Canadian Philosophical Association 2019 Prize Winning Papers

Peer Idealization and Internal Examples in the Epistemology of Disagreement

TIM KENYON *Brock University*

ABSTRACT: The epistemology of disagreement has developed around a highly idealized notion of epistemic peers. The analysis of examples in the literature has somewhat entrenched this idealization, when using cases of extant philosophical disputes between named interlocutors. These examples make it hard to emphasize the ordinary ways in which discussants, as disciplinary colleagues, may be wrong. Overlooking these possibilities is probably made easier by widespread attitudes in philosophy about the importance of genius or raw intelligence in doing philosophy. The use of such internal examples needlessly limits consideration of the full range of epistemically relevant features of disagreements.

RÉSUMÉ : L'épistémologie du désaccord s'est développée autour d'une notion idéalisée de pairs épistémiques. L'analyse d'exemples dans la littérature a quelque peu enraciné cette idéalisation, surtout lorsque les exemples étudiés sont des désaccords tirés du canon philosophique contemporain et qu'ils opposent des interlocuteurs identifiés. Il est difficile, pour des raisons socio-professionnelles, de souligner les manières ordinaires par lesquelles ces collègues disciplinaires peuvent se tromper. Il est probablement d'autant plus facile de négliger ces possibilités que les attitudes concernant l'importance du génie dans la pratique de la philosophie sont fort répandues dans le domaine. L'utilisation de ces exemples limite donc la prise en compte de toutes les caractéristiques épistémiquement pertinentes des désaccords.

Keywords: disagreement, social epistemology, meta-philosophy, epistemology of disagreement, epistemic peers, expertise

***Dialogue* 59 (2020), 69–79**

doi:10.1017/S0012217319000374

I. Smartness and Peerhood in the Epistemology of Disagreement

> Do you get worried when you discover that some of your philosophical colleagues disagree with your philosophical views? What if you know these colleagues to be *smarter*, and, with respect to the relevant issues, *better informed* than you?[1]

The interesting cases of disagreement in contemporary epistemology of disagreement (EoD) are generally taken to be those arising between *epistemic peers*. Peerhood is sometimes presented in terms of an equiprobability of making true assertions regarding a given domain.[2] More often, it includes some specific reference to the possession of the same evidence, and equality of reasoning capacities relative to that evidence.

> In the stipulative sense of "peer" introduced, peers literally share all evidence and are equal with respect to their abilities and dispositions relevant to interpreting that evidence. Of course, in actual cases there will rarely, if ever, be exact equality of evidence and abilities. This leaves open questions about exactly how conclusions drawn about the idealized examples will extend to real-world cases of disagreement.[3]

Epistemic peerhood is an idea critical for contemporary EoD because it defines a puzzle or family of puzzles that has animated a considerable proportion of that literature. The puzzle is that of what to believe in a case where agents standing in an important sort of epistemic symmetry relative to some area of judgement nevertheless disagree about an issue falling within that area. Thinking in terms of epistemic peers could be a heuristic that helps us reason creatively about what, if anything, distinguishes our epistemic position from that of our interlocutors — it could help us to think about symmetry-breakers.[4] But that's rarely how it's been used. In fact, EoD has tended not only to idealize epistemic peerhood in the ways noted by Richard Feldman and Ted Warfield, but to treat the satisfaction of known (sometimes *revealed* or *acknowledged*) epistemic peerhood as a background condition for the core phenomenon of interest. Jointly this raises the prospect that the core phenomenon of interest to EoD is practically non-existent.

There is a lot to say about epistemic peerhood, and I am not going to attempt to say it all here. It is enough to note that the idealized nature of epistemic peerhood in the EoD literature has often been flagged as a concern, even if only to be set aside in many instances.[5] My interest in the following remarks is to

---

1 Bryan Frances, 2013, pp. 121-122.

2 Amir Konigsberg, 2013, p. 97.

3 Richard Feldman and Ted Warfield, 2010, p. 2.

4 Tim Kenyon, 2018, p. 241.

5 For example, in addition to the Feldman and Warfield remark, see David Christensen 2009; and Nathan King, 2012.

https://doi.org/10.1017/S0012217319000374 Published online by Cambridge University Press

follow up on one surprisingly common way to philosophize about peer disagreement through examples. After all, one might expect that looking at examples would be a reliable way to moderate the degree of idealization involved in theoretical discussions of peerhood. Dealing with examples should enable a rich, empirically informed analysis of the many, many ways that symmetry is broken — perhaps inevitably broken — between approximate epistemic peers in actual, messy cases. This is why it strikes me as significant that a surprising number of those examples of peer disagreement are themselves drawn from within contemporary philosophical debates.

