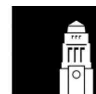
UNIVERSITY OF LEEDS

This is a repository copy of *I—The Presidential Address: Philosophical Scepticism and the Aims of Philosophy*.

White Rose Research Online URL for this paper:
https://eprints.whiterose.ac.uk/183980/

Version: Accepted Version

Article:

Beebee, H orcid.org/0000-0001-7650-3622 (2018) I—The Presidential Address: Philosophical Scepticism and the Aims of Philosophy. Proceedings of the Aristotelian Society, 118 (1). pp. 1-24. ISSN 0066-7374

https://doi.org/10.1093/arisoc/aox017

© 2018 The Aristotelian Society. This is an author produced version of an article published in Proceedings of the Aristotelian Society. Uploaded in accordance with the publisher's self-archiving policy.

Reuse
Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown
If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

eprints@whiterose.ac.uk
https://eprints.whiterose.ac.uk/

AUTHOR-ACCEPTED MANUSCRIPT; PLEASE DO NOT CITE

***The Proceedings of the Aristotelian Society* 118 (2018): 1–24**

Philosophical Scepticism and the Aims of Philosophy[1]

Helen Beebee

Department of Philosophy, University of Manchester

Oxford Road, Manchester M13 9PL, UK

Abstract

I define 'philosophical scepticism' as the view that philosophers do not and cannot know many of the substantive philosophical claims that they make or implicitly assume. I argue for philosophical scepticism via the 'methodology challenge' and the 'disagreement challenge'. I claim that the right response to philosophical scepticism is to abandon the view that philosophy aims at knowledge, and (borrowing from David Lewis) to replace it with a more modest aim: that of finding 'equilibria' that 'can withstand examination'. Finally, I consider what our attitude to our own philosophical theses should be.

I

Introduction. In recent years a lot of attention has been paid to what I'll call the *problem of philosophical scepticism*. Philosophical scepticism, as I'm using the term here, is scepticism *about* philosophy: the claim is that philosophers do not, and cannot, know many of the substantive philosophical claims that they make or implicitly assume. Philosophical scepticism contrasts with what I'll call 'garden-variety' scepticism: the kind of scepticism that claims to cast doubt on more mundane beliefs, such as the belief that the external world exists, the belief that other minds exist, or the belief that the future will resemble the past. A surprisingly large number of philosophers have recently endorsed – or at least come pretty close to endorsing – philosophical scepticism. Thus, for example, Hilary Kornblith says: 'It is just not

[1] I am indebted to the many people and audiences who have commented on versions and ancestors of this paper, especially Chris Daly, Guy Longworth and Tom Smith.

clear that the philosophical enterprise has served as a source of knowledge' (2013, p. 260), and William Lycan says that we have 'only dribs and drabs' of philosophical knowledge 'here and there' – 'nothing to write a song about' (2013, p. 120).

Philosophical scepticism has a variety of sources, of which I shall discuss two: what I'll call the *methodology challenge* and the *disagreement challenge*. According to the methodology challenge – often aimed specifically at metaphysics – our methods can stake no claim to truth-conduciveness; thus we have no particular reason to think that the theories that those methods deliver are true (see e.g. Ross, Ladyman and Spurrett 2007, Ch.1; Ney 2012; Kriegel 2013). The disagreement challenge takes as its starting point the fact of what Sandford Golderg (2013) calls 'systematic peer disagreement'. Pretty much whatever your philosophical view, you'll find plenty of other philosophers who are just as good at philosophy and just as well informed as you are, who sincerely avow a view that's incompatible with your own. And this – so the argument goes – undermines our epistemic entitlement to our own view.

The problem of philosophical scepticism is a problem because we like to think that we *do* have knowledge of substantive philosophical claims: knowledge not just of what follows from what, or of which philosophical theses are compatible or incompatible with other theses or with common sense or whatever, but knowledge of theses about (for example) the nature of reality. Moreover, philosophers tend to assume that the *aim* of philosophy (or least one of its aims) is *also* knowledge – so even if we, individually, don't know very much, we can take comfort in the thought that, somehow or other, we are collectively making progress towards that goal. Philosophical scepticism also puts pressure on this assumption. If we have no reason to think that our methods are truth-conducive, then it's hard to see how we *can* be making progress towards that goal; and if systematic peer disagreement undermines knowledge, only the optimist who thinks that such disagreement might eventually – and by rational means – disappear can hang on to the idea that one day, in principle, knowledge will be achieved. But it's unclear what might warrant such optimism, given philosophy's very long history of failing to resolve major points of substantive disagreement.

In fact, some aspects of contemporary philosophy are quite hard to square with the claim that it is knowledge that we seek. Take David Lewis's thesis of Humean supervenience. When he introduces the thesis (1986, pp. ix-xi), he offers no argument for it. He doesn't even attempt any informal motivation for it. He himself doesn't

even appear to be very interested in whether it's *true*: 'What I want to fight', he says, 'are *philosophical* [and not empirical] arguments against Humean supervenience' (1986, p. xi). And yet, despite the fact that it appeared on the scene as – and to a great extent remains – a mere speculative hypothesis, Humean supervenience has been the subject of an enormous amount of philosophical attention over the last thirty years. That's hard to explain if what we really care about is the pursuit of knowledge. Why did we not, collectively, simply ignore it, pending a passable argument in its favour?

Even in methodological mode, our own accounts of what we are doing sometimes sit uneasily with a self-conception as knowledge-seekers. Consider the way in which the theoretical virtues are normally discussed. Generally speaking, the claim is made that some virtue such as ideological or ontological parsimony or simplicity is a ground for theory *choice* or *preference*. It is very rarely claimed explicitly that they are grounds for thinking the theory in question *true* – that is, that theory choice or preference itself is a matter of choosing the theory that is most likely to be the true one. When that claim *is* explicitly interrogated, the conclusion reached is normally a negative one (Huemer 2009, Kriegel 2013, Willard 2014).

