

OILING THE WHEELS OF THE PHILOSOPHY MACHINE

HELEN BEEBEE

‘Philosophy’, like ‘art’, denotes both an activity and the product of the activity, each of which has value in its own right. I start with the idea that the main role of professional philosophers, qua researchers, is to contribute to the production of the product—philosophical works that are in the public domain—and argue that neither our own doxastic states nor their justificatory status is important when it comes to our ability to fulfil that role. In particular, assertion in philosophical contexts is not governed by a norm that requires us to be in any particular doxastic state with respect to the asserted content.

PAPER FROM THE 2025 JOINT SESSION

I

Introduction. Philosophers: what is the point of them? That’s an annoyingly hard question to answer, especially when it’s posed in a context where the status of philosophy as a public good is what’s at issue. We are paid good money—by universities and research councils, and so, ultimately, by students and taxpayers—to go about our core business of teaching and research; what justifies this use of other people’s money?

That’s not a question I’m going to address in this paper. I’m interested instead in why, when it comes to research, philosophy seems to have a harder time answering the ‘what’s the point of you?’ question than many other disciplines. There are doubtless various reasons for this—famous scientists saying that philosophy is ‘dead’, for example (Hawking and Mlodinow 2010, p. 5)—but I suspect one reason is that we somehow manage to give the impression that we are mostly being paid for pursuing our own intellectual interests for our own edification and enjoyment, and it’s entirely unclear why taxpayers should be funding us to do *that*. The basic aim of this paper is to dispel that impression.

© 2025 THE ARISTOTELIAN SOCIETY.

THIS IS AN OPEN ACCESS ARTICLE DISTRIBUTED UNDER THE TERMS OF THE CREATIVE COMMONS ATTRIBUTION-NonCommercial License (<https://creativecommons.org/licenses/by-nc/4.0/>), WHICH PERMITS NON-COMMERCIAL RE-USE, DISTRIBUTION, AND REPRODUCTION IN ANY MEDIUM, PROVIDED THE ORIGINAL WORK IS PROPERLY CITED. FOR COMMERCIAL RE-USE, PLEASE CONTACT journals.permissions@our.com.

Proceedings of the Aristotelian Society, Vol. XX, Part XX

<https://doi.org/10.1093/arisoc/aoaf014>

Philosophy—in common with art (and science and other fields of human endeavour)—is both an activity and a product. Making art is something that is done by vast numbers of people, a tiny fraction of whom do it for a living or indeed are any good at it: the vast majority of art is sorely lacking in aesthetic value. But we can distinguish between the *instrumental* value of the activity qua means of production—which is low when the result is terrible—and its value qua activity. Its value qua activity is often significant even when the value of the product is low, not just in the sense that the people involved in artistic activity find it enjoyable, but in the deeper sense that it fosters creativity; it provides a means of expressing one's emotions, ideas or predicament; it can be therapeutic; and so on. Call this the *intrinsic* value of the activity.

These distinctions are pertinent to philosophy too. There is the activity of philosophizing, and there is what that activity produces. As with art, philosophizing is done by vast numbers of people, most of whom are rank amateurs, and most of it—if it results in a product at all—gets us a product that's of very low quality: flatly self-contradictory, utterly bizarre, ill-informed, or simply reinventing the wheel. As with art, though, the activity still has *intrinsic* value—not only because people find it enjoyable but because, even when done badly, it is a serious attempt to engage one's intellect in grappling with profound and difficult questions about the world and our place in it. And there is a broader sense in which the activity of philosophizing—or, perhaps better, generally approaching matters with a philosophical mindset, or manifesting philosophical virtues—need not have any identifiable product at all.

The basic point of this paper is that conceiving of philosophy in terms of this distinction between activity and product gives us a way of thinking of our own philosophical activity in a way that avoids the charge that that we are being paid to conduct philosophical research for our own edification and enjoyment. By and large what we're being paid to do, qua researchers, is to make the product; our own edification and enjoyment is just a happy side effect. The distinction also helps us think about what the relationship is between our work qua researchers on the one hand—which is to produce the product—and, on the other, the work we might do that contributes to facilitating and enhancing *other* people's pursuit of the activity of philosophy. Here too, our own edification and enjoyment is just a happy side effect.

By way of a foil, consider Ernest Sosa's recent book, *Epistemic Explanations* (2021). Sosa's aim is to develop and defend an account of what he calls 'a proper epistemology for the liberal arts, including the humanities, and philosophy more specifically' (2021, p. 21). His account is a virtue-epistemological account: it is a normative story about the kinds of competences and environment that are required in order to gain what he calls 'firsthand intuitive insight' (2021, p. 3). Roughly speaking, this is the kind of knowledge you gain by figuring things out for yourself rather than through what he calls 'epistemic deference' (2021, p. 4): while relying on somebody else's testimony may, in the right circumstances, get us philosophical *knowledge*, Sosa claims, it doesn't get us first-hand intuitive insight.

