Synthese (2022) 200:20
https://doi.org/10.1007/s11229-022-03535-y

ORIGINAL RESEARCH

Disagreement in philosophy

Catherine Z. Elgin[1]

Received: 11 May 2021 / Accepted: 4 December 2021 / Published online: 18 February 2022

Abstract
Recent philosophical discussions construe disagreement as epistemically unsettling. On learning that a peer disagrees, it is said, you should suspend judgment, lower your credence, or dismiss your peer's conviction as somehow flawed, even if you can neither identify the flaw nor explain why you think she is the party in error. Philosophers do none of these things. A distinctive feature of philosophy as currently practiced is that, although we marshal the strongest arguments we can devise, we do not expect others to agree. Nor are we dismayed then they do not. Through a survey of familiar professional practices, I argue that philosophy rightly revels in responsible disagreement. This discloses important and perhaps surprising facets of the epistemology of philosophy.

Keywords Philosophy · Disagreement · Understanding · Epistemic acceptance

1 Introduction

Philosophers are a disagreeable lot. By that I don't mean that we are particularly unpleasant. Far from it. Many are lovely people. In saying that philosophers are disagreeable, I mean that we are *able* to *disagree* with one another about virtually anything. Take any philosophical thesis (including the law of non-contradiction), we can probably find someone who will seriously argue for its denial. Take any mundane statement of fact. We can find grounds to challenge it (even if we have to resort to malevolent demons or brains in vats) and we can find venues in which we both do and should take such challenges seriously.

Recent philosophical discussions present disagreement as epistemically unsettling. Evidently, all things considered, we'd prefer that folks agree. Learning that a peer disagrees with you, it is said, should prompt you to suspend judgment, lower your credence, or dismiss your peer's conviction as somehow flawed, even if you can neither identify the flaw nor explain why you think she is the party in error. Philosophers do

✉ Catherine Z. Elgin
catherine_elgin@harvard.edu

[1] Graduate School of Education, Harvard University, Cambridge, MA 02138, USA

none of the above. A distinctive feature of philosophy as currently practiced is that, although we marshal the strongest arguments we can devise, we do not really expect others to agree. Nor should we be particularly daunted or even disconcerted when they don't (see Goldberg, 2013a). That, I suggest, raises questions about what we think we are doing. To underscore the puzzle, let's look at some curious cases.

2 Curious cases

2.1 An anecdote

Shortly after Goodman and I published *Reconceptions* (1988), a workshop on the book was held at the University of Bielefeld. The workshop had the standard format. Someone presented a paper. One of us gave a short reply. Then the floor was open for a general discussion. One issue discussed in *Reconceptions* concerns the identity of literary works. The inscriptionalist criterion advanced—that sameness of spelling is the criterion of identity for a literary work—has the unattractive consequence that translations are not strictly instances of the work they translate. We devised something that, in computer terminology, might be called a patch. It handled the problem, but in a graceless, seemingly ad hoc way. At the workshop, Wolfgang Heydrich presented a paper that respected our nominalist scruples and supplied an elegant alternative (Heydrich, 1993).

For our purposes, the important point is what happened next. I replied: 'He's right. His alternative is lovely. It is much better than ours. I wish we had thought of it.' Everyone apparently agreed with that assessment. No one knew what to do next. Once we said he was right, there seemed virtually nothing left to discuss. The collective befuddlement may have been exacerbated by the technical nature of the issue. In effect Heydrich said: here's a glitch; here's a fix. Probably in other cases, an argument that got (and merited) universal agreement would open avenues for fruitful discussion. 'If, as we all agree, Professor *Z* has conclusively demonstrated that *p*, that opens the door to *q*, *r*, and *s*.' Still, the reactions at the Bielefeld workshop underscored the fact that the standard expectation in philosophy is that we will disagree: we will raise pointed objections, ask challenging questions, identify and spell out the significance of weaknesses. *Philosophers are typically ill-equipped to give or take 'yes' for an answer.*

We might compare philosophy's practice with another genre in which someone tells others what to think or do: the sermon. The topic of a sermon and a philosophy lecture could be the same; they could even be delivered by the same person. Suppose Reverend *X* is also a philosopher. This week he gives both a sermon and a philosophy lecture on, say, the nature of the good. Perhaps he even gives the same talk in both venues. (I have been told that Bishop Butler did this.) Still, there is a difference. In church, there is no Q & A. There is no formal, real-time opportunity for members of the congregation to dispute what is said. They are supposed to sit quietly, take it in, and, if need be, revise their beliefs and behavior accordingly. After the philosophy lecture, however, auditors are expected to raise objections. And the speaker is expected

to reply. Regardless of the esteem with which they hold the lecturer, philosophers do not and should not let his words go unchallenged. Disagreement is expected.

2.2 Lit reviews

Many philosophy papers open with a review of the literature. The practice goes back to Aristotle who began many of his works saying, 'Let us examine the opinions of the many and the wise'(1941). Like Aristotle, we review the important positions on the subject under discussion. (We may ignore the opinions of the many, but typically we consider ourselves obliged to address the opinions of the allegedly wise.) Occasionally we accept and build on what another philosopher has said. But even there, we are apt to intimate that the position we build on regrettably did not go far enough. A literature review in philosophy typically consists of rehearsing extant positions and explaining why they are inadequate. That done, we go on to present our own (which, naturally enough, does not suffer from the flaws we highlighted). The literature review sets the stage; it frames our discussion by explaining how and why we disagree with others who have worked on the same topic.

