

*This is a Pre-Print copy of a paper forthcoming in Synthese. Please cite the published version only.*

Title: Catherine Elgin on Peerhood and the Epistemic Benefits of Disagreement

Abstract: Conciliationism is the view that an agent must revise her belief in a proposition when she becomes aware that there is an epistemic peer who disagrees with her about that proposition. If epistemic peers are anything less than strict cognitive and evidential equals, then even slight differences could explain away why the two parties disagree in the first place. But this strict notion of peerhood never obtains in many, if not most, of real-life cases disagreements between inquirers. One recent account of peerhood which might obtain more frequently in cases of real-life disagreements comes from Catherine Elgin (2018). She argues that two scientists who are epistemic peers can disagree because while they might have the same reasoning abilities, they can have different reasoning styles. While there are merits to Elgin's account, I argue that two people with different reasoning styles are unlikely to be epistemic peers since such differences could serve to explain away why they disagree. I argue that there is a conception of peerhood which can retain the sceptical force of conciliationism without trivializing or dismissing the problem of disagreement. I then conclude that a particularly attractive argument against conciliationism is gestured at by Elgin in one of her earlier pieces on disagreement. This argument is based on the idea that there are epistemic benefits to be gained from remaining steadfast in the face of disagreement (2010). The success of this argument, however, might force an inquirer faced with disagreement to choose between synchronic reasons and diachronic reasons. If Nick Hughes' (forthcoming) recent defense of epistemic dilemmism is correct this may be a feature of the argument, rather than a bug.

## I. Introduction

In 2002, paleontologist Gerard Gierlinski was vacationing on the Greek island of Crete when he accidentally discovered what he took to be human-like footprints. Years later in 2010, he decided to formally study the footprints, and consulted numerous experts from around the world. Gierlinski and his colleagues concluded that the footprints belonged to human relatives and are around 5.6 million years old. This conclusion shocked the scientific community because, if correct, the footprints are 2 million years older than the footprints in Africa, which is widely believed to be the birth place of the human species. If Gierlinski and his team are correct, then the story of human evolution is completely different from the one currently accepted by the vast majority of the scientific community. In light of this, it's not surprising that Gierlinski's work was initially met with almost universal scorn. Eventually, the work was published in a 2017 article, "Possible hominin footprints from the late Miocene (c.5.7 Ma) of Crete?"<sup>1</sup> Since its publication the paper has received attention from the mainstream media, and it seems that the scientific community is more willing to seriously consider the data. Indeed, since its publication most of the scientific community now accept that the alleged footprints are in fact footprints.<sup>2</sup>

---

<sup>1</sup> Per Ahlberg, the member of Gierlinski's team responsible for publishing their findings says it took six and a half years and many rejections to get their research published. Ahlberg has published numerous papers including in prestigious scientific venues such as *Nature* and *Science*. He claims to have never before experienced so much trouble getting his work published. For instance, many referees flatly denied that the footprints could be related to humans, despite expert opinion to the contrary. Ahlberg is sceptical the review process was fair, and claims many referees were just trying to bury the paper.

<sup>2</sup> I discovered this story in a news article. Much of the information is from the article: [https://newsinteractives.cbc.ca/longform/human-footprints-greece?cnis=01d7e7d5-8294-4320-bd25-1463e0aa309b\\*C\\*1195\\*0\\*0\\*A](https://newsinteractives.cbc.ca/longform/human-footprints-greece?cnis=01d7e7d5-8294-4320-bd25-1463e0aa309b*C*1195*0*0*A). See Gierlinski et al 2017 for the published results.

This story raises many questions about the appropriate response to peer disagreement, some of which have been explored by epistemologists in the last few decades. The epistemology of disagreement literature focuses on the following question: Suppose an agent *S* rationally believes proposition *P* and becomes aware that an epistemic peer, agent *W*, believes proposition *not-P*. *Is S rational to continue to believe P in the face of her disagreement with W?* Of course, there are different ways to phrase this question, but the focus is always on whether peer disagreement constitutes a (partial) defeater for an agent's initial belief. For now, let's suppose that agents *S* and *W* are epistemic peers with respect to whether *P* if they are in the same evidential and cognitive situation with respect to whether *P*. If they're in the same evidential and cognitive situation with respect to whether *P* then *S* and *W* are both equally likely to reach the correct verdict with respect to *P*. This definition of epistemic peerhood will suffice for now, but as we'll later discover there are different ways of defining peerhood, and such definitions can have important implications as to the appropriate response to disagreement.

While there have been hundreds of articles published on disagreement recently, it's clear that really only two main answers to the question about the rational implications of peer disagreement have emerged: Conciliationists (conformists, revisionists) claim that when *S* becomes aware of peer disagreement about *P* she possesses a (partial) defeater for *P*. Weak versions of conciliationism require only slightly lowering one's confidence in *P*, such that it might still be rational for *S* to believe *P*. Strong versions of conciliationism say that *S* must either lower confidence in *P* (below the threshold required for rational belief), or suspend judgment altogether about *P*. If strong conciliationism is true, then scepticism looms large for the many topics over which we encounter peer disagreement. Non-conciliationists (non-conformists, non-revisionists, steadfasters), on the other hand, deny that when *S* becomes aware of peer disagreement about *P*

she always has a (partial) defeater for *P*. For the non-conciliationist, then, peer disagreement need not always pose a sceptical threat about the many topics over which we find ourselves in disagreement with peers.<sup>3</sup>

If strong conciliationism (from here on just ‘conciliationism’) is true, then when Gierlinski encounters peer disagreement about his footprint thesis he should lower his confidence in or suspend judgment altogether about his thesis. On the other hand, if non-conciliationism is true then it’s possible Gierlinski need not lower his confidence in his belief that the footprints are human-like; Gierlinski and his team can remain steadfast in the face of disagreement. While a number of reasons in support of non-conciliationism have been offered in the literature, for the rest of this paper I’m going to assume that conciliationism isn’t defeated by any of these arguments (or at least not obviously so) and hence disagreement poses a serious challenge to many of our beliefs.<sup>4</sup>

Some may worry, however, that real-life cases of disagreement in the scientific community (and communities of inquiry more generally), like the one between Gierlinski and his opponents, are too messy for the current literature to be a of any relevance to them. Thus, in the rest of this paper I do the following: (i) I examine the idealized cases used to motivate conciliationism and show why they assume a very strict notion of epistemic peerhood. It’s doubtful this strict notion of peerhood – one that needs evidential and cognitive equality – ever obtains in cases of real-life

---

<sup>3</sup> Some have advocated for hybrid views which sometimes recommend conciliating and sometimes remaining steadfast. However, these are minority views and so won’t be my focus here. Likewise, it’s often difficult with such views to figure out precisely when the criteria is met for either conciliating or non-conciliating. See Lackey 2010 and Kelly 2010.

