



13

Epistemic Circularity and Epistemic Incommensurability

Michael. P. Lynch

1. INTRODUCTION

If I were to believe you are trustworthy just on your say-so, my reasoning would be infected with what is called *epistemic circularity*. I would be supposing a source is trustworthy by relying on that very source. Generally speaking, we tend to think this is a very bad idea. It is why we don't bother asking politicians or salesmen whether they are honest.

A well-known line of reasoning stemming from the Ancients—what we might call the criterion argument—claims that epistemic circularity leads to radical scepticism. Baron Reed has recently put the point this way:

Let F_1, F_2, F_3 , etc., be a subject S 's cognitive faculties, of which S has a finite number. In order to know that F_1 is a reliable source of knowledge, S will have to use either F_1 , or another faculty. But if S uses F_1 his belief that F_1 is reliable will be epistemically circular. So, S must instead use (say) F_2 . But S should not use F_2 unless she knows that it is a reliable source of knowledge itself. In order to come to know this, S will have to use F_1, F_2 , or some other faculty. But S cannot use F_2 , on pain of epistemic circularity. And S cannot use F_1 , without first knowing that it is a reliable source of knowledge, which is still in question. So, S must use some other source-say, F_3 . But it should be clear that the same issues will arise with respect to F_3 , and that S will eventually run out of faculties to which she has not already appealed. Epistemic circularity is inescapable.

(Reed 2006: 186–7)

The invited conclusion: we don't know whether our faculties or sources of belief are reliable. And if we don't know that, then, according to the sceptic, those same sources cannot be sources of knowledge—that is, they cannot produce justified beliefs. Consequently, says the sceptic, I don't—and perhaps can't—know anything.



Epistemic Circularity and Epistemic Incommensurability 263

The criterion argument is disarming in its simplicity. But it is open to an equally simple response grounded in epistemic externalism.¹ The externalist rejects the key assumption of the argument—namely, that one must first know that a source of belief is reliable before it can be said to produce knowledge. From the externalist's viewpoint, all that matters is whether the source *is* reliable, not whether we know or even believe that it is. As long as the source is in fact reliable, it can produce knowledge, including knowledge about whether it, or some other source, is reliable. Epistemic circularity may be pervasive, says the externalist, but it doesn't prevent us from having knowledge or justified belief.

The externalist response to the problem of the criterion continues to be the subject of a wide-ranging controversy.² One thought is that far from providing an easy victory of scepticism, the externalists' response only reveals the inadequacy of their view, by demonstrating that knowledge is too easy to come by if externalism is correct.³ To some, the 'if my sources of belief are reliable, then I can know they are' response simply shows that the externalist is missing the point, or talking past the sceptic, or somehow not addressing the original issue.

Despite these doubts, I'm going to assume for present purposes that some version of the externalist response succeeds in answering the criterion argument. My reasoning, such as it is, is that it is difficult to see how that argument *could* be answered without relying on such a response. Nonetheless, I also think that there is something right about our misgivings about the externalist answer to the criterion argument. What is right about those misgivings is that the externalist response has no traction against a different problem, distinct from scepticism. Like the sceptical argument above, this other problem is rooted in part in the issue of epistemic circularity. But unlike its cousin, it is not a problem about whether we in fact have knowledge or are justified in our opinions. It is about rationally resolving explicit disagreement over the reliability of our most basic methods for forming beliefs. Although I will not try to show it here, this latter problem, what I will call the problem of epistemic incommensurability, is arguably the root worry behind the criterion argument, and, in my view, the reason why Chisholm was right to say that the latter was 'one of the most important and difficult of all problems of philosophy' (1982: 61).

The bulk of this chapter will be concerned with clarifying the problem and its nature. In the final section, I briefly sketch a solution.

¹ See Alston (1989), Sosa (1997a, 1997b). For a novel defence of this strategy that links externalism to Reidian common-sensism, see Bergmann (2004). Reed (2006) offers a criticism of Bergmann's position.

² See e.g. Fumerton (1995), Stroud (1994), and Pritchard (2005).

³ This is one of the lessons that can be drawn from S. Cohen's (2002) well-known discussion in 'Basic Knowledge and the Problem of Easy Knowledge'. Cohen himself is targeting the notion of basic knowledge in general.





2. EPISTEMIC INCOMMENSURABILITY

Cain and Abel, let's imagine, are having coffee and arguing about the age of the Earth. Abel asserts with great confidence that the earth is a mere 7,000 years old. Cain, amazed, points out that Abel's claim is not justified by the evidence of the fossil record, the best explanation of which is that the Earth is far older. 'Inference to the best explanation from the fossil and historical record can work sometimes' Abel concedes, 'but the *best* method for knowing about the distant past is to consult the Holy Book; it overrides any other competing evidence'. Cain scoffs and rejects the book as an unreliable source for knowing about the distant past; the *only* reliable method, he insists, is to employ a combination of abduction and induction from the fossil and historical record.

