

4 Disagreement in Philosophy

An Optimistic Perspective

Many philosophers (and non-philosophers) think persistent disagreement is a characteristic feature of philosophy. This thought is central in much scepticism about philosophy as a discipline: *Surely, the thought goes, if the goal is to find answers to important questions, philosophers have failed: they've been at it for more than 2000 years and they haven't managed to agree on anything!*

This paper is an extended response to that line of thought and also an exploration of the sources of persistent disagreement in philosophy.¹ I argue for three central theses:

1. We have no evidence that there is more persistent disagreement in philosophy than in relevantly similar disciplines.
2. There is *some* persistent disagreement in philosophy, but it is: (a) inevitable (because of the failure of *Evidence Neutrality*) and (b) intellectually valuable.
3. Convergence is irrelevant to the intellectual value of any view – the value of the views in for example, contemporary physics has nothing to do with the convergence on those views within academia.

THE VIEW THAT PHILOSOPHY IS PLAGUED BY PERSISTENT DISAGREEMENT

In 1933 Ferdinand Schiller wrote:

We are all aware that philosophers are even more prone to disagree than doctors, and probably all of us are ready upon occasion to contribute our quota to the disagreements that mark, and scar, the face of philosophy.

(Schiller 1933: 118)

Here is a similar sentiment expressed in 2004 by Peter van Inwagen:

Disagreement in philosophy is pervasive and irresolvable. There is almost no thesis in philosophy about which philosophers agree. If there is any philosophical thesis that all or most philosophers affirm, it is a negative thesis: that formalism is not the right philosophy of mathematics, for example, or that knowledge is not (simply) justified, true belief. That is not how things are in the physical sciences.

(Van Inwagen 2004: 332)

Chalmers (2015) defends a related view. Chalmers' focus is on what he calls 'the Big Questions of philosophy'. These are questions like:

- *What is the relationship between mind and body?*
- *How do we know about the external world?*
- *What are the fundamental principles of morality?*
- *Is there a God?*
- *Do we have free will?*

According to Chalmers there hasn't been what he calls 'large collective convergence' on answers to the Big Questions. Here is his account of *collective convergence*:

We can define collective convergence on an answer over a period as the increase in degree of agreement on that answer from the start of the period to the end of the period. Degree of agreement can be defined using one of various mathematical measures of agreement across a group of people on a set of issues. Collective convergence (*simpliciter*) over a period is defined as the collective convergence on the dominant answer at the end of that period over the period.

(Chalmers 2015: 5–6)

Here is the account of *large collective convergence*:

We can say that *large* collective convergence over a period requires as much convergence as there has been over big questions in the hard sciences in the same period. Here I will take the hard sciences to include at least mathematics and the natural sciences: paradigmatically physics, chemistry, and biology.

(Ibid.: 6)

What is the evidence that there hasn't been large collective convergence on the big questions in philosophy? This is an empirical claim and there isn't much empirical work on this issue. Almost the only effort in this direction is the 2009 PhilPapers Survey (published as Bourget and Chalmers 2014). This was an online survey sent to 2000 professional philosophers in North America, Europe, and Australasia. Respondents were asked about their views on thirty important questions in philosophy. I won't go through the details of the responses here, instead I report Chalmers' summary:

The degree of disagreement here is striking, if unsurprising. Only one view (non-skeptical realism about the external world) attracts over 80% support. Three views (*a priori* knowledge, atheism, scientific realism) attract over 70% support, with significant dissent, and three more views attract over 60% support. On the other 23 questions, the leading view has less than 60% support.

(Ibid.: 9)

CHALLENGES TO THE EMPIRICAL EVIDENCE FOR PERSISTENT DISAGREEMENT

Claims about how practitioners of a discipline like philosophy compare to practitioners of other disciplines raise very complicated empirical questions. This is even more so when such claims invoke difficult notions like 'agreement' and 'disagreement' – and the idea that this can be measured in the relevant comparative way. This section raises some concerns about the data in Bourget and Chalmers (2014). Most of what I have to say here is quite tentative and at best adds up to the rather non-controversial conclusion that we need much more empirical work in order to be confident in making the comparative claims.