<sup>4</sup> Arguments for non-conciliationism include the claim that one should privilege one’s own perspective, conciliating would exhibit a kind of spinelessness, and that conciliationism is self-referentially incoherent (Decker 2014; Sampson forthcoming). Others have suggested that conciliationism should be rejected because it can lead one to accept obviously absurd results (such as one’s doctor telling you to take one hundred thousand pills per day), and that it leads to an unpalatable scepticism (Oppy 2010; Sosa 2010). Finally, some have argued that one should not conciliate because there are epistemic benefits to be gained from remaining steadfast in the face of disagreement (Elgin 2010; Lougheed 2018; Matheson 2015b).

disagreement.<sup>5</sup> (ii) One recent account of peerhood which might obtain in cases of real-life disagreements comes from Catherine Elgin (2018). She argues that two scientists who are epistemic peers can rationally disagree because while they might have the same reasoning abilities, they can have different reasoning styles. (iii) I argue that while there are merits to Elgin's account including her focus on the epistemic benefits of disagreement, two people with different reasoning abilities are unlikely to be epistemic peers since such differences could serve to explain why they disagree. (iv) I thus offer an account of epistemic peerhood which obtains in real-life cases of disagreement and thus preserves the sceptical pressure generated by conciliationism. (v) I conclude that one reason for certain researchers to reject conciliationism is that disagreement is often epistemically beneficial to inquiry. This is an argument first gestured at by Elgin, at least within the disagreement literature (2010). However, this line of argument might ultimately force a researcher to choose between synchronic and diachronic reasons. This isn't a problem, though, if Nick Hughes' (forthcoming) recent defense of epistemic dilemmism is correct.

## II. Epistemic Peerhood

In this section I outline some of the seminal idealized cases in the literature intended to support conciliationism. I then show why such cases, if they are to support conciliationism, require a very strict notion of epistemic peerhood.

### 1. Peerhood and Simple Cases

---

<sup>5</sup> For simplicity I'm not going to worry about the idea that we should account for epistemic inferiors and superiors. Even if we did worry about them there is no practical way to accurately assess an opponent's epistemic status as an inferior or superior in real-life. Without an accurate assessment and hence weighting the same lessons are going to apply here as they would in the case of peers.

Much of the early work in the epistemology of disagreement centres around a number of idealized cases intended to support conciliationism. Consider the following cases:

#### Horse Race

You and a friend are to judge the same contest, a race between Horse A and Horse B. Initially, you think that your friend is as good as you at judging such races. In other words, you think that in case of disagreement about the race, the two of you are equally likely to be mistaken. The race is run and the two of you form independent judgments. As it happens you become confident that Horse A won, and your friend becomes equally confident that Horse B won... Here is the bottom line. When you find out that you and your friend have come to opposite conclusions about a race, you should think that the two of you are equally likely to be correct (Elga 2007, 166-167).

#### Dean in the Quad

Suppose you and I are standing by the window looking out on the quad. We think we have comparable vision and we know each other to be honest. I seem to see what looks to me like the dean standing out in the middle of the quad. (Assume that this is not something odd. He's out there a fair amount.) I believe that the dean is standing on the quad. Meanwhile, you seem to see nothing of the kind there. You think that no one, and thus not the dean, is standing in the middle of the quad. We disagree. Prior to our saying anything each of us believes reasonably. Then I say something about the dean's being on the quad, and we find out about our situation. In my view, once that happens, each of us should suspend judgment. We each know that something weird is going on, but we have no idea which of us has the problem. Either I am "seeing things," or you are missing something. I would not be reasonable in thinking that the problem is in

your head, nor would you be reasonable in thinking that the problem is mine (Feldman 2007, 207-208).

#### Restaurant Case

Suppose that five of us go out to dinner. It's time to pay the check, so the question we're interested in is how much we each owe. We can all see the bill total clearly, we all agree to give a 20 percent tip, and we further agree to split the whole cost evenly, not worrying over who asked for imported water, or skipped desert, or drank more of the wine. I do the math in my head and become highly confident that our shares are \$43 each. Meanwhile, my friend [i.e. epistemic peer] does the math in her head and becomes highly confident that our shares are \$45 each. How should I react, upon learning of her belief (Christensen 2007, 193)?

The intuitive results in these cases are meant to support conciliationism. Such cases require strict evidential and cognitive symmetry between the two opponents who disagree.<sup>6</sup> For any difference in evidence or cognition could justify reasonable disagreement. The kind of evidential and cognitive equality required is possible to achieve in these cases, because they involve relatively encapsulated tasks (Oppy 2010). If I know that you've taken an LSD pill before attempting to do the mental math for the restaurant bill, then I have a dispute *independent* reason to distrust your mental math. At least at the time we're calculating the bill you aren't my epistemic peer because you've taken a drug which makes your cognition unreliable. In other words, you aren't my epistemic peer with respect to the bill total. Likewise, if we have different bodies of evidence about the bill, then we aren't peers with respect to the bill total. Suppose we both receive different bills because the restaurant's receipt printer is running out of ink and mis-prints the bill my friend

---

<sup>6</sup> Or at least they require that one need not be aware of relevant asymmetries.

receives. In such a case not only do I have different evidence than my friend, but I have better evidence. On the assumption that I'm aware of this mistake then I have a reason to distrust my friend's calculation. She isn't my peer with respect to the bill because we have different evidence.<sup>7</sup> The right assessment of these idealized cases seems obvious: they support conciliationism. If these or similar cases arise in the real-world, then we should lower confidence in our initial belief, or suspend judgment, or withhold judgment until we can recheck or confirm regarding the dispute in question. Whatever we do, we aren't rational to simply stand pat and remain steadfast in our original belief once we become aware of disagreement about that belief. Important questions arise, however, when we try to apply these lessons to more complicated real-life cases of disagreements, particularly those that arise in the context of inquiry.

