

IS IT SAFE TO DISAGREE?

Jaakko Hirvelä

Abstract

This paper offers a new account of the epistemic significance of disagreement which is grounded in two assumptions; (i) that knowledge is the norm of belief and, (ii) that the safety condition is a necessary condition for knowledge. These assumptions motivate a modal definition of epistemic peerhood, which is much easier to operate on than the more traditional definitions of epistemic peerhood. The modal account of the epistemic significance of disagreement yields plausible results regarding cases of disagreement. Furthermore, it is able to tap into the intuitions that have motivated the conformist and the nonconformist positions and it locates a fruitful middle-ground between these two conflicting positions. It will be shown that the conformist is correct in that cases of real peer disagreement force us to suspend our judgment. The reason for this is that in cases of real peer disagreement our beliefs fail to be safe. The nonconformist, on the other hand, is right in that disagreement in itself does not have any epistemic power. It is only by the grace of nature that we gain knowledge. The fact that someone disagrees with you does not mean that you do not have knowledge.

1. Introduction

What kind of epistemic significance does disagreement have? Does it have the power to destroy knowledge, undermine the rationality of our disputed beliefs, or rob us of justification, or is it impotent regarding such matters? The epistemology of disagreement is typically framed in terms of reasonableness and rationality. The central questions have been, (i) can there be reasonable disagreements between epistemic peers, and (ii) what ought one do in the face of peer disagreement? These issues have for the most part been studied by examining whether epistemic peers who disagree are rational in holding onto their beliefs in the face of disagreement, or whether they are justified in their disputed beliefs given that they are disclosed to the fact that an epistemic peer disagrees with them.

A relatively neglected approach to the issue has been to inquire whether a subject whose belief amounts to knowledge before the disagreement can retain her knowledge in the face of peer disagreement. The main goal of this paper is to examine the epistemic power that disagreement might have over knowledge. Can disagreement defeat knowledge?

To get off the ground we will make two assumptions; that knowledge is the norm of belief and that the safety condition is a necessary condition for knowledge. While these assumptions are far from uncontroversial, many epistemologists will be inclined to accept them.¹ According to the knowledge norm of belief one should believe that p only if one knows that p . According to the safety condition one knows that p only if one could not easily have erred in one's belief that p . Given these assumptions the question that we seek to answer is this: 'Does disagreement undermine the safety of our beliefs?' If it does, then it undermines our knowledge, and if it undermines our knowledge, our disputed beliefs will fall into the realm of ignorance rather than knowledge, in which case our beliefs fail to comply with the norm of belief. If our beliefs fail to comply with the norm of belief then they fall short of what beliefs should attain.

The conclusion of this paper is that in cases of real peer disagreement both parties lack knowledge, since their contested beliefs are not safe. Full-blown scepticism is diverted, since in cases of merely apparent peer disagreement one's contested belief might be safe, and thus can amount to knowledge. Crucially, however, the fact that one lacks knowledge in cases of real peer disagreement has nothing to do with the fact that one happens to disagree with an epistemic peer. The reason why we lack knowledge in such cases is that our beliefs are unsafe, and our beliefs were unsafe even before someone disagreed with us. Disagreement does not have any epistemic power with respect to knowledge.

To reach this conclusion we will briefly characterize the knowledge norm of belief and the safety condition. Then we will turn to look at the definition of epistemic peerhood. After that we will

¹ The knowledge norm of belief has been advocated by Williamson (2000), Huemer (2007), Sutton (2007), Jackson (2012), Sosa (2011), and Littlejohn (2013) to name a few. For objections against the knowledge norm of belief see McGlynn (2013, 2014). The safety condition on the other hand has been put forward as a necessary condition for knowledge by Luper (2006), Pritchard (2005, 2007, 2012a, 2015), Sosa (1999), and Williamson (2000). For arguments against the necessity of the safety condition see Neta and Rohrbaugh (2004), Comesana (2005), and Bogardus (2014).

examine cases of real peer disagreement and cases of merely apparent peer disagreement. Finally, we will conclude by examining how our modal approach to the epistemic significance of disagreement is situated in the conformism versus nonconformism debate.