2.3 Hiring decisions

Philosophy departments typically ignore the question of whether the candidate's philosophical views are true (Lewis, 2000). Assuming that the goal of the department and the university is the advancement of knowledge, Lewis maintains, this practice is epistemically costly. Ceteris paribus, someone with false beliefs on the topic he plans to devote a considerable portion of his professional life to is less likely to advance knowledge than someone with true beliefs on the topic. The reason, Lewis says, is this: 'To the extent that a researcher is guided by false doctrines, he is liable to arrive at new and different false doctrines, since he will choose them partly to cohere with the doctrines he held before' (2000: p. 189). Although this is rather obvious, hiring committees in philosophy routinely disregard it. Our hiring practice is odd. We purport to seek the truth, yet largely bracket the question of whether the candidate's views are true. If we sought a medical opinion, or advice about how to fix the furnace, we would presumably restrict our informants to those we had reason to think held true beliefs about the topic. Minimally, we would eschew those whose beliefs we thought were definitely false. Why are philosophy hiring decisions different?

Lewis suggests that the practice is the product of a tacit treaty. Since there is no assurance that the majority in any given philosophy department believes the truth, there is a chance that if the truth value ascribed to a candidate's views were considered relevant, the result would entrench falsehood. Hence we tacitly agree that all parties will willingly forego the best outcome (bringing it about that the department be dominated by those with true philosophical beliefs) to block the worst (bringing it about that the department be dominated by those with false ones) (2000: p. 198). Perhaps so. But when the result of the treaty is that the department hires someone whose views a member regards as false, the hiree will be her colleague for the foreseeable future. So

they will disagree. Lewis's argument suggests that this is a regrettable consequence that we all have to live with. Below I will suggest a more positive spin.

2.4 Response to fallibility

Philip Kitcher (1990) argues that when there is a significant chance that a currently popular scientific theory is false, the community has sound epistemic reasons to insure that alternative positions remain viable. This requires that adherents of those positions get some proportion of the jobs in the field, a measure of financial support to enable them to continue their investigations, and opportunities to present their results and be assured of a fair hearing. If the members of the community want to believe only what is true, Kitcher argues, it should not foreclose inquiry prematurely. Presumably such open-mindedness is restricted to scientifically plausible alternatives. Kitcher is not recommending that the National Science Foundation fund the research of contemporary followers of Thales who are bound and determined to demonstrate that, the history of science notwithstanding, everything is water.

Inasmuch as it pertains to science, Kitcher's argument may seem an outlier here. But his argument applies to philosophy at least as much as to science. Maybe more. However confident we may be about the epistemic status of currently accepted scientific theories, we should probably suspect that our favored philosophical position, regardless of its popularity, stands a good chance of being false. As Richard Fumerton points out, 'most philosophical views are minority opinions' (2010: p. 109). On any issue of philosophical significance, there is nothing close to consensus. Not only should we therefore refrain from foreclosing inquiry into currently unpopular but plausible rival views, we should provide venues where the merits of different positions, even implausible ones, can be compared and contrasted—that is, venues where disagreements can be aired and taken seriously. Our practices reflect this policy. Journals and conferences welcome divergent views. To be sure, we don't want to spend our careers fending off kooks and nuts. So we will have to figure out how to exclude them. The point is that the criterion for exclusion should be something other than the kookiness or nuttiness of the theses being proposed.

2.5 Required courses

Many philosophy departments require their students to take courses in the history of philosophy. They strong-arm their students into spending considerable time and intellectual energy studying positions that their mentors reject as clearly false—positions that they both hope and expect the students will also reject. We require our students to study Plato, even if we are convinced that the forms do not exist. We require them to study Spinoza even if we are confident that there's more than one thing in the world. We require them to study Kant even if we consider it obvious that transcendental idealism is a non-starter. You can continue the list, enumerating plainly untenable views that you were required to study, and those that you in turn will require your students to study. I am not suggesting that this is a bad practice. In fact, I think that studying the great philosophers of the past is vital. I will say more about why shortly. My main

point here is that there is something at least odd about expecting ourselves and our students to spend a lot of time studying positions that we reject, and reject for good reason. We actively disagree not only with the views of our contemporaries, but also with those of our forbears. And we can't disagree with those views if we don't know what they are.

Philosophy's curious institutional arrangements and practices seem to indicate that we consider responsible philosophical disagreements valuable. What can that value be? Before addressing that, we should consider what it is to hold a philosophical position.