Of course, philosophers are not unaware of the problem of connecting the theoretical virtues with truth-conduciveness. But, given the prominence of the theoretical virtues in discussions in contemporary ontology in particular, if ontologists really think what they're up to is figuring out the truth then it is surprising, to say the least, that they spend so much time – to use David Armstrong's expression – avoiding a compulsory question on the examination paper.

I consider the methodology challenge in §2 and the disagreement challenge in §3. The upshot is that we know, both collectively and individually, very little by way of substantive philosophical theses. Moreover, the lesson of the methodology challenge is that very many of the systematic peer disagreements that undermine our current claims to know substantive theses are unlikely to be tractable. Hence if knowledge is the *aim* of philosophy, then we should be sceptical of the possibility, even in principle, of ever reaching it. In §4 I sketch a position, which I call 'equilibrism', which, drawing on some well-known remarks of David Lewis's, proposes a more modest aim for philosophical enquiry: to find out 'what equilibria there are that can withstand examination'. Finally, in §5, I address what I call the 'commitment problem': the problem of giving an account of the kind of attitude we need to have towards substantive philosophical theses if we are to engage in the

practice of philosophy, given that in many cases we lack warrant for believing those theses.

II

The methodology challenge. Start from the assumption that metaphysics is to be pursued in a broadly naturalistic vein (in some sense of 'naturalistic'): that is, the methods it deploys are not substantially different from those deployed in the sciences. We construct theories on the basis of our evidence, and we choose between rival empirically equivalent theories on the basis of the theoretical virtues: explanatory power, simplicity, parsimony and the like. (This, I would say, is the dominant conception of the method of metaphysics – and certainly of ontology – in contemporary philosophy.) But now, the worry goes, how does metaphysics fare in comparison to the sciences? Answer: spectacularly badly. Our evidence base is woefully bereft, consisting largely as it does of intuitions, about whose alleged connection with the facts of the matter we lack any good story; nor do we have grounds for thinking that, as deployed in metaphysics, maximisation of the theoretical virtues is truth-conducive.[2]

Responses to the methodology challenge have largely focussed on the status of intuitions, and in particular those intuitions that are elicited by thought experiments. I shall briefly consider a particular response to the challenge, namely that of Timothy Williamson (2005), who argues that the 'intuitions' typically thought to be evoked by thought experiments are merely the upshot of our common-or-garden counterfactual thinking – something that, in general, we are pretty good at. Hence there is no special problem of accounting for the provenance or veracity of our intuitions.[3]

Well, perhaps Williamson is right that 'intuitions' (or at least those that are often conceived as 'data' that philosophical theories must accommodate) are really just the upshot of counterfactual thinking. But how far does that get us, really? Even if we have on the table a plausible story about how we come by the judgements elicited by thought experiments, that story does not help us at all in the very many cases

---

[2] Kriegel 2013 spells out this argument in considerably more detail.

[3] Williamson generally eschews using the term 'intuition' for the judgments elicited by thought experiments, noting that in general the 'main current function' of the term 'intuition' and its cognates 'is not to answer questions about the nature of the evidence on offer but to fudge them, by appearing to provide answers without really doing so' (2007, p. 220). I here persist in using the term merely for the sake of expediency and not to signal any disagreement with him in that regard.

where we disagree with each other: plenty of disputes remain. You might think that Designed Ernie in Alfred Mele's 'zygote argument' is not morally responsible, and conclude that non-historicist compatibilism is false (Mele 2006, p. 189). Or you might not (Fischer 2011, p. 271). You might come up with what you take to be a cast-iron case of the failure of the transitivity of causation (McDermott 1995, p. 524). Someone else will inevitably disagree (Lewis 2000, p. 194). You might think you have described two different possible worlds that agree with respect to the distribution of matters of particular fact but disagree with respect to the laws, thus refuting the claim that the latter supervene on the former (Carroll 1994, pp. 57-68). Or you might not (Beebee 2000, pp. 586-92). And so on.[4]

In the case of run-of-the-mill counterfactuals that we deploy outside of philosophy, disagreements can often (though of course not always) be resolved. I say that if your alarm clock hadn't gone off this morning (as it did at 7.45), you would have missed the bus. You point out that you were already awake and about to get up when the alarm went off. I revise my judgement. Or you point out that, while you were indeed woken up by the alarm clock, you never sleep past 8 o'clock so you would have woken up at most 15 minutes later, and you waited at the bus stop for a good 20 minutes, so – again – you would have caught it anyway. Again, I revise my judgement in the light of this new information. And so on.

We have no such method for resolving *philosophers'* clashes of intuitions. When I say that Designed Ernie is morally responsible for his crime or that in our transitivity case *x* really was a cause of *z*, and you disagree, we have no agreed way of reconciling our differences. There are no empirical facts about which you might correct me, in the light of which I would recognise my mistake and change my mind. I have not illegitimately held fixed some facts that I should not have held fixed, or vice versa – not by my lights, anyway. We might try to convince each other otherwise, of course, and sometimes one of us succeeds; but often we both fail. Often we both fail simply because each of us is holding fixed some element of our own background philosophical theory that the other rejects. But then our thought experiment serves only to provide an example of the different consequences of our respective theories; it

[4] Williamson says: 'Levels of disagreement over thought experiments seem to be significantly lower among fully trained philosophers than among novices' (2007, p. 191). I can only report – from the corners of philosophy where I spend my time – that they do not seem so to me.

cannot adjudicate between them. So even if Williamson is right, the 'intuitions' that are thus vindicated are a relatively circumscribed set – namely the ones that philosophers by and large agree about.

The claim here is not merely a special case of the claim that we cannot, in general, know or claim to know things that other people deny. That claim is false: in general we can of course know things that other people deny. We can even know things that others take themselves have good reasons to deny. We can even, often, legitimately (if fallibly) *claim* to know such things – and typically do, in the face of cranks, conspiracy theorists, and those who are simply misinformed or whose evidence base is inadequate. But what the defender of thought-experiment-induced intuitions as a source of knowledge needs to claim is that we can know things that are denied by people in whom we have no prospect of finding any underlying epistemic fault whatsoever.