2. Two Assumptions

According to the knowledge norm of belief one should believe that p only if one knows that p . But in what sense of 'should'? There are at least two different ways of interpreting the knowledge norm of belief:

- (i) in believing that p one is rationally committed to knowing that p .²
- (ii) in believing that p one is aiming to know that p ; knowledge is the epistemic standard of success for believing.³

Options (i) and (ii) are not mutually exclusive, and many would accept both of them. However, there are some epistemologists, such as Jackson, who explicitly reject (i) while embracing (ii). If (i) is true the conclusions that we can draw are somewhat stronger than the conclusions that we can draw if only (ii) is true. For the sake of the argument, let us accept both (i) and (ii).

Now let's briefly look at the safety condition. The idea behind the safety condition is that in order to know that p one could not easily have erred in believing that p . To know, is to be safe from error. But how should this condition be fleshed out? Here is a mainstream formulation of the principle:

SAFETY: a subject S 's true belief that p amounts to knowledge only if:

- (i) in all the nearby possible worlds in which S believes that p (via the same method M that she uses in the actual world) her belief will continue to be true.

² Huemer (2007, pp. 145-46) and Gibbons (2013, p. 101) understand the knowledge norm of belief explicitly in this way. Williamson seems to be committed to this reading as well since he states that 'If believing p is, roughly, treating p as if one knew p , then knowing is in that sense central to believing. Knowledge sets the standard of appropriateness for belief' (2000, p. 47).

³ Jackson (2012) endorses this reading of the knowledge norm of belief and rejects interpretation (i) on the grounds that it places too strong constraints on rationality.

But this formulation is inadequate for two reasons. Firstly, we need to think about SAFETY in terms of a continuum of tolerance to epistemic risk when it comes to knowledge. Thus, we cannot tolerate false beliefs in the very closest possible worlds, but do tolerate some false beliefs in nearby possible worlds that are further away from the actual world.⁴ Secondly, SAFETY is trivially satisfied in cases where the subject believes in a necessary truth or in a stable contingent truth. The reason why SAFETY is unable to deal with such cases is because it demands only that the belief that the subject *actually* formed has to continue to be true in all the nearby possible worlds where the subject continues to form that very same belief. In order to deal with such cases we simply need to globalize the notion of SAFETY to a set of propositions that the subject could easily have believed in nearby possible worlds. Of course the relevant set of propositions will have to be restricted somehow, since we do not want to demand that in order for a subject's belief that p to be safe the subject cannot form *any* false beliefs by using the same method that she uses in the actual world. I propose that we restrict the set of relevant proposition in terms of subject matters of inquiry. All of the relevant propositions have to be roughly about the same thing, they have to be propositions in which the subject could easily have formed a belief in her inquiries.

For example, if a subject believes that it is snowing outside, then, when evaluating whether she knows that it is snowing outside, we need to check whether she could easily have ended up with a false belief in her inquiry. The subject matter of her inquiry is the weather outside. The relevant set of propositions will therefore comprise of propositions that describe the weather outside, such as [It is raining], [It is windy], [It is sunny], etc. In order to be safe from error, the subject must not believe in any of these propositions in nearby possible worlds where the proposition in question is false. This way of restricting the relevant set of propositions allows the proponent of safety to deal with cases featuring necessary truths. If a subject forms the necessarily true belief that $12 \times 13 = 156$ by using a malfunctioning calculator that generates answers at random, she is not safe from error. Since the calculator is generating answers at random there will be plenty of nearby possible worlds where she ends up having a false belief about the

⁴ See Pritchard (2007, p. 292) for motivation for making the distinction between the very closest possible worlds and nearby possible worlds in this context.

product of 12 and 13. Since the subject matter of her inquiry is the product of 12 and 13, these false propositions must be included in the relevant set of propositions. Taking these points into account we get the following version of the safety condition:

GLOBAL SAFETY: S knows that p (which belongs to a set of propositions P) only if

- 1) in nearly all nearby possible worlds where S believes in a proposition belonging to P (by the same method M that S uses in the actual world) S's belief is true, and
- 2) in all of the very closest possible worlds where S believes in a proposition belonging to P (by the same method M that S uses in the actual world) S's belief is true.⁵

Now that we have laid the groundwork, let us take a look at cases of disagreement.

3. Peer Disagreement

The epistemology of disagreement has focused on cases of peer disagreement and this essay is no exception. But what does it take to be someone's epistemic peer on a given subject matter? Some epistemologists require evidential and cognitive equality. They claim that two persons are epistemic peers relative to the question whether p , if and only if, they are equally familiar with the evidence relevant to the question whether p and they are equally competent and reliable in assessing the evidence relevant to the question whether p .⁶ Others think that two subjects are epistemic peers on the question whether p if and only if, conditional on their disagreement they are both equally likely to be mistaken.⁷ Although I have no argument against these definitions of epistemic peerage, I will provide a new kind of definition that will be somewhat easier to operate on in the context of this essay.

