just shifting the burden of proof: we are trying to understand what is being demanded of us.

Faced with these conflicting pressures, what should we say? Perhaps that making sense is not an all-or-nothing matter. In a general way, we can see what the sceptic is up to: he sees himself as entitled to issue naked challenges, because he has a conception of justification that underwrites their legitimacy. But this does not mean that such challenges are really fully intelligible. They have no content beyond that conferred by a highly abstract model of the structure of justification.

This is correct as far as it goes. But we can say more. Perhaps the sceptic's naked challenges are not quite as stripped down as they first seem. Perhaps his demand for a global legitimation of our beliefs involves further theoretical presuppositions. If so, our diagnosis is not yet complete. We shall return to this question in Chapter 16. For now, I want to turn to a different kind of objection: that we have not escaped from the sceptical trap but merely deferred the sceptic's triumph.

Meta-scepticism?

Agrippan scepticism is supposed to be intuitive: to depend only on ideas that are evidently implicit in our most ordinary understanding of practices of entering and evaluating claims to knowledge. The first task of theoretical diagnosis is to dispel the illusion of intuitiveness by making explicit the ideas that the sceptic takes for granted or passes off as mere common sense. Here we have made a beginning. If my argument is on anything like the right lines, Agrippan scepticism is not intuitive at all but rather depends on a complex theoretical account of the relations between epistemic entitlement, epistemic grounding, and the ability to cite evidence.

A sceptic can reply that this exercise in diagnosis doesn't get us anywhere. What we now have are two competing conceptions of knowledge—Prior Grounding and Default and Challenge—both with some plausibility. But unless we can show that one of them is correct, the sceptic simply triumphs at second order. If the default and challenge conception is the correct theorization of our ordinary practices of epistemic assessment, we have lots of justified beliefs. But if the Prior Grounding conception is the right way of seeing such things, we do not. If we cannot decide the issue, then, for all we know, radical scepticism is the right verdict on our pretensions to hold justified beliefs. This meta-sceptical conclusion is enough for any sceptic.

If scepticism really is to be a verdict on justification, as we ordinarily understand it, it is crucial that the sceptic not be imposing gratuitously severe epistemic standards. Even if they are not immediately intuitive, his standards must be reflectively recognizable as *implicit* in ordinary practices of epistemic assessment. So the question is: do we have any reason to suppose that the game of giving and asking for reasons, as ordinarily played, respects (or even aspires to respect) the demands of the Prior Grounding Requirement? More precisely, does the Prior Grounding conception fit ordinary practices of claiming, conceding, and denying justification better than does the default and challenge conception?

Before answering this question, let me say that even if the phenomenology of everyday justification were to seriously underdetermine the choice between the competing models of justification—in other words, if both models proved to fit everyday epistemic practices more or less equally well—it would still be theoretically reasonable to prefer the default and challenge account. By hypothesis, that model fits the agreed facts equally well and has the added merit of not generating gratuitous, sceptical paradoxes. It is therefore a better account of ordinary justification.

This point is worth expansion. The Prior Grounding conception both generates the threat of scepticism and constrains our responses to that threat: with the Prior Grounding Requirement in place, we are forced to choose

between foundationalism and the coherence theory. If, as I have argued, neither option is satisfactory, the Prior Grounding Requirement leads to scepticism as the final verdict on our pretensions to hold justified beliefs. This means that the Prior Grounding conception of justification represents our ordinary practices of epistemic assessment as self-defeating. Although the whole point of such practices is to make invidious comparisons, there are no grounds for making them. A theory that represents working practices as unworkable is a bad theory.

On theoretical grounds, then, we would be entitled to prefer the default and challenge to the Prior Grounding conception of knowledge and justification, even if the two conceptions were equally faithful to the phenomenology of everyday epistemic practices. However, they are not.

Just as the Default and Challenge model requires, ordinary discussion of epistemic entitlement presupposes a sharing of justificatory responsibilities between claimants and challengers. The Prior Grounding Requirement places all justificatory burdens on claimants and none on challengers. It leaves no room for legitimate challenges to challenges. But such challenges to challenges are a pervasive feature of ordinary reason-giving.

Sceptical arguments, and the traditional epistemological theories to which they give rise, show an evidentialist bias that is far from obviously present in ordinary epistemic practices. It is certainly true that *to justify* a belief is typically to marshal evidence, offer one's credentials, explain away apparent counter-evidence, and so on. *Justifying*, in other words, just is giving grounds. But *being* justified is not always a matter of *having gone through* a prior process of justification.

Connected with this is that fact that ordinary practices of epistemic assessment reflect what Robert Brandom calls 'the social articulation of the space of reasons'.[3] Although knowledge cannot be detached in any general way from the ability to give reasons (when called for), nevertheless in special cases we can attribute knowledge to another person because *we* can defend his reliability, even if he cannot. This social distribution of reason-giving abilities allows us to inherit knowledge by deference to experts. In a complicated society, an enormous amount of knowledge is acquired this way. Here too, the Prior Grounding Requirement, which ties knowledge to an individual's ability to cite adequate evidence for what he believes, is seriously at variance with everyday practice.

These *prima facie* discrepancies are not absolutely decisive. That is, I am not claiming that the Prior Grounding conception of knowledge and justification cannot be squared with the phenomenology of everyday practices of epistemic assessment. Any theory can be squared with *prima facie* counter-evidence if we are willing to take on board enough *ad hoc* hypotheses. For example, while allowing that the Default and Challenge conception gives a good account of what we ordinarily *call* 'knowledge', the sceptic may claim

that this so-called knowledge falls short of what we can reflectively see are our own standards. It is really something less than knowledge: knowledge for all practical purposes, say. However, if the only reason for accepting this claim is that doing so saves the Prior Grounding conception from refutation, we strengthen my earlier claim that the Default and Challenge conception is preferable on theoretical grounds. Since it fits more readily with obvious features of everyday practices of epistemic assessment, that conception does not have to be laden with *ad hoc* hypotheses to fill the gaps. It is therefore a better theory.

A defender of the Prior Grounding conception will reply that the advantages are not all on one side. His conception has virtues—and the Default and Challenge conception has corresponding vices—which have so far not been recognized.

When we first introduced Agrippa's trilemma, we agreed that the sceptic will not dispute the phenomenology of everyday justification. In particular, he will admit that ordinary cases of challenge and response will often come to what the interlocutors regard as a satisfactory resolution: typically, when the claimant is able to cite something that the challenger does not dispute. But the sceptic will claim that such justification is entirely dialectical; and we can all see on reflection that such purely dialectical justification is not sufficient for knowledge. Accordingly, if all the so-called 'Default and Challenge conception' amounts to is the claim that such dialectical justification is sufficient for knowledge, that conception is inadequate.

The answer to this is that the Default and Challenge conception is not committed to the view that dialectical justification is all the justification we need. One of the theoretical strengths of the Default and Challenge conception is that it can accommodate insights from epistemic externalism, thereby allowing for genuine external constraint. I shall develop this point further in Chapter 15. But for the present, there is no reason to suppose that our diagnostic strategy leads only to meta-scepticism.

The Problem of the Criterion

Some readers may feel that the argument of the preceding section is question-begging. I have claimed that the Default and Challenge conception of knowledge and justification meshes more smoothly with the phenomenology of everyday epistemic practices than does its competitor, the Prior Grounding conception. Further, I have claimed that the Default and Challenge conception is theoretically preferable, and would be even if both conceptions fitted the agreed facts equally well. But does all this show that the Default and Challenge conception is definitely correct? Or does my position amount to the claim that we are entitled to accept it because it copes with problems better than does its rival? I seem to be arguing that the Default and Challenge conception is

default-and-challenge justified: acceptable according to its own standards. Why isn't this the Mode of Circularity?

This objection is an instance of a very ancient sceptical problem, the Problem of the Criterion. Get the anti-sceptic to articulate his own standard for knowledge or justification—his 'criterion'—and then argue that he cannot defend it in a non-question-begging way. Are we justified in adopting a Default and Challenge conception? If so, according to what standard of justification: the same standard or a different one? If we say 'different', not only are we threatened with a regress, we are admitting that the standard we want to defend is not ultimate. But if we say 'the same', we are begging the question. Since I have no wish to defend the Prior Grounding Requirement at any level, the charge of circularity is what I must address.