But it matters quite a lot what we take 'a proper epistemology' for philosophy to mean. Setting aside the question whether first-hand intuitive insight is even possible, at least in any great quantity—to which I think the answer is probably 'no' (Beebe 2018)—I have no beef with the idea that it is in principle a philosophical good. Perhaps it's even the Supreme Philosophical Good: the ideal state to be in, qua philosophical thinker. But if it's what we, qua professional philosophers, are *aiming* to achieve, then it does indeed look very much as though we're being paid to pursue philosophy for our own edification. *Our own* first-hand intuitive insight does not, in itself, have any obvious value to anyone apart from us.

What a 'proper epistemology' for philosophy as an academic discipline, undertaken in an institutional environment where one might reasonably expect it to produce *public* epistemic goods, should do is explain what, epistemically, needs to be done in order to produce *those* goods. I'm not going to tell a positive story about that here. But I *am* going to conceive the discipline of philosophy in a way that separates the products of philosophical research from what is, or what would ideally be, in the minds of the producers. And the upshot of that is that whatever a 'proper epistemology for philosophy', qua discipline that delivers public goods, looks like, it need not concern itself with philosophers' *beliefs* or the epistemic standing of those beliefs. And the same goes for other broadly doxastic states such as 'endorsement' and 'acceptance'.

In §II, I draw an analogy between science as conceived by Alexander Bird in his recent book, *Knowing Science* (2022), and philosophy. I characterize the discipline of philosophy as a machine that is apt for producing outputs (the machine's 'product') that have

been subjected to the machine's quality control mechanisms and are in the public domain.

This conception of the 'product' raises the question of what the function of the broadly doxastic states of individual philosophers is when it comes to the effective operation of the machine. In §III, I suggest that, given the obvious sceptical threat posed by pervasive peer disagreement in philosophy, it would be better not to conceive assertion in philosophical contexts as the expression of belief. Instead, when it comes to assertion, steadfastness in the face of pervasive peer disagreement is licensed by the positive role it plays in producing a high-quality product.

In §IV, I argue that the norms of assertion in philosophical contexts can be characterized in terms of 'commitment', conceived as something like a set of dispositions, rather than in terms of belief (or indeed some weaker broadly doxastic state such as 'acceptance' or 'endorsement'). In §V, I consider the broader contributions, beyond the operation of the machine, that we might reasonably want philosophers to make to the public good, the upshot of which is that we do not need to believe any particular philosophical view in order to make those contributions either. In §VI, I briefly sum up.

II

What's Our Product? When it comes to art, the product of the activity, generally speaking, is some artefact out there in the world: a novel, a musical composition, and so on.¹ What is or has been in the producer's head might be relevant to our aesthetic appreciation or evaluation of the product, or indeed to whether the product counts as art at all, but it's not the product. I suggest we should think of the products of philosophy in a similar way: they are artefacts, and they do not reside in anyone's head.

Shifting from art to science: in *The Logic of Scientific Discovery* (1959), Karl Popper argues that beliefs based on induction are never warranted—but that's not bad news, he says, because science in fact relies on the method of conjecture and refutation, and not inductive inference. Propose a theory that's consistent with the available

¹ I am playing fast and loose with 'artefact' here; also, the distinction between activity and product is often murkier than I'm making it sound: obviously a performance is in some sense both activity and product.

evidence, subject it to rigorous testing, and reject or amend it when the evidence contradicts the theory. Theories that have survived rigorous testing so far are highly ‘corroborated’, but corroboration of a *theory* is very much not the same thing as inductive confirmation of the *belief* that the theory is true.

Popper’s view is now not remotely widely endorsed by philosophers, of course. But here I’m interested in his insistence that we should not really care that much about what’s in scientists’ heads, at least when it comes to considering how science works, what its aims are, and how it progresses. A hypothesis is advanced, it’s tested, and at some point it will end up getting modified—or thrown on the scrapheap and replaced with a better one. The mental state of someone who advances a hypothesis is neither here nor there: it’s the hypothesis itself, the predictions it implies, and its status as corroborated or refuted that are important. Of course, it’s scientists themselves who construct and test and reject the hypotheses, so it’s not as though what’s going on in their heads is *causally* redundant. But their *attitude* towards a hypothesis that they’re advancing and testing is just not to the point. It makes no difference as far as the progress of science, or your own contribution to it, is concerned whether you (irrationally, in Popper’s view) *believe* it, or whether you think it’s almost certainly false, or whether you have no idea one way or the other.

Popper’s view here goes hand-in-hand with another view—one that I’m not sure got much traction even at the time—about the nature of scientific knowledge (Popper 1979). For Popper, we should think of such knowledge as ‘objective’, meaning that it is to be found not in the heads of scientists but—as he somewhat mysteriously put it—in the ‘third world’, by which he meant something like: scientific knowledge is the *publicly available content* that scientists produce. It is what’s in the books and journals and libraries (the ‘first’ world being the world of material stuff, and the second being the contents of people’s minds). As Popper put it, we should think of scientific knowledge as ‘knowledge without a knowing subject’.