A natural starting point in trying to find a definition of epistemic peerage is to note that epistemic peers should have roughly the same kind of epistemic standing in matters that they are epistemic peers. If you and I are epistemic peers regarding Roman

⁵ For reformulations of the safety condition that are similar in spirit see Pritchard, (2012b, pp. 273; 2015, p. 102) and Williamson (2009, p. 325).

⁶ See Kelly (2005, pp. 174-75), Christensen, (2007, pp. 188-89), Feldman (2007, p. 201), Lackey (2010, p. 302), and Lammenranta (2011, p. 5).

⁷ See Elga (2007, p. 481).

history, we should have roughly the same amount of true beliefs regarding Roman history and we should be disposed to form true beliefs about that subject matter to roughly the same degree. Furthermore, we should have roughly the same amount of knowledge regarding Roman history. Assuming that GLOBAL SAFETY is a necessary condition for knowledge, we can easily define epistemic peerage in line with these observations.

S and S* are epistemic peers regarding a set of propositions P only if:

- (i) S and S* have true (and false) beliefs in propositions belonging to P to almost the same degree across the scope of nearby possible worlds in a similar distribution.⁸

This definition of epistemic peerhood is not extensionally equivalent with Elga's probabilistic definition, because probabilities do not track perfectly the modal profile of an event. After all, winning in the lottery is an extremely unlikely event, but, provided that you have bought a ticket, it is an event that happens in a possible world that is very close to the actual world, since all that would have needed to change for you to win, is for a few coloured balls to fall in a slightly different configuration.

Furthermore, our definition does not demand that epistemic peers have to be evidential and cognitive equals. This should be seen as a virtue of the definition because it is doubtful whether two subjects can ever share exactly the same evidence. We are often in possession of *personal evidence* about our own epistemic status that is not available to other people (Lackey 2010, p. 309; Lammenranta 2011, pp. 8-9). If you and I disagree about whether p , I will be able to rule out a number of possible mistakes that I could have made in believing that p that I am not in a position to rule out with respect to your belief that $\neg p$. You might be drunk, lying, tired, have something in your eye, or you might be joking with me. I am, however, in a position to know that I am not drunk, lying, tired, etc.

Moreover, those epistemologists who favor the conditions of evidential and cognitive equality often focus on cases where the relevant evidence is fully disclosed in such a way that the disagreeing parties have knowingly shared with one another all the relevant

⁸ The true and false beliefs that S and S* have must be distributed in a similar fashion across nearby possible worlds. Otherwise S could be S*'s epistemic peer even if S had relevant true beliefs only in the very nearest possible worlds while S* had relevant true beliefs only at possible worlds that barely counted among the nearby possible worlds.

evidence on the disputed question (Feldman 2007, p. 201). But the evidence that we have is often so subtle that we cannot cite it or bring it to focus, and thus we are often not able to fully disclose our relevant evidence (Sosa 2010, pp. 290-91). If that is true, then there are far less interesting cases of peer disagreements than we originally thought.

However, I do not want to claim that the definition given above is superior to earlier definitions. Instead I merely think that it is more easily put to use given our assumption that some kind of safety condition is a necessary condition for knowledge. The reason for this is that the modal definition of epistemic peerhood has traction with other modal conditions such as GLOBAL SAFETY. It will be a lot easier to evaluate, whether a subject can satisfy GLOBAL SAFETY, given that she disagrees with her epistemic peer and that we understand epistemic peerhood along the lines of the modal definition of epistemic peerhood.

Now let us take a look at some familiar cases of peer disagreement with our two underlying assumptions and the definition of epistemic peerage in mind.

BILL CALCULATION:

While dining with four of my friends, we all agree to leave a 20% tip and to split the cost of the bill. My friend, Ramona, and I rightly regard one another as peers where calculations are concerned – we frequently dine together and consistently arrive at the same figure when dividing up the amount owed. After the bill arrives and we each have a clear look at it, I assert with confidence that I have carefully calculated in my head that we each owe \$43 and Ramona asserts with the same degree of confidence that she has carefully calculated in her head that we each owe \$45.⁹ (Lackey 2010, p. 315)

What should I do in such a situation? Should I stick to my guns or lower my confidence in my belief that we owe \$43? According to Christensen (2007, p. 193),

it seems quite clear that I should lower my confidence that my share is \$43 and raise my confidence that it's \$45. In fact, I think (though this is perhaps less obvious) that I should now accord these two hypotheses roughly equal credence.

