

WHO IS AN EPISTEMIC PEER?

Axel GELFERT

ABSTRACT: Contemporary epistemology of peer disagreement has largely focused on our immediate normative response to *prima facie* instances of disagreement. Whereas some philosophers demand that we should withhold judgment (or moderate our credences) in such cases, others argue that, unless new evidence becomes available, disagreement at best gives us reason to *demote* our interlocutor from his peer status. But what makes someone an epistemic peer in the first place? This question has not received the attention it deserves. I begin by surveying different notions of ‘epistemic peer’ that have been peddled in the contemporary literature, arguing that they tend to build normative assumptions about the correct response to disagreement into the notion of peerhood. Instead, I argue, epistemic peerhood needs to be taken seriously in its own right. Importantly, for epistemic agents to count as peers, they should exhibit a comparable degree of reflective awareness of the character and limitations of their own knowledge.

KEYWORDS: epistemic peers, peerhood, disagreement, ignorance

1. Introduction

Recent years have seen an explosion of the epistemological literature on peer disagreement. How should we, as individuals, react when we encounter someone who disagrees with us on a point of fact, yet who we have every reason to believe is our epistemic equal? Disagreement is as much part of our everyday epistemic lives as reliance on others for knowledge, and thus merits attention from epistemologists – perhaps especially so at a time when traditional sources of institutional testimony (experts, media, higher education, science, etc.) are widely perceived as becoming more diverse and polarized.

In trying to understand this phenomenon, epistemologists of disagreement standardly operate with an idealized situation, in which two equally well-informed ‘peers,’ on the basis of the same evidence, come to different conclusions regarding a specific subject matter. At issue is the question of whether *mere* acknowledgment of disagreement is sufficient to require us to abandon belief, or at least to revise our corresponding degree of belief downwards. One prominent view, the *equal weight* (EW) view, holds that two epistemic peers, upon learning that they assign different credences to a given proposition *p*, should revise their

credences to be ‘roughly equal’¹; in other words, two epistemic peers who find themselves disagreeing on whether p should ‘split the difference’ and update their respective credences so as to reflect the average of the two initial credences. The equal weight view has been subject to severe criticism, with some critics arguing that, in ‘splitting the difference,’ we are in effect discarding whatever initial evidence we had – ‘pre-disagreement,’ as it were – for our original belief. As one critic puts it, this “does not seem to be a virtue in an epistemological theory.”² On what has been called the *right-reasons* (RR) view, we owe it to our initial reasoning (which led us to conclude that p) that we regard our epistemic peer, who claims that not- p , as “having gotten things wrong” on this occasion.³ Rather than revising our own credence, we treat the fact of disagreement essentially as a reason for *demoting* our disputant from his position as epistemic peer – after all, we already know that one of us must have incorrectly assessed the force of the evidence, and for all we know it is our disputant who is mistaken.

Such, in general outline, is the shape of the debates that have placed peer disagreement at the heart of recent epistemology. Curiously, however, whereas there has been no shortage of discussions of the phenomenon of *disagreement*, comparatively little attention has been paid to a systematic analysis of what makes someone an *epistemic peer* in the first place. The present paper attempts to fill that gap. In Section 2, I survey the notions of ‘epistemic peer’ that have been peddled, often implicitly, by contributors to the epistemology of disagreement. In Section 3, I challenge two key assumptions concerning the character, and alleged frequency, of peer disagreement. Section 4 argues that determining epistemic peerhood is a more important task than adjudicating between different responses to the mere fact of disagreement; importantly, epistemic peerhood comes in degrees. In Section 5, I argue that, beyond being equally knowledgeable, epistemic peers should also display a similar degree of reflective awareness of the limitations of their own knowledge. The paper concludes by suggesting a possible refocusing of the philosophical debate on disagreement.

¹ David Christensen, “Epistemology of Disagreement: The Good News,” *The Philosophical Review* 116 (2007): 193.

² David Enoch, “Not Just a Truthometer: Taking Oneself Seriously (but not Too Seriously) in Cases of Peer Disagreement,” *Mind* 119 (2010): 969.

³ Adam Elga, “Reflection and Disagreement,” *Noûs* 41 (2007): 486. Note that Elga is here *characterizing*, not *endorsing*, the right-reasons view.