While the sceptic's argument sounds impressive, it invites diagnosis. The first point to notice is that there are two ways in which we might be interested in the 'correctness' of an abstract model of justification. One way, which has been our main concern up to this point, is descriptive or theoretical: is the model plausible as an explicit statement of the implicit normative structure of our ordinary practices of epistemic assessment? The other way is normative: is the Default and Challenge conception normatively correct? Is it the right way to think about knowledge and justification? Does it embody standards that ought to inform our practices, whether or not they do inform them, or that we are entitled to adopt, whether or not we have yet adopted them?

The distinction between the descriptive and the normative issue is not absolute. Bringing implicit structures to light may well involve an element of idealization, blurring the distinction between exposing our presuppositions and revising them. But distinctions can be useful without being knife-edged.

Taking the question of the 'correctness' of the Default and Challenge conception descriptively, there is a kind of circularity in defending it according to Default and Challenge standards. But the circularity is not vicious. We can responsibly accept this conception if we can meet the relevant challenges; but setting this as the standard we *must* meet does not guarantee that we *can* meet it. An epistemological theory can fail to be justifiably acceptable by standards that the theory itself sets. Self-vindication is not guaranteed in advance.

The same general answer applies, taking the question normatively. However, in this case matters are complicated by the issue of what considerations bear on the correctness of epistemic standards. Or is talk of correctness, in the sense of truth, out of place with respect to norms? For aren't norms like the rules of a game, mere arbitrary conventions?

The question of how to understand the correctness of epistemic norms is important. Indeed, it is the key to a deeper diagnosis of scepticism than we have achieved so far. I shall therefore explore it further in succeeding chapters. For now, however, I suggest that we look at the issue pragmatically. In saying this, I mean that we take into consideration the interests that are subserved by

practices of epistemic assessment. (Recall from the Introduction my claim that there must be some *value* in knowledge, something that gives us an interest in having it.) Bringing in interests, we can see that likening epistemic standards to the rules of a game (the game of giving and asking for reasons) does not make them arbitrary. After all, the rules of football are not arbitrary. Rather, they are adapted to human physical capabilities, in the light of our interest in playing a game with a certain level of competitiveness. If the goals were forty feet wide, scoring would be too easy. If they were two feet wide, no one would ever score. Similarly with standards of justification: they can be evaluated in the light of our epistemic interests: avoiding errors, coming to believe significant truths, improving our theories, and so on.

Viewed in this pragmatic perspective, the Prior Grounding conception does particularly badly. No proposed normative structure for our epistemic practices is useful if it precludes making any distinction between justified and unjustified beliefs. Unless this distinction can be made, there are no such practices. The Prior Grounding Requirement, which makes scepticism unavoidable, is self-defeating in just this way. The Default and Challenge conception, which heads off scepticism, is normatively preferable.

Of course, this assumes that the default and challenge conception does not also self-destruct. To be sure that it is satisfactory, we need to develop it further. We also need to see how it helps us deal with Cartesian scepticism. In pursuing these projects, we will be led to deepen our diagnosis of both sceptical questioning and the traditional epistemological theorizing it inspires.

Notes

1. The distinction between theoretical diagnosis and direct refutation is not a sharp one. Inevitably, what the sceptic acknowledges and what he leaves unsaid depends on how sceptical arguments are formulated. In practice, however, the distinction is clear enough. This is because the need for sceptical arguments to appear to be natural or intuitive sets limits to what the sceptic can afford to acknowledge. The danger of over-explicitness is that his scepticism will modulate from philosophical scepticism to scepticism about certain philosophical theories of knowledge. He will not show, in an unqualified way, that knowledge is impossible but only that it is impossible given certain adventitious standards, dictated by controversial epistemological ideas.
2. Although the phrase is Brandom's, there are many prior articulations of this general conception. One of the most important is to be found in J. L. Austin's seminal essay 'Other Minds', in Austin (1961). I believe, however, that this conception of justification is very ancient, originating in Academic theories of 'sceptical assent', particularly Carneades' doctrine of the 'tested impression'. Since the ancient sceptics and their opponents contrast sceptical assent with knowledge, ancient attacks on the possibility of knowledge are not necessarily radical in my sense. However, it is notable that the Pyrrhonian sceptics attacked even 'sceptical' epistemologies like that of Carneades. Default and Challenge structures show up in non-epistemological accounting practices

too. Consider a different sense of 'responsibility': accountability for one's actions. Here, again, 'responsibility' is the default position: one is accountable unless in possession of an appropriate excuse. This view of responsibility is the one taken by Austin in his famous paper 'A Plea for Excuses', in Austin (1961). Austin's views on knowledge and freedom are importantly connected.

3. Brandom (1995).

14

Knowledge in Context

Contextualism

The Agrippan problem seems to force us to choose between two conceptions of knowledge: as having either a foundational structure or a coherence (mutual support) structure. Our critical examination of the Prior Grounding Requirement has thrown up a third possibility: that knowledge conforms to a Default and Challenge structure. In this model, questions of justification always arise in a definite justificational context, constituted by a complex and in general largely tacit background of entitlements, some of which will be default. Thus the Default and Challenge model leads to what I shall call (with some misgivings) a *contextualist* picture of justification.[1]

Because the contextualist picture shows affinities with both foundationalism and coherentism, it does not amount to a simple rejection of either of these traditional approaches. However, since it also differs sharply from both, it is best seen as *sui generis* and not as a variant of either traditional theory. But to make good on this claim, I need to clarify what I understand by a contextualist conception of knowledge and justification.

Aspects of context

The fundamental idea of contextualism is that standards for correctly attributing or claiming knowledge are not fixed but subject to circumstantial variation. We can classify the factors that influence the epistemic status of claims and challenges under five main headings.[2]

In the first place, justification is subject to *intelligibility or semantic constraints*. Wittgenstein remarks that, if you tried to doubt everything, you would not get as far as doubting anything. This is not a matter of practicality: one reason we have lots of default entitlements is that holding many true beliefs, or not being subject to certain kinds of error, is a condition of making sense, thus of being in a position to raise questions at all. Unless we routinely get lots of things right, it is not clear what we are talking or thinking *about*, if anything. Wittgenstein makes the point with characteristic flair: 'Suppose a man could not remember whether he had always had five fingers or two

hands? Should we understand him? Could we be certain of understanding him?'[3] The answer is: no, we could not be certain. At some point, 'mistakes' shade off into unintelligibility. Someone who cannot do the simplest calculations, or perform the simplest counting operations, is not making arithmetical mistakes: he does not understand numbers. Of course, there is no sharp boundary here—that is why mistakes shade off into unintelligibility—but the fact remains that one cannot, in just any circumstances, be mistaken about anything whatsoever. A contextualist must therefore be careful about how he states his fallibilism. It is tempting to say that anything can be called in question, but not everything at once. But it is not true even that anything can be called in question in any situation. To be intelligible at all—and not just to be reasonable—questioning may need a *lot* of stage-setting. As we shall see, this is true of the sceptic's attempt to call in question our most ordinary and obvious judgements about the world around us.

Intelligibility constraints have to do with our being able to raise meaningful questions at all. But the exclusion of certain types of doubt can also arise from what is required to raise questions of some specific kind. I shall call these exclusions *methodological constraints*; and I shall call propositions that have to be exempted from doubt, if certain types of question are to be pursued, *methodological necessities*.

The boundary between intelligibility and methodological constraints is not a sharp one. Methodological constraints can be seen as intelligibility constraints specific to certain forms of investigation. We are thus led to our second—broadly methodological—type of factor influencing contextual constraints on justification. Methodological necessities are a source of default entitlements because they determine the *direction of inquiry*. For example, serious worries as to whether the Earth even existed five minutes ago, or whether every piece of documentary evidence is some kind of forgery, do not result in an especially scrupulous approach to historical investigation. Rather, they preclude any engagement in historical research. The need to recognize methodological limitations on doubt is not, as the sceptic has it, a reflection of our practical limitations but a fundamental fact about the *logic* of inquiry. Nor is it a matter of our applying more relaxed standards, lowering what Robert Fogelin calls the 'level of scrutiny' to which claims are subjected.[4] The direction of inquiry has to do, not with the level, but with the *angle* of scrutiny. There is no simple relation between level and angle. Within the practice of doing history, we can be more or less strict in our standards of evidence. But some questions have to be set aside for us to think historically at all.