What about the theoretical virtues? Responses to this aspect of the methodology challenge have been rather thin on the ground.[5] L. A. Paul claims that 'if simplicity and other theoretical desiderata are truth conducive in scientific theorizing, they are truth conducive in metaphysical theorizing' (2012, p. 22). My sense of the matter is that most metaphysicians endorse both this conditional and its antecedent. But both of those claims are contestable, and indeed contested. Starting with the antecedent: as is well known, the claim that the theoretical virtues *are* truth-conducive in the sciences has been contested by Bas van Fraassen (1980). So far as I am aware, that claim has not been refuted.

In fact, as James Ladyman notes, even in the scientific case 'it is widely acknowledged that it is difficult to argue that simplicity or elegance are direct evidence for the truth of a theory; so [scientific] realists have tended towards the view that all the superempirical virtues are subsumed under explanatory power so that the solution to the underdetermination problem is IBE [inference to the best explanation]' (Ladyman 2012, p. 42). So suppose we grant (as of course van Fraassen does not) that IBE is indeed a truth-preserving inference in the sciences. May we infer, as Paul suggests, that the same is true in metaphysics? Well, Ladyman argues that the role of IBE as a means of theory-choice in science has been overstated: 'explanatory power', he says, 'plays the role that it does in [scientific] theory choice because of the

[5] An exception is Cowling 2013.

relationship between theoretical explanation and the empirical virtues of scientific theories. We have inductive grounds for believing that pursuing simplicity and explanatory power in science will lead to empirical success, but no such grounds where we are dealing with distinctively metaphysical explanations, since the latter is completely decoupled from empirical success' (2012, p. 46). In other words, to the extent that – and *only* to the extent that – best explanations in science have turned out to be predictively superior to their rivals, we have *inductive* grounds in the scientific case for thinking that inference to the best explanation is truth-conducive. But no such inductive grounds exist in the case of metaphysics.

Finally, what if we set all this aside and granted that the theoretical virtues are, individually, truth-conducive – or, roughly equivalently, that IBE is a truth-preserving inference in metaphysics? Would we then have a method for deciding between rival metaphysical theories? Well, we would if we had some agreed way of trading the virtues off against one another (correspondingly: a method for determining which of various rival explanations that make different trades is the best). But of course we don't – and plenty of disputes in metaphysics arise from differences in our trade-offs. Paul herself says: 'While choosing one theory over its competitors based on a particular balancing of theoretical virtues is perfectly rational, we will almost always have a range of alternative theories to choose from: this is because there are many different ways of balancing respective theoretical virtues and hence many different ways of arriving at reflective equilibrium' (2012, p. 22). I think she's right about that. But the sceptical worry here is precisely *how* it can be rational to believe one theory rather than another when we have no grounds for thinking that our particular balancing act is more likely to deliver the truth than those of our competitors.

Consider Argle and Bargle's dispute about the nature of holes (Lewis & Lewis 1970). At the end of their discussion about whether or not holes are material objects, Bargle (who thinks that holes are not material objects) accuses Argle (who thinks they are) of violating common sense in his description of a kitchen-roll tube that has a small hole in its side: 'you must pay a great price in the plausibility of your theories', he says. The conversation continues:

> *Argle:* Agreed. We have been measuring the price. I have shown that it is not so great as you thought; I am prepared to pay it. My theories can earn credence by

> their clarity and economy; and if they disagree a little with common opinion, then common opinion may be corrected even by a philosopher.
>
> Bargle: The price is still too high.
>
> Argle: We agree in principle; we are only haggling.
>
> Bargle: We do. And the same is true of our other debates over ontic parsimony. Indeed, this argument has served us as an illustration … of the nature of our customary disputes. (Lewis & Lewis 1970, pp. 211-12)

Argle and Bargle have reached stalemate in their cost-benefit analysis of the claim that holes are material objects. They agree on what the costs and benefits of the view are: on the benefit side we have 'clarity and economy', and on the cost side we have dissonance with common sense. But they disagree over what the costs and benefits are *worth*: they weigh them up and get different results. So – here's the sceptical problem again – it is unclear why Argle is entitled to hold that his theory 'earns credence' by its clarity and economy despite the cost, when he has no response to Bargle's view that it doesn't.

The methodology challenge, then, is that we have no grounds for trusting our intuitions – perhaps at all, or, conceding the point to Williamson, when they conflict with those of our philosophical peers. And we have no grounds for thinking that the theoretical virtues are truth-conducive; indeed, even if we grant that they are, we have no agreed methodological standard for trading them off against one another, and hence, where such trade-offs are what divide us, we have no more reason to believe our own theory than we do a theory whose (alleged) justification comes from making the trade differently.

From here on, I shall shift my attention from metaphysics in particular to philosophy in general, and will assume without argument that a version of the methodology challenge applies to other areas of philosophy as well. That is, admittedly, a big assumption, since the kind of 'naturalistic' approach at which the methodology challenge is aimed is largely (though not entirely) confined to contemporary metaphysics – and not even to the whole of metaphysics, given that not all metaphysicians sign up to this way of proceeding.

On the other hand, intuitions are the stock-in-trade of much of philosophy, and a version of the Argle's and Bargle's situation is also pretty commonplace outside metaphysics. Simpler, more elegant theories – whether in ethics or aesthetics or philosophy of language or wherever – have an annoying tendency to throw up some counter-intuitive consequences, thus generating disputes about whether the solution to that problem is to add epicycles to the theory in order to save the intuitions or rather to attempt to explain away, or bite the bullet with respect to, the recalcitrant intuitions. More generally, at least in many areas of philosophy it is hard to see how other conceptions of appropriate philosophical methods are going to be any less susceptible to worries about the truth-conduciveness of those methods.

III

The disagreement challenge. The disagreement challenge starts from the obvious and undeniable fact of pervasive systematic peer disagreement. (In §2 I briefly considered two sources of such disagreement: divergent intuitions with respect to thought experiments, and differences in how we are to trade costs against benefits.) How – so the challenge goes – can we claim to know or be justified in believing our philosophical views, when we know that equally capable and well-informed philosophers disagree with us – and often continue to do so despite our best efforts to persuade them of the error of their ways?