As this toy example indicates, disagreements usually start with a disagreement over the facts. But disagreement over facts often turn into disagreement over whose view of those facts is best supported or justified. And sometimes we move still further up what we might call the epistemic ladder: we begin to disagree over how we *ought* to support our views of the facts, about the sort of evidence that should be admitted, and whose methods more accurately track the truth. When we do that, we are engaged in a truly *epistemic* disagreement: a disagreement over epistemic principles.

By an epistemic principle, I mean a normative principle to the effect that some source or way of forming beliefs has some valuable epistemic status. The valuable epistemic status I'll be discussing here is reliability. So we might say that an epistemic principle is one that says that some doxastic practice or method is reliable. The reason reliability is valuable is that reliable doxastic methods are those that are likely to produce true beliefs. And having true beliefs, I'll assume for present purposes, is a good; it is what we might call a, or even the, epistemic goal.⁴

Doxastic methods are more or less basic. A less basic method would be a blood test, since the reliability of that test would presumably be justified by employing other methods. Plausible candidates for *basic* methods, on the other hand, include inferential methods like deduction or induction and non-inferential methods, like sense perception. It is hard to see how these sorts of methods could be justified solely by appeal to any other method. Consider just one case: induction. If Hume was right, there isn't much hope in justifying the reliability of inductive inference without relying at some point on induction to do so.

But I will not quibble here about examples. It seems likely that there are basic epistemic methods, and I will assume that there are. Let's say an epistemic principle is fundamental when it is about such a method and derivative when it is not.

⁴ I have defended the value of truth elsewhere, see Lynch (2004, 2009a, 2009b).





Epistemic Circularity and Epistemic Incommensurability 265

Belief-forming methods, whether basic or not, are more or less restricted in applicability. Deduction is clearly of wide scope—since deductive inference from true premises, if it is reliable, is reliable for forming beliefs of any sort. Not so with sense perception—sense perception is not a reliable means for forming beliefs about mathematics. It is, broadly speaking, applicable to ascertaining the contingent facts about the external world. Still more restricted are Cain and Able's principles: both assert the reliability of methods with regard to a single domain of facts—facts about the distant past.

Let us say that A disagrees with B over some epistemic principle just when A does not believe EP and B does. Thus one disagrees with someone over an EP in this sense when one either disbelieves the EP or withholds belief in it (e.g. when one doubts that it is true). A *overtly* disagrees with B over some EP just when A explicitly withholds assent from an EP B asserts. The paradigm case of overt epistemic disagreement is where assent is withheld because the relevant EP is simply denied. An overt epistemic disagreement is *mutual* just when both sides to the dispute deny an epistemic principle the other asserts.

Overt mutual epistemic disagreement can be deep or shallow. I'll say a disagreement is *deep* when it meets the following conditions:

1. *Commonality*: The parties to the disagreement share common epistemic goal(s).
2. *Competition*: If the parties affirm distinct principles with regard to a given domain, those principles (a) pronounce different methods to be the most reliable in a given domain; and (b) these methods are capable of producing incompatible beliefs about that domain.
3. *Non-arbitration*: There is no further epistemic principle, accepted by both parties, which would settle the disagreement.
4. *Mutual Circularity*: The epistemic principle(s) in question can be justified only by means of an epistemically circular argument.

By these criteria, most epistemic disagreements will *not* be deep. In most cases where my methods of belief formation are questioned, Non-arbitration, and hence Mutual Circularity fail—as might happen when one doctor questions another about the reliability of a given test, and the dispute is settled by appeal to other more basic methods. In other cases, subsequent investigation may reveal that the methods in question do not really compete—perhaps because one is more restricted in application than another. Finally, in some cases, the disputants are simply talking past one another—one is aiming at truth, for example, the other at practical success or some such.

Nonetheless, it seems pretty clear that deep epistemic disagreements can and do occur. Cain and Abel's dispute seems a likely example. The two principles involved, again, are:





C: Inference to the best explanation of the historical and fossil record is the most reliable method for knowing about the distant past.

A: Reading the Holy Book is the most reliable method for knowing about the distant past.

We can easily imagine that Cain and Abel share the same goal—to ascertain the truth about the distant past of the planet. Moreover, their principles compete. Both practices cannot be the best for determining what is true of the distant past. And assuming that neither thinks that time travel is an option, there appears to be no shared principle that will settle the matter either.