While particularly clear with respect to the methodological necessities of organized forms of inquiry, such as historical research, the point that the direction of inquiry limits the range of the dubious is entirely general. What we *are* looking into is a function of what we are leaving alone. We can no more inquire into everything at once than we can travel simultaneously in all directions. This point applies as much to the most arcane theoretical inquiries as to practically oriented investigations. Indeed, to speak of it as a 'limitation' is misleading. The point of such constraints is to make focused questioning possible. Such constraints are agents of change, not bulwarks of dogmatism.[5]

Factors of the third type are *dialectical*. Given a certain direction of inquiry, various possible defeaters may or may not be in play. Sometimes, claims may face standing objections, in which case they will not enjoy default status. But default status can be lost as new problems arise, just as hitherto accepted justifications can be undermined by new evidence. The epistemic status of claims and beliefs changes with developments in the dialectical environment.

Factors of the fourth type are *economic*. A defeater does not come into play simply by virtue of being mentioned: there has to be some reason to think that it might obtain. How much reason we require fixes the severity of our epistemic standards or level of scrutiny. If we insist on ruling out even very remote error-possibilities, we are imposing severe standards for knowledge and justification. Costs and benefits—economic factors—figure in here. If it is important to reach some decision, and if the costs of error are fairly low, or if we gain a lot by being right and lose little by being wrong, it is reasonable to take a relaxed attitude to justificational standards. If the costs of error are high, more demanding standards may be in order. The opportunity costs of further inquiry can also be relevant. Of course, in referring to 'economic' factors I do not have in mind only monetary or material considerations. Anything we value is a benefit and anything we would rather avoid is a cost. I call these considerations 'economic' to stress the point that there is typically no purely epistemological answer to the question of what level of epistemic severity is contextually reasonable.

Methodological, dialectical, and economic factors concern primarily the epistemic responsibility dimension of justification. They reveal that the relationship between personal and evidential justification is multiply contextual. In the first place, with respect to maintaining epistemic responsibility, the *existence* of a properly motivated challenge determines whether evidential justification—in the strict sense of citable evidence—is required *at all* to secure personal justification. In the second place, contextual factors fix the adequacy conditions on evidential justification's securing personal justification. Most importantly, they determine what potential defeaters ought to be excluded. These will never amount to every logically possible way of going wrong, but will be restricted to a range of relevant alternatives.[6] However, there are two sources of irrelevance that must not be confused. An error-possibility may be beside the point—strictly irrelevant to the subject in hand—or, while not strictly irrelevant, it may be too remote a possibility to be worthy of serious consideration.

So much for personal justification. But for a person to have knowledge, his belief must be adequately grounded (whether or not he is aware of its grounds

and can cite them as evidence). Here, contextual factors of a fifth type come into play. We can call these *situational* factors. Epistemic contexts are not exhausted by methodological and dialectical considerations: facts about the actual situations in which claims are entered or beliefs held are crucial too. This is because, in claiming knowledge, we commit ourselves to the objective well-groundedness of our beliefs. We are open to the existence of relevant defeaters that we have overlooked or not yet uncovered, even if blamelessly. Our commitment to groundedness is thus an important source of openness to self-correction.

Situational factors highlight the externalist element in contextualism. The 'adequate grounds' dimension of justification has a doubly 'external' character. Because there is default entitlement, a claimant need not always be aware of the grounds for his belief, in order to be epistemically responsible. But the adequacy of his grounds—whether he is aware of them or not—will depend in part on what real-world possibilities those grounds need to exclude. I say 'in part' because standards of adequacy are always standards that we fix in the light of our interests, epistemic and otherwise. Even when considering the objective adequacy of grounds for a belief, questions of epistemic responsibility can never be wholly forgotten.

An example

A simple example will illustrate how contextual constraints on justification operate. Seeing someone drive by in an old sports car, two people engage in the following exchange:

A: Isn't that old sports car an E-Type?
B: Yes, a rare early model.
A: What makes you say that: don't they all look pretty much the same?
B: Sure, but that one had external bonnet latches which you only get on the first five hundred cars.

Here A and B concede to each other various default entitlements—that an old car just passed by, and so on. Without such concessions, their conversation could not take the specific direction it does take (methodological factors). Indeed, if certain things (e.g. the capacity to tell a car from an elephant) could not be taken for granted, the speakers could not have any kind of intelligible exchange about the vehicle passing by (semantic factors). But B makes a claim implying special expertise. Although sometimes this alone may be enough to invite a challenge, here A has a specific reason for querying B's identification. He points out, correctly, that the various versions of this particular car are very similar. If, as A suspects, B's remark was prompted by a casual glance at the passing car, this is an epistemic defeater. Citing it challenges B to back up his claim, altering the dialectical context.

B's response is sufficient to meet A's challenge. Or rather, it is sufficient given the severity of epistemic standards appropriate to a casual conversation. For example, B has not ruled out the possibility that what seem to be external latches are non-functional, added to make a later car look like an early model. Normally he wouldn't be expected to. But if A were thinking of buying the car (an economic factor), this might be worth looking into.

Even if the evidence B cites is not defective this way, it still does not rule out every conceivable way in which his claim might have been false. For example, it does not exclude the possibility that what they have just seen is a *replica* of the model in question. But B has no obligation to exclude this possibility unless there is some reason to suspect that it might obtain. Of course, A could try to provide a reason (appropriate to the operative standard of justificational severity):

A: Maybe it's not the real thing. I read that some firm was making replicas of vintage Jaguars.
B: Not E-Types. There are so many survivors that a replica would cost more than an original.

Unless A can think of some way of challenging B's latest response, B maintains entitlement to his claim.

By showing that his original claim was not epistemically irresponsible, B secures personal justification. Still, we might wonder about his evidential justification. Evidence is supposed to be objectively adequate: to support a knowledge claim, it has to 'establish the truth' of what has been stated. Do B's grounds accomplish this, given that there are ways of going wrong that they do not exclude? After all, B's argument does not rule out the possibility that some wealthy enthusiast has had a replica built regardless of cost, perhaps with updated mechanical components to improve reliability.

It is in connection with the objective adequacy of grounds that situational factors come into play. We can think of these factors as determining a broader informational context or environment constituted by relevant facts, known or knowable, though not necessarily by or to the epistemic agents in question. To see what this comes to, suppose that, in fact, nobody does go in for building replica E-types, or for modifying later cars to look like early models: in this situation, B's grounds for identifying the car as an early model are objectively as well as dialectically adequate though, in a different situation, they would not have been.

There is considerable indeterminacy about the objective adequacy of grounds, resulting from the fluidity of contextual boundaries. Suppose that a few relevant replicas have been made. But all have been sold abroad, so that there aren't any where A and B live. Are B's grounds (the external locks) still objectively adequate? It depends how we draw the boundaries. How we do this always depends at least in part on considerations of epistemic responsibility.

Suppose that B had known that there were a few replicas in foreign hands: would he still have been responsible in making his claim? Surely he would: there is no reason to demand evidence that excludes such (literally) distant possibilities, even when their existence is acknowledged. However, the fact that standards of adequacy can never be fixed in total abstraction from questions of epistemic responsibility (and the interests that guide them) does not compromise the objective character of the grounding requirement. The fact that considerations of responsibility help set the standards for a belief's being adequately grounded does not guarantee us the ability to meet those standards. We can and do fail to satisfy standards that we set for ourselves.

Contextualism and foundationalism

To see how contextualism relates to foundationalism, we must recall that foundationalism comes in two forms. One is structural foundationalism:

> (STF) (i) There are basic beliefs, beliefs that are *in some sense* justifiably held without resting on further evidence. (ii) A belief is justified if and only if it is either itself basic or inferentially connected, in some appropriate way, to other justified beliefs.

The other is foundationalism proper, substantive foundationalism:

> (SUF) (i) and (ii) as above. (iii) There are certain kinds of beliefs (or other forms of awareness—e.g. experiences) that, by their very nature—that is, in virtue of their content—are fitted to play the role of terminating points for chains of justification. These beliefs (or other conscious states) are epistemologically basic because *intrinsically credible* or *self-evidencing*.