2. Epistemic peerhood: the received view

Lack of unanimity is a common experience in social interactions, and in everyday language we readily help ourselves to the term ‘*disagreement*’ and its cognates. By contrast, the expression ‘epistemic peer’ is a technical term, invented by philosophers in order to bring into sharper focus a set of theoretical questions. Interestingly, the term makes its first appearance in the philosophy of religion: If there was only isolated dissent about the (alleged) proper basicity of such claims as “God exists,” we might rationally dismiss disagreement about a proposition, but, as Gary Gutting argues, “the disagreement of substantial numbers of those who, as far as I can tell, are my *epistemic peers* (i.e. my equals in intelligence, perspicacity, honesty, thoroughness, and other relevant epistemic virtues) is surely another matter.”⁴ Jonathan Kvanvig, by contrast, insists that “if disagreement is important to justification, it cannot be disagreement among epistemic peers” since one of the peers might, for purely contingent reasons, be “in a better epistemic situation.”⁵

Early uses of ‘epistemic peer’ rely on an understanding of peerhood that stresses parity with respect to general epistemic virtues, as enumerated by Gutting. This contrasts with the dominant view in the debate, which adds the requirement that, for two epistemic agents to count as peers in a factual dispute, they must be “equals with respect to their familiarity with the evidence and arguments which bear on that question.”⁶ In other words, peers must have “been exposed to the *same* evidence and have worked on it comparably long, carefully, etc.”⁷ It seems reasonable to assume that, barring special conditions, such sweeping similarity in epistemic outlook will lead to epistemic peers being similarly reliable as sources of information. Indeed, this is sometimes seen as the defining feature of epistemic peerhood, as when David Enoch defines an epistemic peer as “someone who is, somewhat roughly, antecedently as likely as you are to get things right (on matters of the relevant kind)”⁸; others have argued for the same view precisely “because it generates the puzzles about disagreement” without entailing any substantive

⁴ Gary Gutting, *Religious Belief and Religious Skepticism* (Notre Dame: University of Notre Dame Press, 1982), 83.

⁵ Jonathan Kvanvig, “The Evidentialist Objection,” *American Philosophical Quarterly* 20 (1983): 54.

⁶ Thomas Kelly, “The Epistemic Significance of Disagreement,” *Oxford Studies in Epistemology* 1 (2005): 174.

⁷ Bryan Frances, “The Reflective Epistemic Renegade,” *Philosophy and Phenomenological Research* 81 (2010): 424. In the contemporary debate, even ‘slight differences’ in evidence are seen as undermining peerhood; see, for example, Nathan King, “Disagreement: What’s the Problem? or A Good Peer is Hard to Find,” *Philosophy and Phenomenological Research* (2011, Early View: doi: 10.1111/j.1933-1592.2010.00441.x): 13.

⁸ Enoch, “Not Just a Truthometer,” 956.

similarities between the agents' evidence or aptitude.⁹ Thus, what started off as an attempt to identify agents with broadly comparable epistemic virtues (though possibly different exposure to the evidence), gradually became assimilated to a probabilistic point about the relative reliability of epistemic agents as sources of information.

At the beginning of this section, I noted that 'epistemic peer' is an epistemological term of art. As such, any definition is bound to have a stipulative element. Nowhere is this stipulative character more evident than in Adam Elga's – at the time, by his own admission, "nonstandard"¹⁰ – definition of (perceived) peerhood in terms of our (pre-disagreement) assessment of how likely each agent is to be mistaken. On this view, "you count your friend as an epistemic peer with respect to an about-to-be-judged claim if and only if you think that, conditional the two of you disagreeing about the claim, the two of you are equally likely to be mistaken."¹¹ While this obviously sits well with the equal weight view – of which Elga is a leading proponent – others have questioned why determinations of peerhood should be concluded *before* a claim is about to be judged: Why "exclude from one's conditionalization process the disagreement itself as reason for demoting" our interlocutor from his position as one's peer?¹² In the remainder of this paper, I shall argue that rather than building normative assumptions about the correct response to disagreement *into* our definition of 'epistemic peer,' we need to disaggregate both phenomena and develop a richer notion of epistemic peerhood.

3. Acknowledgment and the paucity of peer disagreement

Why worry about disagreement in the first place? The intuitive reason is clear: Instances of disagreement are, at least *prima facie*, good occasions to reflect on the fallibility of the methods by which we acquire beliefs (including the belief at issue). It would seem, then, that upon encountering a disagreement with a peer, we should engage in some epistemic soul-searching with the aim of identifying weak links in our reasoning processes and eliminating possible sources of error. But note that this is not how contemporary epistemologists of disagreement typically approach the problem: Their concern is with the *pro tanto* epistemic reason given to us by the brute fact of disagreement alone, irrespective of any

⁹ Alex Bundy, "In Defense of Epistemic Abstemiousness," *Logos & Episteme* 2 (2011): 288.