In his recent book, *Knowing Science* (2022), Alexander Bird conceives of scientific knowledge as collective knowledge, where the ‘knowing subject’ is what he calls the ‘institution of science’. But, in broadly Popperian spirit, he does *not* take it to be a condition on collective knowledge that it reside in anyone’s head. It might be languishing unread in a long-forgotten journal article, say, ready to be

rediscovered at a later date (Bird 2022, ch. 4). I think of science according to Bird's view as a big machine. Its input consists of theoretical proposals, experimental results, and evidence. The machine gets to work on the input via two basic mechanisms: scientific activity (experimentation and inference) and truth filters (peer review and replication of experiments). And its output is 'scientific propositions', which, having been suitably filtered by the machine, get to count as knowledge only if they are true. So the machine is designed to spit out items of collective knowledge, in the form of scientific propositions that are in the public domain. Of course, it's a fallible machine; not everything it spits out is true and therefore a candidate for knowledge. But that's what the machine is *for* and—arguably—it's doing a pretty good job.

Human minds are, of course, essential to the Science Machine: they contribute in a wide variety of ways to its existence and maintenance by inventing and modifying hypotheses, writing papers, running experiments, raising objections at conferences, refereeing journal submissions, and so on. (Though the machine also has other components: equipment for generating experimental results, measuring devices, journals, and so on.) But the output of the machine is not mental states; it is an extra-mental resource for the edification of and deployment by human beings. So—as with art—we have a distinction between the activity of *doing* science (the operation of the machine) and its product (the output), where the activity has instrumental value in so far as it contributes to producing, or improving the quality of, the product.

I suggest that we think of the discipline (or the 'institution') of philosophy as being like the sciences in this regard: as a complex machine with inputs and outputs, which operates via the various processes that characterize philosophical enquiry. The Philosophy Machine is of course different to the Science Machine in various ways: most of us don't go in much for evidence-gathering; we mostly subject our hypotheses to other kinds of critical scrutiny. By and large, then, what mostly goes on inside the machine is the activity of philosophy, as conducted by people—mostly, these days, professional philosophers—with the skills and knowledge that equip them to play a part in producing a high-quality product. The operation of the machine includes a wide range of activities that we're all familiar with: developing and refining theories and arguments, coming up with objections, writing papers, giving talks, editing journals, refereeing book manuscripts and grant applications, and so on.

My claim, then, is that the Philosophy Machine has the same basic structure as the Science Machine: it has as input theoretical proposals and (what we might at least loosely term) evidence; we all do the research part of our jobs; and it spits out books, journal articles, and so on. So in one sense that's the answer to the question, *what's our product?* But of course it's not a very illuminating answer, because what the machine is designed to do is not just to spit out any old philosophy books and journal articles, but ones that meet some pretty rigorous epistemic standards. So what does the application of those standards get us?

At the very least (though of course one might wish for more), one thing we surely want the machine to spit out is: *viable* systematic philosophical theories—theories of a given phenomenon that are reasonably complete, plausible according to standards that are generally tolerated by the community of philosophical enquirers, and consistent with the evidence, such as it is. (That's pretty loose, but it will do for the purposes of this paper.)² The machine is well-designed for doing that—though of course fallible: some of what we now take to be viable theories may turn out in hindsight not to be viable even in broad outline.

Is that all that can reasonably be asked of the Philosophy Machine? For example, is it designed to produce individual or collective *knowledge* of such theories? I'm sceptical about the possibility of achieving much by way of philosophical knowledge (Beebe 2018), but let's set that aside. What I hope is pretty uncontentious is that, for pretty much any phenomenon of philosophical interest, there is *currently* more than one theory on the table that, according to the standards that we collectively tolerate within philosophy, counts as viable—at least as far as we, collectively, are able to tell. Or, to put it another way, very many philosophical disputes over even the approximate truth of a given philosophical theory are not tractable by us, now, even if they are tractable in principle.

So here are the two main claims I'm putting on the table. First, the output of the Philosophy Machine is not mental states; it is (reasonably) *publicly available content*. The product is a public good and, as

² Bird takes the product of the Science Machine to be propositions, which count as knowledge only when they're true. It's unhelpful, I think, to think of the product of the Philosophy Machine in such a piecemeal way, but I don't think anything I'm going to say hangs much on this. Philosophical theories are composed of and imply propositions, so 'propositions that are constituents of or implied by viable philosophical theories' would do just as well.

such, it resides in the public domain, or at least not locked away in the heads of a bunch of philosophers. It is to be found, at least primarily, in books and journal articles. Of course, some of that content also resides in people's heads in the sense that philosophical theories are entertained, believed, grappled with, or whatever. But the *point* of the Philosophy Machine is not to induce mental states—not beliefs or knowledge or understanding or first-hand intuitive insight—in philosophers. Its product is an extra-mental resource, for the edification of and deployment by human beings in general, present and future.

Of course, such edification and deployment will inevitably involve some such mental attitude, and hence the systematic theories we produce must, as far as we can tell, constitute appropriate *resources* for getting oneself into whatever the relevant mental state might be: knowledge, understanding, first-hand intuitive insight, or whatever. But the product itself is publicly available theories.

The other claim is that the Philosophy Machine itself is not a kind of aggregate of what is going on in the minds of individual philosophers; it is a social and intellectual institution. What is going on in our own heads makes a difference to the operation of the machine only in so far as it results in some kind of tangible effect on the rest of the machine. Our own personal mental economy qua people engaged in servicing the Philosophy Machine and producing the outputs is therefore not terribly important, in so far as the value of what is produced is concerned. What we believe, how strongly we believe or are warranted in believing it, whether it counts as (individual) knowledge, and so on are just not of primary importance when it comes to the product. A viable theory, or a good argument, or a well-aimed objection, or whatever, has that status independently of the doxastic state that the person presenting it happens to have towards it. So if said states *are* important, their importance must lie in their necessity for keeping the machine in good working order. The basic point of the next two sections is to argue that they are not, in fact, important for keeping the machine in good working order.

