

# On the Limits of Philosophical Progress

David J. Chalmers

Let's be honest. There is considerable progress in philosophy. But there is less progress than we might like.<sup>1</sup>

One can see the glass as half-full or as half-empty. Many other speakers at this conference have taken the half-full approach, focusing on the many sorts of philosophical progress that there are. I agree with much of what they have said. But here I will take the half-empty approach for a moment, focusing on the sort of progress that there is not.

Here's a central thesis: There has not been large collective convergence to the truth on the big questions of philosophy.

Here the big questions of philosophy are things like: What is the relationship between mind and body? How do we know about the external world? What are the fundamental principle of morality? Large collective convergence on these questions requires a high average degree of agreement among members of the philosophical community on the answers to the questions: to a first approximation, as much or nearly as much agreement as there is in scientific communities on big questions in the sciences. There has been a good deal of collective convergence on the big questions in many areas of science: physics, chemistry, biology, and so on. There has arguably been much less collective convergence in philosophy. Convergence to the truth requires that this convergence lead to true beliefs regarding the answers to these questions.

Here is an argument for the central thesis. It has two premises, an empirical premise and a logical premise.

---

<sup>1</sup>This is a slightly revised version of a 15-minute talk presented in a symposium at the Harvard-Australia conference on Philosophical Progress in September 2011. The paper is something of an exercise in saying the obvious. But on this topic it is worth saying the obvious first so that less obvious things can be said from there. I have mostly abstracted away from my own slightly less obvious views about these matters, although there are some hints about them toward the end.

- (1) Empirical premise: There has not been large collective convergence on the big questions of philosophy.
- (2) Logical premise: If there has not been large collective convergence on the big questions of philosophy, there has not been large collective convergence to the truth on the big questions of philosophy.
- (3) Conclusion: There has not been large collective convergence to the truth on the big questions of philosophy.

I will take the logical premise to be uncontroversial and focus on the empirical premise. I take the empirical premise to be plausible to those with passing familiarity with philosophical practice. We do not have all the empirical data required for a systematic investigation of the premise, but we have some of it.

For example, the 2009 PhilPapers Survey surveyed around 1000 professional philosophers on answers to thirty important questions in philosophy, and typically found that answers to major questions were distributed something like 50-50 or 60-40 or 70-30, once agnostics and other intermediate options were removed. This suggests that at least where these questions are concerned, large collective convergence has not been achieved.

Now, you might say that these are the big questions of the moment and therefore are precisely those that are unanswered, so the result is no surprise. There is correspondingly little agreement on the *current* big questions of physics: the status of string theory, for example. To avoid this worry, it is important that the big questions be individuated not by current debate but by past importance.

To properly address this issue, we would need analogs of the PhilPapers Survey in (for example) 1611, 1711, 1811, 1911, and 2011, asking members of the community of philosophers at each point first, what they take to be the big questions of philosophy, and second, what they take to be the answers to those questions (and also the answers to any big questions from past surveys). We would also need to have analogous longitudinal surveys in other fields: the MathPapers Survey, the PhysPapers survey, the ChemPapers Survey, the BioPapers Survey, and so on. And we would need a reasonable measure of convergence to agreement over time. I predict that if we had such surveys and measures, we would find much less convergence on answers to the big questions suggested by past surveys of philosophers than we would find for corresponding answers in other fields.

Even if we do not have large collective convergence to truth on the big questions, we can have various weaker forms of progress, each corresponding to dropping one of the central requirements.

Drop “large”: It is certainly possible that there has been major convergence on answers to a

small number of the big questions of philosophy (questions about logic or about god, for example?), and/or minor convergence on answers to other questions. Maybe if God looked at survey results knowing the answers to all the questions, she would see that on average 15 answers in 2011 than in 1611. But I suspect that the corresponding figure would be much higher in other areas.

Drop “collective”: The data are consistent with the claims that various individuals or sub-communities have themselves had large convergence to true answers on the big questions. For example, maybe over time I have converged on the truth on most of these questions. But if so, sadly, it has not led to collective convergence to the truth. Likewise, perhaps groups such as the logical empiricists or the Oxford realists have converged on the truth. But again, community-wide convergence has not ensued.