The case plausibly satisfies Mutual Circularity as well. In Cain's case, it is hard to see how he could avoid appealing to a track record argument to justify the reliability of his method. And it is equally hard to see how he could avoid assuming the reliability of that method in accepting the premises of such an argument. For there seems to be no other means (that Cain would accept) which can serve as a check on whether our abductions about the past are accurate. In Abel's case, we can imagine that the circularity might manifest itself in a very different way. He might, for example, give the following argument for his trust in the Holy Book's pronouncements about the past:

God wrote the Holy Book in the distant past.

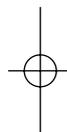
Everything God wrote in the Holy Book is true.

Therefore the Holy Book is the most reliable guide to the distant past.

Suppose that we ask Abel why he is committed to the first premise, and he responds that the book says that God wrote it. Since this is a belief about the past, it assumes the reliability of the book as a guide to the truth about the past. Abel's argument is not premise-circular: the conclusion is not literally one of the premises of the argument, but it is epistemically circular—it assumes the reliability of a method in an attempt to show that the method is reliable.

Some will wonder if there is an asymmetry to the cases Cain and Abel can make for their epistemic principles. One might think that Cain, but not Abel, will be able to independently confirm the reliability of his method. In fact, however, both Cain and Abel can offer independently confirming arguments. But any such argument will either be itself epistemically circular or push the disagreement back to the applicability of a different fundamental epistemic principle.

Take Cain first. First, recall that the issue in question is whether abduction from the fossil record is the best method for learning about the *distant past*. Unless time travel is an option, *direct* confirmation of the reliability of this method by observation is not an option. That is, we can't check to see whether the method has produced accurate results by direct observation. This fact is part of what will make any track-record argument for the reliability of the method epistemically circular. Moreover, even if Cain were to appeal to certain non-inferential methods (such as sense perception) to *indirectly* bolster the case for





Epistemic Circularity and Epistemic Incommensurability 267

the reliability of abduction from the fossil record, these methods can themselves be challenged. And as we've already noted, a defence of the reliability of sense perception, the most likely candidate in Cain's case, will most probably itself be subject to epistemic circularity.⁵ As Alston and many others have argued, it is hard to see how we could show that sense perception is reliable without relying on its reliability. And presumably it will not help Cain to protest that his favoured method is merely a local application of a more general reliable method: abduction. For in the imagined case, Abel does not deny the reliability of that method. What he denies is that abduction from the fossil record is the most reliable method for forming beliefs about the distant past.

In addition, note that Abel too might well try to give some independent argument, in addition to the one I've given above, in defence of his principle, one that does not rely on what the book says. He might appeal to the teaching of various prophets, or other sacred texts. Cain will question the reliability of these sources. So Abel might appeal to still more basic, non-inferential methods of belief formation, such as divine revelation or mystical perception of God's actions and will. But here again the question can (and presumably will) be raised about the reliability of such methods. As in Cain's case, it is difficult to see how Abel will be able to demonstrate their reliability without an epistemically circular argument, for it is difficult to see how one could defend one's claims to reliably speak for, or perceive, God without appealing either to the book, or to mystic perception again.

These points hark back to Reed's argument given at the outset. The general lesson is that where deep epistemic disagreement occurs—and again, it frequently does not—then the disagreement is ultimately disagreement over fundamental epistemic principles.

I emphasize that the Cain and Abel case is an example. You may not agree that it is a deep epistemic disagreement. If you don't, then choose an example you prefer. If you doubt there are any such examples, then reflect on the possibility of an overt debate over the reliability of induction, or sense perception, or any other basic doxastic method.

This is worth stressing. In order for an epistemic disagreement to qualify as deep, there need not be a clash of principles. It only needs to be the case that one side does not affirm a principle that the other side does affirm, and that Mutual Circularity, Competition, Non-arbitration, and Commonality are all satisfied. So overt scepticism about induction, for example, where the sceptic overtly

⁵ This illustrates a general lesson: where the methods in question each claim privileged access to a special domain of facts, appeals to other principles will be apt to lead to non-inferential principles that are themselves subject to epistemic circularity. Consider, to take another example, where Hylas and Philo disagree about the only reliable method for determining what is morally good. Hylas claims it is divine revelation; Philo rational intuition. It is hard to see how mutual circularity can be avoided in such a case, since both principles claim that a certain non-inferential method has privileged access to a restricted domain of facts. Any other method will be rejected as not being capable of accessing those facts.





doubts its reliability, qualifies as a deep epistemic disagreement. Commonality is easily satisfied. Competition is satisfied vacuously since the antecedent will be false. Non-arbitration will be satisfied—presumably because, if Hume is right, there is no more fundamental epistemic principle that can justify the reliability of induction. And Mutual Circularity is satisfied because the principle in play is defensible only by epistemically circular arguments. Consequently, overt scepticism about a fundamental epistemic principle like induction is a default case of deep epistemic disagreement. This is a helpful point to keep in mind if you are doubtful about the prevalence of mutual disagreements like that between Cain and Abel. But it is also helpful because it illustrates that overt sceptical challenges are merely a limit case of a broader phenomenon—epistemic disagreement.