III

Steadfastness in the Face of Peer Disagreement. One might think—and very many philosophers *do* think—that a well-oiled philosophy machine requires that our assertions, as made in philosophical conversation, presentations, published work, and so on, express justified

beliefs. After all, we do want the product to meet pretty rigorous epistemic standards, and that requirement would seem to do the trick. But, unfortunately, it faces the problem of peer disagreement.

The question of what we should do, epistemically speaking, in the face of disagreement from our epistemic peers has attracted a lot of attention recently. More specifically, in the case where you take your belief that p to be justified or warranted, or, alternatively, where you have a high degree of credence in p , what do you do when confronted with an epistemic peer—someone whose level of expertise and evidence is roughly the same as yours—who denies that p , and who thereby poses at least a *prima facie* threat to your happy epistemic state? Hardline defenders of steadfastness hold that you have no reason to revise your epistemic attitude: to suspend judgement on whether p , or to regard your belief as less justified, or to reduce your confidence level. Hardline conciliationists hold that you should suspend judgement. And there is a range of views that occupy positions closer to the middle of the scale.

We see the same range of views with respect to philosophical beliefs in particular (see, for example, Baghramian, Carter and Cosker-Rowland 2024). In the context of philosophy, disagreement is pervasive and often very persistent, and yet our assertoric practice in philosophical contexts clearly manifests steadfastness. (This is a perfectly sensible exchange: ‘I added up the bill and you owe me £13.50.’ ‘No, it’s £12.75.’ ‘Oh, maybe I got my sum wrong then.’ This, in philosophical contexts, isn’t: ‘I’ve thought about it a lot, and I think causation is a matter of counterfactual dependence.’ ‘No it isn’t.’ ‘Oh, maybe I made a mistake then.’) Indeed, conversational norms in philosophy *demand* steadfastness. When X asserts that p and Y responds by asserting $\neg p$, that conversational move normally constitutes, at best, a potentially useful piece of information for the purposes of further dialogue. Depending on the precise nature of the conversation, Y may be signalling that they are about to launch into an objection to p , or that they have now identified the underlying cause of their earlier disagreement over the truth of q and the conversation can now move on, or whatever. What absolutely isn’t expected of X —just on the basis of Y ’s having asserted that $\neg p$ —is that she concede any ground whatsoever.

This is all rather puzzling, given the very obvious *prima facie* sceptical threat posed by epistemic peers, or even acknowledged epistemic superiors, going around asserting just the opposite of what

you assert. What licenses such steadfastness? Answers to that question are thin on the ground (although not entirely absent; see, for example, Kelly 2016; Buchak 2021). But the puzzle only arises if we assume that our assertions in philosophical contexts are expressions of belief—or, to put it another way, if the norm of assertion in such contexts is or requires belief. The rationality of steadfastness is only under threat in the first place if the name of the game here is the justificatory status of philosophical beliefs or the appropriate degree of credence. That's the dominant assumption in the peer disagreement literature.

But there's another strand to the literature—exemplified by, for example, Alessandra Tanesini and Catherine Elgin—that, in effect, pays attention not so much to the question of what a philosopher should *think* in such a situation—how, if at all, they should make adjustments to their mental economy—but what they should *do* (Elgin 2022; Tanesini 2024). These authors note that disagreement is not only pervasive in philosophy but lies at its beating heart. What we actually *do* when someone presents a position that conflicts with our own, they point out, is engage in a conversation. We ask our interlocutor to provide reasons for the claim, we consider and discuss those reasons, and so on. We might, eventually, retract or amend our claim. Or they might amend or retract theirs. The cogs in the Philosophy Machine shift their position. So when, in a philosophical conversation, someone merely denies what you just asserted without any clear signalling of what is supposed to happen next, they've made an illegitimate move in the conversation because they've failed to make any cogs shift. We all implicitly recognize that when *p* is, or is implied by, some claim that is widely contested, merely denying *p* when the speaker has just asserted it, even when done by an epistemic peer or even an acknowledged epistemic superior, on its own achieves precisely nothing.

Viewed from one perspective, then—according to which assertion in the context of philosophy is a matter of the expression of belief—philosophers' general refusal to concede any ground in the face of peer disagreement can make a defence of steadfastness against the sceptical challenge seem like an urgent task. Otherwise our assertions are apt to look like those of, as Keith DeRose puts it, 'overconfident blowhards, displaying absurd levels of confidence in our answers to questions we are in no position to be so confident about' (2017, p. 276). But if we approach the issue of what licenses steadfastness in

our assertoric practice from the perspective of what we, collectively, are trying to achieve in philosophy, and hence focus on how that practice contributes to helping the Machine do what it's designed to do, the question of the epistemic standing of any philosophical beliefs we might have lapses into irrelevance. Steadfastness in our assertoric practice is to be explained by the fact that, by and large, it is required for the smooth operation of the Philosophy Machine.