First note a couple of very obvious limitations of Bourget and Chalmers (2014) (these points are also mentioned in Chalmers 2015):

- a) The study covers only a single point in time so it doesn't tell us anything about convergence over time.
- b) Convergence, as Chalmers and others think of it, is comparative. The claim is that philosophers converge less than some other people (e.g., those who work in 'the hard sciences'). But Bourget and Chalmers (2014) has no comparative component.

So their study is completely silent on the crucial comparative question: is the situation in philosophy significantly different from that in similar disciplines? What we would need would be a series of diachronic studies of several disciplines and then a comparison. That's missing.

So even if we take the study at face value, it is at best extremely weak evidence for Chalmers' empirical claim. Of course, many philosophers have a hunch that there's a lack of convergence in their discipline, but hunches about empirical matters even (or maybe in particular) about ourselves and our own disciplines are horrifically unreliable (this is one of the points Bourget and Chalmers (2014) correctly emphasize). One point all parties to this debate should agree on is that more empirical work is needed. At this point there's simply no genuine evidence for the claim that philosophy is significantly different from other disciplines with respect to the amount of disagreement.

That's not all: there are at least four additional reasons for concern about how Chalmers (2015) uses the data from Bourget and Chalmers (2014).

(i) **Verbal disputes are not real disagreements (and Bourget and Chalmers (2014) doesn't screen for verbal disputes):** The study assumes that if two respondents gave divergent answers to one of the questions asked, they disagree. But that's a naive assumption. It follows only if they interpret the words in the same way. Take, for example, the question, 'Physicalism or non-physicalism?' Suppose one respondent says 'yes' and another 'no'. Do they disagree? Only if they interpret the word 'physicalism' the same way. If they mean different things by that word, then we have no evidence they disagree.³ Do we have reason to believe that participants in the study interpret the words in the question in a uniform way?

No, we don't. What follows is in danger of looking like an *ad hominem* argument, but it's not: I happen to wholeheartedly agree with Chalmers (2011). In that paper, Chalmers argues that a) verbal disputes are pointless, and b) almost all the questions that are mentioned in Bourget and Chalmers (2014) have been involved in verbal disputes. Chalmers says:

In the Socratic tradition the paradigmatic philosophical questions take the form 'What is X?'. These questions are the focus of many philosophical debates today: What is free will? What is knowledge? What is justification?

What is justice? What is law? What is confirmation? What is causation? What is color? What is a concept? What is meaning? What is action? What is life? What is logic? What is self-deception? What is group selection? What is science? What is art? What is consciousness? And indeed: What is a verbal dispute?

(Chalmers 2011: 531–2)

He continues in a footnote: ‘I think that the philosophical literature over almost all of the questions in the last paragraph is beset by verbal disputes, in a fashion that is occasionally but too rarely recognized’ (ibid.: 532).

Suppose Chalmers is right in this. If so, an alternative interpretation of the Bourget and Chalmers (2014) results is that they provide additional evidence of widespread verbal disputes in philosophy. The differences in replies do not reflect differences in substantive views, but instead differences in how crucial terms are used.

I should mention that Chalmers briefly mentions verbal disputes in his 2015 paper. He asks: *Why is there so much disagreement in philosophy?* and one of the replies he considers is this:

there is little convergence because participants are talking past each other. Each side is using key terms in different ways and each is correct where their own use of the term is concerned. In ‘Verbal Disputes’ I argued that verbal disputes are common in philosophy. For example, I think many debates in the philosophy of free will and the philosophy of language have a significant verbal element. And I think that resolving verbal disputes can lead to philosophical progress. Still, often when we clarify the key terms in a partly verbal dispute, we find that a substantive dispute remains. And there is a core of fundamental questions (including many normative questions, as well the mind–body problem and other issues involving ‘bedrock’ philosophical concepts, in the terms of ‘Verbal Disputes’) for which the diagnosis of a verbal dispute seems quite implausible.

(Chalmers 2015: 26)

Two thoughts about this:

- a) Verbal disputes don’t explain disagreement. They undermine the idea that there is disagreement in the first place. Verbal disputes are not genuine disagreements, they are pointless verbal confusions.