⁹ The case is originally presented in Christensen (2007, p. 193).

Christensen's intuition is widely shared.¹⁰ The question we need to ask is whether BILL CALCULATION is a case of peer disagreement, and whether my, or Ramona's, belief can be globally safe if one of us calculated the amount correctly.

It seems clear that BILL CALCULATION can be a case of real peer disagreement given our definition of epistemic peerage. That is to say, it is possible that I and Ramona acquire roughly the same amount of true beliefs when trying to split bills evenly among several people across the nearby possible worlds in a similar distribution and nevertheless disagree about the amount each of us owes. Lackey (2010, p. 315) has noted that the disagreement in BILL CALCULATION is inexplicable if it is understood as featuring two subjects who are evidential equals, since it is difficult to see how the subjects could share all the relevant evidence given their disagreement. Since we do not demand that epistemic peers have to be evidential equals, we are free to view BILL CALCULATION as a case of real peer disagreement rather than as a case of disagreement where two subjects merely think that they are epistemic peers.

Given how BILL CALCULATION is set up, Ramona and I clearly are real epistemic peers. Now suppose that I split the bill correctly and each of us owes \$43. Should I stick to my guns, lower my confidence in my answer, or suspend judgment? To find out, let us test whether my belief could be globally safe in BILL CALCULATION.

If Ramona and I really are epistemic peers regarding calculations, it seems that my belief, even if true, does not satisfy GLOBAL SAFETY. After all, if we are epistemic peers, then we acquire roughly the same amount of true and false beliefs across nearby possible worlds (in a similar distribution), while doing calculations. Ramona ended up with a false belief in the actual world. Thus there is a possible world very close to the actual world where I end up with a false belief by doing a calculation. But if that is true I will not form only true beliefs in the very nearest possible worlds regarding some subject matter and therefore, I fail to satisfy GLOBAL SAFETY. Given our assumption that GLOBAL SAFETY is a necessary condition for knowledge I do not know that each of us owes \$43. What ought I to believe then? The knowledge norm of belief tells us that I ought to believe that p only if I know that p . Therefore I should suspend judgment. My

¹⁰ See Sosa (2010, p. 292) and Lackey (2010, p. 316). The theories advanced by Elga (2007) and Feldman (2007) also mandate suspension of judgment.

belief falls short of the aim of belief. I am rationally committed to knowing that p even though I do not know that p . In short, my belief is impermissible. This of course does not render it irrational, but I would clearly be better off (epistemically speaking) without my luckily true belief. So it seems that it is rational to re-calculate and suspend judgment for the time being.

Our modal approach to disagreement yields the same verdict regarding other cases featuring real epistemic peers.

HORSERACE:

You and I are watching a horserace between horses Ain't Misbehavin, Batteries Not Included, and Cadillac Jack. Suppose that we are epistemic peers when it comes to evaluating which horse won the race and rightly regard each other as such. The race is a close one, but we both have a good vantage point, and as we form our beliefs about which horse won the race, we are fairly confident in our judgments. But to our astonishment we disagree about the outcome. You believe that Ain't Misbehavin won, while I believe that Batteries Not Included won.¹¹

The common intuition regarding this case is that we should suspend judgment on which horse won (Elga 2007, pp. 486-87; Kelly 2010, pp. 151-52; Lamménranta 2011, p. 5). Our modal approach to disagreement gives here the same verdict. Since you are my epistemic peer I am no better off, epistemically speaking, even if my belief happens to be true. Given that we are epistemic peers there are three possible ways that the world might be.

- (i) Either both of us have a false belief regarding the winner of the race, in which case both of us lack knowledge since knowledge requires truth,
- (ii) you have a true but globally unsafe belief since your epistemic peer has a false belief in the actual world in which case we lack knowledge since knowledge requires global safety and truth,
- (iii) I have true but globally unsafe belief since I have an epistemic peer who has a false belief in the actual world in which case we lack knowledge since knowledge requires global safety and truth.

¹¹ Slightly modified from Elga (2007, p. 486).