Contextualism is definitely *not* substantively foundationalist. Default justification is not associated with any particular kinds of belief. Whether a belief enjoys this status depends on a large number of contextually variable factors: the current state of knowledge (including critical responses to received views), the particular inquiry in which we are engaged, and much else besides. Indeed, the very same belief can go from default to non-default justificational status. So, not only is there no need for the contextualist to postulate special 'epistemic kinds' of basic beliefs, it would go against the spirit of contextualism to do so.

While contextualism is definitely *not* substantively foundationalist, it looks *prima facie* like a form of structural foundationalism. Since our default entitlements do not depend on further evidence, don't they amount to a kind of heterogeneous and shifting foundation? If you like. The fact remains, however, that contextualism differs sharply from traditional foundationalism, even when foundationalism is viewed as a largely structural doctrine. As we saw, foundationalism embodies an atomistic conception of justification:

intrinsic credibility ensures that each basic belief is justified all by itself, without regard to further beliefs, the context of inquiry, or our real world situation. The justificational relevance of basic to non-basic beliefs enjoys a similar independence of such contextual factors: justificational relations between beliefs depend on content (meaning) alone. Contextualism rejects both aspects of the foundationalist's epistemological atomism. In any context of justification, explicitly or tacitly, there will always be a large number of beliefs and commitments in play. As we shall see in the next chapter, some of these commitments can be epistemically relevant without playing a direct justifying role. For example, we may need to know certain things in order to be able to recognize what would or would not count as a reasonable challenge to a particular claim we have advanced. But this knowledge will not function as grounds for the claim in question. This means that, in a given context of justification, certain commitments function as fixed points without its being correct to see whatever judgement is the focus of interest as *resting* on them.

This point is worth elaborating. I have said that, in any context of justification, there will always be lots of further beliefs and commitments in play. Exactly what they are will depend on what is at issue. This does not exclude the possibility that certain propositions function as fixed points across a wide range of contexts. Indeed, for some propositions, there may be no contexts in which they are up for grabs: we simply have no idea what would count as a challenge to them. They will, as Wittgenstein puts it, 'lie apart from the route traveled by inquiry'.[7] Nevertheless, such fixed-point propositions need not be seen as foundations in anything like the traditional sense. As Wittgenstein explains:

> I do not explicitly learn the propositions that stand fast for me. I can discover them subsequently like the axis around which a body rotates. This axis is not fixed in the sense that anything holds it fast, but the movement around it determines its immobility.
>
> No one ever taught me that my hands don't disappear when I am not paying attention to them. Nor can I be said to presuppose the truth of this proposition in my assertions etc., (as if they rested on it) while it only gets its sense from the rest of our procedure of asserting.[8]

As this example suggests, ordinary practices of empirical inquiry and justification are not beset by sceptical doubts about the continuing existence of physical objects, such as my hands. But this is not just a matter of making an assumption. These bedrock certainties derive their content—their meaning—from the particular practices of inquiry and justification that hold them in place. To believe in an historical past, or an external world, just is to recognize certain types of error possibility, to demand certain kinds of evidence (in appropriate circumstances), and so on. Because such certainties are *semantically embedded* in our epistemic practices—thus unintelligible apart from

them—it is a mistake to see those practices as justificationally dependent on the 'presuppositions' they embed. They are not assumptions because they are not, in the relevant sense, foundational at all.

These reflections point to a very deep difference between contextualism and traditional foundationalism. As I noted in Chapter 1, our epistemological tradition has tended to focus on propositional knowledge, drawing a sharp distinction between 'knowing that' and 'knowing how'. From a contextualist standpoint, it is wrong-headed either to insist on a sharp distinction or to treat knowing-that as theoretically fundamental. Certainly, we can investigate the contextual constraints and presuppositions governing particular kinds of inquiry: but it is doubtful in the extreme whether we can make them fully explicit in the sense of listing them exhaustively or reducing them to some simple set of rules. Relevant evidence, appropriate objections, and suitable replies are things we have to learn to recognize by projection from examples. No one who lacked this ability could make much of propositionally formulated rules or presuppositions. Possessing it is what we call having 'good judgement', though there is a (vague) limit to how bad a person's judgement can be without his lacking the capacity to judge at all. For a contextualist, there cannot be a sharp distinction between knowing-that and knowing-how because being able make judgements—the precondition of any knowing-that—involves know-how essentially. This is why propositional knowledge and certain kinds of know-how are acquired together. Propositional knowledge is not self-contained: not because it rests on some pre-propositional knowledge-by-acquaintance, but because it is embedded in the practical mastery of forms of discourse and inquiry.

Contextualism and the coherence theory

Contextualism and coherentism have certain affinities. Contextualism embodies a kind of local or modular holism. Thus, like the coherence theorist, the contextualist rejects the thought that particular beliefs have an intrinsic epistemic status and lacks an interest in sorting out our beliefs into *a priori* and *a posteriori*, necessary and contingent, or analytic and synthetic. It is not so much that these traditional distinctions are shown to be completely 'incoherent'. Rather, they begin to look like ways of talking we could usefully do without. From a contextualist standpoint, the empiricist account of *a priori* knowledge in terms of analytic or conceptual truth runs together a number of ideas that should be disaggregated: justification without empirical evidence, meaning-constitutivity, unrevisability, and truth in virtue of meaning alone. The methodological necessities that inform some particular kind of inquiry may be held as default entitlements and may fix what we mean by, say, 'historical evidence'. But this does not make them unrevisable: if the kind of investigation they sponsor runs into trouble, they may need to be changed. And if they do not have to be true, they do not have to be true by virtue of meaning.

Again, however, there are differences to balance the similarities. Principally, while anti-atomist, contextualism is not *radically* holistic. Although it insists that justification always presupposes a *critical epistemic mass* of contextually relevant beliefs, contextualism has no use for the coherence theorist's pseudo-totality: our 'system of beliefs' or 'total view'. It is not that we *cannot* think of the sum total of our commitments and methods of inquiry and argumentation as a system (unless doing so leads us to overestimate the extent to which the 'system' is systematic). The point is simply that radical holism gives a misleading picture of justification and inquiry. For the contextualist, all questions of justification arise against and depend for their full intelligibility on presuppositions that are (reasonably) not currently in question. Some will be standing commitments, held in place across a wide range of epistemic or investigative contexts, whereas others will be specifically relevant to the matter in hand. While in neither case will they be absolutely unquestionable, questioning them will shift the focus of inquiry in ways that take other commitments out of the line of fire. Furthermore, when we justify a claim by producing evidence, that evidence will always depend for its significance on a broader informational environment which is never fully surveyable. If these points are accepted, the idea of 'global' justification must be regarded as chimerical. Thus, by offering a principled rejection of the very idea of 'global' justification, the contextualist obviates the need for either criteria of global coherence or privileged access to the contents and structure of one's supposed total view. Contextualism therefore avoids the coherence theory's collapse into foundationalism.

The coherence theory builds in a sharp distinction between factual commitments and epistemic norms. Our factual commitments are expressed by the beliefs that make up our system or total view. Their epistemic status—how well justified they are—depends on how well the system they constitute satisfies certain general epistemic norms or standards: the criteria of coherence. In the contextualist picture, the norm/fact distinction is more methodological than ontological (to borrow a phrase from Brandom).[9] Exempting certain propositions from doubt fixes the direction of inquiry, making a particular context of inquiry or justification the context that it is. Such propositions thus play a normative role: not questioning them is part of the rules of the game. As Wittgenstein says, 'we use judgments as principles of judgment'.[10] Nevertheless, they may represent substantive factual commitments: situations may arise in which they can and should be subjected to scrutiny. They do not 'lie apart from the route traveled by inquiry' because of some special subject-matter but because of the functional role that we accord them in our investigative practices (and which with effort and imagination we might rescind).

From a contextualist point of view, both foundationalism and the

coherence theory are overreactions. Seeing the need for fixed points to block the indefinite repetition of requests for evidence, the foundationalist postulates a stratum of intrinsically credible or self-justifying beliefs. The coherence theorist sees the problems that the demand for intrinsic credibility gives rise to. In particular, he notices that individual beliefs owe their justificational status and significance to relations with further beliefs. But he goes from this reasonable anti-atomism to radical holism, with its own serious difficulties. Contextualism, because it allows both fixed points and epistemic interdependence, has a good claim to incorporate the best features of its traditionalist rivals.