- b) The appeal to bedrock is unhelpful for two reasons: first, we now have a new empirical assumption: *that there are substantive bedrock disagreements*. We have no evidence of that. The questionnaire in Bourget and Chalmers (2014) certainly doesn't support that, since it is not conducted at the bedrock level. Second, one can question whether there is a bedrock level and, if there is, whether it is purged of verbal disputes. It would go beyond the scope of this paper to explore those issues further, but they are important issues to get clear on for those who want to understand the connection between disagreement in philosophy and verbal disputes in the way Chalmers hints at.

(ii) **Why compare to 'hard sciences' and not instead to the social and cognitive sciences?** Van Inwagen, Chalmers, and others take the relevant comparison class to be 'the hard sciences' and go on to claim that philosophy compares unfavourably to those disciplines. But why is that the relevant comparison? Why not instead compare to the cognitive and social sciences? The remit of philosophy is, taken as a whole, much more closely related to those fields than to theoretical physics and mathematics. Here is a **bold empirical conjecture**: if we were to pick a set of 'core' (or 'Big') questions within economics, political science, sociology, or psychology, we would find as much (if not more) disagreement as in philosophy.

With that comparison in mind, two connected points are worth noting: a) practitioners of these disciplines – for example, economics and psychology – have the highest *social* standing of any disciplines. So-called experts in these fields are given enormous decision-making powers. Economists run the infrastructures of our societies and psychologists are given immense power over individual lives. In this respect they have much higher standing than theoretical physicists and mathematicians. If persistent disagreement is a negative aspect of a discipline, then, if my conjecture is correct, we are no worse than practitioners of those highly respected disciplines. b) It is striking that self-doubt isn't nearly as prevalent in for example, economics and psychology.³ That's one respect (maybe the only) in which we should learn from those disciplines.

In the light of a) and b), here is a psychological/sociological conjecture about why there's so much more self-doubt in philosophy than in economics and psychology: it's not because there's more

progress or more convergence in those disciplines. It's because they have high social standing (and are given lots of responsibility) and philosophers don't. It's this lack of social standing (and role) that triggers rather than any objective features of philosophy.

(iii) **Convergence leads to special sciences:** In the history of philosophy, the following has happened many times: philosophers work on a set of problems and then at some point they converge on precise questions and develop some shared methodological standards. Then, *under certain hard to understand conditions*, that set of issues becomes a new discipline. At some point in history, central issues in what we today describe as economics, psychology, sociology, theology, linguistics, and biology were considered philosophical questions. They then developed into what we today consider separate disciplines. This 'spawning function' of philosophy is undeniable, but very difficult to understand in detail. However, for current purposes the details don't matter much. What matters is that, speaking roughly, *there is some correlation between increased convergence and the spawning of a new discipline*. So when philosophers start to converge on a cluster of issues, then, under hard to predict and explain conditions, that cluster tends to become identified as a separate discipline. It seems a bit arbitrary, then, to not treat that as an instance of philosophical convergence. Insofar as these other disciplines count as converging disciplines, philosophy should share in some of that 'glory' (if you think that is what it is.)

(iv) **There's more consensus on 'small questions', negative claims, and conditional claims (and these are just as important as answers to 'Big Questions'):** Chalmers and van Inwagen focus on the 'Big Questions' and treat the adjective 'Big' in a somewhat normative way. It's as if being 'Big' (with a capital 'B') makes a question more valuable or interesting or central to philosophy. So if we measure disagreement in philosophy, we should measure it with respect to the Big Questions. Here are some reasons to resist that attitude towards the Big ones: first, there's more convergence on the 'smaller' questions in philosophy. There is also more convergence on negative claims (e.g., 'the meaning of a word isn't an associated mental image' or 'disagreement data fails to establish that truth is relative') and conditional claims (e.g., 'if evidence for distinctive de se attitudes

is just evidence of opacity, then there are no good arguments for an essential indexicality thesis'). Chalmers recognizes this. He says:

This is not to deny that we have attained a great deal of collective knowledge in philosophy. As Timothy Williamson (2006) has said, we knew much more in 2004 than in 1964, much more in 1964 than in 1924, and so on. But this collective knowledge typically does not involve answers to the big questions. It is mainly knowledge of the answers to smaller questions, of negative and conditional theses, of frameworks available to answer questions, of connections between ideas, of the way that arguments bear for and against conclusions, and so on. In the absence of convergence on the big questions, collective knowledge of the answers to those questions eludes us.