I will argue that this choice of *internal examples* is apt to entrench, rather than mitigate, a hyper-idealized notion of epistemic peers. These examples skew discussions away from a thorough consideration of the prospects for even experts to be not just mistaken, but outright unreasonable, and the surprisingly fine-grained ways in which expertise can succeed or fail. I submit that underlying meta-philosophical views about the significance of smartness or genius in doing philosophy well are a factor in making it seem natural to recruit internal examples. But that is an unhappy way of thinking about both philosophical expertise and the philosophical project; and internal examples are a foreseeably distortive influence on our theorizing about EoD.

II. Peerhood and Sudden Lapses of Expert Knowledge

First, just a bit more context on peerhood. If we take known or revealed peerhood to define the core cases of epistemically significant disagreement, then we will find it very hard to operationalize the notion in actual cases. In part, this is due to the variations of definition of epistemic peerhood, into which weeds for simplicity's sake I will not descend, and the related difficulties of making precise such notions as 'all the same evidence,' and 'equality of abilities.' Partly, too, it is because the role of epistemic peerhood in the literature often presupposes that we have a way of carving up propositions into domains of disagreement, to allow that a seeming peer with respect to a domain can remain a peer despite being *prima facie* wrong about some assertion that falls within that domain. And this must hold even though being wrong about one thing in the domain is virtually certain to involve being wrong about other closely related things in it, and hence is suggestive of some non-trivial disagreement over elements of the domain. If we disagree about whether olive oil or sunflower oil is better for frying up hash browns, we are at least pretty likely to disagree about which tastes better, what the appropriate maximum frying temperature is, and so forth. Noting this widening circle of potential disagreements, how obvious is it that we are strictly epistemically *symmetrical* with respect to the domain of fry-cooking? The raw materials for disagreements ramify pretty freely, so an account is owing of the degree of differences in domain-specific belief consistent with peerhood within that domain.

But partly, too, peerhood is a vexed notion just because it's really hard to individuate domains of knowledge and predict what false, or even frankly weird, things an expert can believe at the margins of such a domain.

To see this problem in concrete terms, consider prospective disagreements between me and Nobel Laureate Linus Pauling. (Setting aside that he is no longer living, that is.) Pauling was a brilliant award-winning molecular biologist, whereas I am … well, *made out of* biological molecules. It's no surprise that we wouldn't be peers on disagreements in that field. But it might be surprising to consider the full reason for this. Pauling's knowledge and judgement are likely superior to mine on every molecular biology topic — *except* for the effects of large ascorbic acid doses on human health. On that topic, my knowledge and judgement are relevantly superior to his. Pauling became obsessed with what he saw as the health benefits of mega-doses of Vitamin C; but his claims had the character of zealotry even at the time, and have not stood up to empirical investigation. I've read through some of the studies regarding the effects of ascorbate mega-doses, and, though no expert myself, I have the cognitive and affective capacity for appropriate uptake of their key points. Pauling, it seems, had some motivated reasoning or other intellectual barrier that prevented his taking up such information. So he would be my epistemic superior regarding the discipline quite generally, but I would be his superior regarding at least one fairly thinly sliced topic within the discipline (setting aside the problem of just how thinly peer disagreements may be sliced).[6]

Similarly, although Andrew Wakefield is no longer licensed to practice medicine due to his utterly fraudulent medical research claiming a link between vaccinations and incidence of autism in children, he did qualify as a doctor after completing medical school, and was for a long time a Fellow of the Royal College of Surgeons. By contrast, I have watched some episodes of *House*, and once gave chest compressions to a mannequin. If Wakefield and I disagreed about the presentations of cardiac arrest or symptoms of a hernia, we would not be epistemic peers; he would by far be my superior. Between the two of us, he's the one you would want sitting next to you if you started choking on a gummy bear. But if we disagreed (and we do) about whether fear of autism should lead you not to vaccinate your child or yourself, I am without question his epistemic superior. I would give the consensus expert opinion — which I am emotionally willing and cognitively able to incorporate into my psychological economy, at an educated layperson's level — while Wakefield, unable or unwilling to do so, would say things about vaccination that virtually all medical researchers regard as sheer nonsense.