This isn't to say that there aren't plenty of overconfident blowhards in philosophy. It's also not to say that the question of how, if at all, we should adjust our mental economy in the face of peer disagreement is illegitimate or uninteresting. It's just to say that, as far as the *practice* of philosophy—the operation of the Philosophy Machine—is concerned, the question is an idle one.

IV

*What the Philosophy Machine Requires of Our Mental Economy.*³

One way to try and escape the problem of accounting for the assertoric norm of steadfastness in the face of peer disagreement in philosophy would be to weaken the doxastic state that we expect our philosophical assertions to express. There has been a lot of recent philosophical discussion about whether belief is really the appropriate attitude to have with respect to many philosophical claims, with a lot of authors suggesting that, given endemic persistent peer disagreement, something weaker than belief is appropriate: perhaps 'endorsement' (Fleisher 2021, p. 364), 'regarding-as-defensible' (Goldberg 2013), 'acceptance' (Beebe 2018; DeRose 2017, Appendix C) or some such. Such purported states are tailor-made to be appropriate in contexts where disagreement is persistent and where each party to the dispute has what they take to be good arguments for, or broader considerations that tell in favour of, their own position; in other words, in contexts where peer disagreement would threaten one's epistemic entitlement to belief. (I'll henceforth refer to states such as belief, acceptance, and so on as 'broadly doxastic states', or just 'doxastic states' for short.)

³ The discussion of norms of assertion in this section has been especially influenced by my collaboration with Ylwa Sjölin Wirling (cf. note 7 below).

I, however, want to defend a more radical position than this: that *no* broadly doxastic state concerning some disputed philosophical claim is required for legitimate assertion of it, because it's just not required by the smooth running of the Philosophy Machine. Rather, the norm of assertion in philosophical contexts demands of us only what I shall call 'commitment', which is (very roughly) a matter of having certain kinds of dispositions concerning the claim in question.

In her discussion of, as she puts it, 'disagreeing well or badly', Tanesini appeals to epistemic virtues and vices. She defines 'the intellectual virtues of disagreeing well as comprising attitudes, sensibilities, character traits, and dispositions that facilitate epistemically good outcomes in debates and discussions involving seemingly incompatible viewpoints and/or involve epistemically good motivations such as caring about truth, evidence and other people's epistemic agency and autonomy' (Tanesini 2024, p. 211). The corresponding vices include 'arrogance, unwillingness to listen, and testimonial insensitivity' (2024, p. 212). This is a helpful move for my purposes because it gives us an alternative way to think about what it is about the state of one's own mental economy, if not some broadly doxastic state, that plays a central role in the smooth operation of the Philosophy Machine, namely epistemically virtuous dispositions, character traits, and so on, which in turn (by and large) manifest in epistemically virtuous verbal behaviour, including assertion. And such behaviour is epistemically virtuous precisely because it fosters productive philosophical dialogue: it keeps the Philosophy Machine running smoothly. Manifesting these virtues in the various contexts in which philosophical disagreements arise furthers the debate in productive rather than counter-productive ways, the result of which—or so we might reasonably hope—is better product.

Assertion is one of many elements of verbal behaviour that 'disagreeing well'—or badly—encompasses (and indeed non-verbal behaviour: nodding enthusiastically, frowning, looking quizzical, rolling one's eyes, thumping the table, and so on). My claim is that, when it comes to espousing a philosophical view, we should think of the norm of assertion in terms of the role assertion plays in productive philosophical dialogue, that is, in functional terms: in terms of its efficacy in promoting philosophically useful debate and discussion.

A good starting-point is Will Fleisher's account of 'endorsement' (2020, 2021), which he proposes as the norm of assertion of what he calls 'advocacy role claims'—that is, claims 'that aim (or function)

to promote productive debate and disagreement' (2020, p. 242)—in the context of philosophical publishing. Fleisher's conditions on 'endorsement' include, *inter alia*, some dispositions—to 'assert the content of p or otherwise express commitment to p ', to 'treat p as a basis for further reasoning', and to 'shape [one's] research program ... in part based on p '—and the requirement that the asserter 'takes herself to be obligated to defend p ' (2021, p. 369). He takes endorsement to be a mental state that is relevantly similar to the kinds of mental state that other authors have recently touted as a suitable replacement for belief in the context of concerns about the sceptical threat posed by persistent peer disagreement. Endorsement, as Fleisher sees it, is a mental state that is in some ways more demanding than belief (one can believe all sorts of philosophical claims that one has no intention to 'shape' one's research program around, for example), but in other ways less demanding. In particular, one can reasonably endorse p while remaining agnostic about p . But I suggest that assertion of p in philosophical contexts need not constitute the expression of some broadly doxastic state with respect to p *at all*.⁴

So what, in terms of their mental economy, is required of someone who makes an assertion in philosophical contexts? I propose to call the answer to that question 'commitment'. Consider someone who is committed to a relationship or to being a good neighbour, or is a committed vegetarian. Broadly speaking, someone who has such a commitment takes themselves to have—or has voluntarily taken on—a range of responsibilities or obligations, and has (and manifests in appropriate circumstances) a range of dispositions to act in certain ways. Someone who is committed to being a good neighbour might be disposed to learn her neighbours' names, to keep the noise down, to drop Christmas cards through doors. She might take on responsibility for checking in on the elderly man next door if she hasn't seen him for a while, or for mowing his lawn; indeed, she might take herself to be obliged to do those things. And we reasonably expect these features to be stable over time: someone who keeps up such behaviour for a couple of weeks and then stops was,