(Chalmers 2015: 15–16)

For the sake of argument, let's assume that the last sentence is right – without convergence on the Big Questions we don't get *collective* knowledge⁴ of the answers to the Big Questions. That, however, doesn't address the question: why focus on the Big Questions? Why does only the convergence on those questions matter? Again, comparison with economics might be useful. There's no more consensus on 'Big Questions' in economics, but to the extent that there's a sense of progress it's connected to the little questions – they get lots of little questions worked out. It's important for us philosophers to emphasize that we too get lots of little questions worked out. We know much more now about how semantics for relativist treatment of modals would have to go, or how scoping arguments can try to respond to the modal argument, or what Lewis' triviality argument tells us about the interaction between probabilities and conditionals. We just are less likely to count that as 'progress', because we're taking progress to mean solving the Big Questions. That bias is a mistake – the 'smaller' questions are just as important (and maybe what we need is to give up the entire category of 'Big Questions').

In sum: so far no one has made the case that there's persistent disagreement in philosophy that distinguishes the field from other relevant disciplines. That said, I don't take these doubts to undermine the plausible idea that there is a set of philosophical questions that there's persistent and genuine disagreement about. The next section provides an account of why there will always be a set of questions like this.

FIRST SOURCE OF PERSISTENT DISAGREEMENT IN
PHILOSOPHY: FAILURE OF EVIDENCE NEUTRALITY

There are some philosophical questions there's persistent disagreement about. One important reason for this is the failure of what Williamson calls *Evidence Neutrality*. Here is Williamson's helpful introduction to the idea:

As far as possible, we want evidence to play the role of a neutral arbiter between rival theories. Although the complete elimination of accidental mistakes and confusions is virtually impossible, we might hope that whether a proposition constitutes evidence is in principle uncontentiously decidable, in the sense that a community of inquirers can always in principle achieve common knowledge as to whether any given proposition constitutes evidence for the inquiry. *Call that idea Evidence Neutrality*. Thus in a debate over a hypothesis *h*, proponents and opponents of *h* should be able to agree whether some claim *p* constitutes evidence without first having to settle their differences over *h* itself.

(Williamson 2007: 210, my emphasis)

If Evidence Neutrality were true, then we should expect those debating philosophical issues to be able to converge on answers, given time and good will. If, however, Evidence Neutrality is false, then we should not expect to reach agreement – we should expect those who disagree about central philosophical questions to be deadlocked.

Next step: *Evidence Neutrality is false*. Here, again, is Williamson:

Having good evidence for a belief does not require being able to persuade all comers, however strange their views, that you have such good evidence. No human beliefs pass that test. Even in principle, we cannot always decide which propositions constitute evidence prior to deciding the main philosophical issue; sometimes the latter is properly implicated in the former. Of course, we can often decide whether a proposition constitutes evidence prior to deciding the main issue, otherwise the notion of evidence would be useless. But the two sorts of question cannot be kept in strict isolation from each other.

(Ibid.: 203)

This paper will assume, with Williamson, that Evidence Neutrality is false. It would take us too far afield to engage in a full-blown

argument for denying Evidence Neutrality, but briefly: first-order and second-order philosophy are intertwined. So in many cases, to take a stand on a core philosophical issue is also to take a stand on what the correct way to do philosophy is and also on what counts as evidence in philosophy.

The assumption that Evidence Neutrality is false provides a partial explanation of persistent disagreements in philosophy. Often the following will be the case, call this a Paradigmatic Irresolvable Philosophical Dispute:

Paradigmatic Irresolvable Philosophical Dispute: A believes that p and B believes that not-p. If p is true, then q counts as strong evidence for p. If p is not true, then q doesn't count as evidence for p (i.e., the question of what counts as evidence for p is in part settled by the truth or falsity of p). A and B might agree that q, but that won't help resolve their disagreement. B will discount q as evidence (since she doesn't endorse p) and A will cite q as evidence (since she endorses p).