3 Philosophical stance

Many philosophers assume that to hold a philosophical position is to believe it. This strikes me as wrong. To believe a theory is to believe that it is true. Some philosophers—notably Williamson and Stanley—take this view about their own theories. Most, I suspect, do not. If one believes that a theory is true, one ought to believe that it will never justifiably be rejected. Any objections raised against it are misleaders. *Never?* Is it remotely plausible that 200 years from now even our best philosophical theories will be accepted, exactly as they stand? Even if you are, for example, so convinced of externalism that you believe that in 200 years some version of externalism will dominate epistemology, is it plausible that it will be the very version you, or Alvin Goldman, or Hilary Kornblith worked out? Here, I suggest, the pessimistic meta-induction is sound. If we look at the history of philosophy, we do not find a body of received truths that were never subject to revision. It seems unjustifiably arrogant to think that *my* theory and *my* arguments for it (or Goldman's or Kornblith's) are so powerful that they will escape the fates that befell Aristotle, Descartes, Kant, Frege, and whatever intellectual heroes you want to add to the list. Nor is the pessimistic meta-induction the only ground for doubt. Fumerton maintains that he probably shouldn't believe his own theory given how many of his very smart friends and colleagues do not (2010: p. 109). He is not alone in this.

This is not to deny that philosophers are and should be committed to their theories. It is to say that their commitment is not a matter of belief. We are, and think we should be, committed to our philosophical views. What is the nature of that commitment?

> Goldberg suggests that our attitude should be
> *Attitudinal speculation*: 'One who attitudinally speculates that p regards p as more likely that $\sim p$, though also regards the total evidence as stopping short of warranting belief in p.' (2013b: p. 284).

Attitudinal speculation, Goldberg holds, is appropriate when the evidence supports p as opposed to $\sim p$, but not so strongly that the threshold for full belief is reached.

Theories are complex. They consist of a multiplicity of interwoven commitments. There are any number of dimensions along which a philosophical theory could be wrong. Holding that a theory is more likely to be true rather than false is probably excessively optimistic. To think, for example, that Korsgaard's constructivist ethics, or Sosa's virtue epistemology, or Chakravartty's scientific realism is more likely to be

true than to be false is to think that it is unlikely that there is a mistake anywhere in the theory.

Another problem can be seen if we assess a view against its rivals. Suppose a consequentialist recognizes that the plausible rivals to his theory are deontology and virtue theory. (To keep things simple, we will ignore his attitudes towards rival consequentialist theories.) He might hold that consequentialism is more likely to be true than either deontology or virtue theory. But he might still assign the probabilities as follows:

Consequentialism—40%
Deontology—30%
Virtue theory—30%.

In that case, he does not think the theory he is committed to is more likely to be true than to be false. He does think it is more likely to be true than any of its competitors considered alone (Barnett, 2019: p. 114). Barnett suggests that Goldberg would do better to accept

*Attitudinal Speculation**: One who attitudinally speculates* that *p* regards *p* as the likeliest option (given some set of options), though also regards the total evidence as stopping short of warranting belief in *p* (2019: p. 115).

Barnett (2019) argues that even with this modification, Goldberg faces a problem. Peer disagreement provides higher-order evidence that the position an agent espouses is unwarranted. So the vast number of philosophers who disagree with consequentialism supplies enough evidence against it to bring its probability to below 50%. And the vast number of philosophers who think that some other view, say deontology, is correct diminishes the probability that consequentialism is the likeliest alternative. The bad news is that this is not just a problem for moral theories. It holds throughout philosophy.

Barnett suggests that our attitude toward our favored philosophical position is what he calls '*disagreement insulated belief*' (2019: p. 121). The position you should hold, he argues, is the one that is the likeliest option given your evidence, once you have set aside the evidence of disagreement. It is, in effect, what you would believe if you were insulated from information about what your peers believe.

There are a couple of difficulties with this suggestion. The first is that it ignores the pessimistic meta-induction. That well supported theories in philosophy have been rejected or seriously revised in the past should give me pause if I'm tempted to think that my theory will have a happier fate. A second worry is that insulating ourselves from disagreement deprives us of information. Even if I would justifiably believe a theory if I was unaware that others disagree, the news that they disagree seems germane. Once I learn of their disagreement, it seems irresponsible to ignore it. This is especially so if those who disagree with me are people I regard as having relevant expertise.

4 Reorientation

To appreciate the importance of disagreement for philosophy, I suggest, a reorientation is in order. Standardly, discussions about disagreement focus on differences of opinion

as to the truth value of a particular proposition. This carries over to the discussions of disagreement in philosophy. Such discussions are cast as differences of opinion as to whether to believe, for example, externalism or internalism, consequentialism or deontology, scientific realism or constructive empiricism, as though the issue is what proposition to endorse. But such generic positions admit of multiple, divergent theories. A confirmed reliabilist, presumably holds that a belief qualifies as knowledge only if it is reliably connected to the facts. Still, opinions differ over which connections are reliable, and how reliable they have to be. There is plenty to disagree about within the reliabilist camp. Theories are intricately interwoven networks of commitments. Disagreements in philosophy concern what network of commitments is best on balance; they are not typically reducible to disagreements as to the truth value of a particular proposition. I'll have more to say about this shortly.