Regardless of whether option (i), (ii), or (iii) is the correct description of the situation, our epistemic standing is the same. We are in the realm of ignorance, not in that of knowledge. According to the knowledge norm of belief we should believe that p only if we know that p . Since in all cases of real peer disagreement we lack knowledge we should not continue to hold our beliefs in the face of peer disagreement. Is this 'sceptical' result untenable? Is peer disagreement as widely spread as we suppose? Moreover, does merely apparent peer disagreement have similar consequences? In the next section we will argue that merely apparent peer disagreement does not have the same epistemological power that real peer disagreement has.

4. Apparent Peer Disagreement

In cases of apparent peer disagreement two subjects believe falsely that they are epistemic peers and disagree about the truth value of some proposition, say p . Such cases are easy to construct. For example in an altered version of BILL CALCULATION you think falsely that Ramona is a competent calculator, while in fact she is poor at such tasks. In an altered version of HORSERACE you believe that I am good at judging horseraces and that my perceptual abilities are in working order, while I actually suffer from bad sight and have forgotten my spectacles at home. Does the modal approach to disagreement give the same kind of results regarding apparent cases of peer disagreement and real cases of peer disagreement? In order to be able to tell, we need more details. Is your disputed belief true? If it is, can it satisfy GLOBAL SAFETY?

There is no reason to suppose that your belief could not satisfy GLOBAL SAFETY if it is true. You and your apparent peer do not share the same modal profile regarding the relevant set of propositions, so the fact that her belief is false does not mean that you could easily have formed a false belief. Merely apparent peer disagreement does not defeat the safety of your contested belief.¹² Therefore, you can have knowledge in the face of apparent peer disagreement, and thus you can remain steadfast in your belief if

¹² Maria Lasonen-Aarnio (2010) defends the view that beliefs retained in putative defeat cases are not always unsafe. She argues that this does not undermine the claim that safety is a necessary condition for knowledge. I am sympathetic to this line of thought, though it should be admitted that the view is controversial. Baker-Hytech and Benton (forthcoming) argue also for the conclusion, that knowledge can be retained in the face of cases of putative defeat.

it is globally safe. However, if we fill the details of such cases in a different way, the modal approach will give a different answer. If your disputed belief happens to be false, or globally unsafe, you should abandon your belief.

The modal approach to disagreement gives different results in cases of real peer disagreement and in cases of merely apparent peer disagreement. This is not all that strange. Even though many epistemologists tend to think that it does not matter whether the disputants are real epistemic peers as long as they believe that they are (Sosa 2010, p. 283; Christensen 2007, pp. 188-89), it is important to note, that many epistemologists think it does. For example, Kelly (2005, pp. 174-75; 2010), Feldman (2007, p. 201) and Elga (2010, pp. 175-76) frame the discussion in terms of subjects who are real epistemic peers. It would not be altogether surprising if cases featuring merely apparent epistemic peers and cases featuring real epistemic peers would mandate different doxastic attitudes. Such cases might be indistinguishable from each other from the subject's perspective, but so are Gettier cases and genuine cases of knowledge, and yet the knowledge norm of belief tells us that our beliefs in Gettier cases are impermissible.¹³

Moreover, it is interesting to note, that by giving a different diagnosis regarding cases featuring real epistemic peers and cases

¹³ But should this not be seen as a strike against the knowledge norm of belief? After all, surely the beliefs of Gettier subjects are just as rational and permissible as the beliefs of those subjects who have not been Gettiered. The defenders of knowledge norm of belief can deal with such an objection by appealing to a distinction that has been drawn by defenders of the knowledge norm of assertion. This is the distinction between the primary and secondary propriety of assertion (DeRose 2002; Williamson 2000, ch. 11). An assertion is primarily proper if it conforms to the norm of assertion. An assertion is secondarily proper if the subject can reasonably believe that she has met the norm of assertion. In Gettier cases the assertions that the subjects make are primarily improper, since they lack knowledge but they are nevertheless secondarily proper since the subjects can reasonably believe that they know what they assert. Thus the proponent of the knowledge norm of assertion can explain the intuition that the assertions of Gettier subjects are proper in some sense. Now, since it is overwhelmingly plausible that belief is the inner analogue of assertion (Williamson 2000, pp. 255-56), it seems that the same distinction can be put to use when defending the knowledge norm of belief. Gettier subjects fail to conform to the knowledge norm of belief and thus their beliefs are improper in the primary sense. However, they might reasonably believe that they have met the norm of belief and thus their beliefs can be proper in the secondary sense. How does this reflect to cases of real and apparent peer disagreement? If you remain steadfast in your belief in cases of real peer disagreement then your belief will be improper both in the primary and secondary sense. If you remain steadfast in your belief which amounts to knowledge in the face of apparent peer disagreement, then your belief will be proper in the primary sense, since it conforms to knowledge norm of belief, but might be improper in the secondary sense, since you might not be in position to reasonably believe that you know what you believe.