Drop “on the big questions”: There may certainly have been large convergence on smaller questions. There has been large convergence on the thesis that knowledge is not justified true belief, or that conditional probabilities are not probabilities of conditionals. We are especially good at ruling out certain specific views and at making a case for conditional connections between views. But I take it that these are not really answers to the big questions of philosophy. Philosophers like to note that various subfields have split off from philosophy once philosophers answer certain key questions well enough: logic, psychology, formal semantics. Maybe in a few cases (perhaps logic is the best candidate?) this involves answers to a fairly big questions. In the other cases, though, it is unclear whether the fields that have split off have really answered the original big questions in philosophy (e.g. “What is consciousness?” or “What is meaning?”) that predated their formation.

Drop “convergence to the truth”: There are many forms of philosophical progress that do not involve convergence to the truth. For example, there is increased understanding (Bengson), or explorations of possibilities in philosophical space (Feldman, Hellie/Wilson), or the development of beauty and beautiful views (Pasnau), or philosophical pedagogy (Bauer), or making a difference to the world (Stanley). Still, I take it that reaching true belief, and indeed knowledge, is a primary aim of philosophy, probably *the* primary aim of philosophy. After all, most philosophy, or at least most analytic philosophy, seems to consist in putting forward theses as true, and arguing for their truth. I know that at least in my own case, the primary motivation is to figure out the truth about the relevant subject areas: What is the relation between mind and body? What is the nature of reality and how can we know about it? And so on.

Let me reiterate the glass-half-full perspective for a moment: there has been enormous progress in philosophy. In the last century alone, in the philosophy of mind and language alone (to pick

two areas I know well), the field has been transformed, leading to tremendous understanding and sophistication that did not exist before. But this transformation has not been accompanied by a large degree of convergence on the big questions: say, the mind–body problem and the question of what is meaning. Instead, progress has taken one of the four forms above. There has been large collective convergence on some smaller questions and large collective increases in understanding on all the questions, and perhaps reasonably large convergence on the big questions among various subcommunities. But across the philosophical community as a whole (and even across the analytic philosophy community as a whole), convergence on the big questions has at best been small.

We have become enormously sophisticated in addressing philosophical questions. In the last hundred years, we have seen the development of many sophisticated philosophical methods and programs: linguistic philosophy, empirical or naturalistic philosophy, formal philosophy, feminist philosophy, crosscultural philosophy, and experimental philosophy, to name a few. Each of these has led to significant new insights on the big questions of philosophy. But none of them has led to much in the way of convergence. Instead, they have led to more and more sophisticated versions of the same old disagreements, along with some new and interesting disagreements alongside them.

To me, this is the largest disappointment in the practice of philosophy. Perhaps once one has been doing philosophy for a while, it is no surprise that there is not large collective convergence to the truth. But antecedently to doing philosophy, one might have hoped that something more was possible.

In particular, one might have hoped that philosophical argument would have the power to produce widespread philosophical agreement. But in practice, while we can recognize powerful philosophical arguments when we see them, these arguments produce less agreement than one might have expected. Typically, the arguments simply force opponents to elaborate their positions, ruling out some versions of a view and lead to more development of other versions. Some areas of philosophical space get ruled out this way, but usually a lot is left open on both sides of a major issue. Even through continued fractal elimination of regions of space, the area that is left has something of the character of the Mandelbrot set, with large areas of infinite possibility all over the map. This leads to a rich and fascinating philosophical culture. But one could reasonably hope for more.

A final question is: if the thesis is true, why is it true? That is, why is there as little convergence as there is on the big questions of philosophy? A number of diagnoses are possible. It is quite possible that different diagnoses apply to different philosophical domains.

