



OXFORD JOURNALS
OXFORD UNIVERSITY PRESS

Is Epistemology Dead?

Author(s): David Papineau

Source: *Proceedings of the Aristotelian Society*, 1981 - 1982, New Series, Vol. 82 (1981 - 1982), pp. 129-142

Published by: Oxford University Press on behalf of The Aristotelian Society

Stable URL: <https://www.jstor.org/stable/4544984>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <https://about.jstor.org/terms>



The Aristotelian Society and Oxford University Press are collaborating with JSTOR to digitize, preserve and extend access to *Proceedings of the Aristotelian Society*

JSTOR

X*—IS EPISTEMOLOGY DEAD?

by David Papineau

1. Introduction 'Nothing is more easily or manifestly perceptible to me than my own mind.' If you believe that then the problem is to justify our knowledge of other things. The various epistemological strategies are familiar. Other things come in as the causes of sense experience. 'Other things' are in fact mental. We have to recognise other things if we are to be aware of ourselves as having sense experience in the first place.

But suppose we lose faith in the idea of a self-intimating mind. What then becomes of the epistemological enterprise? Is there still a problem of accounting for our knowledge of the 'external' world, if there is no incorrigible knowledge of the mental world to contrast it with?

It might seem that the whole business of 'accounting for' knowledge, of showing how our beliefs about the external world can be justified, must presuppose that there are certain 'privileged representations', certain beliefs not themselves in need of justification. If our aim is the philosophical validation of doubtful beliefs, then surely we need some indubitable beliefs to serve as premises in the validations, to provide a foundation for the superstructure we shall then go on to justify.

But is epistemology essentially committed to the metaphor of foundation and superstructure? It has seemed to a number of modern writers, most notably Quine, that many of the functions of traditional epistemology can be served if we 'naturalize' the enterprise, and conduct it from within accepted scientific theory, rather than as a propaedeutic to science. The aim then is not to stand outside science and show how we can construct it starting with sense experience alone, but rather to take existing science and its view of the world as given, and use that to show how we manage, in our interaction with the world, to acquire reasonable beliefs.

A natural objection to this 'naturalized epistemology' is that

*Meeting of the Aristotelian Society held at 5/7 Tavistock Place, London WC1, on Monday 10 May 1981, at 6.00 p.m.

it puts the cart before the horse. On any account, surely we accord authority to the findings of science insofar as they have been corroborated by observational evidence. If now we are worried about the authority of such evidence, is it not straightforwardly circular to turn to accepted scientific theory for further support?

However this objection to naturalised epistemology is less than conclusive. It is not obvious that the circle involved is at all vicious. Why should we not use existing perceptual beliefs to adjudicate scientific theory, and then subsequently use that scientific theory (in particular our scientific understanding of the relevant sense organs) in turn to check the reasonableness of those perceptual beliefs? By this kind of iterative process we could hope to arrive at a state where our general scientific theories and our particular perceptual beliefs were in optimal accord. In Neurath's celebrated metaphor, we could rebuild our boat while afloat, without any part of the construction serving the role of unquestionable starting point. This process is not of course going to give us certainty, for it does not aspire to logical derivation from incorrigible foundations—but it does at least seem to share the traditional epistemological concern to evaluate our beliefs, to consider whether the judgements our perceptual and cognitive dispositions have led us to do indeed reflect the world as it is.

There is however a more sophisticated objection to naturalized epistemology, which queries not so much its circularity but whether it has the basic ability to play this traditional evaluative role. The difficulty is that while scientific theory (physiology, optics) might well *explain* the occurrence of certain mental *events*, it is not at all clear what this has to do with the *justification* of *judgements* about the world. In the traditional Cartesian context, with those magical 'ideas' which were at the same time both *things*, which could interact causally, and *awareness*, for which questions of representational correctness arose, it was understandable, if unfortunate, that explanation and justification got run together. But once we give up on the 'given', there is no escaping the conceptual gap between explaining a mental event and justifying a judgement. To do the former is to place the mental event in the space of facts, to see it as a thing (physical, physiological, or functional) which inter-