featuring merely apparent epistemic peers, we have effectively found a fruitful middle-ground between the conformist and the nonconformist. The conformist claims that we should give equal weight to our own beliefs and to those held by our epistemic peers.¹⁴ Thus the conformist thinks that in cases of peer disagreement significant doxastic revision is called for. The nonconformist on the other hand thinks that the mere fact that you disagree with an epistemic peer does not mandate any doxastic revision on either side.¹⁵ According to the modal approach the conformist gets the right result in real cases of peer disagreement, but delivers the wrong result in cases of apparent peer disagreement, while the nonconformist gives the wrong verdict regarding cases of real peer disagreement, but gives the correct verdict regarding cases of apparent peer disagreement. Thus by adopting the modal approach we can tap into the both conformist and nonconformist intuitions.

This observation gains additional support when we examine how the nonconformist and conformist have argued for their respective positions. The nonconformists often claim that there is some symmetry breaker (not necessarily independent of the disagreement itself) that allows one of the disagreeing parties to demote the other, so that she no longer counts as an epistemic peer. Kelly writes that:

Whether your demoting me is reasonable will typically depend on such things as whether my best attempts to parry objections are weak and unresponsive as you take them to be, or whether your conviction that they are weak and unresponsive is due (for example) to your being so dogmatically committed to the opposite conclusion that you fail to appreciate the merits of what I say. (2010, p. 165)

¹⁴ Christensen (2007), Elga (2007), and Feldman (2007) argue for the conformist position.

¹⁵ Kelly (2005, 2010), Wedgwood (2010) and Sosa (2010) argue for the nonconformist position, which is nowadays the default position in the literature. Wedgwood supports a version of epistemological egoism, according to which it is rational for us to place greater trust in our own intuitions, while Kelly advocates a view which he calls the total evidence view. According to the total evidence view what is rational to believe in the face of peer disagreement is determined by one's total evidence. The total evidence view therefore mandates steadfastness in the face of peer disagreement in some cases, whereas in others it requires us to suspend judgment since in some cases of peer disagreement the total evidence will not favor either side of the dispute.

What matters for Kelly is whether your reasons for demoting your opponent are true. In other words; what matters is whether she *really* is your epistemic peer or not. The cases that nonconformists appeal to, are usually ones where the disagreeing parties consider one another to be epistemic peers, but in the face of the disagreement end up demoting the epistemic status of their opponent. If the reasons for reasonable demotion have to be true, as Kelly suggests above, then you never really were epistemic peers to begin with. Therefore it seems that nonconformists have actually focused on cases of apparent peer disagreement.

Lackey thinks that this is indeed the case. She argues that the disagreement literature has actually focused on cases of apparent peer disagreement (which she labels ordinary disagreement) and that the nonconformist in particular draw support from such cases. According to Lackey the conformist should insist that we ought to focus on cases of real peer disagreement (which she labels idealized disagreement), since in such cases the conformist view gives the correct verdict. After all, if we focus on such cases, the possible symmetry breakers will be ruled out, and there will be no reasonable ground for demoting your epistemic peer. However, since Lackey defines epistemic peerhood in terms of evidential and cognitive equality she thinks that cases of idealized disagreement are almost impossible, or at least very rare, and that the conformist is therefore ill-advised to focus on cases of idealized disagreement (2010, pp. 310-11). However, since we do not require that epistemic peers have to be evidential and cognitive equals, we need not think that real peer disagreements are next to impossible or even rare. Therefore we are in a position to claim that the conformist gets the right result in cases of real peer disagreement, and that such cases do not have to be rare, but might very well be quite common. If we accept the modal definition of epistemic peerhood, we do not have to think that real cases of peer disagreement happen only in hyperidealized conditions, as we would have to if we accepted the conditions of evidential and cognitive equality.