The beliefs of A and B could both be the best supported by the evidence given what their theories treat as evidence. 'The role of evidence as a neutral arbiter is undermined' (ibid.: 213).

The details of how an irresolvable philosophical dispute is grounded in different views of what counts as evidence will always be complex. A full presentation of a single case would require a level of detail that takes us beyond the scope of this contribution. However, it's not too hard to get a rough sense of what I have in mind. Here are five simple illustrations of junctures where first- and second-order philosophy intersect in the relevant ways:

- Different views of the nature of knowledge will affect what you think counts as evidence (and even more obviously: different views of what evidence is will affect what you think counts as evidence).
- Whether you think there is an external world or not will affect how you think about evidence (and how to collect it).
- Whether you are a dualist, panpsychist, or a reductive materialist will affect how you think about evidence.
- What you think concepts are will have consequences for how you think about philosophical methodology – in particular it

will have consequences for whether you think philosophers can and should be engaged in *a priori* conceptual analysis.

- Different views about the semantics of 'intuition' and the metaphysics of intuitions can have direct consequence for how you think of philosophical evidence and methodology – in particular it can have consequences for whether you think intuitions can and should serve as evidence in philosophy.

Often the connections between first-order and second-order issues won't be quite as obvious as in these cases, but I suspect that in many of the cases Chalmers lists as 'Big Questions', such connections can be unearthed with a bit of work.

A final point to note before leaving the topic of Evidence Neutrality and Irresolvable Disagreements. Not only does the falsity of Evidence Neutrality explain persistent disagreements, but it also tends to exacerbate those disagreements: 'both the good fortune of being right and the misfortune of being wrong are magnified' (ibid.: 213). Those who are right are also right about what counts as evidence and so are in a good position to improve their theory: they know where to look for additional evidence. However, those who are wrong are also wrong about what constitutes evidence and so are not in a good position: their further theorizing will be contaminated by their false beliefs about what counts as evidence.

SECOND SOURCE OF PERSISTENT DISAGREEMENTS: INSTITUTIONAL IMPARTIALITY

The second explanatory component for disagreement is institutional/sociological. It's simply this: as a matter of fact, philosophy departments don't tend to hire just people who agree with each other. We hire people who fundamentally disagree. We also educate and supervise students who fundamentally disagree with us. There's no doctrinal or methodological entry-ticket to becoming a professional philosopher. Throughout our 2000-year history, we see a spectacular array of different views represented across the world among those who would self-describe as professional philosophers.

It could have been different. We can imagine an alternative development in which one set of answers and one methodology became institutionally dominant. That would mean we had a

form of institutional partiality, where hirings required allegiance to a specific conception of philosophical methodology and also agreement on substantial issues.

In the light of the failure of Evidence Neutrality, institutional impartiality is a very good thing for at least two reasons: (i) non-experts are not in a position to choose between competing philosophical views. The people building up academic institutions are, in large part, non-experts. They are also not in a position to adjudicate between competing experts. So it's good that they stay impartial and don't try to force philosophy into one direction. (ii) Pluralism is good even for those who are right. When they have around them people who disagree with them, they are constantly challenged to sharpen their arguments. If, by some miracle, those who are right got to be in a position to make all philosophy hires and just hired other people who were right, that wouldn't be good for them: they need people who challenge them fundamentally.

CONVERGING DISCIPLINES, INSTITUTIONALLY
CONSTRUCTED CONSENSUS, AND THE VALUE OF
PERSISTENT DISAGREEMENTS

In the light of the points made in the previous section, the concern should be with the converging disciplines rather than the non-converging disciplines. Structurally, what happens in converging disciplines is often something like the following: they start with fundamental assumptions, both methodological and substantive. Those assumptions then provide a framework for a research programme and the basis for assessing both contributions and participants. What is not questioned or investigated are those fundamental assumptions. Typically, practitioners are not in a good position to justify their fundamental assumptions: when challenged about these, they can't provide reasons that aren't just internal to the framework. When challenged about the framework as a whole, they would have no clue what to say.