acts causally with other such. To do the latter is to enter the space of reasons, to evaluate an item with content for truth or falsity, for rationality or lack of support. And there is no obvious way of getting from the former space to the latter: to mention just one difficulty, what in the explanatory story about the physiology of perception is supposed to correspond to the evidence which in the justificatory tale might or might not lend support to our perceptual beliefs? To pick on the stimuli reaching our sensory peripheries (rather than, say, the light emitted by the objects themselves, or the neurones firing in our visual cortex) is surely quite arbitrary, explicable only as some half-remembered residue from our lost Cartesian past.

Considerations such as these have suggested to some (even more modern) writers, most notably Richard Rorty, that once we reject the given we should therewith abandon altogether the search for ultimate philosophical justifications for knowledge. Of course there are our actual intellectual procedures, the rules whereby we in practice adjudicate the acceptability of our knowledge claims. And of course we can chart the structure of these procedures, and use them to pass judgement on specific beliefs. But this scarcely gives us a philosophical justification for our beliefs. For the intellectual procedures themselves have simply been *described*, as in an anthropological report, and this in itself does nothing to show that they will enable us to get reality right. On the other hand, as soon as we set off in the most plausible direction for such a justification, and start theorising naturally about the relation between mind and the rest of reality, then we seem to step out of the world of reasons and into the world of facts, and all questions of justification seem quite to disappear.

In this paper I want to show that this pessimism about the possibility of epistemology is unjustified. I shall not of course be resuscitating Cartesian incorrigibility. My main concern is only to find a place for naturalized epistemology, to show that our theories about ourselves as perceiving beings can indeed play an evaluative role in our intellectual practices. This will occupy the next four sections. Having then defused the immediate argument for killing off epistemology, I shall briefly consider the prospects for what then survives.

2. *The Malleability of Observations* I want to press the anti-epistemological position at what I take to be its weakest point—namely, its account of observational judgements.

What account can the opponents of epistemology give of such judgements? They will of course recognise that there are particular statements about local matters of fact that we accept straight off, and not as the result of inferences from other statements. But they take the authority of such observation reports to derive, not from infallible access, but simply from the fact that it is part of our social practice, part of our language game, that if a normal person in normal circumstances reports the presence of, say, a tree, indicating, if appropriate, that their warrant for this claim is that they saw it to be so, then any requests for further justification are quite out of order.¹ On this conception it is not because observation reports have some special kind of authoritative provenance that we accept them unquestionably, but vice versa—we accord them authority because it is part of our practice not to question them. (Not that such reports are incorrigible: if it turns out that the circumstances, or the observer, were abnormal (the lighting was bad, he or she was drugged) then of course their acceptance can be revised. But still, if we have no reason for any such defeating suppositions, then what the observer says goes.)

The trouble with this story is that the range of matters on which an observer's word goes is something about which we change our minds over time. Our observational practices have a history. True, the earliest stages of this history can only be conjectured. We do not know when men first learnt to distrust immediate judgements about the shape of objects half in and half out of water, or about the pitch of the sounds emitted by moving sources. But if these episodes are lost, there are plenty of similar cases within recorded history. For as well as illusions recognised by our prehistoric ancestors, there are also illusions that we have come to recognise as a result of relatively recent theoretical advances. Thus with judgements about the sun's motion, or the simultaneity of distant events, or the colour of fast receding stars, . . .

What I am claiming here should not be particularly contentious: simply that on occasions it happens that existing procedures generating the acceptance of observation reports under-

go revision. (In some such cases it may be that we come to accept different judgements observationally in the relevant circumstances—‘the stick is straight’, ‘the sun is coming into view’. Alternatively we might decide that the matter in question cannot be decided directly and immediately on the basis of observation, but only by means of some more complicated process of inference.)