It is merely psychologically odd, and not theoretically paradoxical or puzzling, that Pauling was a great biochemist, but deeply unreliable on the biochemistry of ascorbic acid in human aging processes. And I can reason clearly and write frankly about the symmetry-breaking factors explaining my latent disagreement

[6] See Nathan Ballantyne, 2019, for the view that Linus Pauling was an 'epistemic trespasser,' making claims outside his discipline, with respect to the functions of ascorbate. I think mine is a permissible individuation of disciplinary boundaries.

with him. But that's largely because Pauling is deceased, and he and I don't share a discipline, and he's sort of a public figure. By contrast, when EoD has focused on applied cases, it has often resorted to examples of philosophical disagreements between named interlocutors, often living, often involving the author of the example personally. I call these *internal example*s because they are disagreements internal to the discipline considering the significance of disagreement. Here, I use the term to mean the richer notion of internal disagreements just canvassed: not merely disagreements internal to philosophy as an entire discipline and history — say, Bertrand Russell disagreeing with Thomas Aquinas — but disagreements between named living or recent living peers.

The use of internal examples makes EoD less generally applicable, and less theoretically robust, because it encourages silence about factors crucial to our understanding of epistemic peerhood and the significance of expertise.[7] To some extent, this is just because case studies of philosophers talking to philosophers about philosophers talking to philosophers are not very obvious first steps towards generality. More importantly, though, these examples comprise a methodology that skews the analyses of the phenomena, discouraging empirically open contextual and theory-based explanations of disagreement that would be available for less professionally and personally entangled cases.

Candidate explanations for disagreement between rough and ready peers ought to include epistemic asymmetries arising from such factors as ignorance, bias, and self-deception, perhaps specific to the case at hand. But the idealizing spirit and fondness for puzzles that drive the literature towards considering strictly symmetrical scenarios as the defining cases of philosophical interest get a boost from the assumption that some actual philosophical disputes count as instances of this symmetry (and sometimes even purport to find an asymmetry in the counterintuitive 'wrong' direction, creating a still sharper sense of puzzlement). This idea, that philosophical interlocutors count as sufficiently strict peers on the topic in question to render the disagreement itself theoretically interesting, is facilitated in part by professional courtesy, personal modesty, and even friendship, rather than open empirical imagination about the cases.

Of course, from a collegiality perspective, it's quite right that professional courtesy and so forth should influence the conjectures one is willing to publish regarding colleagues' inclinations to believe particular things one regards as demonstrably mistaken. But an unconstrained set of such symmetry-breakers is essential to robust reasoning and fertile theory-building about the significance of disagreements in general. I don't hold that philosophers should be

[7] For current purposes, I set aside the prospect of scepticism regarding philosophical expertise and the benefits or differential abilities it might confer (for example, Jennifer Nado, 2014; Jimmy Alfonso Licon, 2012; Eric Schwitzgebel and Fiery Cushman, 2012).

less pro-social in their treatment of examples, but that EoD should choose examples so that pro-social treatment does not risk distorting the theoretical project.

III. Internal Examples as Vexed Empirical Method

Once internal examples are taken to involve unimpeachable interlocutors, the upshot tends to be diagnostic perplexity as to how such agents can persist in holding opposing views.[8] Thus, Peter van Inwagen is driven to extreme lengths to explain why David Lewis was unpersuaded by his arguments for incompatibilism about the will.