Contextualism and the sources of knowledge

From the outset, I have treated epistemological questions as normative. As Kant made clear to us, the philosophical question is not one of simple *fact* (Where do our beliefs come from?) but one of *right* (What entitles us to hold them?).[11] Historically, however, most philosophical discussions of knowledge have focused on identifying the sources of knowledge. The question has been 'Where does knowledge come from: Reason, the Senses, Revelation?' Indeed, it is because of this stress on sources that Kant is so insistent on distinguishing between *Quid facti?* (the question of fact) and *Quid juris?* (the question of right). He thinks that talk of sources of knowledge is apt to obscure the essentially normative character of epistemological questions. I agree. Such talk encourages us to think that asking 'Does knowledge comes from the senses?' is like asking 'Do diamonds come from South Africa?', whereas these questions are quite different in character.

Now, as I noted in Chapter 7, traditional talk of the sources of knowledge need not entail ignoring the normative dimension of epistemological assessment. This is because the faculties that are identified as the sources of knowledge (as opposed to opinion) are conceived in partially normative terms as recognizably reliable and therefore authoritative. Nevertheless, Kant has a point. Talk of the sources of knowledge is not harmless.

To begin to see why, we should take note of two further features of traditional 'sources' of knowledge. In the first place, they are generic: the sources of knowledge are 'the senses' or 'Reason'. Secondly, they are *ultimate*. Such authoritative faculties are sources of what Richard Rorty calls 'privileged representations', beliefs that are the basis for all further inference because they themselves possess a special credibility, derived from their pedigree. In identifying such beliefs, we reach rock bottom: questions of justification cannot be pressed any farther. The picture of knowledge as belief that derives from authoritative sources thus creates a strong prejudice in favour of substantive foundationalism. It also builds in a kind of meta-epistemological foundationalism. Once questions about justification are raised to the level of epistemological theory, citing a generic source, the pedigree that gives basic beliefs their special status, is the last word.

All this bears on questions raised in the previous chapter but not fully resolved there: whether the sceptic's apparently naked challenges are fully intelligible, and whether they are really naked. If I make a claim, implying that I know what I say to be true, the sceptic asks me how I know. Obviously, the sceptic's 'How do you know?' is asked, as Austin says, *pointedly*, the implication being that perhaps you *don't* know. His question is a challenge to me to defend my right to make or accept the claim I have made. Taking for granted the Prior Grounding Requirement on knowledge, the sceptic does not think that this challenge itself needs any licensing or even further explanation. Since I am the one laying claim to knowledge, I am the one who needs to provide grounds, which is as good as to say that naked challenges are entirely in order. By contrast, contextualism insists on further explanation. If a challenger implies that we might be making a mistake, we are entitled to ask how. If the challenger has nothing to say—if his challenge is genuinely naked—then no real challenge has been entered. We have no idea what sort of defence is being demanded of us.

Why doesn't the sceptic see this? And why doesn't he see the emptiness of his naked challenges as reflecting on the conception of justification that licenses them? I think we have the answer. The sceptic's challenge is not altogether naked. Rather, it is generic. The sceptic meets a claim, any claim, with a challenge. If we accept it, our immediate reaction will be to cite evidence specifically relevant to the claim under challenge. But the sceptic's willingness to renew his challenge at every stage indicates that this is not the sort of answer he is looking for. Indeed, my ability to give specific evidence is liable to run out fairly quickly. If you ask me how I know that the car that just went by is an early model E-Type, I can say that I recognized it by its external bonnet locks. But if you ask why I suppose it had bonnet locks, I am not sure what to say, other than that I had a good look at the car and could see them. This is the answer the sceptic wants. In giving it, I raise the debate to the epistemological level. As the sceptic interprets my reply, I commit myself to the authoritative character of a certain generic source of knowledge, in this instance observation or 'the senses'. The real, though typically unavowed, function of the sceptic's apparently naked challenges is to raise the level of the debate in just this way. From the word go, they are implicitly generic. As indefinitely renewable, the very first challenge, if simply accepted, is already a challenge to any and all specific responses it may give rise to. What makes this procedure reasonable—and intelligible, to the extent that it is intelligible—is the presupposition that there are and must be generic sources of knowledge, if there is to be such a thing as knowledge at all.

What we know of ancient scepticism comes mostly from the writings of Sextus Empiricus. In Sextus' writings, the Agrippan strategy is used almost

exclusively in connection with the Problem of the Criterion: the problem of validating a proposed source of knowledge. If what I have been arguing is anything like correct, this is no accident. We ascend to the epistemological level to escape the sceptic's infinitely renewable challenges. But once we have agreed to trace knowledge to some ultimate, generic source, there really will seem to be no alternative to either refusing to defend it (thus making an ungrounded assumption) or appealing to that source itself (arguing in a circle). The Agrippan trilemma seems inescapable. However, from a contextualist standpoint, this entire way of thinking is deeply suspect. Many factors influence epistemic status; default credibility has no connection with one particular source or faculty; and since anything we presuppose is potentially criticizable, there is no point in thinking of any source of knowledge as ultimate. Given appropriate stage-setting, erstwhile default entitlements can be called in question. But this is just the open-endedness of inquiry, not a vicious regress of justification. The possibility of error does not imply the impossibility of knowledge.

Epistemological realism

The sceptic raises the Problem of the Criterion by getting his interlocutor to articulate some standard source of knowledge. He then argues that no such standard can be defended in a non-question-begging way. Applied to my diagnosis of scepticism, this problem can be used to suggest that neither the default and challenge conception, nor the contextualist picture of justification to which it leads, can be shown to be correct. Or at least, cannot be shown to be correct by any neutral standard. We may in some sense be entitled to adopt this picture, but we cannot know that it is true.

We are now in a position to see that this charge either ignores the normative character of justification, or treats it in an implausible way. To be sure, the Prior Grounding and Default and Challenge conceptions set different standards for epistemic responsibility, hence for epistemic entitlement; but it is a bad first move to ask, in a flat-footed way, which conception is true. This is to proceed as though there were some fact of the matter—some fact about what the correct standards of epistemic justification are, or ought to be—that holds quite independently of what we take them to be. This is not how things are. Norms, including epistemic norms, are *standards that we set*, not standards imposed on us by 'the nature of epistemic justification'. A belief is no more justified, wholly independently of human evaluative attitudes and practices, than a certain kind of tackle in football is a foul, wholly independently of our practices of judging certain types of tackle to be against the rules.

In saying that epistemic norms are standards we set, I am not supposing that we ever got together to set them. As I have already argued, the constraints

that govern particular forms of inquiry exist, in the first instance, implicitly in practice rather than explicitly as precepts. But we can make them (partially) explicit should the need arise; and if it seems like a good idea, we can modify them.

The view I am recommending can be considered a pragmatic conception of norms. The alternative to pragmatism is *epistemological realism*. I briefly introduced the idea of epistemological realism back in Chapter 7, in connection with the distinction between structural and substantive foundationalism. To repeat, epistemological realism is not a metaphysical or ontological position within epistemology: the view that there is a real world out there, which we want to know about. It is a form of extreme realism about the objects of epistemological theory: the view that we have some fixed 'epistemic position' determined by facts about the nature of knowledge or the structure of justification. With the contextualist picture partially sketched in, we see how strong a theoretical commitment this is. We also see that epistemological realism is by no means forced on us by evident features of ordinary justification. It represents a definite choice of theoretical orientation and is, in that sense, optional.

By encouraging us to assimilate normative questions about epistemic standards to causal-factual questions about origins, traditional talk of the sources of knowledge tempts us into epistemological realism. If there are certain ultimate, generic sources of knowledge, which fix our 'epistemic situation', epistemically responsible believing must pay these sources due respect. If it turns out that our ultimate sources are not really up to the job, our epistemic situation is intrinsically defective and we cannot know anything. Or if we cannot show that they are up to the job, then for all we know, we know nothing. One way or another, the sceptic wins the day. By tacitly invoking epistemological realism, the sceptic implies that we are *stuck* with his conception of knowledge: that we cannot responsibly change it, unless we know (and can prove) that it is false. After all, the sceptic is merely opening our eyes to how things are. The right response is: so much the worse for epistemological realism.