Finally, even if deep epistemic disagreements never occur, it is clear that they could. And that is enough to raise the questions with which we will be concerned. The basic issue is simple enough. Where the debate is over basic methods for finding the truth, common ground is peculiarly elusive. It is not as if the parties to the dispute are sharing evidence and drawing distinct conclusions from it. Rather, they are disagreeing over what should be counted as evidence.⁶ As a result, there is a symmetry of argumentation in epistemic disagreements. This is because while we don't share the same evidence in cases of deep epistemic disagreement, our epistemic principles are themselves subject to evidence that has *the same structural character*. Neither side can fully justify their principles without Circularity; their principles are epistemically incommensurable.

3. THE THEORETICAL PROBLEM OF EPISTEMIC INCOMMENSURABILITY

For all I've said so far, one could be excused for thinking that epistemic incommensurability is anything but a regrettable fact about the human condition. After all, I've already conceded that epistemic circularity doesn't bar us from knowing which epistemic principles are true. Given externalism, I can give an argument for my epistemic principles. If so, then so long as induction is reliable, I can give an inductive argument for its reliability. And that means, that so long as my epistemic principles are true, I can know they are. So epistemic incommensurability

⁶ For this reason, deep epistemic disagreements are not properly classed as what some have come to call 'reasonable disagreements'. A reasonable disagreement in this sense is one where each side recognizes that they are drawing equally reasonable but incompatible conclusions from shared evidence. Deep epistemic disagreements don't fit this model. For I can't say of someone who does not believe my FEP that we 'draw distinct conclusions from shared evidence' for the simple reason that I won't believe that we share evidence if we don't share *what counts as evidence*. Think of a situation where an advocate of divine revelation tells me that *p*. I doubt that *p* is the case because I doubt that reliability of divine revelation. I may claim we have the same evidence, my opponent will not.



Epistemic Circularity and Epistemic Incommensurability 269

may be unfortunate, but if we've already agreed that scepticism doesn't follow, then what about this poses a distinctly *philosophical* problem?

In fact, it poses two problems. The first looks formidable but, is, I believe, actually not so. The second may seem less deep but is in fact more so. I'll take the easy one first. It is meta-epistemic. I'll lay it out in the simplest way I can and then discuss the premises.

1. Deep epistemic disagreements are rationally irresolvable (epistemic incommensurability).
2. The best explanation for why deep epistemic disagreements are rationally irresolvable is that there are no objectively true fundamental epistemic principles.
3. If there are no objectively true fundamental epistemic principles, there are no objectively true derivative epistemic principles.
4. All epistemic principles are either fundamental or derivative.
5. Therefore, there are (probably) no objectively true epistemic principles.

The argument parallels a familiar argument for moral anti-realism. That argument moves from the premise that deep moral disagreements are rationally irresolvable to the conclusion that there are no objectively true moral principles. Such arguments are widely used to undermine moral realism.⁷

As I see it, the first premise of the argument is reasonable, but only when qualified in a certain respect. But once qualified in that respect, the second premise is false.

Let's start with the first premise. What is meant by 'rationally irresolvable' and why believe that epistemic disagreements are so? One relevant sense of 'rational' here is presumably *epistemic rationality*. Epistemic rationality trades in epistemic reasons. An epistemic reason is a reason for thinking that some belief or principle is true. Let us say that a debate is *epistemically resolvable* just when it is possible that an epistemic reason could be *given*, as I'll call it, for resolving it. A debate is epistemically irresolvable when this is not the case.

What do I mean by 'giving a reason'? Suppose we are wondering which of two plans for building a bridge is the safest, and we consult an engineer to find this out. In so doing, we not only expect the engineer to be able to discriminate whether the one plan is safer than the other, but to be able to articulate, or otherwise make clear, the rational basis of her discrimination. In short, we expect her to *give us a reason* for favouring one plan over the other—to articulate the rational basis of her discrimination so that *we can recognize it as such*. For only if we can, from our evidential standpoint as it were, recognize her reason for favouring one plan *as a reason* can it serve to rationally underwrite our subsequent plan of action. Barring a recognized reason that one or the other plan is safest,

⁷ Note the 'probably' in the conclusion. The argument is deductive, but its conclusion inherits its content from the abduction in the second premise.



we will either not build the bridge at all, or we will judge them equally safe, and choose between the plans in some other way.