## 2. The Problem with Peerhood in More Complicated Cases

We've seen that even slight evidential or cognitive asymmetries can justify an agent remaining steadfast in the face of disagreement. But in the real-world such relevant asymmetries are always bound to exist. If one's evidence and cognition partly depends upon and are bound up with one's total life history, then there are no true epistemic peers in the real-world. So, neither conciliationism nor non-conciliationism can apply to cases of real-world disagreements between inquirers because there simply are no epistemic peers who disagree with each other. This observation about strict peerhood not obtaining in real-world cases often goes unnoticed. Many advocating for conciliationism appear to assume that lessons from simple idealized cases apply to real-world cases. It's often assumed that the cases are relevantly and sufficiently similar for the lessons to apply seamlessly across cases. For instance, David Christensen writes:

---

<sup>7</sup> Peerhood is indexed to particular times and situations because the key question is whether it is rational to believe P *right now* once one becomes aware of peer disagreement about P.



The restaurant case is designed to be simple in two ways: in the evidential situation and in the evaluation of the general capacities my friend and I exercise in reacting to that sort of evidential situation. This makes our intuitions about the case particularly clear. *But the same lessons emerge, I think, from cases involving a bit more complexity.* (Christensen 2007, 193; emphasis mine).

Jonathan Matheson agrees with this assessment:

Disagreements like that involving the restaurant check can occur in a wide variety of cases. Such disagreements seems to occur frequently concerning matters of politics, religion, ethics, and philosophy in general. What I have said about the restaurant check [that both agents should conciliate] seems to apply equally well to these other cases of disagreement. (Matheson 2009, 271)

While we might be able to say the lessons from the idealized cases would apply to real-life cases if in such cases the disagreement was between two opponents who were genuinely epistemic peers, with respect to real-life disagreements between inquirers there are no epistemic peers.<sup>8</sup> Perhaps one way of avoiding this worry is by thinking of idealization and application in terms of degrees. In the three completely idealized cases above we discover the strict requirements of conciliationism (or in other cases of non-conciliationism). Such requirements can still be applied to less-than-ideal cases, just not as strictly. Thus, as cases become less ideal, the requirements of conciliationism (or non-conciliationism) apply less strictly. Therefore, even if it's technically true that in real-life cases of disagreement there are no precise epistemic peers, it's not difficult to see why in a case like Gierlinksi's that conciliation is still required.<sup>9</sup>

---

<sup>8</sup> For more on why the idealized cases don't apply to real-life cases see Lougheed 2018.

<sup>9</sup> Thanks to an anonymous referee for prompting me to address this concern.

In order for this type of response to succeed, however, it has to be shown that idealized cases and non-idealized cases are sufficiently and relevantly similar. Not only hasn't this been shown, but we have reason to think just the opposite; the cases in question aren't relevantly and sufficiently similar. To see why, consider two detectives, Smith and Jones. They are equally reliable when it comes to solving cases (suppose they have identical truth-tracking records). Let's even suppose that they are exact cognitive peers in reasoning their way through cases. Suppose that Smith and Jones are examining the same case separately from one another. Suppose that 99% of the evidence they gather is identical. Further suppose, however, that 1% of the evidence is different. The case as described is almost ideal, Smith and Jones have identical cognition and almost identical evidence. Yet if they disagree with each other it is *not* clear that they should conciliate. Why? Because that difference in 1% of the evidence might represent *decisive* evidence. If Smith possessed decisive evidence, he would be irrational to conciliate. Real-world cases of disagreement are even less idealized than this one. In real-world cases of disagreement two parties will always have much different cognition and bodies of evidence.

Likewise, consider that we are rarely, if ever, in a position to know how much the differences in cognition and evidence turn out to matter in particular cases. That is, we are often not in a good epistemic position to tell whether one's cognition or evidence is better (or worse) than our own. Without being able to tell the epistemic significance of these differences it is doubtful that the lessons from idealized cases of disagreement should be applied to real-world non-ideal cases.

All of this leaves us in a somewhat puzzling position. An inquirer can remain steadfast in the face of epistemic peer disagreement, but only because there are no epistemic peers in real-

world cases of inquiry. In making similar observations about peerhood not obtaining in real-life cases of disagreement Nathan King writes:

It doesn't follow from the claim that there's no wide-ranging problem about [strict epistemic] *peer* disagreement that there's no widespread problem about disagreement of any kind. Moreover, even upon learning that acknowledged peer disagreement is rare, one may have a lingering sense of puzzlement or dismay in the face of disagreement. Could there really be *no* problem here? That seems too good to be true. But if [strict epistemic] *peer* disagreement isn't the problem, what is? (King 2012, 266)

### III. Elgin on Epistemic Peerhood

In a recent article, Catherine Elgin (2018) provides one potential way to address King's worry. On her account, there can indeed be epistemic peers in cases of real-life disagreements between inquirers. In this section I explain why Elgin believes this is the case. I then conclude by explaining why she is mistaken.

#### 1. Elgin on Epistemically Acceptable Alternatives

Elgin appears to reject the notion of strict epistemic peerhood. She claims that “[b]ecause judgment involves choices from within ranges of epistemically acceptable alternatives, not all disagreements, and not all peer disagreements, about matters of fact rest on mistakes” (2018, 10). This view doesn't commit Elgin to the claim that two competing views can both be true, it only means that the two different views in question could both have been arrived at by legitimate reasoning. She offers an example of a disagreement between two scientists to help demonstrate her point:

Suppose an animal bone with holes pierced in it is found among the artifacts in a Neanderthal settlement. It looks enough like a primitive flute that had it been found in a prehistoric *Homo sapiens* settlement, paleoanthropologists would have no reservations about calling it a flute. But according to current theory, Neanderthal brains were not complex enough to create music, nor were their hands dexterous enough to finger a flute. Jen and Ken, both eminent paleoanthropologists, disagree about whether the artifact is a Neanderthal flute. Jen thinks it is a Neanderthal flute; Ken thinks it is not. Their disagreement may be irresolvable. Given the paucity of evidence about Neanderthals, there is no expectation that anyone will ever come up with sufficient evidence to settle the matter conclusively. Moreover, something significant is at stake here. The verdict as to whether the artifact is a Neanderthal flute will contribute to the ongoing debate about how and in what respects Neanderthals differed from early *Homo sapiens*. If Neanderthals were capable of making music, a host of related hypotheses are seriously off the mark (Elgin 2018, 12).

Elgin continues to explain that there are no relevant epistemic asymmetries between Ken and Jen that can explain away why they disagree in the first place. There's no dispute independent reason to think one of them has made a mistake. Here again we are faced with sceptical pressure from conciliationism. If strong conciliationism is true, then Ken and Jen are probably required to suspend judgment about whether the animal bone is a flute. Elgin claims, however, that both Ken and Jen may have arrived at their respective conclusions by perfectly acceptable reasoning, and hence both are permitted to remain steadfast. She writes:

Disagreements occur in a multi-dimensional space. Along a variety of axis, epistemic requirements fix ranges within which acceptable verdicts must lie. But the

requirements are not sufficiently fine-grained to assure uniqueness or to provide a decision procedure for differentiating among judgments that fall within the range of acceptability. If parties disagree because they make different choices within a range of epistemically acceptable options, neither has made a mistake. Where an agent comes down depends to a considerable extent on which alternatives she had chosen. Had she made other choices within the acceptable range, she would have reached a different verdict. *Epistemic peers are not epistemic clones* (Elgin 2018, 15; emphasis mine).