One kind of response to the disagreement challenge is to try to explain how a given philosopher *S* can legitimately claim to know or be justified in believing some proposition *p* even in the face of persistent disagreement from *S*'s peers about whether or not *p*. Broadly speaking, such responses try to argue that whatever evidence that the belief in ¬*p* by one's philosophical peers might in principle provide for the falsity of *p*, that evidence can, at least in some contexts or circumstances, legitimately be ignored.

Thus for example David Lewis, in a letter to Richard Cartwright, says:

> Subject to a qualification, I … say that the one who's right in a philosophers' impasse typically knows that he's right ...
>
> The qualification is that the knowledge may temporarily vanish in the course of philosophical discussion because 'alternatives' are made 'relevant' that normally are not … [Lewis refers to his 1979, p. 355 here; see also Lewis 1996].

> Maybe I cannot truly say, in the midst of a discussion of far-fetched sceptical hypothesis, that I know I have hands. Attending to these hypotheses temporarily raises the standards for what may be called knowledge. But I can truly say, even then, that under more ordinary circumstances I can truly say that I know I have hands. (1989, p. 3)

The context of the discussion here is Graham Priest's view that contradictions can be true; Lewis, of course, is firmly of the view that they cannot. And he says that he knows that he is right (and that his view is reasonable) and Priest is wrong (and that his view is unreasonable), even though there is no serious possibility of establishing this to Priest's satisfaction – it is a genuine 'philosophers' impasse'. So Lewis is not merely interested here with garden-variety sceptical scenarios – evil demons and the like. Rather, he is claiming that his contextualist account of knowledge applies equally to philosophers' impasses, where the alternative hypothesis is something that is actually believed, and by their own standards reasonably so, by actual philosophers.

Another example: in his book *The Strife of Systems*, Nicholas Rescher argues that persistent disagreement is an inevitable feature of philosophy because philosophers inevitably have different cognitive values – and even when they share the same cognitive values, they won't all get weighed against each other in the same way. Nonetheless, he holds that one can perfectly well be said to *know* philosophical theses (1985, p. 208); 'from the vantage point of one's *own* cognitive-value perspective – the only one that one has – only one optimally adequate position on philosophical issues is rationally warranted' (1985, p. 265). Again, then, philosophical knowledge is entirely possible: one can know *p* by one's own lights, while fully appreciating that by different lights *p* does not constitute knowledge (and indeed is not even rationally warranted).

Suppose for now that some such response to the sceptical problem is viable. What these kinds of response have in common is that the challenge they are attempting to meet is that of vindicating the knowledge claims of individual philosophers. But those responses fail to address a broader problem. We often take the bearer of knowledge to be not an isolated individual but a community of some sort. When, at the point in the TV crime drama where the detective gathers her team in the incident room and asks, 'what do we know?' she is asking for a pooling of resources: her aim is to compile a list of (purported) facts about the suspect that can be used as a

basis for further investigation. That list will generally consist of individuals' items of knowledge, of course, but it would be of no help at all to the investigation if it were to turn out that different team members' items of knowledge only *counted* as knowledge in the absence of salient alternatives, or from their own 'cognitive-value perspective'. Were PC Smith to convey her (purported) knowledge – based on excellent evidence – by saying 'the suspect lives in her mother's house', and PC Jones to convey *her* (purported) knowledge – equally well grounded in a different set of evidence – by saying 'the suspect shares a flat with her sister', the detective would quite rightly conclude that the team collectively does not know what the suspect's living arrangements are. Similarly, when we ask 'what do we know about climate change?', or 'how much do we know about the causes of arthritis?', or whatever, any coherent answer to our question will be a set of purported facts whose members are consistent with one another. No proposition *p* that is a matter for persistent peer disagreement can make it onto the list, since there will be no grounds for having *p* rather than ¬*p* on the list.

The TV detective and her team have an aim: to catch the perpetrator. Philosophy, too, has an aim, or perhaps several aims; one such aim is often thought to be knowledge. That aim, if it is an aim, is clearly a collective aim: it is an aim of philosophy *qua* intellectual endeavour, pursued by a broad epistemic community that spans not only the globe but several millennia. So even if we could save *individual* philosophical knowledge from the disagreement problem by somehow insulating our individual beliefs from the challenge posed by the disagreement of our peers, such knowledge would not, in itself, constitute progress towards our collective aim.

Consider what would happen if the philosophy detective were to gather the philosophy team – the epistemic community of philosophers – in the incident room and ask the question, 'what do we know?'. The list of purported facts – at least when it comes to substantive philosophical theses – is likely to be rather short. That is not because we contain within our midsts garden-variety sceptics intent on raising their sceptical hypotheses, thereby challenging our claims to know that the external world exists, or that there are other minds, or whatever. The garden-variety sceptic does not claim to know her sceptical hypothesis, so she comes to the room not with a list that conflicts with ours but rather – concerning whatever domain she is sceptical about – no list at all. Doubtless, for any claim that the philosophy detective might want to put on her list of what we collectively know, a garden-variety sceptic is likely to complain

on the grounds that she has up her sleeve a sceptical scenario that cannot be ruled out; after all, very many philosophical claims presuppose the existence of the external world or other minds or whatever. But we are here assuming for the sake of the argument that *individual* philosophical knowledge is possible, and hence that we are entitled to ignore such scenarios. But even with the garden-variety sceptic banished from the room, the situation in the philosophy incident room still leaves us with little by way of collective knowledge of substantive philosophical theses. Argle and Bargle are in the room. So are the compatibilists and the incompatibilists; the metaphysical realists, the idealists and the pragmatists; the consequentialists and the deontologists; the semantic externalists and internalists; the physicalists and the dualists and the panpsychists; the scientific realists and the constructive empiricists; the Humeans and the dispositional essentialists; and so on. Grant that, individually, the occupants of the room are entitled to claim to know the items on their lists. Even so, collectively they know very little.