I suggest that philosophy is a quest for understanding rather than knowledge. If we want to understand the moral realm or the epistemic realm or the metaphysical realm, we want to grasp how a variety of considerations interweave to constitute a tenable take on the topic. Grasping is not believing. 'When you grasp a relationship between... propositions, you have that relationship under your control. You can manipulate it. You have a set of abilities or know-how relevant to it, which you can exercise if you choose' (Hills, 2016: p. 663). You understand how the various components hang together and provide support for one another. To grasp the Pythagorean theorem is to see how the various axioms and inferences together establish its truth. The question then is not whether externalism is true, but whether, and if so how and how adequately the network of externalist commitments collectively afford insight into the epistemic realm.

Rather than belief (or belief-minus-something), we should focus on *epistemic acceptance*. To accept a theory or network of commitments is to be willing and able to use it as a basis for assertoric inference or for action when one's ends are cognitive (see Cohen, 1992; Elgin, 2017). Epistemic acceptance does not require belief. Someone might accept Newtonian premises in calculating the force of a baseball on the catcher's mitt even though he knows that Newton's laws are strictly inaccurate. He considers $f = ma$ a good basis for inference in a cognitive context where relativistic effects, although real, are negligible. He might design and deploy a device to measure the force with which the ball impacts the glove, again relying on $f = ma$. His actions and inference further his cognitive ends even though, in relying on Newton's laws, he departs from truth.

A theory or network of commitments is acceptable when it is at least as good as any available rival. The goodness in question is a matter of reflective equilibrium: the commitments that comprise a tenable network are reasonable in light of one another, hence mutually supportive; and the network as a whole is at least as reasonable as any available alternative in light of the antecedent commitments we deem relevant. Because I have argued for these theses elsewhere, I won't offer a detailed defense of them here (see Elgin, 1996).

The commitments that comprise an epistemic network are not all judgments of fact. They include norms, standards, methods, rules, and prospects. Nor is reflective equilibrium supposed to insure truth. It only insures that the system is as reasonable as any available alternative in the epistemic circumstances. As a result, someone can

justifiably accept a theory while recognizing that it is likely eventually to be superseded. She can justifiably accept a theory while recognizing that it has rivals that are equally tenable. Indeed, she can even justifiably accept a theory while recognizing that it is a bit of a long-shot. Then its claim to equal tenability would be based on an assessment of its promise, as compared with the promise of its rivals. An evidentialist, for example, might consider the current support for her theory somewhat worse than the support for reliabilism, but deem its prospects better. She might think this on the basis of an assessment of the trajectories of the improvements in the two positions over time. In that case, she would hold that even if reliabilism has a slight edge today, evidentialist theories are improving at a faster rate. (Improvement would be measured, presumably, in terms of increasing ability to solve epistemological problems.) This reorientation puts us in a position to see how philosophical disagreements can be valuable.

5 Bet hedging

Mill (1978), Kitcher (1990), Lewis (2000) make a strong case that epistemic diversity is valuable. By keeping alternatives to a received view alive we hedge our bets. The issues under investigation are difficult; the available resources evidently afford considerably less than conclusive reasons. So it makes sense to incorporate into our disciplinary practices a recognition that despite the evidence in favor of a particular position, it still might be wrong.

A division of cognitive labor is an efficient way to sustain diversity. Rather than expecting each of us to master the pros and cons of each of the alternatives and fair-mindedly argue 'on the one hand,...; on the other hand...', it may be preferable for one party to develop and, to the best of her ability, argue for one position, while another develops and argues for an alternative. This strategy provides an incentive to tolerate and even encourage the development of positions at odds with your own. It does not, by itself, provide any reason to value disagreement per se. To reap the benefits of cognitive diversity, it might be best to simply allow for parallel tracks: One group champions dualism; another, monism. One champions nominalism; another, platonism. One champions internalism, another, externalism. Each can proceed in isolation from the other. The downside is that such a strategy would impede the possibility of learning from one another. Ceteris paribus, a better approach would be to encourage communication between groups with rival commitments.

Recognizing our fallibility is not as helpful as we might hope. To acknowledge that the position put forth in a 30 page article or a 250 page book might be wrong somewhere does not give any indication of where or how or why it might be wrong. Disagreement can play a diagnostic role. It provides focus by pitting positions against one another. This is why the disputational style of philosophy's Q & A's (or the style they should display) pays epistemic dividends. A would-be questioner who simply muses about what he consider related issues is unhelpful. So is one who takes the opportunity to spout off about his own pet theory rather than addressing the position that has just been presented. So is one who announces that the position just has to be wrong, but can't pinpoint any defect. Such so-called questioners distract from the goal of the practice—to put the author of the paper and the auditors in a better

position to assess the claims that are being made. But a questioner who carefully articulates a worry—something of the form, 'if you say that *a,* doesn't that commit you to *b*?'—pinpoints a locus of concern. She does not simply announce that something *might* be wrong somewhere or even that it *must* be wrong somewhere. We already knew the first and may suspect the second. She indicates where she thinks the problem lies. She might, of course, make her case by appealing to her own theory. But the appeal should be relevant and apt. If it is, then by looking at the argument from her perspective, we may discern something worrisome.