¹⁰ Elga, "Reflection and Disagreement," 499.

¹¹ Elga, "Reflection and Disagreement," 499.

¹² Enoch, "Not Just a Truthometer," 977.

changes that we would make to our belief systems, were we to independently reassess our individual processes of belief formation.¹³

3.1. Acknowledgment and the fact of disagreement

For the fact of disagreement to take on a compelling degree of urgency, it must be acknowledged by at least one party to the dispute.¹⁴ If both peers were blissfully unaware of their doxastic differences, the mere counterfactual observation that they *would* disagree with one another *if* quizzed on the point in question, can hardly be expected to have an effect on their beliefs and credences. (Indeed, it is doubtful whether we could legitimately speak of a ‘disagreement’ at all in this case.) Exactly what follows from the acknowledgment of the fact of disagreement, is of course the bone of contention among contemporary epistemologists of disagreement. For the defender of the EW view, acknowledgment of the fact of disagreement – combined with the recognition (taken for granted by the EW view) that our interlocutor is our epistemic peer – immediately requires us to adjust our credence downwards and ‘split the difference.’ By contrast, the RR view considers recognition of the fact of disagreement sufficient for demoting our interlocutor from his position as peer; it takes more than *mere* disagreement to sway what was, after all, our initial considered judgment.

3.2. The paucity of (acknowledged) peer disagreement

Both EW and RR theorists agree on the crucial role of acknowledgment and the purely auxiliary function of judgments of epistemic peerhood – where the latter are seen either as *preceding* the disagreement (as in the EW view) or as *being superseded* by it (as insisted by the RR view). At the same time, parity – both in terms of general epistemic virtues and concerning equal exposure to (and consideration of) the evidence – does remain an important background assumption: After all, disagreements with obvious epistemic inferiors would hardly inspire the sense of urgency that fuels the debate about disagreement among peers. Yet it is worth questioning just how often we do, in fact, find ourselves in genuine situations of *acknowledged* peer disagreement.¹⁵ Parity in terms of the reliability of getting it

¹³ On this point, see also Enoch, “Not Just a Truthometer,” 957.

¹⁴ Other authors prefer to speak of ‘revealed’ disagreement, e.g. the editors in their Introduction to the volume *Disagreement*, ed. Richard Feldman and Ted A. Warfield (Oxford: Oxford University Press, 2010), 4.

¹⁵ This is not to say that acknowledged peer disagreement, even if found to be rare, does not pose a genuine theoretical challenge, but only that the contemporary debate derives much of its force from the alleged *pervasiveness* of peer disagreement.

right on a specific point – say, the question of whether p – requires considerable fine-tuning of the specific epistemic histories of both agents (in terms of their overall experience, relevant background beliefs, etc.), which may often outweigh similarities in general epistemic virtues. As Nathan King notes, “the path toward equal reliability is not straightforward,”¹⁶ and even subtle differences in the dispositional response to evidence undermine the fine-grained notion of peerhood as ‘equal familiarity with the evidence’ and with relevant arguments.¹⁷ Moreover, many longstanding disputes – for example in science – are driven by (often axiological) disagreements about which findings should count as evidence in the first place¹⁸; in such cases, “equal familiarity with relevant evidence and arguments” as a criterion for determining peerhood is itself under dispute, and parties to such disputes – no matter how objectively well-acquainted with the facts – will rarely acknowledge each other as epistemic peers.

4. The primacy of peerhood

What the preceding discussion suggests is that determining whether two disputants are indeed epistemic peers – and, if they are not, identifying in what ways their relationship falls short of peerhood – has primacy over the question to what extent the mere fact of disagreement offers a *pro tanto* reason for each party to adjust their credences. In this section, I wish to take a first step towards developing a richer notion of epistemic peerhood, before arguing, in the next section, for a plausible connection between peerhood and questions of epistemic value, which so far appear to have been ignored in the philosophical debate about disagreement. For now, I wish to argue that our understanding of peerhood can be enriched by recognizing that peerhood comes in degrees.

Early definitions of ‘epistemic peerhood’ lacked specificity, insofar as they emphasized the overall similarity in epistemic character, while discounting the important influence of contingent differences (e.g. in exposure to relevant evidence). By contrast, recent attempts to reduce peerhood to *mere equal likelihood to be right on a particular occasion* are overly narrow, insofar as they offer no guidance as to how best to judge whether someone is an epistemic peer or not. Instead of oscillating between these two extremes, what is needed is a ‘middle ground’ that can account for two sorts of scenarios: First, the possibility that, due to contingent features in their epistemic histories, two equally virtuous agents may fail to be

¹⁶ King, “Disagreement,” 13.