When aimed at the claim that we *collectively* know a lot of substantive philosophical theses, then, the problem of systematic disagreement cannot be met by appealing to the idea that, in the face of such disagreement, we can salvage our *individual* claims to know such theses. But I suggest that the difference between garden-variety sceptical scenarios on the one hand and the kinds of 'sceptical scenario' which constitute rival substantive philosophical theses that are a matter for systematic peer disagreement on the other serves to put pressure on even our *individual* claims to know any substantive philosophical theses. Consider the stalemate that is the outcome of Argle and Bargle's dispute about the nature of holes again. According to Lewis's contextualist solution to the problem of philosophical scepticism, insofar as Argle can claim to know that holes are material objects, Bargle's contention that they are not constitutes a sceptical challenge: a relevant alternative to Argle's view that, once made salient, Argle cannot rule out. Conversely, Argle's contention that holes are material objects constitutes a sceptical challenge to Bargle's view that they aren't. Collectively Argle and Bargle do not know whether or not holes are material objects. But do they even *individually* know – when not discussing holes with one another – whether or not holes are material objects?

Consider the standard zebra/painted mule case (Dretske 1970, pp. 1015-16). Cargle looks at the animal in the zebra enclosure and quite reasonably concludes on the basis of her evidence that the animal is a zebra. Dargle points out that Cargle

cannot rule out the possibility that she is looking at a mule cleverly painted to look just like a zebra, and concludes that Cargle does not *know* that the animal is a zebra. There are of course various anti-sceptical moves one might make at this point. We might take the contextualist route, and claim that while Dargle's sceptical scenario is on the table she is right that Cargle does not know; but in normal circumstances, when Dargle is not around, the sceptical scenario is not a relevant possibility and hence, in those circumstances, Cargle does know. Or we might go externalist and claim that, while Cargle may not know that she knows the animal to be a zebra, nonetheless – assuming it really is a zebra and her perceptual faculty is operating normally – she does know that it's a zebra. Or whatever. Let's assume some such story succeeds.

Now compare Dargle with Fargle. Unlike Dargle, Fargle claims to *know* that the animal is not a zebra but a painted mule. Fargle is a sensible, rational, generally trustworthy interlocutor. When pressed, Fargle can cite some evidence. Fargle has evidence from the local newspaper that heard the zoo has fallen on hard times recently, and has heard from a reliable source that it has, where possible, taken to disguising cheap animals as more exotic, expensive animals when the expensive ones die. Also she was here last month – before the zoo's financial crisis – and the animal in the zebra enclosure then looked decidedly unwell and, from memory (and her memory for the appearance of animals is, she claims, pretty reliable), dissimilar in various other respects to the animal now in front of her. She points out that it is actually quite easy for someone with the requisite skill to paint a mule to look just like a zebra; she shows Cargle an authentic-seeming Youtube clip of someone doing just that, and Cargle really can't tell the difference.

Some of this evidence seems somewhat dubious to Cargle, though not egregiously so; he is certainly not in a position to decisively rule any of it out. Still, on balance he thinks Fargle has not given sufficient evidential weight to the brute implausibility of her claim, nor to the fact that it would be extremely risky, to put it mildly, to the zoo's reputation to replace its deceased animals with disguised cheaper doppelgangers. Cargle does not think Fargle is entitled to claim to know that the animal is not a zebra. Fargle, for her part, thinks that Cargle is being too naïve. He is too unwilling to take into account the zoo's perilous financial situation and the lengths that one might go to, in those circumstances, to keep it afloat; and he is too unwilling to take on board Fargle's testimony about the change in the animal's appearance or

the reliability of her source. Fargle does not think Cargle is entitled to claim to know that the animal is a zebra.

Either Cargle and Fargle are both right about the other's entitlement to claim to know, or they are both wrong. (If I haven't told the story in such a way as to convince you that this is so, tweak their respective evidential situations so that it turns out to be so.) Which is it? I think they are both right: neither of them is entitled to claim to know; and indeed whichever of them happens to be right about the zebra does not know. In particular, even if we grant that Cargle was entitled to ignore Dargle's sceptical painted-mule scenario once she had wandered off to get some chips, knowing what he now knows about the *Fargle's* evidential situation he is not entitled to ignore Fargle's claim that the animal is not a zebra when *she* wanders off to get some chips. And Fargle is in the same position with respect to Cargle's claim that it *is* a zebra.

So it is with Argle and Bargle. Argle knows – not just while discussing holes with Bargle but all the time – that Bargle's position is entirely sensible by her own lights (valuing, as she does, common sense over theoretical economy), and Bargle knows that Argle's is entirely sensible by *his* lights. Their situation is analogous to that of Cargle and Fargle, and not to that of Cargle and Dargle. And so it is with systematic disagreements in philosophy more generally. We differ from Cargle and Fargle in that the sources of our reasons to believe our philosophical claims are generally more complex and abstract and, correspondingly, we have less reason than they do to think that our disagreements can be resolved. (Cargle is surely right that if the zoo really has substituted a painted mule for a zebra, it will get found out soon enough. By contrast, we know that many falsehoods have continued to be endorsed by philosophers for centuries.) Indeed, given the methodology challenge of §2, I claim we have *no* reason to think that many of our disagreements will be resolved. But of course these dissimilarities only serve to make our position worse, and not better, than that of Cargle and Fargle.

IV

Equilibrism. What to do, then, if we do not know – collectively or individually – very much at all when it comes to substantive philosophical theses, and if there are no serious prospects of realising the aim of collective knowledge of such theses even in the long run? Well, we could aim for something else instead: something achievable in

principle, and – ideally – something that does not require us to change the way we go about our business in order to make progress with respect to that aim.