Actual disagreement is not required to perform this function, though. Often valuable discussions involve challengers who play the role of devil's advocate—asking, 'What would you say if *p*?' when they do not for a moment countenance *p.* Inasmuch as philosophical positions are supposed to be defensible against bizarre as well as plausible alternatives, this is a common and valuable practice. Still, actual disagreement is apt to dig deeper. Someone who responsibly disagrees with a position is likely to have thought more deeply about the issue than a mere devil's advocate. She has probably subjected her alternative to scrutiny. She is intellectually invested in it. So, ceteris paribus, her objection is apt to have more traction than one a devil's advocate can formulate on the spot.

Much of the literature on disagreement poses the problem as though the believer knows nothing more than the brute fact that a peer disagrees. But normally we know who disagrees, what they disagree about, and why. Even just knowing who disagrees is informative. If, for example, I learn that Miranda Fricker disagrees with me, I worry about whether I've been insensitive to epistemic injustice; if I learn that Sophie Horwitz disagrees, I worry about whether I have violated Bayesian constraints. Once I am privy to the content of the disagreement, I can consider why they hold an opposing view. What do they see that I do not? What are they focusing on that I am not? If I insulate myself from disagreement, as Barnett advocates, I may never know. One epistemic contribution of responsible disagreement then is to pinpoint errors, vulnerabilities, or potential problems in a position. That is not the only benefit.

6 Advancing understanding

A critical question concerns the basis for the disagreement. Philosophical disagreements are typically not like Christensen's restaurant case (2007). If they are at all serious, no one has made a stupid mistake or overlooked something obvious. Sometimes, although I suspect rather rarely, there is a disagreement about a simple matter of fact or logic. Maybe one party to the disagreement just overlooked a fact that ought to have been considered, or made an invalid inference. More often though, the grounds are different. Elsewhere I argued that 'peers can disagree along a variety of epistemically relevant axes: the weight they assign to evidence, the standards of acceptability, the identity of misleaders, the relevance and importance of various bits of background information, the favored styles of reasoning' (2018: p. 17). If, for example, a consequentialist maintains that the good is prior to the right, while a deontologist maintains that the right is prior to the good, those priorities are likely to figure in the assessment of arguments, in judgments about particular cases, in whether or to what extent

an unforeseen outcome tells against a moral judgment, and so forth. The consequentialist and the deontologist adopt different orientations toward the moral realm. Each perspective highlights some seemingly relevant factors and occludes others (see van Fraassen, 2008). Each provokes questions that the other does not. A deontologist probably need not worry much about unforeseeable consequences; a consequentialist has to. A scientific realist has to concern herself with the condition of Schrödinger's cat; a constructive empiricist might not.

Disagreements then are not just additional reasons to think that you might be wrong. They expand your epistemic range, disclosing previously unrecognized powers and limitations of your own theory. They underscore that the truth is rarely the whole truth. They reveal that there are potentially important factors that have been sidelined, omitted, unappreciated, or ignored. They sensitize you to weaknesses or limitations of your position. Even if your argument is solid, it may be vulnerable if the background assumptions are not precisely as you take them to be. Disagreements acquaint you with relevant alternatives and the considerations that might favor them. They put you in a position to recognize the range of considerations your position depends on—for example, how strong your modal commitments are; what types of evidence you can draw on, and what weight that we should attach to it; perhaps what boundaries you set on your theory. (Is it an ideal theory, or is it supposed to reflect real world constraints? Should you be concerned about the fact that it holds only under a limited range of actual circumstances?) These features are epistemically valuable because they enable you to better understand your own theory, the alternatives to it, and the topic it bears on.

A theory that diverges from yours may have attractive features that yours lacks. Consider David Lewis's realism about possible worlds. A nominalist does not think there is any chance that it is correct. But if he temporarily sidelines his reservations and seriously entertains it, he may recognize that it has the resources to solve a variety of problems that his metaphysically austere nominalism cannot. He learns something about the problem space he is working within, the questions he would like to answer, and the resources he has for answering them. He can recognize that Lewis's theory highlights the importance of modal distinctions that his theory elides. By taking it seriously, he can discern perhaps regrettable shortfalls in his own position, the costs that have to be paid to eschew metaphysical extravagance. Even though he comes away convinced that nominalism is more tenable overall, he also recognizes that it doesn't do everything he might want. Maybe, by studying Lewis's position, he can identify features that he can export. Perhaps, for example, with a suitable theory of fiction, a fictionalist about possibility could adopt some of Lewis's insights or strategies. To figure out if this is so, he needs to locate his disagreement with Lewis. What exactly does he disagree with? What features of the theory are peripheral to their disagreement?

A theory, even a correct theory, is not an intensional replica of reality. In theorizing, we decide what factors are important, how the domain is to be partitioned into kinds, at what grain and in what vocabulary the data are to be characterized, what factors can be set aside or treated as of only marginal significance, and so on. We should recognize that there are tradeoffs. One theory favors precision; another, breadth. One

regards epistemic accessibility as paramount; another cheerfully allows that that the truth may be unknowable.[1]

When we structure a domain, we draw lines, stipulating that some things will be counted as alike, others as unlike. Metaphysically, our stipulations are bound to be correct. Every collection of objects consists of items that are somehow alike. If nothing else, they are alike in being members of that particular collection. Most likenesses, although real, are utterly idle. There is no reason to recognize them. The critical issue in structuring a domain is to draw lines that group together items whose likenesses matter and to segregate from one another items whose differences matter. This holds, not just for objects, but also for principles, norms, and standards. It is unsurprising that there should be controversy about where lines should be drawn. Often critical disputes in philosophy come down to disputes about what issues are central.