But while my point should be uncontentious, it does raise a difficulty for the opponents of epistemology. What are they to say about such changes in our observational practices? According to the story they have told so far, it seems that all such changes will be on an evaluative par. Each set of observational rules will define a different language game, each will set up its own standard of rationality, but there will be no possibility of any questions about the relative merits of different games.

But this is not a comfortable position to end up in. When observational procedures are altered it is generally because of a conflict between some new understanding of how the world works and certain hitherto acceptable observations. Observations are ‘theory-dependent’, at least to the extent that it is perfectly possible for a conflict between a theory and an observation to be won by the theory. But if theory sometimes wins, what is to stop it always winning? What is to stop scientists always sticking to whatever theories they favour, continually dismissing awkward observations as illusory and unreliable?

Of course something does stop scientists doing this. Even if it is sometimes acceptable for observations to succumb to theory, it is clear that they are not always allowed to do so. But if the story told by the opponents of epistemology were the right one, if all observational procedures were really on an evaluative par, then there would be nothing to be said against the theoretical dogmatist’s wanton dismissal of unwanted observations. Something needs to be added to our account of observation, if we are to avoid granting a licence to extreme theoretical relativism.

(Some readers might have been tempted to resist the argument of this section by denying its starting point, the changeability of observational procedures. They might feel that the examples offered are superficial, that what they illustrate are not changes in our most immediate observational judgements, but simply in the further inferences we draw therefrom. Perhaps

there is something in this line of objection. But it is scarcely one that can happily be made by my anti-epistemological opponent. For the only plausible candidates for unchanging observational procedures in the examples mentioned are introspective processes giving rise to reports of subjective psychological states ('It seemed bent to me'). Even if we allow that the idea of such processes makes sense, what is then supposed to be so special about them? If other observational procedures can in principle change, then why not these? And if we need an account of what governs such changes (or lack thereof) then why not in this case too? The only possible answer seems to be that there is something especially authoritative about the relationship between introspective awareness and its objects not present in other cases of observation. Which of course is just what the opponents of epistemology start off by denying.)

3. *Constraints on Rejecting Observations* What then does constrain the changing of observational procedures? The remarks of the last section showed that conflict with a favoured theory is not by itself enough for the rejection of an observational report. Not even the most enthusiastic proponents of theoretical immunity to observational falsifications, such as Paul Feyerabend, suggest that counter-exemplary observations should simply be *ignored*. Rather the idea is that the defender of the theory should attempt to *reinterpret* the awkward observations. He needs to explain, if things are indeed not as those reports specify, exactly what is going on, what the real significance of those reports is. Only if he succeeds in finding such a reinterpretation is he entitled to hang on to his theory. Otherwise he should accept the observations and reject the theory. (Feyerabend's excesses are not so much a matter of recommending that awkward observations be ignored as of urging that the search for reinterpretations never be abandoned.)

Up to a point it is useful here to compare our observing selves with scientific instruments. With a thermometer, say, we have a reading (the height of the mercury) which, together with a general understanding of how thermometers of that kind work, enables us to infer a conclusion (about the temperature). If (say because it conflicts with a favoured scientific theory behind some prediction) we are inclined to dismiss the conclusion about

the temperature indicated by the thermometer, then we have some work to do. For (supposing that we don't go so far as to deny the height of the mercury) this conclusion is forced on us by our interpretational theory of the workings of thermometers. So to reject it we have to reject that interpretational theory. Now we could simply plead agnosticism here, simply say that we did not after all know how thermometers like this work. But this would save the favoured predicting theory merely at the cost of dilapidating another bit of our overall view of the world—our interpretational theory of thermometers. In Imre Lakatos' phrase, it would be a theoretically degenerate step. Thus there is an obligation, if we are indeed to save the predicting theory, to come up with some alternative account of thermometers and therewith a reinterpretation of the significance of the height of the mercury.

Perhaps then the same story goes for human observers. Our acceptance of observational judgements depends on our general understanding of how we work in producing reports to that effect, in a way that requires us to revise that understanding if we are to reject those observational judgements. And this then gives us the constraint on the wanton dismissal of awkward observations: if the revision of our interpretational theory of our observational workings is not to be degenerate, it will need to go beyond simple rejection, to replacement by an alternative theory which reinterprets the old reports.