This suggests a general constraint: where *A* gives a reason (in the sense intended) of some type to *B* for some *p*, it must be possible for *B* to recognize, from his standpoint, that it is a reason. This constraint is consistent with some familiar distinctions. One might be justified in believing *p* without having a reason for the belief. The belief may simply be produced by a reliable method. Likewise one's belief might be justified by a reason without you recognizing that it is a reason. But that is not what is at issue in the case of the engineer, or in the case of an overt epistemic disagreement. In both cases, there is a standing demand to *justify* one's beliefs or commitments. And to actively justify (epistemically) a belief essentially requires *giving an epistemic reason* for the belief, and one gives a reason only if it can be recognized as such by oneself and others.⁸

So a debate is epistemically irresolvable just when no epistemic reason can be given for resolving it. Overt deep epistemic disagreements are marked by epistemic circularity. Together, these facts support the first premise, at least when it is understood in the epistemic way we've been developing. This is because epistemically circular arguments cannot be used to give a reason for why some epistemic principle is true.

This is not to deny that an epistemically circular argument might *be an epistemic reason* for believing a principle (the belief might be justified by such an argument). But it cannot be used to effectively give an epistemic reason for a belief in the face of a standing demand to actively justify the belief. For it will fail to be recognized as a reason by anyone who doesn't already share the principle. As the bridge-plan case illustrates, in a case where live demands for justification are on the table, we demand reasons we can recognize as such. Epistemically circular arguments are hopeless for doing this. If I am wondering whether I can trust you, the fact that you say I can doesn't give me a reason for why I should—even if, in fact, I can.

For this reason, epistemic externalism, even if it true, is of little use here. We may well know (via an epistemically circular argument perhaps) which basic methods are reliable. But that fact has absolutely no traction when one is trying to justify employment of a method in the face of disagreement. When posing a challenge, it doesn't help to be told that 'if you adopt the right method, you'll be able to know that you did'. That answer offers stone instead of bread.⁹

⁸ Compare Thomas Nagel: 'To reason is to think systematically in ways anyone looking over my shoulder ought to be able to recognize as correct' (1997: 5). Here Nagel oversteps: he suggests that reasoning itself (and therefore possessing a reason presumably) requires that the reason so possessed be recognizable by anyone as a reason. Not so: I can have a belief about my own mental states that may be a reason for some other belief that could not possibly be recognized by you as a reason for the latter belief. Nonetheless, the point holds where we are concerned with giving a reason, as opposed to being or having a reason.

⁹ It would be like telling someone faced with a moral dilemma, where e.g. the demands of utility clash with respect for persons, not to worry because, 'if you act on the correct moral principles, you'll do the right thing'. Thanks a lot.



Epistemic Circularity and Epistemic Incommensurability 271

Nor, I think, is it any more help to claim that one is *entitled*—in the sense intended by Crispin Wright (2004)—to one's fundamental epistemic principles. To be entitled in Wright's sense is, roughly, to have epistemic warrant for a proposition without having any epistemic reasons to believe it. One is entitled to a proposition in this sense when (a) that proposition is presupposed by one's cognitive projects; (b) one has no reason to think it is not true; and (c) any attempt to justify it would rest on propositions whose epistemic standing is no better off. Perhaps we are entitled in this sense to our fundamental principles. And perhaps, as Wright suggests, entitlement is epistemic and not pragmatic in character. The problem is that even if these points are true, citing that one is entitled to believe some proposition is not going to help rationally resolve a deep epistemic disagreement.

Suppose, for example, that Abel claims to Cain that he is entitled to his fundamental epistemic principles. This may well be so. We can imagine that treating the Holy Book as the most reliable guide to the past *is* presupposed by one of his basic cognitive projects—like understanding the will of God. And he sees no reason to think it is unreliable. Cain of course will hardly be moved, nor should he be. He doesn't accept Abel's cognitive project (he may even think it is incoherent). And from his standpoint he has offered a reason to believe that Abel's methods are unreliable, and therefore that his fundamental epistemic principles are false. But of course, whether this is so is the very issue they are disagreeing about. If their epistemic disagreement is truly deep, Cain and Abel aren't going to recognize each other's epistemic reasons in the domain of belief in question precisely because those reasons themselves are produced by methods that are the very methods in question. In short, 'I'm entitled to *p*' cannot be used to give a reason for *p* in a case of deep epistemic disagreement, *even if* the person who is entitled is warranted in holding *p*.