Elgin says that while peers are taken to have the ‘same reasoning abilities’ this doesn’t entail that they have the ‘same reasoning styles’. For instance, “Ken tends to reason analogically and to credit analogical arguments. Jen, although adept at analogical reasoning prefers inference to the best explanation, considering analogical arguments rather loose” (Elgin 2018, 16). Sometimes peers disagree because they give different weightings of emphasis to different pieces of evidence. Likewise:

A subject’s background information about a topic consists of all the information she has that directly or indirectly bears on that topic. Although epistemic peers by stipulation have the same background information, they need not draw on the same bits of information or assign the bits they draw on the same weight. If they do not, background information plays different roles, in their reasoning (Elgin 2018, 16-17).

For example, Ken believes that since Neanderthals had thick phalanges they would have been physically incapable of the fine motor control require to finger a flute. Jen disagrees because “[s]he considers the anatomical theory sketchy and its bearing on the case slight” (Elgin 2018, 17). Elgin concludes by pointing to the epistemic benefits of these sorts of disagreements between researchers. I will say more about that later, but for now I want to examine Elgin’s claim that

epistemic peers aren't clones and thus can choose different options among a range of epistemically acceptable alternatives.

## 2. Problems for Elgin on Peerhood

One of Elgin's main claims is that because epistemic peers aren't clones, they can choose from a range of epistemically acceptable options. This explains why epistemic peers can reasonably disagree, even while it remains the case that one of them is necessarily wrong.<sup>10</sup> After all, peers have the same reasoning abilities but different reasoning styles. But it's difficult to parse exactly what this means. It might be tempting to understand Elgin as claiming that reasoning styles are something like methods for assessing evidence, or weighting of evidence, along with relevant truth-tracking records. Perhaps it's enough that as long as two agents have the same truth-tracking record in relevantly similar cases, then they are epistemic peers. For if two parties gave all pieces of evidence exactly the same weightings they would be clones, not peers, at least so it seems according to Elgin.

However, if two parties have the same truth-tracking record in previous like cases, then it's puzzling why they disagree in the case *right now*. Why do Ken and Jen disagree about this particular case if they have agreed in the past about like cases? Even if we allow, with Elgin, that they've given different weightings to evidence, Ken and Jen must have consistently offered the same answers to similar questions in order to be considered peers. If Elgin counters that it's a mistake to think the two peers gave the same answers to similar questions in the past (because they have the same reasoning abilities but not styles), it's difficult to understand why we should consider them epistemic peers in the first place. If Ken and Jen have been consistently giving different answers to similar questions in the past, then they will have completely different truth-

---

<sup>10</sup> This doesn't exclude from possibility cases where with further explanation it could be shown that both parties are correct (or wrong). The cases I have in mind are ones in which both parties cannot be correct.

tracking records. This is true on the assumption that only one answer can be correct in such cases, even if we never discover who was correct. But we don't need to know who was correct. The past disagreements indicate that they have different truth-tracking records even if we don't know the specific content of those records.

The problem raised here is precisely why Elgin's characterization of epistemic peerhood is incorrect. While it's rarely said explicitly in the literature, the notion of peerhood oft assumed is indeed what Elgin refers to as epistemic clones. At the very least, even if different reasoning styles between peers is permitted, peers will have equivalent truth-tracking records to like questions. But it seems unlikely that peers consistently using different reasoning styles will have similar – let alone identical – truth-tracking records. So we begin to see that two agents with different reasoning styles probably aren't peers. To be peers, they need to have similar (if not identical) truth-tracking records. To have similar (if not identical) truth-tracking records they need to use similar (if not identical) reasoning styles. To be peers, then, two agents actually need to be what Elgin refers to as clones.

Asymmetries at the level of what Elgin calls reasoning abilities can show why two agents aren't epistemic peers. They can justify reasonable disagreement. But asymmetries pushed back to the level of what Elgin refers to as reasoning styles do just the same. They show why two agents aren't really peers, and hence justify reasonable disagreement. For slight differences in reasoning styles can explain why two parties disagree, but they also explain why the parties in question aren't really peers. But now we can again ask why this is bad news? If epistemic peers are really what Elgin calls epistemic clones, then it's pretty easy to see that there are no epistemic peers in cases of real-life disagreement between inquirers. This brings us right back to the previous worry expressed by King. This is too easy. We can explain how inquirers can reasonably disagree with

one another, but only by denying that there are epistemic peers in contexts of inquiry. The disagreement is made reasonable only by completely dissolving the sceptical pressure of conciliationism. The disagreement is made reasonable only by altogether denying that epistemic peers exist.

An objection to my criticisms of Elgin's account, however, can be formulated by adding slightly more detail to her initial example. Suppose that Ken and Jen having identical truth-tracking records is perfectly consistent with them using different reasoning-styles. Suppose Ken successfully predicts P while using method M1 (her preferred method), while Jen correctly predicts P using method M2 (her preferred method). Further suppose that they do this for numerous cases and thus generate identical truth-tracking records, even though they use different reasoning-styles. It's when Ken and Jen arrive at the artifact in the Neanderthal settlement that they first disagree with one another. Assuming that no asymmetry has suddenly arisen between the two (e.g. Jen hasn't slept in five days), there is no reason to think that Ken and Jen aren't still epistemic peers upon discovering they disagreement about the significance of the bone. While it's true that if they were to continue to have these sorts of disagreements their truth-tracking records would diverge (and hence they'd lose peerhood status), we can stipulate that the Neanderthal case is the first (and only) over which they disagree.<sup>11</sup>

The key claim in this objection is that it's possible for two parties to have identical truth-tracking records while using entirely different reasoning-styles for arriving at their judgments (something I denied above). First, it's worth noting that this is a case in which Ken and Jen cannot both be correct (indeed, disagreement isn't much of an epistemic problem in such cases). Second, more importantly, the key claim is actually something closer to: it's possible for two parties to

---

<sup>11</sup> This objection is taken, almost verbatim, from an anonymous referee.