> Consider, for example, the body of public evidence that I can appeal to in support of incompatibilism (arguments and other philosophical considerations that can be expressed in sentences or diagrams on a blackboard or other objects of intersubjective awareness). David Lewis "had" the same evidence (he had seen and he remembered and understood these objects) and was, nevertheless, a compatibilist. If I know, as I do, that David had these features (and this feature, too: he was a brilliant philosopher), that he had these features is itself evidence that is (or so it would seem to me) relevant to the truth of incompatibilism. Should this new evidence not, when I carefully consider it, lead me to *withdraw* my assent to incompatibilism, to retreat into agnosticism on the incompatibilism/compatibilism issue?
>
> … The difficulty of finding anything to say in response to this argument, taken together with my unwillingness to concede either that I am irrational in being an incompatibilist or that David was irrational in being a compatibilist, tempts me to suppose that I have some sort of interior, incommunicable evidence (evidence David did not have) that supports incompatibilism.[9]

Van Inwagen is driven to propose an *incommunicable* body of evidence, possessed by him and lacked by Lewis, to make sense of his persisting conviction that he is right and Lewis is wrong. Why? Here's one reason: so many other less exotic explanations are denied us by the choice of example.

For example, why move so quickly past the possibility that the relevant evidence was not merely communicable, but already communicated? Maybe it was just under-appreciated by Lewis upon uptake, initial conditions that gave rise to reasoning that was slightly path-dependent thereafter. Or maybe Lewis was, on this specific question, very slightly less reflectively thorough than van

---

[8] Thomas Kelly, 2005, p. 193, does say this much, to be fair: "Although I tend to find it somewhat unsettling that many disagree with my view, I am inclined to regard this psychological tendency as one that I would lack if I were more rational than I in fact am."

[9] Peter van Inwagen, 2010, p. 26.

Inwagen. These prospects are left unsaid as a matter of social grace and collegiality.[10]

Nor do *we*, as critical audience to the example, see van Inwagen taking very seriously the prospect that other potentially controversial commitments of his own are implicated in the disagreement. What philosophical labour might van Inwagen want his theory of free will to do, and how could this be motivating his own reasoning? He raises the example in a paper about religious disagreement, after all. A theory of praise and blame; a theodicy; an apologetic for the Problem of Evil; a particular conception of persons — these or other open questions on which van Inwagen has preferred views might have him resting a thumb on the scales of his rational evaluation of the free will disagreement with Lewis, in just the way that anyone would be unlikely to notice from their own perspective.

These are clear possibilities, bright with implication, from the perspective of an empirically informed diagnosis of specific disagreements. The choice of example makes these possibilities hard to entertain, though, and harder still to discuss collegially. The case, van Inwagen writes, is one "in which neither philosopher labors under the burden of any cognitive deficiency from which the other is free. I know that David labored under no such deficiency. I like to think that I do not."[11] Notice, though, that these are *empirical* claims of freedom from relevant biases. And as empirical methodology, it is wretched: the impossibly blunt notion of cognitive deficiency; the gallant assumption of freedom from any such deficiency in one's colleague; the similarly optimistic assessment of oneself. Yet having raised the question of cognitive deficiencies, broadly construed, how else could van Inwagen possibly have handled it in the context of an internal example?

Catherine Elgin also uses an internal example of a disagreement with a named philosopher in sketching her own view in EoD. Again, it is a disagreement with David Lewis, this time over the reality of possible worlds. Elgin describes Lewis as "incredibly smart, philosophically gifted, and intellectually responsible," adding that

> [i]t is no false modesty for me to say that David Lewis was a far better philosopher than I am. Nevertheless, I think he was wrong. I cannot refute his position; it is admirably well-defended. But, despite Lewis's intelligence and arguments, I do not believe that there exist real possible worlds … But David Lewis thought otherwise. He was not my epistemological peer; he was my epistemological superior. So should I not revise my opinion to agree with him?[12]

---

[10] I am, of course, not saying that van Inwagen wrote this *merely* as a matter of showing social grace; I expect it is a sincere but mistaken thought that the evidence at his disposal bears out a judgement of true symmetrical peerhood on the communicable elements of the free will debate.

[11] Van Inwagen, 2010, p. 26.