I want to propose an aim that fits the bill by drawing on some well-known remarks of Lewis from the Introduction to Volume I of his *Philosophical Papers*:

> The reader in search of knock-down arguments in favor of my theories will go away disappointed. Whether or not it would be nice to knock disagreeing philosophers down by sheer force of argument, it cannot be done. Philosophical theories are never refuted conclusively. (Or hardly ever. Gödel and Gettier may have done it.) The theory survives its refutation—at a price. Argle has said what we accomplish in philosophical argument: we measure the price. Perhaps that is something we can settle more or less conclusively. But when all is said and done, and all the tricky arguments and distinctions and counterexamples have been discovered, presumably we will still face the question which prices are worth paying, which theories are on balance credible, which are the unacceptably counterintuitive consequences and which are the acceptably counterintuitive ones. On this question we may still differ. And if all is indeed said and done, there will be no hope of discovering still further arguments to settle our differences.
>
> It might be otherwise if, as some philosophers seem to think, we had a sharp line between 'linguistic intuition', which must be taken as unchallengeable evidence, and philosophical theory, which must at all costs fit this evidence. If that were so, conclusive refutations would be dismayingly abundant. But, whatever may be said for foundationalism in other subjects, this foundationalist theory of philosophical knowledge seems ill-founded in the extreme. Our 'intuitions' are simply opinions; our philosophical theories are the same. Some are commonsensical, some are sophisticated; some are particular, some general; some are more firmly held, some less. But they are all opinions, and a reasonable goal for a philosopher is to bring them into equilibrium. Our common task is to find out what equilibria there are that can withstand examination, but it remains for each of us to come to rest at one or another of them. (1983, p. x)

I propose that we conceive what Lewis sees as our common task as our *aim*: that is, our collective aim is to 'find out what equilibria there are that can withstand examination'. Intractable disagreements, on this picture, make for different equilibria:

Argle and Bargle agree on a great many things, we may suppose, and to that extent their respective equilibria will overlap. But they disagree about the nature of holes – and, we can assume, on other issues besides, where their different cost-benefit trade-offs deliver different views – and those differences fall outside the intersection of their respective equilibria.

With our aim thus conceived, philosophical scepticism, conceived as a threat to the possibility of knowledge of many or most substantive philosophical theses, turns out not to pose a threat to our ability to make progress with respect to philosophy's aim. The fact that we bring to the table different sets of substantive assumptions and different methodological principles – the kinds of differences that generate intractable disagreements – is no threat to our ability to contribute to the aim of philosophy. On the contrary: such differences are precisely what generate the plurality of equilibria that it is our collective aim to uncover. And if our individual aim is to find an equilibrium position of our own – or at any rate a partial one, given the vastness of the task – then it no longer really *matters* if our substantive philosophical views are unjustified. Argle need not hide under his desk when Bargle is around, fearing that Bargle's insistence on reminding him of his view of holes will temporarily rob him of his knowledge that holes are material objects. Insofar as Argle and Bargle are convinced that their argument has run its course and each is occupying a stable position, neither poses any kind of threat to the other.

I'll call the combination of this broadly Lewisian conception of the aims of philosophy on the one hand, and philosophical scepticism – understood as the thesis that knowledge of at least very many substantive philosophical theses is impossible – on the other, *equilibrism*. Equilibrism is, I hope, a relatively conservative position with respect to the practice of first-order philosophy. I take it that at least a very large part of what most individual philosophers do just *is* a matter of trying to find a (partial) point of equilibrium, at which one is happy to – as Lewis puts it – come to rest. (Whether or not that is what they think of themselves as doing is another question, of course; I don't expect equilibrism to be conservative with respect to *that*.) In particular, equilibrism does not in the least undermine the philosopher's stock-in-trade of argument and counter-argument (see §5 below). Mostly, then, equilibrism is merely an alternative account of why it is we're doing all of this: of what purpose we are serving in going about our business in the way that we generally do.

It's important to stress that even the aim of bringing all of one's own opinions into equilibrium is an exceptionally tall order – one that surpasses the lifespan and cognitive powers of any actual philosopher. For one thing, our own positions are only ever partial; they never span the whole of philosophical space. You might have your view in some area of metaphysics nicely buttoned down, or so it seems – but there is always the possibility (indeed, likelihood) that when you attempt to bring your guiding assumptions and methodological principles to bear on some new area – the philosophy of language, say, or meta-ethics or the philosophy of physics – you find yourself forced to accept some unpalatable claims: equilibrium is disturbed and a rethink is required.

Nor should the range of 'equilibria that can withstand examination' be overstated. To the extent that there can be said to be some common knowledge and a common set of very general methodological standards in philosophy, there will be overlap between equilibria. Nor does just any old philosophical disagreement over whether *p*, no matter how persistent, necessarily engender competing equilibria that differ with respect to whether *p*. Some persistent disagreements do get resolved eventually: it really does turn out that one side was right and the other was wrong. It is of course hard to tell, while disagreement persists, whether or not it will eventually be resolved; and trying to resolve such disagreements is no less recommended by equilibrism than it is by a conception of philosophy as aiming at knowledge. Equilibrism recommends only that – when the occasion demands – stopping the argument and moving on is a legitimate move to make in pursuit of our collective aim.

V

The commitment problem. Lewis himself conceives the aim of the individual philosopher as a matter of bringing all of one's 'opinions' into equilibrium. But if knowledge of substantive theses can be claimed neither by us collectively nor even by each of us individually, doesn't that leave us – firstly – without any rational entitlement to so much as hold opinions at all? And in that case, how can our individual aim be to bring all of our *opinions* into equilibrium? Second, *argument* is of course central to the philosophical enterprise. And to present an argument is, surely, to present a purported justification for one's view. If no justification is in the offing, how can argument be preserved as the central implement in the philosopher's toolkit?

Let's start with the second question, about the purpose of argument. Suppose that you offer an argument for *p*. It looks very much as though – and indeed you might explicitly claim that – you are offering a justification for *p*, or a reason to believe *p*. Surely, one might think, the point of argument is precisely to figure out which claims are justified and which are not – otherwise, what *is* the point?

Well, your argument for *p* constitutes a challenge to an audience that endorses ¬*p* – and also an invitation to endorse *p* to an audience that has no view one way or another. Philosophical debate ensues. Your audience may attempt to find fault with your argument: it is a bad argument (perhaps, according to them, there are clear cases of arguments of that form that have true premises and false conclusions), one or more of its premises are false (according to them – though of course you may disagree), or whatever. They may or may not succeed by their own lights, or indeed by yours. Adjustments ensue, wherever they need to be made, in order for equilibrium to be restored to whoever's position was challenged by the original argument itself or by the ensuing debate. We move a tiny step closer to philosophy's aim by filling out a small corner of the space of equilibria. All of this is entirely rational.