This is not the end of the story. The contextualist/pragmatic outlook provokes anxieties of its own. If we set the standards, why can't we believe anything we like, so that objective knowledge is an illusion? Or allowing that, given a particular set of rules, there are correct and incorrect ways of proceeding, aren't different groups free to play by different sets? In other words, does contextualism imply relativism? And if standards can change, is there such a thing as progress?

We have already seen that, in the contextualist picture, inquiry and justification are constrained in several ways. But in allaying anxieties about objectivity, the most important constraint is that provided by observation. Accordingly, observational knowledge will be our next topic. Having

developed contextualism further, and having applied its lessons to Cartesian scepticism, we will take up questions about relativism and progress.

Notes

1. My misgivings have to do with the fact that there are several epistemological views described by their advocates as 'contextualist'. See n. 2 below.
2. All contextualist epistemologies involve the basic idea that the standards for attributing knowledge, or justified belief, are in one way or another subject to contextual variation. However, contextualists differ over the sources and effects of this variability. For alternative contextualist epistemologies see Annis (1978), Cohen (1988), DeRose (1995), and Lewis, 'Elusive Knowledge', in Lewis (1999). I think that Fogelin (1994) is a kind of contextualist, though he repudiates the label. For a response to Fogelin that is also relevant to the views of Cohen and DeRose, see Williams (1999*b*).
3. Wittgenstein (1969), para. 157.
4. Fogelin (1994), ch. 5, p. 93f.
5. Similar remarks apply to the suggestion that, even if practically excusable, it is always *epistemically* irresponsible to believe anything on less than adequate evidence. Fogelin calls this view 'Cliffordism', after W. K. Clifford, and claims that traditional theories of knowledge and justification result from the attempt to live up to (unreasonable) Cliffordian standards. See Fogelin (1994), ch. 6. I have learned a great deal from Fogelin's insightful discussion of these matters.
6. The relevant alternative account of knowledge has its roots in Austin. It is further developed in a number of important papers by Fred Dretske. See Dretske (1972) and (1981).
7. Wittgenstein (1969), para. 88.
8. Ibid., paras 152–3.
9. Brandom uses this phrase in connection with Sellar's distinction between observational and theoretical terms. See Sellars (1997), 163. Sellars himself says that this distinction is methodological rather 'substantive' (ibid. 84). Brandom's reformulation is an illuminating gloss.
10. Wittgenstein (1969), para. 124.
11. Kant (1964), p. 120.

15

Seeing and Knowing

Observational constraint

Although foundationalism must be rejected, it can be misleading to say without qualification that knowledge has *no* foundations. As Sellars remarks, to put the point this way

> is to suggest that [empirical knowledge] is really 'empirical knowledge so-called', and to put it in a box with rumours and hoaxes. There is clearly some point to the picture of human knowledge as resting on a level of propositions—observation reports—which do not rest on other propositions in the same way as other propositions rest on them.[1]

Foundationalism, especially in its empiricist form, is an attempt to respond to a legitimate demand: that our beliefs be responsive to observational evidence, which must in turn be intelligible as a reliable source of information about the world. Moreover, the requirement of observational constraint is not merely epistemological in a narrow sense. Unless we can see how our reason-giving practices hook up with the world, they are going to look like games played with meaningless counters.[2] Our problem has a semantic as well as an epistemic dimension.

How should we think about observational constraint? The need to see observation as a reliable source of information about our surroundings suggests that we must incorporate insights from externalist reliabilism. At the same time, pure reliabilism is not an option. If *de facto* reliability is to link up with reason-giving, we must make room for reliability-knowledge: knowledge of the extent to which we are reliable observers of the passing scene. But in allowing for reliability-knowledge, we must not relapse into the coherence theory. To find one's way through this maze is no simple matter.

Reasons and responses

The account of observational knowledge I shall present is due in all essentials to Sellars.[3] Following Sellars's lead, I shall begin by not worrying about 'inner' experiences. Instead, I shall discuss observational knowledge as if it were solely

19

Relativism

Relativism and scepticism

Relativism is much in the air. Various intellectual movements, such as 'post-modernism' and 'social constructivism', are associated by their critics (and sometimes their proponents) with a relativistic outlook, though how far and in what sense the programmes and positions that go under these labels are relativistic is not always clear.[1]

Relativism takes two main forms, subjective and cultural. The subjectivist says that nothing is justified *simpliciter*: things are only justified-for-me. The cultural relativist is less individualistic: he thinks that beliefs are justified for particular 'cultures'. Cultural relativism sometimes leads to the embrace of a 'standpoint epistemology', according to which ethnic, class, gender, or other 'cultural' differences are associated with distinct 'ways of knowing'. But in general, relativism, whether subjective or cultural, can encompass not only substantive beliefs about the world but criteria of justification, methods of investigation and inference, styles of explanation or understanding, and so on. There is no common ground, substantive or methodological, on which argument between those who see the world from different standpoints can proceed. Clashes between them are thus power struggles, 'objectivity' being no more than the dominant world-view's preferred self-description.

I have stated relativism as a thesis about justification. But relativism can also be presented as a doctrine about truth: that nothing is true *simpliciter* but only true-for-me or for some culture. However, I doubt that there is a deep difference here. Relativists reject the idea of objective truth because they have grave doubts about the existence of neutral methods or criteria for establishing the truth. Relativists are suspicious of objective truth because they are suspicious of objective justification, not the other way around. An 'objective truth' that was beyond any human being's capacity for knowledge would be no threat to relativism.

No one needs to be told that human beliefs vary widely across time and space; that differences of opinion can extend to fundamental methodological ideas; and that when they do, disputes can be difficult, or even practically impossible, to resolve. However, relativism is a philosophical doctrine that goes far beyond such obvious facts. The relativist does not hold merely that differences of outlook are often so profound that disputes involving them can be intractable, but rather that there are principled, across-the-board objections to such ideas as objective truth or objective justification. In this respect, the relativist is like the philosophical sceptic. The sceptic does not hold merely that we know a lot less than we like to think, but that knowledge is impossible.

Many philosophers distinguish between scepticism and relativism on the grounds that, whereas the sceptic denies the possibility of knowledge, the relativist holds that knowledge, in its properly relativized form, is not just possible but actual. We might even think of relativism as a defence against scepticism. The thought would be that scepticism results from a hankering after an impossible form of objectivity. Give this up and the sceptic has nothing to say to us.

While this is a fair point, we must not press it too hard. The fact is that, pre-theoretically, we are inclined to accept the possibility of objective justification and objective truth. Of course, we may not take this attitude towards all areas of inquiry, in which case we commit ourselves to significant demarcational views, worthy of further exploration. But the relativist rejects the idea of objectivity completely. In this respect, relativism is close to scepticism.

We should not place too much faith in prepositions. The relativist allows me to speak of things being justified (or true) *for* me or for my culture. How does this differ from their being justified (or true) *according to me*, hence perhaps not really justified at all? Isn't the point of justification—for example, the collection and weighing of evidence—to help us *resolve* differences, not to enshrine them? Like idealism, relativism is a reaction to scepticism that is difficult to distinguish from scepticism itself.

Limits to relativism

So far I have treated relativism as a doctrine about justification or truth. But there is also conceptual relativism, the doctrine that different cultures 'carve up the world' in different ways: that is, think in terms of fundamentally different categories or 'conceptual schemes'. According to conceptual relativists, 'reality itself is relative to a scheme: what counts as real in one system may not in another'.[2] Since everyone must think in terms of some conceptual scheme—there is no neutral standpoint from which rival schemes can be compared—different schemes are 'incommensurable'.

Conceptual relativism is a dramatic doctrine, but as Davidson has famously (or notoriously) argued, not one that it is easy to make clear sense of. What do different conceptual schemes 'carve up' differently. The standard answer is 'the world' or 'experience'. But what is 'the world' here? If 'the world' is something beyond all description, an ineffable thing-in-itself, no real answer to the question has been given. On the other hand, if it is the familiar world of

trees and rocks, it is the world as we already know how to describe it; and while different people may have strikingly different views about how this familiar world operates, any sense of pervasive conceptual disparity has evaporated. Parallel remarks apply to 'experience'. Indeed, the notion of experience operative here is the one we already rejected: experience as a kind of 'non-propositional' knowledge. In other words, the Myth of the Given.