⁴ I'm not entirely sure how much I really disagree with Fleisher here, although he himself describes 'endorsement' as a 'doxastic attitude' (2021, p. 364). He defines endorsement in part in terms of an undefined notion of 'commitment' (2021, p. 369), and I am unsure whether he conceives 'commitment' as itself a broadly doxastic attitude or as constituted by the kinds of disposition just described.

temporarily, a good neighbour, but not a committed one (perhaps unless she has had a major change of heart for some reason and has decisively renounced her prior commitment). She will, of course, have a range of beliefs and desires that underpin all of this: the belief that the neighbours don't want to hear her practising the piano at all hours; the desire to be appreciated by her neighbours, perhaps, or to foster similar behaviour in them. But no *specific* beliefs or desires seem to be constitutive of the commitment; someone might be a committed vegetarian because she believes it is morally wrong to kill animals for food or that eating meat is a major contributor to the climate crisis, or because she promised never to eat meat again, or simply because avoiding meat makes her physically feel better.

Similarly—or so I claim—for philosophical commitments. One might be committed to Humean metaphysics, or to consequentialism or libertarianism or scientific realism. Having a commitment to some such broad-brush philosophical view consists in having taken on responsibility for acting in certain ways, and in actually having the disposition to act in those ways in appropriate circumstances: to assert and deny, in relevant situations, propositions according to whether they accord with or conflict with the view; to defend the view when confronted with objections; to articulate, when asked, the basic arguments for or the attractive features of the view; and so on. And, as with other kinds of commitment, we reasonably expect these features to be stable over time: someone who gives a talk defending dualism on Monday and another defending physicalism on Tuesday clearly is not and was not—absent a Damascene conversion in the interim—committed to either of those views. The same goes for specific philosophical claims as opposed to broad-brush views, although often more specific claims go along with less durable commitments: a consequentialist's exact position might well evolve over time in response to objections to earlier versions, but consequentialists (or compatibilists, or whatever) rarely renounce *that* commitment.

And similarly, I claim, as far as beliefs go: one can be committed to some broad philosophical view or some specific philosophical claim without believing it, and indeed without having any of the various candidate doxastic states of the kind mentioned earlier. That's not to say that there are *no* doxastic constraints on commitment. If I take responsibility for mowing Mr Jones's lawn, I am rationally required to believe myself to be *capable* of mowing Mr Jones's lawn.

Similarly, if I am committed to dualism, I am rationally required to at least believe that dualism is defensible, since such commitment involves, *inter alia*, taking on responsibility for defending dualism against objections. But we need not see that belief as *constitutive* of commitment.

Commitment, thus conceived, delivers steadfastness in our philosophical behaviour—as it should, because, as discussed in the previous section, that behaviour *is* steadfast, and rightly so, given that steadfastness is required by the smooth operation of the Philosophy Machine. If I assert p and an epistemic peer responds by asserting $\neg p$, I am under no obligation to rescind my commitment to p and hence retract—or even defend— p , since no objection to p has been raised. The mere fact that someone else has taken on responsibility for defending $\neg p$ gives me no reason to think myself any less capable of defending p . This explains why merely denying what one's interlocutor has just asserted is idle.

One worry one might have at this point—and which has been expressed by some authors in the debate about whether published assertions in philosophy are subject to the norm of belief—is that, absent such a norm (or perhaps a norm of acceptance or some such), we fail to capture the legitimate demand for *sincerity* in people's philosophical assertions (Buckwalter 2023; Basu forthcoming). But it's not obvious to me that the notion of commitment, as just sketched, fails to do that job. We need to make sense of ourselves and others as philosophical enquirers, and doing that may well involve taking a view on why I, or someone else, might have taken on the relevant commitments. But the psychological explanation might vary significantly from person to person. Philosopher X might take on responsibility for asserting p , for defending it in the face of objections, and so on because they're an overconfident blowhard and they jolly well think they're *right* and everyone who disagrees with them is an *idiot*. Philosopher Y might take on the commitment because they really *like* the view in question, or because they find the world according to p more aesthetically pleasing than the alternatives. Philosopher Z might take on the commitment because, hey, this is the latest hot topic and devoting myself to defending p is going to get me some publications in big journals. Intuitions may differ, but mine are that in each case it is appropriate to describe both the commitment to and the assertion of p as sincere in all three cases. By contrast, philosopher W might

publish a newspaper article arguing for some egregious claim p , wilfully ignoring—and not knowing how to respond to—the many existing arguments against p . W is insincere because she advertises herself as having a commitment to p that she lacks: she knows she cannot, and hence lacks the disposition to, defend p against objections.

At any rate, it's not clear that it *matters*, at least as far as the business of keeping the Philosophy Machine in working order is concerned, what the psychological explanation of commitment is in any given case. You might want to avoid Philosopher X, the overconfident blowhard; and you might suspect that careerist Philosopher Z isn't going to do you any favours unless it results in some improvement in their CV. You might also have legitimate concerns about whether their respective attitudes are symptomatic of a character that will negatively affect the smooth running of the Philosophy Machine in other contexts; overconfident blowhards tend not to make very sympathetic journal referees, for example. Fine. But their contributions to the Machine, as far as their philosophical assertions are concerned, are not criticizable on the grounds that they lack sincerity.