So understood, the argument's incommensurability premise seems reasonable. Deep epistemic disagreements are rationally irresolvable in the epistemic sense: the principles involved require defence by an epistemically circular argument, but epistemically circular arguments can't be used to give reasons. But of course, it also leaves room for the claim that epistemic disagreements might be rationally resolved in another sense of rationality. I'll return to just that suggestion shortly. For now, I'll bracket premise one and turn to the second premise.

4. A BETTER EXPLANATION

When qualified as I've done above, the first premise of the above argument is reasonable. Deep epistemic disagreements are rationally irresolvable in the epistemic sense. But once qualified in this way, the second premise appears to be false.





The second premise is a straightforward appeal to best explanation. In order for it to be plausible, the advocate of the theoretical argument has to rule out the other alternatives –or at least the most reasonable ones. This I don't think can be done. Indeed, I think there is a much simpler explanation available. But first let's briefly consider a few other alternatives.

A *sceptic* will claim that the best explanation for why deep epistemic disagreements are epistemically irresolvable is not that there are no objectively true epistemic principles, but that we can't know which epistemic principles are true. But if externalism about knowledge is correct, as I've been assuming, we *can* know which epistemic principles are true. So given externalism, scepticism is not the best explanation of the rational irresolvability of epistemic disagreements.

This consideration also underlines an important fact already noted but worth noting again: namely *that the argument is consistent with externalism being true*. For an externalist will claim that if our epistemic principles are true, then we can know they are true (by way of those same epistemic principles). But this conditional can be true even if the antecedent is false. So externalism is not ruled out by the argument.

Of course, if externalism is false, then it may well be that the best explanation for why deep epistemic disagreements are epistemically irresolvable is that we can't know which epistemic principles are true. But that is cold comfort surely.

Another possibility is that premise two is false because the epistemic irresolvability of deep epistemic disagreements is due to defective concepts. That is, we might think that our basic epistemic concepts are simply too vague or imprecise, or otherwise mangled, and that is why disagreements which essentially employ such concepts can not be epistemically resolved. This too would be cold comfort. For if our basic epistemic concepts are radically defective, it may be that many of our epistemic principles which employ such concepts will lack a determinate truth value. We end up reaching the same conclusion by other means.

And of course, there are the familiar possibilities of non-factualism and relativism about epistemic principles.¹⁰ But far from being rival explanations, non-factualism and relativism presuppose the general negative claim (that there are no objectively true epistemic principles) allegedly supported by the above argument. Consequently, it remains our first task to see whether there is a simpler explanation of the epistemic irresolvability of epistemic disagreements.

There is a simpler explanation. The simpler explanation for the epistemic irresolvability of deep overt epistemic disagreements is that it is a direct consequence of the nature of epistemic rationality itself together with the relevant facts about the kind of disagreement in question. More precisely: once one understands what it means to give an epistemic reason, the epistemic irresolvability of *deep* epistemic disagreements is just what one should expect. It is a consequence of the

¹⁰ For a non-factualist account, see H. Field (2001). Steven Hales (2006) discusses a relativist solution to a similar problem.



Epistemic Circularity and Epistemic Incommensurability 273

human cognitive condition; what Pritchard has memorably called our ‘epistemic angst’.¹¹

Where there is deep epistemic disagreement over some fundamental principle, the disagreement has hit bedrock, the spade has turned. It has hit bedrock precisely because the disagreement obeys Mutual Circularity, Non-arbitration, and the other constraints constitutive of deep epistemic disagreement. And when bedrock is reached in this way, the process of giving epistemic reasons—of engaging in the *activity of justifying in the epistemic sense*—loses its point. But the best explanation for this is not some deep metaphysical fact about the objectivity of our epistemic principle. The best explanation is that epistemically circular arguments for the reliability of some *method won’t be recognized as a reason to accept that method by those challenging its reliability in the first place*. In other words, the very nature of the debates themselves, together with facts about what is to give a reason, already explain the epistemic irresolvability of deep epistemic disagreements.

So it is plausible that epistemic disagreements are not epistemically resolvable. But that fact should not incline us to think that there are no objectively true epistemic principles. Nor does it entail that we can’t know which epistemic principles are true. It merely demonstrates that, like all things, the process of giving epistemic reasons has its limits, limits imposed by epistemic circularity and the nature of epistemic disagreement.