Different conceptual schemes are presumably embodied in different languages. Where languages express different concepts, there must be problems of translation. So perhaps the best way to make sense of conceptual relativism is in terms of the impossibility of translating one language into another. It is a commonplace that some languages have words with no precise equivalent in others: there is no English word for *Schadenfreude*. Perhaps, then, we want to say that English lacks the concept although, since even in English the meaning of *Schadenfreude* can be explained in a more roundabout way, perhaps we don't. Either way, such failures of easy translatability are far too localized to encourage talk of different conceptual schemes. For that, we need languages that are globally non-intertranslatable.

On the suggestion we are exploring, beings with a conceptual scheme different from our own would have to speak a language that we could not translate. So our question becomes: can we make sense of the idea of a language that is (in principle) untranslatable? It is not clear how. What would justify us thinking of an untranslatable 'language' as a language? What would distinguish the 'words' of such a 'language' from mere noise?

If, like Descartes, we held that thought is prior to language, we might see no problem here. But if we think of language as the vehicle of thought, and if we think of the meaning of linguistic utterances as immanent in patterns of use, we cannot be so cavalier.

First of all, if we think of meaning as immanent in use, we cannot attribute massive illogicality (or wildly different logicality) to other speakers.[3] This is not a matter of being indulgent: a large measure of consistency is necessary if we are to find patterns, or uses, at all. But where we cannot find patterns, we cannot find meanings; and where we cannot find meanings, we cannot find language. Thus, we can find language only where we can find a common sense of the logical.

The second point is that language needs to be tied to the world. As we saw in Chapter 15, the tie is effected by observation-sentences, keyed by training to circumstances that come and go. But this means that we can find meaning, hence, language, only where we can find reactions to worldly conditions that we ourselves can recognize. This means that finding language requires more than finding a common sense of the logical: it involves finding lots of common beliefs too. Davidson calls the need to find massive agreement, wherever we can find language at all, the 'Principle of Charity'. But as I have stressed, being 'charitable' is not a matter of being indulgent. The meaningful is the interpretable; and in the methodology of interpretation, charity is not an option. (These points recapitulate our earlier discussion of intelligibility constraints on demands for justification.)

Supposing that this argument is broadly correct, how far does it get us? Not very far. As Davidson himself emphasizes, his aim is not to eliminate disagreement but to delineate the conditions that make meaningful disagreement possible. This being so, it seems to me that the global conceptual relativism Davidson argues against is something of a red herring. The kind of variability in belief that draws people to relativism exists within the bounds of Davidsonian possibility. The animist thinks that trees are the homes of spirits; the scientifically minded person does not. In order to disagree about trees in this way, there is a lot the parties have to agree about. But their world-views remain deeply at odds, even though they agree at the level of 'That's a tree'. For all that the argument from charity shows, when it comes to disputes like that between an animist and a physicalist, relativism could be the right view.

We should also remind ourselves that not all disagreement is contradiction. Sometimes whole ways of talking go by the board. The profoundest form of atheism is not the one that involves strenuously denying the existence of God but the one that lets theistic ways of talking fall into desuetude. Or to take a case central to the concerns of this book, consider certain traditional distinctions: the analytic versus the synthetic, the *a priori* versus the *a posteriori*, and the necessary versus the contingent. For foundationalists, it is vitally important to understand how the distinctions line up and what beliefs fall into what categories. But from the standpoint of the fallibilist, contextualist, and inquiry-centred epistemology I have been defending, the usefulness of such ways of thinking is much less obvious. So when a thoroughgoing fallibilist expresses doubts about *a priori* knowledge or necessary truth, he is not claiming, in a straightforward spirit of contradiction, that everything is *a posteriori* or that everything is contingent. He is questioning the theoretical utility of these traditional classifications.[4]

Roots of relativism

I said that, for all that the argument from charity shows, relativism could be the right epistemological outlook. But how can a relativist claim to be right? It is often suggested that he cannot, since, like scepticism, relativism is self-defeating. When the relativist claims that nothing is 'absolutely' true (or justified), he doesn't intend his conclusion to be taken as merely true (or justified) for him. If he did, I could dismiss his conclusion by saying that relativism isn't true for me. If arguments for relativism are meant to make a general epistemological point, they must invoke the very notion of truth (or justification) they reject. However, like its counterpart in the case of scepticism, this argument suffers from being too purely dialectical. As in the case of scepticism, the

problem is not simply to silence an awkward opponent but reach a deeper understanding of how relativistic arguments work.

One of the founding distinctions of Western philosophy is that between what the Greeks called *physis*, nature, and *nomos*, custom or convention. While far from unproblematic, the observation that leads to this distinction is simple and (at first sight, anyway) compelling. Some of what we believe seems to be true in virtue of facts that hold independently of human wish and will. Stones fall to earth always and everywhere, whether we like it or not. Such facts belong to nature. But other beliefs, particularly those having to do with values or right conduct, seem to vary widely from place to place and time to time. They reflect the custom or conventions of particular groups of people. As human creations, customs and conventions do not belong to nature. Of course, if we never venture far from home, we may not realize this. That certain actions are wrong may seem as 'natural' as the fact that stones fall to earth. But the Greeks were traders and travellers. It was soon brought home to them that some things are accepted everywhere, others not.

There is no doubt that this perception tends to raise the spectre of cultural relativism: *autres pays, autres moeurs*. However, the mere fact of variation does not imply the impossibility of invidious comparison. The roots of relativism lie not in empirical data but in certain epistemological and metaphysical preconceptions. In the remainder of this chapter, I want to look at relativism primarily from an epistemological angle, leaving more metaphysical considerations for the next.

Epistemologically, the first and most important root of relativism is the idea that all justification takes place in a 'framework' of 'ultimate' commitments. What makes commitments 'ultimate'? The answer is that commitments that make justification possible must of necessity themselves be beyond rational assessment. In other words, what makes commitments 'ultimate' is the (alleged) impossibility of defending them in any non-circular way. But this makes it clear that relativism is not at bottom an empirical thesis at all. Relativism springs from Agrippan scepticism and the traditional conceptions of knowledge with which it is associated. In effect, the relativist accepts the foundationalist picture of the structure of knowledge while denying that there are any (or enough) foundational elements that are universally valid. The Agrippan argument thus functions in a double role: first to enforce the idea of ultimate commitments as necessary for knowledge and then to deprive those commitments of any objective significance. Relativism, we might say, is pluralistic foundationalism.

In connecting relativism with foundationalism, I do not mean to suggest that the coherence theory is wholly innocent. As we saw, the distinction between coherentism and foundationalism is far from clear-cut. The coherence theorist's radical holism pressures him to assign to his criteria of coherence a foundational role, an idea easily adapted to 'standpoint' epistemology. Instead of modelling itself on the atomistic foundationalism of traditional empiricism, standpoint epistemology can postulate systematic methodological orientations that colour all input and determine its relevance. In this way, variant 'total systems' become incommensurable.

I said that relativism, however much it has been associated with cultural anthropology, results more from epistemological preconceptions than from anthropological data. But I would go farther. I think that the 'data' easily become shaped by the preconceptions. Writers on exotic cultures can be tempted to present them as embodying hermetically sealed total views, able to accommodate any objections, with no loose ends and no invitations to criticism or further inquiry. Such bodies of belief may also be represented as shared by all members of the culture in question. The aliens are taken to accept the local tenets down to the last detail, as if critically or sceptically minded individuals are unknown in foreign parts. Views like this smack more of *a priori* epistemological ideals than of empirical findings.[5]

Contextualism, which is hostile to both foundationalist and coherentist pictures of justification, discourages the reification of all-encompassing 'frameworks' or 'standpoints'. For the contextualist, such ways of thinking are all too reminiscent of epistemological realism with its talk of 'our epistemic position'. Pluralizing such talk effects no improvement. No one occupies a single framework or standpoint: there are as many 'frameworks' or 'standpoints' as there are contexts of inquiry. None amounts to a hermetically sealed 'total view'; none rests on 'ultimate' commitments, beyond rational criticism. Furthermore, because contextualism has no use for the radically holistic idea of a total view, it is not under pressure to be dismissive of the idea of external constraint. Contextualism is thus an antidote to relativism.