But this still depends on the idea that it is legitimate to 'endorse' a position – to have an 'opinion', in Lewis's terms. It is crucial to the philosophical endeavour, I think, that we *take a view*. So – to return to the first question posed above – does equilibrism require us to rescind from taking a view *at all*? I say it doesn't – but a story is needed.

There has been a fair amount of discussion recently concerning whether the agnosticism that is (allegedly) recommended by the phenomenon of systematic peer disagreement is compatible with taking a view. Sandford Goldberg argues that the appropriate doxastic attitude in the face of systematic peer disagreement is that of 'attitudinal speculation', which is 'a matter of having a degree of confidence above .5 but below the threshold warranting outright belief' (2013, p. 285). Goldberg argues that attitudinal speculation does all the work that we might have thought we needed belief to do, including making sense of the idea that, when we assert philosophical claims, we do so sincerely. And he takes 'having' a view, or regarding a view as defensible, to amount to attitudinal speculation plus a certain kind of attachment: 'One who endorses and defends a philosophical view', he says, 'is typically more motivated to persist in defense of the view when challenged, than is one who merely

speculates that *p*. (We are more committed, and perhaps more emotionally attached, to our philosophical views, than we are to our speculations.)' (2013, p. 284).

Zach Barnett objects that the requirement of a degree of confidence above 0.5 is too demanding, and argues for the view that 'a person's philosophical views should be her *disagreement-insulated inclinations*' (forthcoming, p. 13), where one's 'inclination' towards *p* is insulated, specifically, from the evidence against *p* that arises from the fact there is systematic peer disagreement concerning *p*. Barnett argues that this somewhat weaker condition on having a view nonetheless satisfies the sincerity requirement.

Goldberg's and Barnett's proposals are broadly friendly to equilibrism – after all, they both seek to make sense of the idea that sincere commitment to a philosophical view is compatible with the absence of full-blown belief (and hence *justified* belief). However, from the equibrilist perspective Goldberg and Barnett are overly concerned about the evidential status of the fact of systematic disagreement. Recall our common aim on the equilibrist view – to 'find out what equilibria there are that can withstand examination' – and the corresponding individual aim of 'coming to rest' at one of them. From the point of view of those common and individual aims, the fact that one's own position is subject to systematic peer disagreement is not something that presents a *problem*: the fact that other people disagree with us (and for good reasons, by their own lights) on matters of philosophical substance is exactly what we should expect. Argle and Bargle disagree over whether holes are material objects because they have an irresolvable disagreement about how to trade consonance with common sense against theoretical economy. For each of them, at this point, the fact of this disagreement is simply not a threat to their own individual aim. Argle does not need to establish that his own view is worthy of degree of belief of > 0.5; nor does he have any need to insulate himself from Bargle's disagreement. He may get on with the business of pursuing his own individual aim of finding an equilibrium at which he can 'come to rest' while being fully aware of the fact that Bargle is aiming to find a different equilibrium of *her* own that enshrines claims that he, Argle, disagrees with. And from the point of view of the collective aim of philosophy, too, this is all well and good.

In rejecting Goldberg's and Barnett's proposals, however, I have of course just made life harder for myself: equilibrism still needs an account of what it is to endorse

a philosophical claim – in effect, an account of what it is to *disagree* about anything, as Argle and Bargle disagree with each other over the nature of holes. I suggest – tentatively – that something like van Fraassen's view about the 'acceptance' of scientific theories can be made to solve the problem. Constructive empiricism faces a similar problem to equilibrism: given that science does not aim at the truth, and hence knowledge, of scientific theories – it only aims at empirical adequacy – how can we make sense of the fact that scientists *do* (and indeed *must,* for the purposes of pursuing that aim) make assertions that apparently express belief in claims about unobservables for which they have no justification? Van Fraassen's answer, in short, is that 'acceptance' of and belief in such claims are two distinct phenomena – and that only acceptance is required.

Here is what, for van Fraassen, acceptance amounts to:

> Acceptance of theories (whether full, tentative, to a degree, etc.) is a phenomenon of scientific activity which clearly involves more than belief [in empirical adequacy]. One main reason for this is that we are never confronted with a complete theory. So if a scientist accepts a theory, he thereby involves himself in a certain sort of research programme. That programme could well be different from the one acceptance of another theory would have given him, even if those two (very incomplete) theories are equivalent to each other with respect to everything that is observable – in so far as they go.
>
> Thus acceptance involves not only belief [in empirical adequacy] but a certain commitment. Even for those of us who are not working scientists, the acceptance involves a commitment to confront any future phenomena by means of the conceptual resources of this theory. It determines the terms in which we shall seek explanations. If the acceptance is at all strong, it is exhibited in the person's assumption of the role of explainer, in his willingness to answer questions *ex cathedra*. Even if you do not accept a theory, you can engage in discourse in a context in which language use is guided by that theory – but acceptance produces such contexts. There are similarities in all of this to ideological commitment. A commitment is of course not true or false: The confidence exhibited is that it will be *vindicated*. (1980, pp. 12-13)

Roughly, then, the idea is that in 'accepting' a scientific theory that is ontologically committed to unobservables, the scientist does not (or at least need not) adopt the attitude of *belief* towards what the theory says about those unobservables – or indeed any attitude that is a kind of watered-down belief-substitute such as Goldberg's 'attitudinal speculation' or Barnett's 'inclination'. Rather, the scientist takes on a kind of practical commitment – a commitment to 'confront any future phenomena by means of the conceptual resources of this theory', a 'willingness to answer questions *ex cathedra*' and to assume 'the role of explainer', and so on – in other words, to speak and write and act as though the theory is true (or at least, insofar as the theory is incomplete, true or close to true in broad outline).