5. THE PRACTICAL PROBLEM OF EPISTEMIC INCOMMENSURABILITY

So deep epistemic disagreements are epistemically irresolvable; but that does not warrant an irrealist conclusion. It may seem tempting to some to leave it at that. But that would be a mistake, for to do so would be to avoid the obvious—and real—remaining problem. For we’ve just conceded that deep epistemic disagreements are not rationally resolvable in the epistemic sense. That leaves it open that they might be rationally resolvable in some other sense. But what other sense, and how, in that sense, are such disagreements to be rationally resolved?

This seems, to me anyway, a live and pressing question. It is quite different from the question of whether we can answer the Cartesian sceptic. There are very few, if any, radical sceptics. And there are unlikely to be any. But disagreements over basic epistemic methods are, in my view, quite possible. Moreover, the issue at hand is not whether one side or other of an epistemic disagreement can *persuade* the other that their principles are the true ones. Persuasion is a

¹¹ See Pritchard (2005: 245–9).



psychological phenomenon, and can be accomplished in a variety of ways. (A big club is a good method; so is an expensive advertising campaign.) What we presumably want here is not a psychological but normative explanation. We want to make sense of what ought to be done in the face of deep epistemic disagreement.

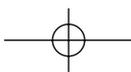
The fact (if it is a fact) that deep epistemic disagreements can't be epistemically resolved means we have to look elsewhere than epistemic reason. An obvious alternative is practical reason. Indeed, one might think that this would have been a reasonable starting place anyway since there is a clear sense in which the problem at issue is practical. Considerations given above, after all, already suggest that the underlying issue isn't a matter of what we know or don't know, but of what we should or shouldn't do. It is not even best conceived as a problem about which principles we 'ought' to believe (because one might say we ought to believe what we know to be true). After all, the root issue at the heart of an epistemic disagreement—that which makes the dispute an 'epistemic' one—is the question of which methods we ought to *employ*. What we want is a reason for employing one method over another. That's a practical matter. No wonder epistemic disagreements seem irresolvable when conceived of as epistemic matters. Epistemic reasons are simply not the right tool for solving them. It is not a matter of trying to justify our belief in our epistemic principles, it is a matter of trying to justify our actions—our employment of a method.

So I suggest that we reframe the problem of epistemic incommensurability as a practical problem. In saying this, I don't mean that it is a moral problem. Not all questions about what to do are moral questions. The question we are faced with in the face of an overt deep epistemic disagreement concerns our ability to justify our doxastic methods. Since these methods are aimed at producing (true) beliefs, the issue at hand is hardly divorced from epistemic matters. It is practical, but not morally practical. Call it a matter of *epistemic practicality*.

6. EPISTEMIC PRACTICALITY

The main goal of this essay has been to get clear on the problem caused by the phenomenon of epistemic incommensurability. That problem, I've argued, is an epistemically practical problem. It concerns how we can justify our employment of our basic epistemic methods in the face of overt disagreement. I've argued that this problem is distinct from the issue of whether we can know that our methods are reliable, and moreover, distinct from the question of whether it is true that they are. Moreover, it is not a problem that will be resolved by appeal to epistemic reasons. It will be solved, if at all, by appeal to practical reasons.

I'll conclude with a suggestion for how to arrive at such reasons. I say the suggestion is about 'how to arrive at' such reasons because one cannot foresee in advance what specific practical reasons might be relevant to cite in favour of



Epistemic Circularity and Epistemic Incommensurability 275

a method for any disagreement over that method. What is needed is a general framework—a way of identifying the types of practical reasons relevant for settling on one method or another.

In setting up the framework, we must recall the general constraints about rational resolution of disagreements uncovered above. Rational resolutions require the giving of reasons. And reasons are given only when they are recognizable as reasons. Together, these constraints suggest that no practical reason for resolving an epistemic disagreement will be successful if it doesn't employ a type of reason—or if, you like, a value—shared by both sides of any such debate. Briefly put, we need common ground.

Of course it is an open question whether such ground can be reached. But rational self-interest provides an obvious initial suggestion. How might we employ such a consideration to argue for one method over another? One possibility is to play what I'll call the *epistemic method game*.¹²