¹⁷ See Section 2.

¹⁸ See Larry Laudan, *Science and Values. The Aims of Science and Their Role in Scientific Debate*. (Berkeley: University of California Press, 1984).

‘epistemic peers’ on a specific question; and second, the possibility that two epistemic agents who happen to be equally likely to get things right *on this occasion*, may nonetheless fall short of peerhood, due to more fundamental differences in epistemic outlook.

Acknowledgment of peerhood is at least as important an element in true cases of peer disagreement as acknowledgment of the disagreement itself. Reflecting for a moment on how peer status – whether objectively warranted or not – is actually accorded by one group (or individual) to another, it seems plausible that judgments of epistemic peerhood, like those of trustworthiness, will often be bound up with social markers of similarity (e.g., indicators of social background, professional affiliation, or academic credentials). Benjamin Wald,¹⁹ following a suggestion by Mark Vorobej, develops a useful distinction between *close* peers (who not only assess the evidence pertaining to a particular topic in similar ways, but also have good reason to believe that they both have similarly good track records in forming true beliefs on the basis of evidence), *distant* peers (who fail to meet one of these conditions), and *remote* peers (who fail both of these conditions). Importantly, distance – in the sense discussed here – can be the result of lack of familiarity with what would constitute a good track record for the other party, and such lack of familiarity can in turn be a side effect of social distance. This is especially pertinent in cases of disagreement among experts from different disciplines (say, disagreement between nuclear engineers and radiation ecologists on the safety of nuclear power plants). As Wald notes, “members of different epistemic communities can count as remote peers to one another”, and persistent disagreement between such communities “need not be due to any failure of rationality.”²⁰ The limits of rational disagreement thus need not coincide with the limits of (acknowledged) peerhood.

5. Peerhood and Socratic ignorance

The example of disagreement among experts, on matters of public concern, supports the observation that judgments of peerhood depend, at least in part, on what is at stake. When the stakes are high (and the choices among possible courses of action are stark), it is reasonable to put time and effort into determining the relative epistemic status of disagreeing parties. By contrast, when the claims in question are inconsequential, reserving judgment in the face of disagreement may well be the most prudent thing to do.

¹⁹ See Benjamin Wald, “Dealing With Disagreement: Distinguishing Two Types of Epistemic Peers,” *Spontaneous Generations* 3 (2009): 113-122.

²⁰ Wald, “Dealing With Disagreement” 121.

By extension, we demand of epistemic peers (and, even more so, of experts whom we entrust with policy advice) not only that they be as reliable and well-informed as us, but also that they share, by and large, our commitments as to *what it is important to know*. Epistemic peers should not only get their facts right, but should also agree on which facts it is important to get right.²¹ Or, if this seems too strong, two epistemic peers – beyond being equally knowledgeable – should at the very least be equally aware of the limitations of their own knowledge. As Philip Kitcher puts it, *reflective* ignorance – that is, being ignorant about the truth value of a first-order proposition, but believing, correctly, that one does not know – constitutes an improvement on *mere* ignorance, since the former “can be the start of something better – of an inquiry that can lead to valuable knowledge.”²² A further improvement would be a state of *Socratic* ignorance, when an epistemic agent is reflectively ignorant with respect to a given claim *p* and also holds a correct belief as to the relative importance of knowing the correct answer to the question of whether *p*.²³ In order for epistemic agents to count as peers, they should exhibit a similar degree of reflective awareness of both the character and limitations of their own knowledge, as well as of the extent to which it speaks to live issues of concern.

The significance of being aware of one’s own epistemic predicament – not least with respect to one’s larger epistemic environment – is not adequately reflected by traditional definitions of epistemic peerhood in terms of either (individual) epistemic virtues or mere reliability on a given occasion. In most cases of persistent disagreement, the relative epistemic standing of the disagreeing parties is far from self-evident. Rather than taking epistemic peerhood for granted and battling over the correct normative response to *prima facie* instances of disagreement, epistemologists would be well-advised to pay greater attention to the causes of disagreement and its persistence, and to the many ways in which peerhood can be undermined by tacit commitments or failure of reflective awareness of one’s own epistemic predicament.

²¹ Even when the fact in question is itself one concerning the relative importance of, say, different empirical findings.

²² Philip Kitcher, “How Ignorant? Let Me Count The Ways,” (Manuscript of a talk delivered at ZiF, University of Bielefeld, 31 May 2011), 5.

²³ Kitcher, “How Ignorant?” 5.