[12] Catherine Elgin, 2010, p. 58.

https://doi.org/10.1017/S0012217319000374 Published online by Cambridge University Press

Elgin demurs from van Inwagen's incommunicable evidence hypothesis, writing, "I cannot speak for van Inwagen. But speaking for myself, I think it is exceedingly unlikely that I enjoy any sort of philosophical insight that Lewis lacked."[13] This observation helps to motivate Elgin's *doxastic involuntarist* perspective on disagreement. According to Elgin, she doesn't conciliate her view with Lewis's simply because she can't *choose* to do so. If she could so choose, then rationally she would have to conciliate, because Lewis is at least her epistemic peer, and probably her epistemic superior on the question.

IV. Smartness Disguises Idealization in Internal Examples

In Elgin's case, too, I submit that irrelevant features of the example are apt to distort the theorizing drawn from it. But now we can perhaps start to see something more general at work, something characteristic of the discipline itself. Part of what is happening in examples of this sort, I suggest, is the unusual tendency of academic philosophy to be heavily driven by assumptions about the role of a monolithic philosophy-specific intelligence, or even *genius*, in doing philosophy well.[14] The idea that philosophical insights flow from pure philosophical intellectual horsepower (plus scholarly knowledge, of course) makes it much less natural to think that someone could deliver amazing insights on some philosophical issues, and yet be significantly wide of the mark on closely related issues. Because how could undifferentiated philosophical smartness or genius be so unpredictably and fine-grainedly gappy in its effects? But this conception of philosophical insight, and the associated conception of what it is to *do philosophy* more generally, are deeply flawed metaphilosophical over-simplifications, and here they are spilling over into social epistemology methodology.

Of the many things that comprise doing philosophy, it might be that Lewis did some better than Elgin (thus far). Maybe there's a notion of 'better' that would allow us to say this of the data. I am confident, though, that Elgin does some of the many things that comprise doing philosophy better than Lewis did. Just think of the number of things that comprise doing philosophy, and the number of different ways of doing each of those things felicitously or effectively! Philosophy is too many things — topics, sub-topics, problems and subproblems, scholarly bodies of knowledge and cross-classified interdisciplinary bodies of knowledge, research techniques, discursive techniques, application techniques, cognitive skills, social skills, linguistic skills, teaching skills, emotional skills, abilities to track formal implications, abilities to recognize analogical patterns, abilities to perceive global significances, and meta-skills yoking all this stuff together — for any one competent and experienced philosopher to do uniformly better in all domains and at all times than another competent and

[13] Elgin, p. 59.

[14] Sarah-Jane Leslie et al., 2015.

experienced philosopher. Elements of any of these could inform the credences adopted on one side or the other in a specific extremely fine-grained philosophical disagreement, taken in context. So blanket claims about who is stronger in The Force in some blunt sense, and therefore who is epistemically superior overall for current purposes, are *at best* unrevealing.

Moreover, Elgin does think that she enjoys a philosophical insight that Lewis lacked: she thinks he was just plain wrong about possible worlds. And this is — in the relevant sense! — much like my thinking that Pauling was just plain wrong about Vitamin C, even though he knew a lot more about biochemistry than I do; or, for that matter, that Wakefield is just plain wrong about vaccines, even though he trained as a doctor and was for many years a Fellow of the Royal College of Surgeons.[15] Again, people who are generally epistemically reliable, even expert, in a topic at one level of granularity can be definitely, even notoriously, wrong on some matter at a more fine-grained level within that field. Even luminous, gifted scholars can be victims of their own biases, outsized commitments, failures of attention, and ignorance when it comes to specific propositions or issues falling squarely within their fields of expertise — perhaps for broadly holistic reasons stemming from beliefs or commitments far beyond their fields of expertise. I'm not at pains to say this is true of Lewis, regarding any of his views in particular. I do say, though, that we cannot do sound social epistemology by ruling out the prospect, or depicting it as mysterious, via general ascriptions of genius or concomitant philosophical ability.

Whether this sort of highly local reversal or failure of expertise is happening to me or to my interlocutor in a disagreement is a critical consideration in real contexts. Our theorizing about disagreement must take seriously the factors that indicate failures of otherwise fairly general intellectual excellence, helping explain them rather than constructing them as paradoxical or puzzle-defining. Theory-informing examples that discourage our reviewing those factors explicitly are barriers to sound theorizing. In practice, this means that examples of disagreements in philosophy, between named living (or recently living) philosophers, are liable to distort our thinking about EoD.