Given the number of dimensions along which there are choice points, we should expect a variety of positions to emerge. They will highlight different aspects of the phenomena. Consequentialist theories, for example, highlight the fact that in acting we aim to produce a particular outcome; they maintain that the act should be assessed in terms of the outcome. Deontologists recognize that we are hostage to fortune. We may act with the best will and the best plan in the world, and still our actions come a cropper. That being so, they maintain, since ethics is concerned with blameworthiness, it should focus on the moral character and the intentions of the agent, not on the outcome. Such divergences can provide a foothold for potentially fruitful disagreements. Rather than simply hurling invectives at one another, we might learn from one another. We can ask, 'What is to be gleaned by seeing things from their point of view?' Even learning that a seemingly plausible theory won't do can be informative. The discovery that the appeal to natural kinds won't solve the new riddle of induction, absent an argument—rather than a mere intuition—that 'green' rather than 'grue' is the more natural kind, reveals something about the depth and complexity of the problem.

7 Responsible disagreement

For disagreement to be valuable, it must be responsible. I said earlier that we must block nuts and kooks if we are going to get anywhere. How are we to do that, if not by announcing that a disagreement grounded in realism about possible worlds or in the prospect that all emeralds are grue is too kooky to be taken seriously? I suggest that rather than assessing the content of the theses, we need to consider the epistemic contours of the disagreement.

Is the epistemic agent who disagrees competent? This is not a matter of what degrees he has, but of whether he displays and draws on an understanding of the topic and the ways it might be approached. This involves an assessment of the tenability, appropriateness, and use of background assumptions, evidence, and modes of reasoning. Is he conscientious? That is, does he take his epistemic responsibilities seriously in

[1] See Samuel Elgin (2015) and James Lenman (2000) for arguments that show that consequentialists are never in a position to know or reasonably believe that a given action is good. They know that it is good if and only if it maximizes utility. Because causal chains are endless, we cannot know which action satisfies that requirement.

this matter? Has he done his homework? If he neglects information that is readily accessible or overlooks relevant available evidential or logical resources, there is reason to doubt that his disagreement will reveal anything epistemically significant. The requirements of competence and conscientiousness provide sufficient reason to, for example, dismiss the objections of anti-vaxxers, since they either do not understand or do not respect the evidence that has discredited the contention that vaccinations cause autism or implant microchips. Since their position has already been justifiably rejected, unless and until they can provide new, more plausible evidence, their disagreement can, for epistemic purposes, be ignored. On the other hand, an epidemiologist who maintains that this year's flu vaccine will be relatively ineffective because the flu virus has mutated probably should be taken seriously. The basis for rejecting the views of anti-vaxxers is not that those views are kooky or extreme. It is that ample, available evidence discredits them. Should new evidence in favor of their position emerge, the position will merit reconsideration. But the evidence must be genuinely new; it must not be simply a repackaging of considerations that were previously, justifiably found wanting. Similarly, someone who dismisses skeptical arguments out of hand on the grounds that there are no malevolent demons need not be taken seriously. For such a dismissal misunderstands the philosophical function that demon arguments perform. Someone who argues that inference to the best explanation blocks demon arguments deserves a hearing.

Responsible disagreement needs to be relevant and focused. Pliny the Elder is said to have ended every speech he gave in the Roman Senate saying 'Carthage must be destroyed'. For all I know, he was right about Carthage. But there were no doubt many debates in the Senate where his point was irrelevant. If he voiced his disagreement over, say, a plan to construct a temple or fix the pot holes on the Appian Way, on the grounds that it did not address the danger that Carthage posed, then his disagreement with the proposal before the Senate should probably have been disregarded. Similarly, a philosopher obsessed with the hard problem of consciousness who disagrees with every speaker who fails to solve the problem—even if that speaker is discussing truthmaker semantics or the role of beauty in aesthetics—should probably be disregarded. For him insist 'But you haven't solved the hard problem of consciousness!', although true, would not be particularly useful. Unless he can show that to say something important about the role of beauty in aesthetics requires taking a stand on the hard problem of consciousness, the news that a particular aesthetician did not solve or even address the problem is not significant.

Responsible disagreement should be respectful. It should conform to the principle of charity. The reason is not just that doing so would be polite. By giving the strongest available reading of a position before disagreeing with it, you disclose something that cuts to the core of the issue. A disagreement that bears only on a shallow or uncharitable construal does not shed much light on the central issues under discussion. The problem it points to either is peripheral or is easily solved.

Fruitful disagreements illuminate the phenomena, our current understanding of the phenomena, and our current resources for addressing the relevant issues. They reveal something worthwhile that we were unaware of or insufficiently attentive to. They may sensitize us to weaknesses or vulnerabilities in our approach. They prompt re-thinking.