V

Philosophy as an Activity. I distinguished in §1 between philosophy qua activity and qua product, and claimed that both have intrinsic value. This being so, the activity of philosophizing that we engage in when contributing to that public good itself has intrinsic as well as instrumental value. But we ourselves contribute very little to the potential sum of that intrinsic value by doing philosophy ourselves—even if we're doing philosophy a lot better than any schoolchild or retired accountant and doing it better has more intrinsic value than doing it worse. We are, after all, a tiny and insignificant minority of human beings. So, to the extent that we care about the intrinsic value of philosophy qua activity, we should not pretend to ourselves that engaging in our own research actually contributes very much at all to that, and we should care about improving the reach and quality of the activity more widely, by teaching philosophy not just to our students but in schools and prisons, making podcasts, giving public lectures, publishing articles and books for a general audience, perhaps engaging in

Socratic dialogue (Sheffield 2025)—but of course not just amongst ourselves—and so on.⁵

Doing those kinds of thing is required for achieving something else as well. I have taken it for granted that the product—the publicly available content—has intrinsic value qua resource for the deployment and edification of human beings at large. I've said nothing about *why* it has intrinsic value, and I'm not going to. But, as noted in §II, 'deployment' and 'edification' do of course require some actual humans to get themselves into some kind of mental state or other with respect to that content. Maybe it's knowledge, or understanding, or even first-hand intuitive insight, or something else; a story about why the product has value will have to say something about what the relevant state is, and why it in turn is a valuable state for human beings to be in. But whatever that state is, getting oneself into it, as we all know from personal experience, is no easy matter. Perhaps such a state is only realistically available to those with significant training in philosophy. That raises a worry, of course, about whether and to what extent the product of the Philosophy Machine constitutes a public good, if its intrinsic value can only be cashed in, as it were, by a tiny minority of human beings. But at least philosophy is no worse off in that regard than algebraic geometry or string theory or advanced game theory.

Finally, philosophy can be *useful*—not just in the sense that it provides people with tools to think more clearly about matters that concern them (your elderly parent has dementia or your oldest friend betrayed you or you're confronting a really thorny moral dilemma or you find discussions about race or gender *really* confusing, or whatever), but in a more practical sense: it is apt for—as Philip Kitcher puts in his recent book, *What's the Use of Philosophy?*—'guiding human practices through its achievements in introducing concepts, proposing lines of reasoning, suggesting standards and rules, posing questions, offering striking comparisons, opening up possibilities, and so on' (2023, p. 150). Those practices include,

⁵ Michela Massimi (2025) has recently written on the normative implications of interpreting the right 'to enjoy the benefits of scientific progress', as enshrined in Article 15.1(b) of the UN's 1966 International Covenant on Economic, Social and Cultural Rights, as the right to *participate* in science. It is an interesting question how one might approach this issue if one interpreted 'science' in a broader sense that encompasses other academic disciplines including philosophy.

for example, policy-making in a wide range of areas and legal proceedings.⁶ Putting philosophy to such uses will often involve what Quassim Cassam (2024a, 2024b) calls ‘co-creation’: ‘problem solving by diverse stakeholders with different assumptions, interests, skills, and thinking styles ... It is not a matter of one person *revealing* a solution to others but of two or more people working together to solve problems. It requires openness to diverse perspectives and interests, a willingness to listen, and to see things from other points of view’ (2024b, p. 130).

These kinds of ways in which philosophers can contribute to what are arguably public goods (though, again, I haven’t made the argument) inevitably involve an interplay between (as I’ve put it) activity and product: the ‘achievements’ that Kitcher lists in the passage quoted above are, after all, mostly to be found in the philosophical literature; on the other hand, merely directing people to that literature is not going to achieve much, if anything, on its own. What they require of the philosopher is a variety of skills and capacities—for example those mentioned by Cassam in the case of ‘co-creation’—as well as some knowledge of the relevant literature (the product). What they do not require, however, is the philosopher who is making the contribution to believe or endorse or accept—or even be committed to—any particular philosophical view. In helping an individual to think philosophically about friendship or guilt or personal identity, or an organization to think philosophically about the ethics of its practices or risk or evidence or whatever, the philosopher deploys their skills and the resources that are available to them. What the individual or organization concerned wants—or ought to want, if they are genuinely interested in thinking philosophically—is to know *how* to think about the issue at hand, and not *what* to think about it.

VI

Conclusion. I have argued that our own broadly doxastic states concerning philosophical claims, as people engaged in the various activities that contribute to the discipline of philosophy qua public good, are not terribly important. What we believe (or ‘accept’ or ‘endorse’),

⁶ For a sense of the very broad range of uses to which philosophy can be put, see the database of impact case studies submitted for the 2021 UK Research Excellence Framework at <https://results2021.ref.ac.uk/impact> (filter by ‘unit of assessment’ and search for Philosophy).

how strongly we believe or are warranted in believing it, whether it counts as knowledge or ‘first-hand intuitive insight’, and so on play a surprisingly modest role when it comes both to ensuring that the Philosophy Machine produces a high-quality product and to facilitating and improving the activity of philosophy—and making useful philosophical contributions—in the wider world.