It helps here to distinguish two rough models of how to generate convergence: *institutionally constructed convergence* and *genuine intellectual convergence*. Let's take contemporary physics as an illustration. In what sense is there consensus about their views? First notice that there isn't broad consensus across the population

of humans. Focusing just on the USA, a recent Gallup poll shows that 'More than four in 10 Americans continue to believe that God created humans in their present form 10,000 years ago'.⁵ This contradicts what contemporary physics tells us about the origin of the universe. So if 'communal agreement' requires agreement across society as a whole, then there's no consensus on core views in contemporary physics. Those who take physics as a paradigm of a converging discipline will reply: *what we mean is that people who are hired as physicists in universities agree that the universe is more than 10,000 years old*. That's true, but the restriction is problematic. Here is why: institutional consensus is very easily generated in intellectually irrelevant ways. You just make sure that people who are hired in the institution agree with each other. So, for example, the people who are hired by the Church of Scientology all agree that humans have an immortal spirit who used to live on other planets. That there's consensus among those working in the Church of Scientology about this is of course entirely unsurprising and gives no kind of epistemic boost to those views. It's an institutionally constructed consensus. That, in itself, has no intellectual value.

What we are looking for is some form of intellectual consensus – consensus generated just by the force of the arguments. However, when we focus on institutions such as universities and their physics departments, it is very hard to make the empirical case that it is the force of arguments that has generated consensus. Disentangling the institutional pressures from the intellectual force is practically impossible.

What's the upshot of this? I take the upshot to be that what we should focus on are the arguments and the evidence for a theory. There are great arguments and impressive evidence for general relativity. That's what's important. Whether a majority of people hired by a particular kind of institution happens to endorse the view doesn't matter. Here is a thought experiment to make that clear: suppose horrific political developments result in religious fanatics taking over physics departments. As a result, the consensus among those employed as physicists ends up being that the world was constructed 10,000 years ago. That's unfortunate, of course, for many reasons, but it doesn't undermine any of the evidence for physics as we know it. Its intellectual value remains, no matter how many of those hired in physics departments believe the world

was constructed 10,000 years ago. What matters isn't consensus or convergence. That's irrelevant to the intellectual value of the view.

Moreover, given the failure of Evidence Neutrality in philosophy, some persistent disagreement is unavoidable. To aim for convergence is to aim for something impossible. To have an impossible goal is irrational. Moreover, there are good institutional reasons for having this reflected in the people hired as professional philosophers. In other words, we have good reason to not let academic institutions try to force convergence – that would be a dreadful system and an enormous intellectual risk. In sum, it seems rational to aim for an institutional structure that preserves divergence (even when we know that means employing some people who are wrong – i.e., paying people to work out views that eventually turn out to be blind alleys).

I now go on to consider two objections to the views just expressed. According to the first objection, peer-disagreement undermines knowledge. According to the second, disagreement undermines 'collective knowledge' and collective knowledge is valuable.

DOES LACK OF CONVERGENCE UNDERMINE
KNOWLEDGE?

I have just argued that convergence is something we shouldn't care about. Chalmers disagrees – he thinks it's very important. But why? Chalmers considers two replies and endorses the second. The first reply, that he rejects, is this:

One obvious answer is that we value knowledge, agreement is required for knowledge, and convergence goes along with increases in knowledge. A strong version of this view, suggested by van Inwagen's discussion, is that where there is sufficient disagreement among experts, no individuals can be said to know the truth. Even if some individuals have hit on good arguments for true conclusions, how can they have justified confidence that these are good arguments, when so many of their peers disagree?

(Chalmers 2015: 14)

Chalmers doesn't endorse this view. He says, even though lots of his colleagues deny the existence of consciousness, he still knows that he is conscious. So widespread disagreement about a view doesn't undermine knowledge. I agree. Given the very extensive literature

on peer disagreement, the ideal way to address this would be to go through each and every account of peer disagreement currently on the market to see what it has to say on this issue. That would make this paper far too long and would also go beyond my area of expertise. Instead, I here simply want to flag the issue and to refer readers to other literature on this topic. The best and most thorough paper on this topic that I know of is ‘Disagreement in Philosophy: Its Epistemic Significance’ by Thomas Kelly. Kelly, in my view convincingly, argues that ‘there is *no* plausible view about the epistemology of disagreement, on which philosophical agnosticism is compelling’ (Kelly 2016: 375).⁶ Those who want to explore this question further should refer to Kelly’s paper and the literature he discusses.