have identical truth-tracking records while using different reasoning-styles when it comes to *answering relevantly and sufficiently similar questions to the one in which they currently disagree about*. It is this specification which makes it difficult to see how Ken and Jen could possibly have the same truth-tracking record when it comes to answering similar questions. Remember Elgin stipulates that “Ken tends to reason analogically and to credit analogical arguments. Jen, although adept at analogical reasoning prefers inference to the best explanation, considering analogical arguments rather loose” (Elgin 208, 16). This difference in reasoning-style is going to lead them to disagree about a whole range of related questions. It’s not possible they’re going to yield the same verdicts for the very reason that they use different reasoning-styles. The questions they base their track-records on need to be similar to the one about the Neanderthal in order to be any indication as to whether they’ll be right about the bone. For Ken and Jen could have equally good track-records at predicting the results of hockey games, but such predictions aren’t relevant to the current question about the bone. If this isn’t right, then Elgin needs to explain to us how Ken and Jen can be peers when they consistently arrive at different verdicts. Or she needs to explain why Ken and Jen disagree about the current case and yet agree on all others, even though they consistently use different reasoning-styles. Finally, if she denies they have very similar (if not identical) truth-tracking records with respect to similar questions she needs to explain why we should consider them peers in the first place.

#### IV. Sceptical Epistemic Peerhood

It’s true that even slight asymmetries at both the levels of evidence and cognition, including weighting pieces of evidence differently may very well justify a researcher remaining steadfast in

the face of disagreement. But here's the rub: in most real-life cases of disagreement it's going to often be unclear (i) where and to what extent those asymmetries exist and (ii) which party the asymmetries favour even if they have been identified. Recall the case of Gierlinski and his disagreement with much of the scientific community. It might be clear enough that there are asymmetries between Gierlinski and his opponents. For instance, maybe they place different weightings on the evidence that the first homo sapiens lived in Africa. Perhaps they also place different weightings on the importance of the subsequent theories that are developed from and dependent upon the assumption that human life arose in Africa. However, it's not clear which weightings are correct (i.e. more truth conducive). This just pushes the disagreement between Gierlinski and his opponents back one step to disagreement about weightings, and other issues, etc. In other words, one has to think the asymmetries, once identified, favour a specific party with respect to the truth. And such justification has to be offered *independently* of the dispute itself. In cases of real-life disagreements within communities of inquirers, it's unlikely that such independent reasons are often, if ever, available.

Elgin is right that epistemic peers need not be epistemic clones. They need not even have equality at the level of evidence and cognition. Rather, it simply needs to be the case that neither party to the dispute has independent reasons to elevate their epistemic status or downgrade their peer's epistemic status. This formulation of peerhood can be understood as being supported by a kind of scepticism. If an agent isn't sure she is her opponent's evidential or cognitive superior, then she should treat her opponent as if she is her epistemic peer. We can standardized this notion of epistemic peerhood as follows:

Sceptical Epistemic Peerhood

Two agents who disagree over the truth of proposition P are epistemic peers with respect to whether P iff:

1. Each agent lacks a dispute independent reason to think that each of their bodies of evidence is superior to that of each of their opponents.
2. Each agent lacks a dispute independent reason to think they have assessed their body of evidence more accurately than each of their opponents.
3. (1) and (2) obtain and there are no other relevant epistemic asymmetries between the agents.

It's worth pausing here to make a few observations. This account of epistemic peerhood is able to generate the sceptical pressure oft attributed to conciliationism. The problem of what to do when one becomes aware of peer disagreement is alive and well on this account. And it's a problem that's going to appear with respect to many cases of real-world disagreements between researchers. We've thus successfully addressed King's worry about too easily avoiding the problems associated with peer disagreement.

## V. The Epistemic Benefits of Disagreement

We've now established an account of epistemic peerhood which is able to generate the sceptical pressure typically attributed to conciliationism, but one that also obtains in cases of real-life disagreements. We are left with our original questions about the rational demands placed on us once we become aware of peer disagreement. What should Gierlinski do when he becomes aware of peer disagreement about his thesis that the footprints are from a human-like ancestor? There are a number of arguments against conciliationism in the literature. Here I want to discuss a relatively

underdeveloped argument for remaining steadfast in the face of disagreement which is especially relevant to disagreements between researchers. This is the argument that there are epistemic benefits to remaining steadfast in the face of disagreement. Such benefits would be impossible, especially in a scientific context, if researchers always conciliated in the face of peer disagreement (Elgin 2010, 2018; Lougheed 2018).

This argument is of particular interest because Elgin's explanation of peers using different reasoning-styles is one that will be relevant in contexts where researchers disagree with each other. As we've seen, there are problems for thinking that on Elgin's account two researchers using different reasoning-styles are actually peers. On my account, Sceptical Epistemic Peerhood, peerhood will obtain in many cases of real-life disagreements. Suppose that Sceptical Epistemic Peerhood obtains with respect to Jen and Ken. This means that neither Jen nor Ken is aware of what method the other uses and/or evidence the other possesses. More specifically, neither is aware that they have better evidence or a more reliable method for assessing the evidence than the other party. So, the question remains, what are the rational requirements for Jen and Ken once they become aware of their disagreement? The reason for exploring the argument below is that it is particularly relevant to contexts of inquiry where inquirers disagree with one another. In fact, it's the sort of argument which is highly context sensitive and may not be applicable beyond the scope of inquiry and inquirers.

### 1. The Epistemic Benefits of Disagreement

Here's a standardized version of the argument against conciliating that I have in mind:

1. If agent *S* reasonably believes that there are future epistemic benefits to be gained from continuing to believe proposition *P* in the face of epistemic peer disagreement within a

research context  $R$ , then  $S$  is rational to be a non-conciliationist about  $P$  in the context of  $R$ .

2.  $S$  believes  $P$  within the context of  $R$ .
3. There is at least one epistemic peer of  $S$ 's who believe *not- $P$*  within the context of  $R$ .
4.  $S$  reasonably believes that there are future epistemic benefits to from continuing to believe  $P$  within the context of  $R$ .

Therefore,

5. Non-conciliating is true with respect to  $S$ 's belief that  $P$  within the context of  $R$  (Lougheed, 2018, 265-267).