8 Back to our curious practices

This may shed light on the epistemic value of some of philosophy's curious practices. Why don't philosophers easily take 'yes' for an answer? Agreement typically does little to advance understanding. That like-minded folks agree with me may assuage my insecurities, giving me confidence that I haven't made an obvious, stupid mistake. Maybe it affords evidence that I have a relatively stable base to work from. But it does not do much to push my thinking further. It does not point me in any particular direction. A fruitful disagreement, on the other hand, advances understanding. It can highlight aspects of my theory and its relation to the phenomena that I had not properly attended to. Suppose, for instance, I construed scientific knowledge as a conjunction of propositions each of which has a high probability, and construed the growth of knowledge as a matter of adding new conjuncts. Finnur Dellsén disagrees, pointing out that it is a principle of probability theory that the probability of a conjunction of mutually independent claims decreases with each additional conjunct (2019). As a result, each additional conjunct with a probability of less than 1 lowers the probability of the whole. My account, as it stands, faces a serious problem. Maybe I'll decide to scrap it; more likely I will modify it—perhaps by ceasing to construe a scientific theory as a conjunction, or by rejecting the idea that probability is the measure of tenability.

The advantages of fruitful disagreement sometimes also arise in cases of agreement. I said above that agreement among like-minded people does not advance understanding much. Surprising agreement is different. If, for example, a scientific realist and a constructive empiricist find themselves agreeing about the best way to interpret Newton's bucket, or find themselves with the same reasons for being bewildered about Einstein–Podolsky–Rosen, they may come to understand that (and with luck why) the issues raised by these thought experiments transcend their respective stances toward unobservables. Such agreement, because it was unanticipated, discloses something new and potentially epistemically valuable.

Why do we require our students to study the history of philosophy? Why do we turn to our predecessors as often as we do? Not surprisingly, Aristotle provides the answer. Recall his appeal to the many and the wise. In the lead-up to his definition of eudiamonia in the *Nicomachean Ethics*, he lists a number of views and says, 'Now some of these views have been held by many men and men of old, others by a few eminent persons; and it is not probable that either of these should be entirely mistaken, but rather that they should be right in at least one respect or even in most respects' (1941, §1098b25-30). We look to our forbears not merely for inspiration, nor to wholly endorse or wholly reject what came before, but to cull from their accounts ideas that strike us as worth expropriating, at least in part. Although we think they are wrong, we suspect that they are not entirely wrong. Mill concurs. 'When there are persons to be found who form an exception to the apparent unanimity of the world on any subject, even if the world is in the right, it is always probable that the dissentients have something worth hearing to say for themselves, and that truth would lose something by their silence' (1978: p. 46).

A similar point holds for literature reviews. Because philosophers often frame their positions by juxtaposing them with presumptively plausible alternatives, we use peer

disagreement as a way to triangulate. By sketching the current state of play, perhaps with acknowledgement of the historical trajectory that got us here, we map the terrain of the territory within which we will work. We position ourselves by reference to those we disagree with. This would be fruitless if we thought they were dead wrong. In *De Anima* Aristotle says, '[I]t is necessary... to call into council the views of those of our predecessors who have declared any opinion on this subject, in order that we may profit by whatever is sound in their suggestions and avoid their errors' (1941, §403b20-23). We look at the works of our (partly mistaken) predecessors and peers, with an eye to deciding what might be worth importing, perhaps with modifications, into our own theory. Rather than just dismissing our opponents as wrong, we seriously investigate what is wrong with their views. By figuring out where and why we agree and where and why we disagree, we benefit from the efforts of others. We don't then have to reinvent the wheel; nor, if we are lucky, do we fall into the same traps they fell into.

Hiring practices foster understanding in ways that Lewis does not acknowledge. This, I suspect, is because he focused on truth and knowledge, and did not consider how the quest for understanding might be different. Suppose you wanted a department that would advance understanding in metaethics.[2] It would make very good sense to hire Peter Railton—a confirmed moral realist—and Allan Gibbard—an equally confirmed expressivist. The reason is not because you wanted to hedge the department's bets or diversify its intellectual portfolio. Rather it is the expectation that over many years, Railton, Gibbard, and the students they draw to the department will fruitfully disagree. They will push one another to respond to objections they never would have thought of on their own. The department would also have an incentive to hire Liz Anderson, who will raise objections if, for example, the metaethical theorizing departs too far from everyday moral life. Even if one of them arrives at the truth about metaethics, it is doubtful that he would have achieved such a refined, well-supported, comprehensive, and well-defended understanding without the years of fruitful disagreements with colleagues and students. We learn from one another. We benefit from being pushed to rethink, strengthen, and refine our views. In philosophy, the push often comes in the form of disagreement.

I have argued that peer disagreement is an asset in philosophy; it is not something to shy away from. This is reflected in our practices. Still the question arises, when confronted with a case of peer disagreement, should we suspend judgment, lower our credences, or hold fast to our positions?

I've suggested that most philosophers do not actually believe their theories, so strictly speaking, lowering our credences is not an available option. Still, it might seem, we could lower our level of confidence in our network of commitments, which would amount to much the same thing. This sleight of hand will not work. Most of us probably expect that our views will be supplanted. This seems to indicate that whatever we are confident about, it is not that our views, just as they stand, are correct.