However, it might still be the case that the most interesting cases of disagreement are cases of apparent peer disagreement. Controversy abounds in politics, morality, science, religion, and, of course, in philosophy. Are the disagreements that arise within these fields between real or merely apparent epistemic peers? That of course depends on who disagrees with whom and about what, no unified answer can be given here, nor should it be expected. Nevertheless the modal approach advocated here does

not rule out the possibility that we could retain some of our controversial beliefs in these fields. It really depends on whether we can gain knowledge within these fields, and I think that there is no reason to suppose that we cannot.

But perhaps the reader will not be satisfied with such a vague response. After all, are we not interested in knowing when an apparent peer disagreement is merely apparent? According to the account sketched above, a subject cannot know without gaining further evidence that she is in an apparent peer disagreement with someone. If *S* is in an apparent peer disagreement with *S**, she will believe (albeit falsely) that she disagrees with her epistemic peer. Of course *S* might gain evidence in the future that makes her realize that *S** is not in fact her peer. This could of course happen in various ways.

That said, it is possible that the attributor of knowledge knows that an apparent peer disagreement is merely apparent. As an outside perceiver, the attributor of knowledge could have more information than the disagreeing parties. She could, for example, know that *S* is less reliable than *S** in assessing the relevant evidence, and hence know that *S* and *S** do not have true and false beliefs in the relevant set of propositions to almost the same degree across nearby possible worlds.

But discerning whether a given disagreement is between real or apparent peers is a difficult task. In order to successfully evaluate whether two subjects are epistemic peers we have to have a lot of information, which we often lack. Since the modal account of the epistemic significance of disagreement allows us to hold onto our beliefs in cases of apparent peer disagreement, full-blown scepticism is nevertheless avoided. Knowing whether a disagreement is between apparent or real epistemic peers is then not as crucial as we might have originally thought.

5. Conclusions

The modal approach to disagreement does not lead to outright scepticism. It rather allows us to hold onto some of our cherished opinions regarding moral, philosophical and political matters, provided that the world is such that our cherished opinions are safe. If they are globally safe, they might amount to knowledge, and if they amount to knowledge we are entitled to hold onto

them even if someone whom we take to be our epistemic peer disagrees with us. This, I think, is a most welcome result.

But we are not yet home and dry. We have not yet answered our main question, 'what kind of epistemic power does disagreement have over knowledge?' True enough, we have shown that only real cases of peer disagreement undermine knowledge, but this way of putting the issue is somewhat misleading. It is misleading because in real cases of peer disagreement knowledge was never attained by the disagreeing parties. Remember that their beliefs were globally unsafe, and thus did not amount to knowledge. What prohibited them from knowing that p or that $\neg p$ was simply the way the world was. Their beliefs could easily have been false. The fact that they disagreed had nothing to do with the fact that they lacked knowledge. They lacked knowledge even before they realized that they were in disagreement with an epistemic peer.

Disagreement, then, does not have the power to destroy our knowledge. The nonconformist is right in this respect. Disagreement in itself does not have any epistemic power over our knowledge. Or at the very least, given the framework that we are operating in, we have found no reason to suppose that disagreement has such power. Of course, if GLOBAL SAFETY does not give the necessary and sufficient conditions for knowledge it might be that some other condition, which is necessary for knowledge, is undermined by the mere fact that one disagrees with someone whom one takes to be one's epistemic peer. But in order to examine whether such a condition is undermined by apparent peer disagreement we would have to know what that condition is. Alas, we have not yet succeeded in giving a satisfactory analysis of knowledge.¹⁶

Jaakko Hirvelä

Department of Philosophy, History, Culture and Art Studies

University of Helsinki

Unioninkatu 38

00014 Helsinki

Finland

jaakko.hirvela@helsinki.fi

¹⁶ I would like to thank the audience at the University of Tampere where this paper was presented in October 2015 for helpful feedback. Special thanks to Jani Hakkarainen, Raul Hakli, Markus Lammenranta, and Vili Lähteenmäki for insightful comments on an earlier version of this paper. Thanks also to an anonymous referee at *Ratio* for a number of very helpful suggestions.