Space constraints prevent me from offering an extended defense of this argument, but I can briefly say a few things in its defense.<sup>12</sup> The most controversial premise in need of explanation and defense is (1).  $S$  might reasonably believe that there are future epistemic benefits to be gained from persisting in her belief by reflecting on previous like cases. There are many cases in the history of science where disagreement ultimately proved to be truth-conducive. It's likewise the case that a research environment in which disagreement amongst practitioners exists is truth conducive. The argument thus appeals to past like examples and claims to be making a recommendation based on these past examples. For instance, consider one real-life example I appeal to in order to defend a similar argument:

Consider the case of Ignaz Semmelweis and bacteria. In the 19th century, hospitals in Europe and America were disease and death-filled places. Childbirth in hospitals by male doctors replaced home births with midwives (Downs 1982, p. 227). But this greatly increased cases of puerperal fever and while it "had long been known in the

---

<sup>12</sup> For more on type of argument see Lougheed 2018; Elgin 2010.

medical profession...its causes remained a dark mystery” (Downs 1982, p. 228). Ignaz Semmelweis (1818-1865) was a physician who became concerned with discovering the cause and prevention of puerperal fever which contributed to mortality rates of over 20% among pregnant women. Against the medical consensus of his day, Semmelweis became convinced that there was a way to prevent puerperal fever (Zoltán 2014). He began to observe a connection between rates of the puerperal fever and medical students who had been involved in medical procedures, including examining corpses, before child delivery. When he ordered that all medical students wash their hands in chlorinated lime before child delivery, cases of puerperal fever dropped in his division to 1.27%. Semmelweis’ ideas were, however, generally met with scepticism and hostility from the medical community. Semmelweis published his findings in a book called *The Etiology, Concept, and Prophylaxis of Childbed Fever* (Semmelweis 1861). But though [h]e sent it to all the prominent obstetricians and medical societies abroad...the general reaction was adverse. The weight of authority stood against his teachings. He addressed several open letters to the professors of medicine in other countries, but to little effect (Zoltán 2014). Eventually, Semmelweis experienced a mental breakdown, partly due to his frustration with not being able to convince any of his peers about the truth of his ideas (Lougheed 2018, 266-267).

While I disagree with Elgin about the right definition of epistemic peerhood, we both agree that there are epistemic benefits to be gained from disagreement. In an earlier piece on disagreement she observes that if a philosopher refrains from using *P* or not *P* as a premise in an argument, simply because there are peers who disagree about *P*, this greatly limits the premises that a philosopher can utilize in argumentation. Given the vast amount of peer disagreement on

numerous philosophical topics, this is a very high epistemic cost to adopting conciliationism (Elgin 2010, 67). While Elgin's argument is with specific reference to how in philosophy it might be epistemically beneficial to continue to defend certain claims despite peer disagreement, this idea clearly generalizes to *many* other fields of inquiry.<sup>13</sup> Recently, she builds on these earlier ideas:

If the parties to the disagreement entertain one another's position seriously and respectfully – if they keep a genuinely open mind to the possibility that there is something significant to be said for it – they may come to appreciate the weaknesses in their own position as well as strengths in their adversary's. This may lead one to concede that the other party is correct. Alternatively, it may lead her to shore up her own position so that it can deflect his objections. A third possibility is that together they craft a position different from the ones with which either of them started, or revise their methods, standards, or criteria of acceptance. There is no assurance that they will ever agree (Elgin 2018, 20).

It's true that in non-empirical fields that disagreement tends to be persistent. We don't have close to agreement or even slow convergence on the seminal questions in philosophy. It's also the case that in more empirical fields there sometimes won't ever be enough evidence to point clearly to one answer.<sup>14</sup> For instance, Elgin notes that there might never be enough information about the Neanderthals to decide the dispute about whether the artifact in question is a flute. Likewise, without new discoveries which would bear on the current set of footprints, there appears to be no

---

<sup>13</sup> To be precise, Elgin's argument is about acceptance rather than beliefs. But the point I make here about remaining steadfast in the face of disagreement need not change in light of it. For a book-length discussion of acceptance see Elgin 2017, where she uses this idea to justify the rationality of scientific inquiry.

<sup>14</sup> Note that this is consistent with the truth of the Uniqueness Thesis. See White 2005.

clear way to settle the dispute between Gierlinski and his opponents. However, even if this is the case Elgin is right to suggest that there are epistemic benefits to remaining steadfast.<sup>15</sup>

### 1. Objections

One worry for this line of argument is that it's going to be difficult to tell if epistemic benefits will obtain in particular cases of persistent disagreement. While the history of science is filled with success, it's also filled with many failures. When faced with disagreement how can a scientist tell whether her theory will prevail? Elgin's point about the benefits of disagreement even if no winner is ever declared is instructive. But also consider that a research environment with disagreement itself is more helpful to inquiry than a more homogenous research environment. If this is right, then the epistemic benefits don't just obtain if the researcher in question turns out to be right, the very existence of disagreement itself (even unresolvable disagreement) helps to create an epistemically successful research environment.

Another objection is not just that it's going to be difficult to tell whether epistemic benefits are in offing, but that such benefits even when achievable can lead one to endorse obviously irrational propositions. In other words, there can be cases where achieving positive epistemic consequences are in direct conflict with epistemic rationality. Consider, for example, that a consequence of trusting a strange and unknown Guru is one gets many true beliefs. Suppose that the Guru's beliefs always turn out to be confirmed at a later time  $t_2$ . But suppose that at time  $t_1$  when the Guru is outputting beliefs there is no independent evidence that supports the truth or falsity of the Guru's beliefs. But the Guru makes fairly radical claims about the world in the absence of any apparent available evidence (though he isn't making such claims in the face of contradictory evidence). The

---

<sup>15</sup> Though it's worth noting Elgin is careful to distinguish her view from dogmatism. A dogmatist might remain steadfast in the face of disagreement, but she doesn't sincerely try to understand arguments opposing her view.



problem is that while trusting the Guru has the epistemically positive consequence of leading to true beliefs, it doesn't follow that it's rational to trust the Guru (and to join his cult and trust only what he says about the world). In a case like this one, the forward-looking consequential focus on rationality is absurd. It seems to retroactively justify beliefs at  $t_2$  that were intuitively epistemically unjustified at  $t_1$ .<sup>16</sup>

There are, however, at least five important differences between trusting the Guru about what to believe because of positive epistemic results (i.e. gaining more true beliefs) and remaining steadfast in the face of disagreement. First, the Epistemic Benefits of Disagreement Argument doesn't recommend that a researcher continue to believe *P* *absence of any apparent available evidence* just because she stands to gain some epistemic benefits. The researcher is justified in believing *P* before she becomes aware of peer disagreement about *P*. She already has good first-order evidence for *P* before discovering the worrisome higher-order evidence of peer disagreement. In the Guru case as described above the person doesn't have appropriate (independent) first-order reasons to believe what the Guru says (remember that there is just an absence of evidence pointing in either direction). Second, the Epistemic Benefits of Disagreement Argument is intended only to apply to researchers within research environments (and hence about beliefs directly related to their research). Admittedly, there are issues of vagueness when it comes to identifying who counts as a researcher and what counts as a research environment. But it is intuitively obvious that the Guru in question is not a researcher and not in a research environment. Third, in the case described above the person who is trusting the Guru has no access to the methodology the Guru employs. They don't know how the Guru is arriving at their beliefs. The researcher, however, does know the methodology she used to arrive at her belief that *P*. Remember