Should we then suspend judgment? Similar problems arise here. In being committed to a philosophical position, what are we supposed to be judging? What is it we are supposed to suspend? Acceptance is a matter of being willing to use as a basis for assertoric inference or for action when our ends are cognitive. To suspend acceptance,

[2] This example was suggested to me by Samuel Elgin.

then, is to deprive ourselves of a resource. Sometimes this is entirely reasonable. Once you become convinced, say, that naive set theory leads to paradox, you should withdraw your commitment to it, and refuse to use the contention that sets can take anything as members as a premise in cognitively serious assertoric inferences. In less alarming cases, it may be reasonable to suspend acceptance, but still be willing to use a contention in hypothetical inferences. Perhaps then we should think of philosophical reasoning as largely hypothetical.

The third alternative is to remain steadfast. As it is usually put, remaining steadfast requires not only retaining your commitments, but also believing that your peer has made an error in opposing them. I suggest that we can remain steadfast without making a negative assessment of our peer. On my view, to remain steadfast is to continue to consider one's commitments worthy of being used as a basis for assertoric inference and for action when one's ends are cognitive. It does not require thinking that those who do otherwise are therefore wrong.

Since we already expected disagreement, our level of commitment need not be diminished by finding out that others disagree. What matters is the content of the disagreement. In some cases, it ought to convince us to withdraw a commitment. Russell's objection to Frege's Basic Law V was conclusive (1971). Logicism is untenable. In others, it might prompt serious rethinking. Vogel's car theft case might lead you to rethink your commitment to closure (1990). In yet others, we are justified in holding fast. For now I remain committed to non-factivism about understanding even though a non-trivial number of very smart philosophers disagree with me.

Although I have focused on philosophical practices and philosophical disagreements exclusively, I suspect that my point holds for systematic inquiry more generally. Mill thinks so (1978). He considers disagreement valuable. Except perhaps in mathematics, a large part of understanding a position is knowing what the significant objections to it are, and how to rebut them. If so, in the context of systematic inquiry, responsible disagreement, rather than being grounds for dismay, should be recognized as an epistemic asset.

Acknowledgement I am grateful to Greta Turnbull and participants at the 2018 Ryerson University Philosophy Conference for comments on an earlier draft of this paper, and to two anonymous reviewers for comments. In particular I am grateful for the recommendation that I recognize the ways that Mill's views align with mine.

References

Aristotle. (1941). *The basic works of Aristotle*, R. McKeon (Eds). Random House.

Barnett, Z. (2019). Philosophy without belief. *Mind, 128*, 109–138.

Christensen, D. (2007). Epistemology of disagreement: The good news. *Philosophical Review, 116*, 187–217.

Cohen, L. J. (1992). *An essay on belief and acceptance*. Oxford University Press.

Dellsén, F. (2019). Rational understanding: Toward a probabilistic epistemology of acceptability. *Synthese, 198*, 2475–1494.

Elgin, C. (1996). *Considered judgment*. Princeton University Press.

Elgin, S. (2015). The unreliability of foreseeable consequences: A return to the epistemic objection. *Ethical Theory and Moral Practice, 18*, 459–466.

Elgin, C. (2017). *True enough*. MIT Press.

Elgin, C. (2018). Reasonable disagreement'. In C. R. Johnson (Ed.), *Voicing dissent* (pp. 10–21). Routledge.
Fumerton, R. (2010). You can't trust a philosopher'. In R. Feldman & T. A. Warfield (Eds.), *Disagreement* (pp. 90–110). Oxford University Press.
Goldberg, S. (2013a). Disagreement, defeat, and assertion'. In D. Christensen & J. Lackey (Eds.), *The epistemology of disagreement: New essays* (pp. 167–189). Oxford University Press.
Goldberg, S. (2013b). Defending philosophy in the face of systematic disagreement. In D. Machuca (Ed.), *Disagreement and skepticism* (pp. 277–294). Routledge.
Heydrich, W. (1993). A reconception of meaning. *Synthese, 95*, 77–94.
Hills, A. (2016). Understanding Why. *Noûs, 50*, 661–688.
Kitcher, P. (1990). The division of cognitive labor. *Journal of Philosophy, 87*, 5–22.
Lenman, J. (2000). Consequentialism and cluelessness. *Philosophy and Public Affairs, 29*, 342–370.
Lewis, D. (2000). Academic appointments: Why ignore the advantage of being right? *David Lewis: Papers in ethics and social philosophy* (pp. 187–200). Cambridge University Press.
Mill, J. S. (1978). *On liberty*. Hackett Publishing Co.
Russell, B. (1971). 'Letter to FREGE' *From Frege to Gödel*. In J. Van Heijenoort. Harvard University Press, (pp. 124–125).
van Fraassen, B. (2008). *Scientific representation*. Clarendon Press.
Vogel, J. (1990). Are there counterexamples to the closure principle? In M. D. Roth & G. Ross (Eds.), *Doubting* (pp. 13–27). Kluwer.