Bibliography

- Baker-Hytch, M., and M.A. Benton. (forthcoming). 'Defeatism Defeated', *Philosophical Perspectives*.
- Bogardus, T. (2014). 'Knowledge under Threat', *Philosophy and Phenomenological Research*, 88, pp. 289–313.
- Christensen, D. (2007). 'Epistemology of Disagreement: The Good News', *Philosophical Review*, 116, pp. 187–217.
- Comesana, J. (2005). 'Unsafe Knowledge', *Synthese*, 146, pp. 395–404.
- DeRose, K. (2002). 'Assertion, Knowledge, and Context', *The Philosophical Review*, 111, pp. 167–203.
- Elga, A. (2007). 'Reflection and Disagreement', *Nous*, 41, pp. 478–502.
- (2010) 'How to Disagree About How to Disagree', in R. Feldman and T. Warfield (eds.), *Disagreement* (Oxford University Press: Oxford), pp. 175–86.
- Feldman, R. (2007). 'Reasonable Religious Disagreements', in L. Antony (ed.), *Philosophers without Gods: Meditations on Atheism and the Secular Life* (Oxford University Press: Oxford), pp. 194–214.
- Gibbons, J. (2013) 'Knowledge Versus Truth', in C. Littlejohn and J. Turri (eds.), *Epistemic Norms* (Oxford University Press: Oxford), pp. 97–114.
- Huemer, M. (2007). 'Moore's Paradox and the Norm of Belief', in S. Nuccetelli and G. Seay (eds.), *Themes from G. E. Moore: New Essays in Epistemology and Ethics* (Clarendon Press: Oxford), pp. 142–157.
- Jackson, A. (2012). 'Two Ways to Put Knowledge First', *Australasian Journal of Philosophy*, 90, pp. 353–69.
- Kelly, T. (2005). 'The Epistemic Significance of Disagreement', in T. Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology* (Oxford University Press: Oxford), pp. 167–96.
- (2010). 'Peer Disagreement and Higher-Order Evidence', in R. Feldman and T. Warfield (eds.), *Disagreement* (Oxford University Press: Oxford), pp. 111–74.
- Lackey, J. (2010). 'A Justificationist View of Disagreement's Epistemic Significance', in A. Haddock, A. Millar and D. Pritchard (eds.), *Social Epistemology* (Oxford University Press: Oxford), pp. 298–325.
- Lammenranta, M. (2011). 'Disagreement, Skepticism, and the Dialectical Conception of Justification', *International Journal for the Study of Skepticism*, 1, pp. 3–17.
- Lasonen-Aarnio, M. (2010). 'Unreasonable Knowledge', *Philosophical Perspectives*, 24, pp. 1–21.
- Littlejohn, C. (2013). 'Xv-the Russellian Retreat', *Proceedings of the Aristotelian Society*, 113, pp. 293–320.
- Luper, S. (2006). 'Restorative Rigging and the Safe Indication Account', *Synthese*, 153, pp. 161–70.
- McGlynn, A. (2013). 'Believing Things Unknown', *Nous*, 47, pp. 385–407.
- (2014). *Knowledge First?* (Palgrave-Mcmillan: Basingstoke).
- Neta, R., and G. Rohrbach. (2004). 'Luminosity and the Safety of Knowledge', *Pacific Philosophical Quarterly*, 85, pp. 396–406.
- Pritchard, D. (2005). *Epistemic Luck* (Oxford University Press: Oxford).
- (2007). 'Anti-Luck Epistemology', *Synthese*, 158, pp. 277–97.
- (2012a). 'In Defence of Modest Anti-Luck Epistemology', in T.B. K. Becker (ed.), *Sensitivity Principle in Epistemology* (Oxford University Press: New York), pp. 173–92.
- (2012b). 'Anti-Luck Virtue Epistemology', *Journal of Philosophy*, 109, pp. 247–79.
- (2015). 'Anti-Luck Epistemology and the Gettier Problem', *Philosophical Studies*, 172, pp. 93–111.
- Sosa, E. (1999). 'How to Defeat Opposition to Moore', *Philosophical Perspectives*, 13, pp. 141–54.
- (2010). 'The Epistemology of Disagreement', in A. Haddock, A. Millar and D. Pritchard (eds.), *Social Epistemology* (Oxford University Press: Oxford), pp. 278–97.

- (2011). *Knowing Full Well* (Princeton University Press: Princeton).
- Sutton, J. (2007). *Without Justification* (MA: Bradford Books: Cambridge).
- Wedgwood, R. (2010). 'The Moral Evil Demons', in R. Feldman and T. Warfield (eds.), *Disagreement* (Oxford University Press: Oxford), pp. 216–44.
- Williamson, T. (2000). *Knowledge and Its Limits* (Oxford University Press: Oxford).
- (2009). 'Reply to John Hawthorne and Maria Lasonen-Aarnio', in P. Greenough and D. Pritchard (eds.), *Williamson on Knowledge* (Oxford University Press: Oxford), pp. 313–29.