---

<sup>16</sup> This objection is taken, almost verbatim, from an anonymous referee.

that the researcher has first-order reasons for believing *P* in the first place. Fourth, there is a long history of the existence of disagreement leading to epistemic benefits. This isn't to say that remaining steadfast hasn't sometimes resulted in epistemic harms, or that we don't need criteria to decide when it's rational to remain steadfast in particular cases. But there is still a difference here, because there is no such history that can be appealed to when it comes to trusting Gurus. Trusting Gurus is not a methodology commonly used to arrive at true beliefs. Fifth, remaining steadfast about *P* because there will be epistemic benefits is one way of responding to awareness of peer disagreement about *P*. It's an argument about how to respond to higher-order evidence. Trusting the Guru is not a response to higher-order evidence, and given the differences mentioned here, it should be clear that the Benefits to Disagreement Argument does not commit one to trusting the Guru.

## 2. Practical Reasons, Synchronic Epistemic Reasons, and Diachronic Epistemic Reasons

Another objection to this line of argument says that it conflates practical reasons with epistemic reasons. David Christensen and Richard Feldman, two prominent conciliationists, are clearly aware that there could be benefits from pursuing a line of inquiry in the face of disagreement. They simply believe that this constitutes a practical reason for belief, not an epistemic reason. For instance, Christensen writes that:

It's quite plausible that knowledge is best advanced by people exploring, and attempting to defend, a variety of answers to a given question. Perhaps, human psychology makes this easier to do when investigators actually have a lot of confidence in the hypotheses they're trying to defend. Certain sorts of inquiry might well work best when a variety of investigators have irrationally high levels of confidence in a variety of pet hypotheses. So there may well be important

epistemic benefits to certain patterns of irrational belief. But I would argue that the patterns of belief are no more epistemically rational for all that (Christensen 2007, 216).

Likewise, Feldman shares Christensen's worry:

This skeptical conclusion [that we should suspend judgment in the face of peer disagreement] does not imply that people should stop defending the views that seem right to them. It may be that the search for the truth is most successful if people argue for the things that seem true to them. But one can do that without being epistemically justified in believing that one's view is correct (Feldman 2011, 157).

Here a rarely mentioned distinction is worth exploring. There's a difference between *synchronic* epistemic reasons and *diachronic* epistemic reasons (Lougheed 2018; Lougheed and Simpson 2017; Matheson 2015a, 2015b). The reasons aren't practical in the scientific cases I have been discussing because the benefits in question are epistemic. But there does seem to be a genuine difference between questions about what to believe *right now* versus what to believe *given future considerations*. What should a researcher believe *right now* in the face of disagreement about whether *P*? Conciliationism rightly puts sceptical pressure on a researcher's belief that *P*, once she becomes aware of an epistemic peer who believes *not-P*. The question of what a researcher should believe about the future epistemic status of whether *P* is less clear. We've just seen reasons to think there are future epistemic benefits in the offing when researchers remain steadfast in the face of disagreement. The disagreement literature, and epistemology more broadly, is almost exclusively concerned with synchronic epistemic reasons. However, a researcher may well have good diachronic epistemic reasons for remaining steadfast in the face of disagreement.

It's important to recognize that science and other fields of inquiry need diachronic epistemic reasons. Given widespread disagreement amongst scientific researchers both past and present, it's doubtful that science can be considered a rational enterprise without claiming that the diachronic reasons in question are indeed epistemic (in the sense that they impact what it's epistemically rational to believe). It's tempting here to offer an account of all-things-considered epistemic rationality which gives equal weight to both synchronic and diachronic reasons. This is because principles about how to weigh one type of reason against another are not obviously forthcoming. But reflecting on peer disagreement seems to show us, at least in part, why such principles are hard to find: There can be clear conflicts between the demands of synchronic rationality and diachronic rationality. It may well be that a researcher like Gierlinski is (epistemically) rationally required to conciliate in the face of disagreement when it comes to his synchronic belief that the footprints are human-like. However, this doesn't entail that Gierlinski isn't rational to remain steadfast about his diachronic belief about the footprints. As Matheson notes, "it may be that there simply is no all-epistemic-things-considered perspective, but only the disambiguated epistemic perspectives" (Matheson 2015b, 144). If this is right, then at the very least we can say that scientists like Gierlinski are diachronically epistemically rational to remain steadfast in the face of disagreement. It also means that if we're going to understand science as an epistemically rational enterprise, we likely won't be able to do so without the use of diachronic epistemic reasons. Or at least this is the case if conciliationism is true.

### 3. Dilemmic Epistemology

Choosing between synchronic and diachronic reasons may present what Nick Hughes (forthcoming) has recently referred to as an epistemic dilemma. He offers an example where there is a genuine dilemma between the goals of believing only what is true (and not false), and believing

only what is rational (and not irrational). Both of these are pretty basic epistemic norms (or at least I will take them as such). Here's the example Hughes offers:

MUG: You have just woken up and are in the kitchen making breakfast. You want to make a cup of tea, and you'd prefer to drink it from your favourite mug. Remembering that you left it in the dishwasher last night, you see in there and see what appears to be the mug sitting on the rack, just where you left it. On the basis of visual experience you form the belief that your favourite mug is in the dishwasher. Alas, you are deceived. Last night while you were sleeping a thief broke in, stole your mug, and replaced it with an identical-looking replica. (Hughes forthcoming)

Hughes suggests that the person in this scenario finds themselves in a genuine epistemic dilemma that is not open to rational resolution. He writes:

I say we should take them [i.e. epistemic dilemmas] at face value. My view – dilemmism – is that [believing what is true and believing what is rational] both express full-blown, bona fide, epistemic requirements. So you're required to believe that P. And that's pretty much that. Neither of these requirements outweighs or takes priority over the other, and there is no hope for resolving the conflict by appealing to the idea that there are different senses of 'ought' at work or anything like that. All things considered you both ought to believe that P and ought not to believe that P. Sometimes, through no fault of your own, you can stumble into a situation in which you're damned if you do and damned if you don't. Just as there are moral dilemmas, in which you fall short of living up to the demands of morality whatever you do, so too are there epistemic dilemmas. Life is hard (Hughes forthcoming).