This connects, I think, with Goldberg's and Barnett's concern about sincerity: if we are not entitled to *believe* the claims of our own theories, in what sense can they truly be said to be *our* theories? How can we sincerely endorse the claims those theories make? Acceptance, I take it, is supposed to deliver sincerity. The attitude of acceptance does not, of course, constitute sincere *belief*, but it is sincere nonetheless. The working scientist adopts a theoretical standpoint, works hard to accommodate the existing evidence and explore further consequences of her theory, makes adjustments where necessary, and so on. And she can do all of this entirely sincerely while yet merely accepting rather than believing her own theory, according to van Fraassen. All of this, I suggest, amounts to the scientist's *taking a view* in as much of a sense of 'taking a view' as is required of her for the purposes of playing her part in the progress of science.

If something like van Fraassen's notion of acceptance really can constitute a legitimate sense in which one might 'take a view', then it can, I think, be applied to the working philosopher no less than to the working scientist. The aims of science (according to van Fraassen) and the aims of philosophy (according to me) differ, of course: the aim of empirical adequacy in science is very different to the pluralist aim in philosophy of discovering the equilibrium positions that can withstand examination. But in each case the acceptance of a theory that one cannot rationally believe serves a purpose relative to that aim. In the case of science, the aim of empirical adequacy demands that theories that posit unobservables are developed and tested, and in the case of philosophy the aim of the discovery of equilibria demands that we take on

board a set of core assumptions and methodological prescriptions in order to develop and scrutinise an equilibrium position of our own that can withstand examination.[6]

References

Alston, William P. 2007: 'Audi on Non-doxastic Faith', in M. Timmons, J. Greco & A. Mele (eds), *Rationality and the Good*: *Critical Essays on the Ethics and Epistemology of Robert Audi.* New York: Oxford University Press, pp. 123-40.

Audi, Robert 2008: 'Belief, Faith, and Acceptance', *International Journal for Philosophy of Religion,* 63, pp. 87-102.

Barnett, Zach forthcoming: 'Philosophy Without Belief', *Mind*, available online: https://doi.org/10.1093/mind/fzw076

Beebee, Helen 2000: 'The Non-governing Conception of Laws of Nature', *Philosophy and Phenomenological Research,* 61, pp. 571-94.

Carroll, John W. 1994: *Laws of Nature.* Cambridge: Cambridge University Press.

Cowling, Sam 2013: 'Ideological Parsimony', *Synthese,* 190, pp. 3889-908.

Dretske, Fred 1970: 'Epistemic Operators', *The Journal of Philosophy,* 67, pp. 1007-23.

Fischer, John Martin 2011: 'The Zygote Argument Remixed', *Analysis,* 71, pp. 267-72.

Goldberg, Sandford 2013: 'Defending Philosophy in the Face of Systematic Disagreement', in D. E. Machuca (ed.), *Disagreement and Skepticism*. New York: Routledge, pp. 277-93.

Huemer, Michael 2009: 'When is Parsimony a Virtue?', *The Philosophical Quarterly,* 59, pp. 216-36.

Kornblith, Hilary 2013: 'Is Philosophical Knowledge Possible?', in D. E. Machuca (ed.), *Disagreement and Skepticism*. New York: Routledge, pp. 260-76.

Kriegel, Uriah 2013: 'The Epistemological Challenge of Revisionary Metaphysics', *Philosophers' Imprint,* 13, pp. 1-30.

Ladyman, James 2012: 'Science, Metaphysics and Method', *Philosophical Studies,* 160, pp. 31-51.

[6] A somewhat different notion of 'acceptance' is discussed in the philosophy of religion in the context of non-doxastic faith (see Alston 2007 and Audi 2008). I lack the space to consider whether Alston's conception of 'acceptance' might work out as an alternative to van Fraassen's in the context of having a philosophical view.

Lewis, David K. 1979: ‘Scorekeeping in a Language Game’, *Journal of Philosophical Logic*, 8, pp. 339–59.

-- 1983: *Philosophical Papers,* Vol. I. New York: Oxford University Press.

-- 1986: *Philosophical Papers,* Vol. II. New York: Oxford University Press.

-- 1989: Letter to Richard Cartwright, 4 Sept. 1989, David Lewis Papers, C1520, ‘Cartwright, Richard’, Box B-000662 Folder 7-8, Princeton University Library. Available online at URL = http://www.projects.socialsciences.manchester.ac.uk/lewis/letter-month-august-2017/, accessed 30.8.17.

-- 1996: ‘Elusive Knowledge’, *Australasian Journal of Philosophy*, 74, pp. 549-67.

-- 2000: ‘Causation as Influence’, *The Journal of Philosophy*, 97, pp. 182-97.

-- & Stephanie Lewis 1970: ‘Holes’, *Australasian Journal of Philosophy*, 48, pp. 206-12.

Lycan, William G. 2013: ‘On Two Main Themes in Gutting’s ‘What Philosophers Know’, *Southern Journal of Philosophy*, 51, pp. 112-20.

McDermott, Michael 1995: ‘Redundant Causation’, *The British Journal for the Philosophy of Science,* 46, pp. 523-44.

Mele, Alfred R. 2006: *Free Will and Luck*. New York: Oxford University Press.

Ney, Alyssa 2012: ‘Neo-positivist Metaphysics’, *Philosophical Studies*, 160, pp. 53-78.

Paul, L. A. 2012: ‘Metaphysics as Modeling: The Handmaiden’s Tale’, *Philosophical Studies,* 160, pp. 1-29.

Rescher, Nicholas 1985: *The Strife of Systems*. Pittsburgh: The University of Pittsburgh Press.

Ross, Don, James Ladyman & David Spurrett 2007: ‘In Defence of Scientism’, in J. Ladyman & D. Ross, with D. Spurrett & J. Collier, *Every Thing Must Go*, Oxford: Clarendon Press, pp. 1-65.

Van Fraassen, Bas 1980: *The Scientific Image.* New York: Oxford University Press.

Willard, M. B. 2014: ‘Against Simplicity’, *Philosophical Studies,* 167, pp. 165-81.

Williamson, Timothy 2005: ‘Armchair Philosophy, Metaphysical Modality and Counterfactual Thinking’, *Proceedings of the Aristotelian Society,* 105, pp. 1-23.

-- 2007: *The Philosophy of Philosophy*. Oxford: Blackwell.