If this is right then we seem to have found a place for naturalized epistemology after all. Our acceptance of observational reports seems to depend on our theories (optical, physiological, computational, etc.—henceforth I shall call these all together our 'perceptual' theories) of how those reports get produced. And this dependence certainly seems to be a normative one—people with different such theories will be led to adopt different observational practices.

4. Naturalized Epistemology as Coherence But are things that simple? Anybody with initial doubts about the possibility of naturalized epistemology is unlikely to have been much persuaded by the brief remarks of the last section.

To start with, there is clearly something suspicious about the analogy between scientific instruments and human observers.

Consider the thermometer story again. What it actually comes to is this: starting with an observation of the height of the mercury on some occasion, together with a general interpretational theory relating heights of mercury to temperatures, we infer a conclusion about the temperature. Constraints on theoretical revisions then imply a constraint on rejecting the conclusion about the temperature: we cannot drop that conclusion without dropping the interpretational theory, and we cannot reasonably drop that without an alternative. But note that this story only works because we started with, and took as quite unproblematic, the initial observation about the height of the mercury.

Similarly with humans. If the idea is that a conclusion such as, say, *the stick is bent* was forced on our ancestors by some such naïve ‘theory’ as *when someone says that something is bent, then it is*,² and that they were justified in rejecting this conclusion only because they could come up with an alternative to the naïve theory, then again we are simply taking for granted the ability to come up with the initial observational judgement *X said that the stick was bent*.

But once we spell out the examples in this way, it becomes unclear whether the arguments of the last section cast any light on the acceptance of observations after all. The judgements whose acceptability is at issue turn out not to be observations accepted straight off, but rather conclusions inferred from observations. The role of our perceptual theories is not to validate observations but to serve as premises in inferences. And the only observations in sight are taken quite for granted, without any attempt to evaluate their acceptability.

Those impressed by the gap between the space of facts and the space of reasons are unlikely to be surprised by this. While it is clear enough how perceptual theories can serve as premises in inferences, there was nothing in the last section, they will say, to explain how they could perform the far more tricky task of altering our direct dispositions to respond observationally to the world.

However the opponents of epistemology would be ill-advised to make too much of the irrelevance of perceptual theories to our observational procedures at this point. For without anything to explain what does govern changes in such procedures, they

would be in danger once more of having nothing to say against the extreme theoretical relativist.

Still, what they might well argue here is that the principle governing the rejection of observations is simply that of maximal coherence. The suggestion would be this. At any time we are faced with a given corpus of accepted observational statements. The generalisations we accept express the patterns we have found amongst such observations. Sometimes new observations will conflict with existing generalisations. And on some such occasions it may well be that overall simplicity is best achieved by rejecting the observation. But, further, it needs to be noted that such conflicts cannot be represented as a simple fight between counter-exemplary observation and predictive theory. For characteristically the counter-example will also be derivable, as in our examples, from a report of an utterance (or a report of an experience) and some interpretational perceptual theory. So saving the predictive theory by rejecting the counter-example will at the same time require rejecting the interpretational theory—hence the obligation, imposed by the original aim of overall coherence, to take this path only when we can find some alternative to that interpretational theory.

On this account our perceptual theories do not somehow stand behind our observational judgements supporting them. Their role is not to tie down the individual points at which our theoretical net is moored, to make sure those observational fastenings are independently firm. Rather they are simply themselves parts of that net.

And so on this account there is no problem of explaining how we get from the facts given by perceptual theories to new reasons. For no such step is taken. There is no serious sense in which our intellectual practice depends upon our perceptual theories. Whatever perceptual theories we have, our basic procedure remains the same—aim for maximal coherence. (The only reason the perceptual theories get into the story at all is that they are part of the network we are trying to make coherent.)