Let us say that an epistemic principle is privileged when it is worthy of teaching in the schools, used in evaluating research, and seen as trumping other, possibly conflicting methods. In the epistemic method game, players are charged with cooperatively coming up with reasons for privileging some epistemic methods for forming beliefs over others in some world w —a world distinct from their own. Candidate methods include methods like deduction, induction, sense perception, reading palms, and consulting sacred texts. Players operate, in good Rawlsian style, under the following rules or constraints. First, they cannot, in their deliberations, presuppose that any method is more reliable for producing true beliefs than others in w . That is, they must operate under the assumption that consulting the sacred texts may well be as reliable a method of belief formation for the inhabitants of w as sense perception. In effect, their deliberations must take place under the assumption that scepticism is true—all methods for forming beliefs are as reliable as others because none are reliable. Second, just as they cannot, in their deliberations, assume that one method of belief formation is more reliable than any other, so they cannot assume that one metaphysical picture of the world is any more accurate than others. By a 'metaphysical picture' I mean—admittedly somewhat loosely—a view about the ultimate structure and nature of reality. Thus, players cannot, in their deliberations assume that, for example, naturalism is true, nor that, for example, Christian theism is true. Third, players know that they will each eventually inhabit w themselves. But finally, fourth, they don't know all the methods they will—because of upbringing, education, religion, and so forth—wish to employ themselves in w .

Were we to play the method game, it would seem in our self-interest to favour privileging those methods that, to the greatest degree possible, were *repeatable, adaptable, public, and widespread*. Repeatable methods are those that in like cases

¹² What I am calling the epistemic method game was first suggested to me by Eberhard Herrmann; I am not sure he would approve of its use here, however.





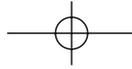
produce like results. It would be in our interest to favour repeatable methods because such methods could be used over and over again by people with different social standings. Adaptable methods are those that can be employed on distinct kinds of problems and which produce results given a variety of kinds of inputs. It would be in our interest to favour such methods because we don't know what sort of problems we'll face in w . Public methods are those whose effectiveness could, in principle, be judged publicly—that is, it is not the case that only one person is its sole judge of effectiveness. It would be in our self-interest to favour public methods because we don't know if we'll be lucky enough to be that one person in w . And finally, widespread methods are those that many people can in fact employ. It seems rational that we would privilege methods with these features simply because by doing so, we would maximize each of our chances to both use and assess the use of the privileged methods. In this sense, such methods could be called democratic.

But aren't players of the epistemic method game using various methods (such as a priori and causal reasoning) to arrive at their conclusion about w ? Obviously they are, and they may be using other, more idiosyncratic methods as well. The rules do not in any case prohibit this. The situation described is not one where the participants are without epistemic methods, nor is it one where the players must share epistemic methods. It is one where they are asked to decide—using whatever methods they have available, and acting under the relevant constraints—which methods should be politically *privileged in w*. And the methods that are so privileged are those that will form the content of their subsequent epistemic standards and principles.

One might object that Abel was asking Cain for some reason to believe that certain standards of reason are true. Our suggested form of argument gives something else. It gives practical, self-interested, reasons for adopting some standards over others. Why isn't this stone instead of bread?

Two answers suggest themselves. The first is that stone soup is sometimes the best item on the menu. Given the dialectical situation forced upon us by the nature of deep epistemic disagreements, giving practical reasons is the best we can do. Moreover, practical reasons—and this has been my point—are still reasons. They are better than big sticks.

There is something to be said for the first answer. But there is a better answer available. Namely, not only is giving practical reasons the best we can do, *it is what we should do*. For if Cain and Abel are to engage in real discussion, they should respect each other as fellow judgers—as being capable of arriving at the truth. And where we respect each other as fellow judgers, we *should*, where possible, defend our claims with reasons that we each recognize as reasons. And that means, if we are to satisfy the demand to respect each other as fellow judgers, we *should* appeal to the common currency of our self-interest to justify why we should privilege certain methods. For only by doing so can we give reasons that can be recognized as reasons. Put most provocatively: if we wish to respect each



Epistemic Circularity and Epistemic Incommensurability 277

other as fellow judges of the truth in cases of deep epistemic disagreement, we should not appeal to reasons for judging something to be true.

I have not tried to argue *which* specific methods would emerge from the epistemic original position, although I very much doubt that ‘reading the sacred text’ would be one. My point is that *even in the face of disagreement*, we can give reasons for our methods. The reasons are practical, not theoretical, but they are reasons all the same. Moreover, we should try to give such reasons. If we are to treat each other as autonomous judges worthy of equal respect, we must engage in the process of giving and asking for reasons, even when the question at hand concerns the reliability of our most basic methods for reaching the truth.¹³

¹³ This chapter has benefited from discussions at the University of St Andrews, the University of Edinburgh, the Bled Epistemology Conference, and the University of Connecticut. My thanks to Don Baxter, Tom Bontly, Matthew Chrisman, Patrick Greenough, Terry Horgan, Scott Lehman, Alan Millar, Baron Reed, and especially, Duncan Pritchard for additional comments.