(1) *Problematic method*. There is no convergence because philosophy does not have a method

that makes such convergence possible. Mathematics has proof, the sciences have experiments, history has archival scholarship. In philosophy we mainly have argument. Arguments are a much less powerful method for producing convergence than the methods above. As suggested above, the primary reason for this seems to be that anyone who antecedently denies the conclusion of a philosophical argument can usually deny one of its premises without enormous costs. It is rare to find an argument for an important philosophical conclusion whose premises are all neutral ground. Even when some premises are relatively uncontroversial, as with premises from the sciences or formal premises, one almost always needs further contentious premises to reach a philosophically interesting conclusion. These premises are typically not hard for an opponent to reject. And even for agnostics, doubts about a conclusion tend to easily infect one's confidence in a premise.

I think that this diagnosis helps to explain why convergence in philosophy is more difficult than convergence in other fields. But one would like a deeper explanation of just why philosophical argumentation on the big questions has the property discussed above. After all, argument does sometimes produce convergence on smaller questions. And one would also like to know just where this leaves the possibility of philosophical knowledge. Here various other diagnoses of the lack of convergence become relevant.

(2) *Anti-realism*. There is no convergence to the truth because there are no objective truths to be had in the relevant domains. There is simply an unruly body of opinion out of which we should not expect convergence to emergence. I have sympathy with anti-realism in some domains (ethics and ontology, for example), but this still leads plenty to be realist about. And even accepting moral anti-realism leaves open the question of why there is so little convergence to the truth over the question of moral anti-realism itself.

(3) *Verbal disputes*. There is no convergence to the truth because participants are talking past each other. Each side is using key terms in different ways and is correct where their own use of the term is concerned. I think this is quite common in philosophy—many debates in the philosophy of free will and the philosophy of language have a significant verbal element, for example. And I think that resolving verbal disputes can lead to philosophical progress. But this still leaves a core of fundamental questions to which the diagnosis of a verbal dispute is implausible.

(4) *Sociological factors*. There is no convergence to the truth because although the problem has been solved and some know the answers, sociological factors have kept others from agreeing. For example, perhaps others are professionally rewarded for disagreement, or are professionally or emotionally invested in alternative answers to the question. And in philosophy as opposed to science, the absence of proof and experimental falsification makes it much easier to keep a

defeated research program alive.

(5) *Unknowability*. The problems are unsolvable in principle. While there are true answers to the questions, conclusive grounds for these answers are unavailable to rational thought, so the answers are not knowable. Even if we were superintelligent and superrational (as, perhaps, after the singularity), philosophers would disagree over the answers to these questions.

(6) *Human inadequacy*. There is no convergence to the truth because while the problems are solvable in principle, they are too hard to be solved by us (Colin McGinn's view). Perhaps humans are not wired to do philosophy, or at least to answer the relevant philosophical questions conclusively. More intelligent beings might do better: perhaps after the singularity philosophers will have converged on the answers to all of our big questions.

(7) *Current nonideality*. There is no convergence to the truth because we are not yet reasoning as well as we can, or because we have not yet had insights that are within human grasp. The problems are solvable in principle, and are solvable by humans, but we have not yet gotten there. Perhaps new methods, more disciplined reasoning, or new empirical discoveries are required. There is a curve of increasing philosophical sophistication such that past a certain point on the curve, major progress is possible. We are not there yet, but we are working our way toward it.

I think that most of these answers apply to some philosophical questions. I have argued elsewhere (especially in *Constructing the World*) that all truths are a priori entailed by certain fundamental empirical truths. So where there is a fact of the matter, ideal reasoning from the fundamental truths ought to get us there. Sometimes there is no fact of the matter, sometimes we have already gotten there but sociology and verbal disputes get in the way, and sometimes we need more empirical knowledge. But on the remaining hard cases the choice comes down to the last two options: solvable but not humanly solvable, or humanly solvable and merely unsolved.

If we do not know which of these options is correct, I think that we have to assume that the last option describes our position. Philosophy is still young and we are still learning how to do it well. It might turn out that we are on the cusp of an answer, or it might not. All one can do is keep doing philosophy as well as one can and see where it leads. That is my own glass-half-full perspective on philosophical progress.