5. Naturalized Epistemology as Practical Reasoning I think that the coherence account of our practice in revising observations given in the last section is demonstrably wrong. In order to show this I shall first present an alternative account.

On my view, the role our perceptual theories play in directing our observational practice does not hinge on their licensing inferences from premises *about* observations to observable conclusions about the world. We do not appeal to such theories in *particular* cases to get some conclusion from some already given observational data. The importance of such theories is rather that they enable us to 'stand back' from our practices and consider the way that we in *general* arrive at a certain kind of observational judgement, and so to decide whether our practice in this respect is indeed reliable. Thus we might use our knowledge of refraction and related matters in deciding whether or not we are any good in general at judging shapes of sticks in water. And clearly we could do this without bringing in any view we might have on the shapes of particular sticks on particular occasions.

Now on this account there is an obvious answer to the question of how we get from the space of facts to the space of reasons. For the upshot of our excursion into naturalized epistemology will not be that we accept some further conclusion as to the facts, but rather that we *do* something, namely change our observational practice. There is no question about what is supposed to play the role of evidence in our perceptual story, for the parts of that story are not supposed to mirror the premises and conclusion of an argument in the first place. Instead the whole story is supposed to move us to an action. The point of our perceptual theories is not that they give rise to interesting theoretical inferences, but that they allow important *practical* inferences. (Not that changes in our observational practice can be achieved at will: since we cannot decide what to believe at will such changes will require periods of self-correction and retraining.)

Now contrast this account with the coherence analysis of the last section. There is an obvious enough difference in intention: while on the coherence account the aim is to make all our observations tally with each other, on the practical account the idea is to make each observation tally with the world.

Moreover, this difference in intention implies differences in execution. For one thing, there seems to be nothing in the coherence account that requires any retraining in our observational dispositions. Since the coherence account only starts

biting given a set of observations to work with, and says nothing about where they should come from, it seems to allow that we should keep on getting observations from procedures which on the practical account would long ago have been discredited as unreliable. (However there is perhaps room here for an advocate of the coherence account to argue that even if such bad observations kept on coming in, they would also keep on getting thrown out—which could itself provide a motive for retraining.)

Perhaps a more conclusive argument against the coherence account is that it does not explain why, in the kind of example I have been considering, coherence should not be saved by rejecting the initial premise about an observation (*X said that the stick was bent*) along with the observation itself (*the stick is bent*) in cases where we cannot find any alternative to the interpreting theory (*when people say something is bent, it is*). If coherence were all we were after, this would seem a perfectly satisfactory way of getting it. But of course this would be a quite inappropriate response in the situation, say, where our ancestors' predictive understanding of the behaviour of rigid bodies started indicating that things which seemed bent were in fact straight. Simply to deny that anybody even took such sticks to be bent would simply be a dishonest evasion of the issue.

On the practical account things come out differently. For there we need to consider utterances not because descriptions of them are premises in arguments for conclusions which threaten incoherence, but simply because they are instances of the practice we are deciding whether to discontinue or not. Correspondingly, if we quit the practice we will stop describing *sticks*, not describing utterances about sticks. (We could, it is true, raise the question of whether to stop this latter practice if we really wanted to. But that would be a quite different matter, and one which would involve a quite different bit of perceptual theory.)

6. Final Remarks So far I have shown that it is necessary (and, moreover, that it is possible) to tailor our observational practice to our perceptual theories. This then removes the first section's objection to the philosophical evaluation of knowledge claims: namely that a naturalized epistemology, the only plausible possibility once the given has gone, is of necessity normatively impotent.

But there remains the possibility of a rather deeper objection to normative epistemology. This too stems from the demise of Cartesian incorrigibility, and in particular from the thought that once we give up on certainty, then we ought to give up also on the idea of an objective reality existing quite independently of our theoretical projections. (What part can such a reality play in our thinking, if we have given up any hope of access to it?)