Contextualism takes fallibilism seriously. Justification, in matters where serious issues of justification arise, is always provisional, never watertight. And it often involves less-than-algorithmic procedures, such as inference to the best explanation. This brings me to the second epistemological preconception underlying relativism: the lingering influence of the demonstrative conception of knowledge and the 'quest for certainty'. Despairing of universal, objective certainty, the relativist settles for personal or group certainty. Perhaps, in the end, the two roots of relativism are really one: the relativist, like the sceptic, is a disappointed foundationalist.

A thoroughgoing fallibilism, hence a contextualist epistemology, is once more the proper antidote. Apparently intractable disagreement no more implies dogmatic relativism than our vast contingent ignorance implies radical scepticism. Just as contingent ignorance invites further inquiry, intractable disputes invite a search for common ground. We might not be able to find it, just as we cannot always fill gaps in our knowledge. But then again we might. Or we might be led to some different views altogether. Whether we

succeed or fail depends on ingenuity and luck. We cannot predict the future of inquiry. Success is not guaranteed, but neither is failure.

Contextualism and relativism

Some philosophers will find implausible the claim that contextualism is the cure for relativism. They think that contextualism itself is a form of relativism.

It is not hard to see why. It is tempting to think of the presuppositions that are held in place by the direction of inquiry and other contextual factors as constituting a kind of 'framework' within which justificational questions can be raised and answered. I think that this way of talking is best avoided. It is almost bound to lead to contextualism's being confused with the relativist thesis that all justification is 'framework-relative', in a way that places frameworks themselves beyond justification or criticism. For example, Fogelin represents the contextualist as claiming:

(FR) S is justified in believing that p if p is justified within the framework in which S is operating.

However, as Fogelin points out, there are at least three ways of resisting the move from a claim to be justified 'within a framework' to a claim to be justified *simpliciter*:

(a) I may reject S's justificatory framework. (S may be using astrological tables.)
(b) I may accept S's justificatory framework, but think S has not used it correctly.
(c) I may grant that S has been epistemically responsible, but think his grounds have been defeated.

Fogelin rejects contextualism because he takes it to exclude critical reactions like these.[6] This is a mistake: a contextualist can and should accept the potential legitimacy of all three critical moves.

A contextualist view of knowledge and justification does not commit one to holding that a reference to context is part of the *content* of a knowledge-claim. We must recall the *sotto voce* proviso. A knowledge-claim commits one to holding that all significant potential defeaters—possibilities which, if realized, would make one's belief either false or inadequately grounded—have been eliminated: the contextual element comes in to fix what defeaters should be counted significant. But presuppositions as to what is significant are themselves open to criticism, which can be informationally or economically triggered.

More precisely, a contextualist will hold:

(C1) All justification takes place in a context of presuppositions (e.g. relating to which potential defeaters need to be excluded) and other circumstances which are not currently under scrutiny.
(C2) These presuppositions and circumstances can themselves be articulated and challenged, but only by a recontextualization of the original justificatory procedure, a recontextualization that will involve presuppositions of its own.
(C3) Recontextualization can go on indefinitely. But this is the open-endedness of inquiry, not a vicious regress of justification.

Bringing to light questionable background commitments is an important strategy for raising motivated challenges to accepted claims. This is a crucial element in contextualist epistemology, not a criticism of it.

I said that talk of 'frameworks' is harmless provided that we do not take it too seriously. But such talk is best avoided. Almost inevitably, it encourages us to think of contexts of justification as insulated from external criticism, a view that contextualism is simply not committed to. Intelligibility constraints guarantee the existence of a wide range of cross-contextual commitments and entitlements. Furthermore, observational evidence operates cross-contextually. Of course, such evidence is not mechanically determinative of what we ought to think, for it is always potentially subject to considerations of relevance and reliability. But it is always there and is not simply to be dismissed.

Another unfortunate feature of framework talk is that it invites us to think in terms of an overly sharp distinction between norms and facts. We may be tempted to conceive different 'ways of knowing' as constituted by divergent epistemic norms which, since they govern the conduct of inquiry, cannot themselves be the objects of critical examination. This temptation should be resisted. A proposition that serves as a 'methodological necessity' has a special normative status within a particular type of inquiry, in the sense that exempting it from doubt is a precondition of engaging in inquiry of that type. A methodological necessity thus owes its normative status to its functional role in a particular investigative practice. But the norm/fact distinction is here methodological rather than ontological. Viewed from another angle—that is, recontextualized—a methodological necessity can appear as a substantial empirical commitment, open to scrutiny and revision.

The new angle of scrutiny can be provided by theoretical advances outside the 'framework' in question. Fogelin's example of astrology is a good one. What killed astrology was not detailed empirical problems within the 'framework' of astrological inquiry (though there were always lots of failed predictions) but a new conception of the physical universe, a conception that made the astrologer's geocentric world of stellar and planetary influences incredible. The astrological 'framework' was never in principle insulated from such external undermining.

Framework talk encourages relativism by leading us to think of contexts of inquiry as more rigid than they are. Contextualists had best not go in for it.

Reason and tradition

It might be argued that my argument is question-begging in a different way. It is conducted entirely within a broadly critical-rationalist perspective on belief. However, rationalism itself is a specific cultural tradition: the 'Western' tradition. In fact, rationalism is just another faith: faith in reason.

Baldly stated, this claim misses something very significant, which is that a broadly rationalist outlook is reflexive, so that adopting it is not a matter of uncritical commitment, still less uncritical commitment to our favourite ideas about rationality. Once we embark on the path of critical inquiry, epistemological goals, norms, and procedures themselves become potential objects of critical scrutiny. We can ask what we are pursuing in our quest for knowledge (the analytic problem), whether the quest has its limits (the problem of demarcation), how to go about it (the problem of method), whether the quest is fundamentally flawed or self-defeating (the problem of scepticism), and even whether the whole thing is really worth the candle (the problem of value). Epistemology is a third-order tradition of inquiry created by the capacity of a broadly rationalist outlook to reflect critically on its own presuppositions. The fact that norms, including epistemic norms, are instituted by attitudes does not exempt them from criticism.

If 'traditionalism' entails blind adherence to inherited views, a fallibilist epistemology is obviously anti-traditionalist. But in another way, contextualism is anything but hostile to tradition. All inquiry takes place in a rich informational context. Such contexts are never the creation of a sole inquirer: they are the legacy of past co-operation. Tradition—the inheriting of results and methods—is the prerequisite of investigation, thus of self-correction.

This obvious feature of inquiry has been obscured by modern epistemology's radical individualism. This individualism is different from classical epistemic individualism, which is a reflection of the moral-practical significance classically accorded to knowledge. Modern epistemic individualism results from modern epistemology's subjectivism, which is the legacy of Cartesian scepticism. Taken at face value, Cartesian scepticism forces the individual inquirer to suspend belief in the objective world and everything in it. It forces the individual to rely entirely on his or her own resources. This result was welcomed by Descartes, who dreamed of personally reconstructing science on a secure foundation. But it is a wholly unrealistic hope. Because contextualism grows out of a theoretical diagnosis of scepticism, it allows us to take the social-historical character of knowledge seriously. This is one of contextualism's virtues.

We are now in a position to see why the contextualist will not be moved by

the claim that his broadly critical-rationalist outlook is an unjustified presupposition. The objection is self-defeating: anyone who can even raise it is already a fellow rationalist. This looks like the 'relativism is self-defeating' argument earlier criticized for being too dialectical. But it isn't. Or rather, this sort of objection takes on a different character when conjoined with an independently defended contextualist epistemology. For a contextualist, questions of justification only arise in contexts of motivated challenges. For views to which there are no serious alternatives, such questions do not arise. This is our situation with respect to a broadly rationalist outlook. Being who we are and knowing what we know, blind traditionalism is no longer a live option. Once released, the rationalist genie cannot be put back in the bottle. Our problem is to say what knowledge and rationality amount to in the situation we are actually in, not to defend them against imaginary alternatives.

Notes

1. For a sympathetic but critical account of social constructivism, see Hacking (1999).
2. Davidson (1984), 183.
3. Quine (1960), 58 f. criticizes the notion of a 'prelogical mentality'.
4. The defence of Rationalism in BonJour (1998) seems to me to be vitiated by a failure to consider this possibility.
5. Reasons for being suspicious of 'descriptive cultural relativism' can be found in Moody-Adams (1997), chs. 1–2.
6. Fogelin (1994), 95–8.