Even to discuss the matter in this way, naming names and considering specific people, has an unmannerly feel to it — perhaps not in the original, very polite examples, but once we abandon the imperative to attribute unbroken expert reliability and practical immunity to bias to the people in question. But we do have to talk about the cases because they are influential examples leveraged to some effect in the literature. We must reflect realistically on the prospect that the interlocutors in these cases, their usual remarkable expertise notwithstanding, could be wrong for entirely unremarkable, even potentially embarrassing, symmetry-breaking reasons — the kind of reasons for which we

[15] Naturally there are many key differences between these cases as well.

https://doi.org/10.1017/S0012217319000374 Published online by Cambridge University Press

know experts can be wrong about things *all the time*. In thinking about disagreement and its epistemological significance, we must be able to speak directly of the prospects of ignorance, bias, carelessness, insensitivity to evidence, or otherwise infelicitous reasoning on the part of otherwise reliable or expert agents. Examples that discourage that sort of open diagnostic discussion predictably lead to over-idealized theorizing.

Acknowledgements: I am grateful to the following people for their helpful conversations or comments on various versions of this work, including the longer piece (2018, in the References) from which its ancestor was excised for reasons of length. My thanks to Nathan Ballantyne, James Chase, Shannon Dea, Mathieu Doucet, David DeVidi, Catherine Elgin, Kirk Lougheed, Dima Mohammed, Erich Rast, and participants in the 2017 ArgLab Research Colloquium, Universidade Nova de Lisboa. Alex Worsnip in particular gave very helpful remarks in his commentary on the paper at the 2019 Canadian Philosophical Association Annual Meeting.

References

Ballantyne, Nathan
2019 "Epistemic Trespassing," *Mind* 128 (510): 367–395.

Christensen, David
2009 "Disagreement as Evidence: The Epistemology of Controversy." *Philosophy Compass* 4 (5): 756–767.

Elgin, Catherine
2010 "Persistent Disagreement," in *Disagreement*, edited by Richard Feldman and Ted Warfield. Oxford: Oxford University Press, pp. 53–68.

Feldman, Richard, and Ted Warfield
2010 "Introduction," in *Disagreement*, edited by Richard Feldman and Ted Warfield. Oxford: Oxford University Press, pp. 1–9.

Frances, Bryan
2013 "Epistemic Renegades," in *The Epistemology of Disagreement*, edited by David Christensen and Jennifer Lackey. Oxford: Oxford University Press, pp. 121–166.

Kelly, Thomas
2005 "The Epistemic Significance of Disagreement," in *Oxford Studies in Epistemology*, *Volume 1*, edited by Tamar Szabó Gendler and John Hawthorne. Oxford: Oxford University Press, pp. 167–196.

Kenyon, Tim
2018 "Disagreement," in *The Routledge Handbook of Applied Epistemology*, edited by David Coady and James Chase. New York: Routledge, pp. 233–246.

King, Nathan
2012 "Disagreement: What's the Problem? *or* A Good Peer Is Hard to Find." *Philosophy and Phenomenological Research* 85 (2): 249–272.

Konigsberg, Amir

2013 "The Problem with Uniform Solutions to Peer Disagreement." *Theoria* 79 (2): 96–126.

Leslie, Sarah Jane, Andrei Cimpian, Meredith Meyer, and Edward Freeland

2015 "Expectations of Brilliance Underlie Gender Distributions across Academic Disciplines." *Science* 347 (6219): 262–265.

Licon, Jimmy Alfonso

2012 "Sceptical Thoughts on Philosophical Expertise." *Logos and Episteme* 3 (3): 449–458.

Nado, Jennifer

2014 "Philosophical Expertise." *Philosophy Compass* 9 (9): 631–641.

Schwitzgebel, Eric, and Fiery Cushman

2012 "Expertise in Moral Reasoning? Order Effects on Moral Judgment in Professional Philosophers and Non-Philosophers." *Mind & Language* 27 (2): 135–153.

Van Inwagen, Peter

2010 "We're Right. They're Wrong," in *Disagreement*, edited by Richard Feldman and Ted Warfield. Oxford: Oxford University Press, pp. 10–28.