COLLECTIVE VS. INDIVIDUAL KNOWLEDGE

In response to the question, ‘Why is *convergence* to the truth important, and why should we be concerned about its absence?’ Chalmers says:

even if agreement is not required for individual knowledge, some degree of agreement is plausibly required for *collective* knowledge. If the community of experts on a question has serious disagreement over the answer to that question, then that community cannot be said to collectively know the answer to that question, and nor can the broader community of which they are a part. (Chalmers 2015: 15)

He adds:

Furthermore, we value collective knowledge. One reason that progress of the hard sciences has been so impressive is that it has plausibly enabled us – the community of inquirers – to collectively know the answers to those questions. But in the absence of sufficient agreement on the answers to philosophical questions, we cannot be said to have collective knowledge of those answers.

(Ibid.)

I think there are two questions worth disentangling here: (i) Can we describe a situation in which there is collective knowledge without large-scale collective convergence? (ii) If the answer to (i) is ‘no’, why should we care?

With respect to the first question, I think the issue is murky. Here is an illustration of the issue as I see it. Consider theories of truth. There are a number of alternative theories on the table (correspondence theories, coherence theories, pragmatic theories, deflationary theories, etc.), all of them very well worked out, but no consensus on the truth about truth. Suppose, for simplicity, that we have eight candidate theories, and suppose one of them is correct, say the deflationary theory. Is there any sense then in which we philosophers know the answer to the question: *What is truth?* I'll assume that if the answer is 'yes', then in some sense we have collective knowledge. If the answer is no, then we don't have collective knowledge. Here are some analogies to help you think about that kind of situation:

Analogy 1: Suppose I'm looking for a golden coin together with 10,000 other people. I can say that *we* have found the coin even if just one of us found it and many are still looking for the coin in the wrong places.

Analogy 2: As in 1, I've found the golden coin, but seven other people found fake-gold coins, and they think theirs is the golden coin and I can't argue them out of their false belief. I can still say that *we* have found the golden coin (though unfortunately, some people don't recognize it).

Analogy 3: As in 2, I have found the golden coin and seven others have found fake coins. Now emphasize that there's no consensus among the 10,000 collaborators about who has found the genuine coin (they are evenly split between the eight of us). We can still, I think, say that *we* have found the golden coin – there's just the complication that we have no consensus about how to pick it out from some fake coins.

The relevance of these analogies is that we have a form of collective achievement by virtue of an individual achievement even when other members of the group fail to recognize that achievement (or even dispute it). In cases involving attitude verbs, we have something similar:

Analogy 4: We can say that Apple knows how to improve Siri when some of the employees have figured it out, even if there's

extensive disagreement within the company about how to do it. It might even be that the disagreement is irresolvable (in that one group can't convince another).

My own judgement about the Apple case is clear: Apple knows how to improve Siri (and it is also true *that Apple knows that Siri can be improved by doing D* (where D is the correct improvement procedure)). However, some people's judgement about this case vacillates somewhat depending on how the details are worked out (e.g., who makes the final decision, what are the practical implications, etc.).⁷ Much will depend on what the point of the collective attribution is. In the Apple case, maybe what matters is the production end: can they as a matter of fact change the software in the relevant ways or do they at least have the capacity to do that? There's no clear analogy to that in the case of philosophy. That said, my view is that widespread disagreement is in principle no obstacle to collective knowledge. However, the details of the conditions under which we can make such attributions are no doubt complex and a full discussion would go beyond the scope of this paper (for some recent discussion of related issues see Bird 2010, Lackey 2014).