If we accept this line of argument, then it seems that epistemology cannot help collapsing into mere anthropology after all. We can describe the procedures which guide the construction of our theories, but without the concept of an independent reality there seems no way of raising the further question of whether they are the right ones to have or not. If all we can seriously mean by 'reality' is the pictures our procedures come up with (or will eventually come up with), then we are left without any further court of appeal to assess those procedures themselves.

What then about our naturalized epistemology? This certainly looked like a way of evaluating beliefs for a match to reality. However the kind of anti-realist I have just introduced might well refuse to take it at face value. He could remind us that a naturalized epistemology does not yield certainty—the theories it invokes are after all just fallible theories, and the reality we make our observational judgements reflect is only the reality they portray. And he could argue that therefore our naturalized procedure for revising out observational procedures is itself just more procedure, not a way of ensuring that our beliefs match some abstract reality.

But is this conclusion at all plausible? It seems to make it quite arbitrary that we adjust our observational practice as we do. But surely it is not. Surely the reason we adjust our observational procedures as we do is that we wish our observational judgements to get reality right, and while our current perceptual theories might not be perfect guides as to how to bring this about, they are the best we've got.

Suppose we *didn't* so adjust our observational procedures (while keeping the rest of our methodology the same). Then we would end up using observational procedures which our perceptual theories showed us to be unreliable. And surely this

would be downright irrational, not just an alternative practice we happen not to have adopted.

This is not intended as an argument against anti-realism itself, but only against the view that anti-realism implies the absence of evaluative epistemology. For the anti-realist ought to find the idea that our naturalized epistemology is arbitrary as implausible as the realist does. Even anti-realists, with their methodology-dependent 'reality', will presumably still think of us as perceiving beings responding to other bits of the world. And presumably they will still prefer our perceptions to be accurate rather than inaccurate representations of those bits.³ So it would be as irrational for them as for anybody else to condone our continuing in a perceptual practice that we believe to be unreliable.

That is, while the anti-realist refuses to recognise a REALITY, in Putnam's typography, beyond all theoretical projections, he still recognises a 'reality', the world as we find it. And this in itself is enough to yield a normative justification for his observational practice.

Still, even if evaluative epistemology is possible for anti-realists as well as realists, there is a sense in which it is *more* possible for realists (a sense which they are unlikely to welcome).

In examining the possibility of naturalized epistemology I have throughout started from a point where we already have a set of observational practices and accepted scientific theories: my concern has been to show that those theories then operate with normative force back on those observational practices to either validate or repudiate them. But of course there is a further normative question that might be raised—namely, why go along with any observational practices in the first place? Nothing I have said will move the resolute sceptic who simply starts by refusing to accept any observational judgements or any theories whatsoever (though note that without the given this sceptic ends up with absolutely nothing at all, not even beliefs about his own mental states). If he has no theories to show him he is silly not to accept any observations, then on my account he has no reason to start doing so.

But note that it is only the realist who need recognise a task of persuasion here. He wants to defend an independent reality which he claims the extreme sceptic is missing out on. For the

anti-realist, on the other hand, 'reality' is in the end a spin-off from the adoption of certain procedures. He might allow that once we have started off with certain procedures then there is a question of how best to continue in pursuit of the 'reality' they have given us (and should, I think, answer it in the way I have suggested). But he need not accept that any corresponding question of justification is raised by extreme scepticism, simply because according to him the sceptic has taken away the materials necessary to construct any such question. In the end the important issue for the standing of epistemology is probably not so much the possibility of naturalized epistemology as whether the anti-realist has here stolen something that ought properly to be earned by hard realist toil.

NOTES

¹ See, for instance, Richard Rorty, *Philosophy and the Mirror of Nature*, (Blackwell, 1980), ch. 4.

² Apart from some irrelevant extra complexities it would come to the same thing if we took the 'readings' on humans to be their experiences (*it seemed to X that the stick was bent*) rather than their utterances.

³ Perhaps the anti-realist need not recognise any relation of representation? I would maintain that since he has 'reality' he will have at least 'representation'. Given more space, I would have said something about the notion of representation my overall argument needs. But any such comments would have gone for realists and anti-realists alike.