This is not to say that the epistemology of philosophy, as it applies to individuals, is a waste of anyone’s time. Questions about what our individual attitudes should be towards the philosophical theories we spend our time developing and criticizing are perfectly legitimate philosophical questions that merit philosophical attention. But my claim is that such epistemology is not an epistemology ‘for’ philosophy, as Sosa puts it, in the sense of telling us anything about what we, individually, should be aiming to achieve for ourselves in our role as contributors to the public goods that philosophy delivers.⁷

School of Philosophy, Religion and History of Science
University of Leeds
Leeds LS2 9JT
UK

REFERENCES

- Baghramian, Maria, J. Adam Carter, and Rach Cosker-Rowland (eds.) 2024:
The Routledge Handbook of Philosophy of Disagreement. New York: Routledge.
- Basu, Rima forthcoming: ‘Against Publishing Without Belief: Fake News, Misinformation, and Perverse Publishing Incentives’. To appear in Sanford Goldberg and Mark Walker (eds.), *Philosophy with Attitude*. New York: Oxford University Press.

⁷ This paper was written during my tenure as Jane and Aatos Erkkö Visiting Professor in Studies on Contemporary Society, Helsinki Collegium of Advanced Study. Thanks to Alexander Bird, Jessica Leech, Michela Massimi and audiences at the Philosophy Society of Finland and the Universities of Turku and Vienna for helpful discussions. Some of the material in this paper overlaps with the content of book I am writing with Ylwa Sjölin Wirling, which has been in progress for so long and involved so much discussion that it’s no longer obvious to either of us whose ideas started out as whose. So I am not at all confident that all of the ideas expressed in my paper are entirely my own. But in any case, I definitely couldn’t have written this paper without those discussions, and I am hugely grateful to Ylwa both for them and for letting me attach my name to some ideas that may well have originated with her.

- Beebee, Helen 2018: 'Philosophical Scepticism and the Aims of Philosophy'. *Proceedings of the Aristotelian Society*, 118(1), pp. 1–24.
- Bird, Alexander 2022: *Knowing Science*. Oxford: Oxford University Press.
- Buchak, Lara 2021: 'A Faithful Response to Disagreement'. *Philosophical Review*, 130(2), pp. 191–226.
- Buckwalter, Wesley 2023: 'The Belief Norm of Academic Publishing'. *Ergo: An Open Access Journal of Philosophy*, 9(54). <https://doi.org/10.3998/ergo.3117>.
- Cassam, Quassim 2024a: 'Liberation Philosophy'. *Proceedings of the Aristotelian Society Supplementary Volume* 98, pp. 1–26.
- 2024b: 'Review: Philip Kitcher, *What's the Use of Philosophy?*' *Society*, 61(1), pp. 128–30.
- DeRose, Keith 2017: *The Appearance of Ignorance: Knowledge, Skepticism, and Context*, Volume 2. Oxford: Oxford University Press.
- Elgin, Catherine Z 2022: 'Disagreement in Philosophy'. *Synthese*, 200(1), pp. 1–16.
- Fleisher, Will 2020: 'Publishing Without (Some) Belief'. *Thought: A Journal of Philosophy*, 9(4), pp. 237–46.
- 2021: 'Endorsement and Assertion'. *Noûs*, 55(2), pp. 363–84.
- Goldberg, Sanford C. 2013: 'Defending Philosophy in the Face of Systematic Disagreement'. In Diego E. Machuca (ed.), *Disagreement and Skepticism*, pp. 277–94. New York: Routledge.
- Hawking, Stephen, and Leonard Mlodinow 2010: *The Grand Design: New Answers to the Ultimate Questions of Life*. London: Transworld Publishers.
- Kelly, Thomas 2016: 'Disagreement in Philosophy: Its Epistemic Significance'. In Herman Cappelen, Tamar Szabó Gendler, and John Hawthorne (eds.), *The Oxford Handbook of Philosophical Methodology*, pp. 374–94. Oxford: Oxford University Press.
- Kitcher, Philip 2023: *What's the Use of Philosophy?* New York: Oxford University Press.
- Massimi, Michela 2025: 'Local Knowledges and the Right to Participate in Science'. Forthcoming in *Philosophy of Science*, pp. 1–25. <https://doi.org/10.1017/psa.2025.4>.
- Popper, Karl R. 1959: *The Logic of Scientific Discovery*. London: Hutchinson.
- 1979: *Objective Knowledge: An Evolutionary Approach*. Oxford: Oxford University Press.
- Sheffield, Frisbee C. C. 2025: 'Socrates and Dialogue as the Greatest Good'. *Proceedings of the Aristotelian Society*, 125(3), <https://doi.org/10.1093/arisoc/aoaf016>.
- Sosa, Ernest 2021: *Epistemic Explanations*. Oxford: Oxford University Press.
- Tanesini, Alessandra 2024: Disagreement and Intellectual Virtues and Vices. In Baghrarian, Carter and Cosker-Rowland 2024, pp. 210–22.