I turn now to the second question: why should we care about collective knowledge attributions? Some people think individual knowledge attributions are important because they endorse some version of the knowledge norms of assertion, belief, or action.⁸ Many of us reject all such views. For example, Cappelen (2011) and Pagan (2015) argue that there are no such norms at all (so in particular, the knowledge norm is wrong). Even if you're not on board with the anti-norm view, an account is needed of the value and significance of making collective knowledge attributions. Whatever you think about the significance of the individual knowledge attributions doesn't transfer to the collective case without argument (i.e., work is needed to show that what we say about the individual cases applies to the collective case).

My view is this: what matters in this context (i.e., understanding the development and nature of disciplines) is not to get clear on whether it's okay for members of a group to say 'We know that ...'. What matters is to understand the details of the interaction between the participants in a non-converging discipline. To focus on a general question such as 'Can we say that they have collective knowledge?',

isn't helpful for understanding anything important. What we should focus on instead are questions such as these:

- What's the detailed source of the disagreement?
- How much agreement is there on conditional claims, negative claims, methodology, space of possible answers, and quality of arguments?
- How do the groups who disagree cluster and interact?

The answers will be sensitive to details and be messy. In a discipline like philosophy, the answer will vary between sub-disciplines, between academic communities, and over time. So trying to answer the very general question 'What is the value of collective knowledge in philosophy' isn't particularly helpful. What's helpful is to try to answer detailed questions such as those just listed for specific sub-disciplines, and specific academic communities at particular times.

Again, it's important to keep in mind that there's not much that's distinctive about philosophy here. Compare again to a discipline like economics. There's no more consensus about Big Questions in economics than in philosophy, but we don't find nearly the same level of handwringing and agonizing about that fact. Why not? Well, they have sub-groups, often centred around specific institutions (e.g., around so-called Freshwater and Saltwater schools). Within each grouping there is a higher degree of convergence (than in the discipline as a whole) and complex theories are developed based on the consensus within each group. So to understand contemporary economics, we need to look at the work within those clusters, the relationship between the clusters, and the sources of the disagreement between them. Just to try to answer the question 'Can economists say that they know that ...' (where '...' is some thesis that there's disagreement over) isn't very illuminating.

In sum: Chalmers asks 'Why is *convergence* to the truth important, and why should we be concerned about its absence?' The answer is twofold: a) that's the wrong question, and b) if you insist on focusing on the very general question, then the important observation is that in some important sense (or in some contexts) it's true to say: 'We have collective knowledge of the answers to all the Big Questions'.⁹

NOTES

- 1 Much of this paper can also be seen as a reply to (Chalmers 2015).
- 2 We might have some evidence that they disagree about what the words mean or should mean, but that's not what Bourget and Chalmers were testing for.
- 3 This is not to deny that there's often internal criticism and debate within economics and psychology. What they have significantly less of is a group of people (internal and external to the discipline) who practically make careers out of criticizing the discipline as a whole.
- 4 At this point I'm bracketing problems with the notion of collective knowledge – more on that topic below.
- 5 www.gallup.com/poll/170822/believe-creationist-view-human-origins.aspx
- 6 Chalmers agrees: 'I think that at least in some cases, a good argument can ground an individual's knowledge of a conclusion even when peers reject it. For example, I think that the presence of any number of peers who deny the existence of consciousness would not undermine my knowledge that I am conscious' (Chalmers 2015: 14–15).
- 7 For empirical evidence that people judge that a group can know how to do something without actual agreement about how to do it, see (Jenkins *et al.* 2014).
- 8 For example, those who endorse the knowledge norm of assertion think it is constitutive of assertion that one should assert *p* only if one knows that *p*. For those who endorse this view, the question of whether *A* knows that *p* is important when assessing *A*'s assertion that *p*. For more discussion of and criticisms of this view see (Cappelen 2011) and (Pagin 2015).
- 9 In January 2016, Nancy Bauer, Paul Horwich, L. A. Paul, Patrick Greenough, Mark Richard, and Bjørn Ramberg participated in a workshop on Progress in Philosophy at the University of Oslo. I learned a lot about these issues from discussions during that workshop. I also got useful feedback from Olav Gjelsvik, Joshua Habgood-Coote, Torfinn Huvenes, Øystein Linnebo, Knut Olav Skarsaune, and Rachel Sterken.