Hughes is referring to an unresolvable dilemma between the requirement to believe what is true and the requirement to believe what is rational. I think the dilemma can also apply to the following epistemic norms that we've been discussing:<sup>17</sup>

*Synchronic Epistemic Rationality:* One ought (epistemically) to be synchronically epistemically rational.

*Diachronic Epistemic Rationality:* One ought (epistemically) to be diachronically epistemically rational.<sup>18</sup>

While Hughes is aware that many epistemologists will reject epistemic dilemmism as counterintuitive or outright absurd, he suggests that arguments against it aren't as easily available as some believe. Again, given space constraints I won't provide a detailed outline of his arguments for this conclusion. Rather, I'm concerned with possible applications that epistemic dilemmism has with respect to choosing between synchronic and diachronic epistemic reasons. Dilemmism is able to preserve the idea that *both* synchronic epistemic reasons and diachronic epistemic reasons express genuine epistemic norms about what agents ought to believe (Hughes forthcoming). Likewise, Hughes' position supports the idea that the requirements to be synchronically epistemically rational and diachronically epistemically rational are both non-optional norms. In other words, it could be that we face a genuine dilemma when it comes to choosing between whether to be synchronically epistemically rational or diachronically epistemically rational.<sup>19</sup>

---

<sup>17</sup> Hughes does briefly say that dilemmism may apply to disagreement if we treat conciliationism and non-conciliationism as norms. It's worth noting that if these different positions are normative, it may well be because of the difference in synchronic and diachronic reasons. This warrants further discussion, but I won't explore it here. See Hughes forthcoming.

<sup>18</sup> This is modified from Hughes forthcoming.

<sup>19</sup> Objections to dilemmism Hughes addresses include that: "It must be rejected because it leads to contradictions and explosions in deontic logic; that it must be rejected because it fails to give agents useful guidance; and that it must be rejected because it requires one to do the impossible, and thereby violates the principle 'ought-implies-can'." (Hughes forthcoming).

## VI. Conclusion

We've seen that in real-life cases of disagreement no two opponents are ever strict evidential and cognitive equals. Elgin suggests that this is because the requirement of strict equality describes clones, not peers. Epistemic peers must have the same reasoning abilities, but need not have the same reasoning styles. However, differences at the level of reasoning styles can lead to epistemic asymmetries which are able to explain how the two parties came to disagree. While Elgin's account explains how they can rationally arrive at different beliefs, it does so at the cost of disregarding the sceptical pressure of conciliationism. If two opponents use different reasoning styles, they will have different truth-tracking records with respect to similar like questions. This makes it difficult to understand how they could be epistemic peers. Rather, we should think of two agents as opponents who are unsure who, if anyone, has the evidential or cognitive advantage. This is surely unclear in many real-life cases of disagreement. On this sceptical account of peerhood the normative force of conciliationism exists in many cases of real-life disagreement. At the level of synchronic epistemic rationality the existence of peer disagreement may very well require conciliationism. But from the perspective of diachronic epistemic rationality it can be rational to remain steadfast and persist in a line of inquiry in the face of peer disagreement. It remains to be seen whether there is a plausible account of an all-things-considered epistemic rationality, one that is able to offer principled reasons regarding how we should weigh different synchronic reasons against diachronic reasons. It could be that there is no all-things-considered perspective and sometimes an agent is faced with a genuine dilemma between synchronic and diachronic epistemic reasons. Whether a scientist in Gierlinski's position is rational to remain steadfast in the face of

disagreement could depend on whether one has the synchronic or diachronic epistemic perspective in view.<sup>20</sup>

### Bibliography

- Christensen, David. (2007). "Epistemology of Disagreement: The Good News." *Philosophical Review* 116:2, 187–217.
- Decker, Jason. (2014). "Conciliation and Self-incrimination." *Erkenntnis* 79.5: 1099-1134.
- Elgin, C. (2010). "Persistent Disagreement." In *Disagreement*, Richard Feldman and Ted A. Warfield (eds). Oxford University Press, 53–68.
- \_\_\_\_\_. (2017). *True Enough*. Cambridge MA: The MIT Press.
- \_\_\_\_\_. (2018). "Reasonable Disagreement." In *Voicing Dissent: The Ethics and Epistemology of Making Disagreement Public*, Casey Rebecca Johnson (eds). New York: Routledge, 10-21.
- Feldman, Richard. (2006). "Epistemological Puzzles about Disagreement" In *Epistemology Futures*, Stephen Hetherington (eds). Oxford: Oxford University Press, 216-236.
- Hughes, Nick. (forthcoming). "Dilemmic Epistemology." *Synthese*.
- Kelly, Thomas. (2010). "Disagreement and Higher Order Evidence." In *Disagreement*, Richard Feldman and Ted A. Warfield. Oxford University Press, 111-174.
- Lackey, Jennifer. (2010). "A Justificationist View of Disagreement's Epistemic Significance." In *Social Epistemology*, Adrian Haddock, Alan Millar, and Duncan Pritchard (eds). Oxford University Press.
- Lougheed, Kirk. (2018). "The Epistemic Benefits of Deep Disagreements." *Informal Logic* 38 (2): 263-292.
- Lougheed, Kirk and Robert Mark Simpson. (2017). "Indirect Epistemic Reasons and Religious Belief." *Religious Studies* (53) 2: 151-169.

---

<sup>20</sup> I am very grateful to two anonymous referees for extremely helpful comments on this paper. This paper was made possible, in part, by funding from the Social Sciences and Humanities Research Council of Canada.



- Matheson, Jonathan. (2015a). *The Epistemic Significance of Disagreement*. New York: Palgrave Macmillan.
- \_\_\_\_\_. (2015b). "Disagreement and the Ethics of Belief." In *The future of social epistemology: a collective vision*. Edited by James H Collier. USA: Rowman and Littlefield, 139-148.
- Oppy, Graham. (2010). "Disagreement." *International Journal for Philosophy of Religion*, 68: 183-199.
- Sampson, Eric. (forthcoming). "The Self-Undermining Arguments From Disagreement." *Oxford Studies in Metaethics* 14.
- Sosa, Ernest. (2010). "The Epistemology of Disagreement." In *Social Epistemology*, Alan Haddock, Adrian Millar, and Duncan Pritchard (eds). Oxford University Press, 278-297. [First published in *Episteme*, 6:3 (2009): 269–279.]
- White, Roger. (2005). "Epistemic Permissiveness." *Philosophical Perspectives* 19, 445-459.