Matriculation Number: 060011386

Module Name: Dissertation for M.Litt. Programme

Module Number: PY5099

Supervisor: Jessica Brown

Word Count: 13,450

Title: The Epistemic Significance of Disagreement

I hereby declare that the attached piece of written work is my own work and that I have not reproduced, without acknowledgement, the work of another.

I hereby declare that I composed the attached dissertation, that the work of which it is a record has been done by me alone, and that it has not been accepted in any previous application for any degree.

The Epistemic Significance of Disagreement

1. Introduction

"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% tip, and we further agree to split the whole cost evenly ... I do the math in my head and become highly confident that our shares are $43 each. Meanwhile, my friend does the math in her head and becomes highly confident that our shares are $45 each."[1] We know that there is a single answer to the calculation and so that one of us has made a mistake. Without reason to think that I am in a better position than my friend to carry out the calculation (I do not have a clearer view of the bill; or a better track-record for this sort of calculation; or any reason to think that my friend is addled by tiredness or wine) I cannot simply assume that the mistake was hers. At the very least, it seems that I am not justified in retaining my belief with undiminished confidence in the face of the disagreement.

Cases of disagreement such as this are of at least *prima facie* epistemic significance, in that they seem to provide those involved with evidence that can affect the confidence that they can justifiably invest in the proposition they disagree over and thereby effect the doxastic attitude that they are justified to take towards it. This paper will assume three broad doxastic attitudes that one can take towards a given proposition: belief, disbelief (that is, belief in the negation of the proposition), and suspension of judgement (that is, the attitude of neither believing the proposition nor its negation). Further states are possible;

[1] Christensen, D. 2007. 'Epistemology of Disagreement: The Good News'. *Philosophical Review*, Vol. 116. p. 193. The example also appears in Christensen, D. 2009. 'Disagreement as Evidence: The Epistemology of Controversy'. *Philosophy Compass*, Vol. 4. p. 757.

that state one is in, for example, with respect to propositions that are totally outside one's cognitive grasp, but this is less a doxastic attitude and more a total absence of doxastic attitude.

Kelly opts instead for "the standard Bayesian convention according to which the credence that one invests in a given proposition is assigned a numerical value between 0 and 1 inclusive, where 1 represents maximal confidence that the proposition is true, 0 represents maximal confidence that the proposition is false, 0.5 represents a state of perfect agnosticism as to the truth of the proposition, and so on" because he worries about less nuanced notions of belief that treat belief as "an all-or-nothing matter";[2] belief must instead be treated "as a matter of degree".[3] This observation is entirely correct but does not threaten the present work because our tripartite division of belief states allows for variation in degree of belief. One agent may have a weak belief that p, while another has a strong belief that p. Both take the same doxastic attitude towards the proposition, but the second agent's belief will withstand further reduction in confidence than the first agent's before becoming an attitude of suspended judgement or of disbelief. This paper allows itself notions such as that of strong belief and weak belief, but does not consider its goals furthered by precise quantification of confidence.

Theories of disagreement can be constitute a spectrum, the two ends of which are *steadfast* views, according to which finding oneself in a disagreement *never* provides evidence that affects the confidence that one is justified to invest in the proposition at the centre of a disagreement (and so can *never* provide evidence that can alter the doxastic attitude that one is justified in taking towards it) and *conciliatory* views, according to which finding oneself in a disagreement

[2] Kelly, T. 2010. 'Peer Disagreement and Higher-Order Evidence'. *Disagreement*, Feldman and Warfield (eds), Oxford: Oxford University Press. p. 117
[3] Kelly. 'Peer Disagreement and Higher-Order Evidence'. p. 118

always provides evidence that affects the confidence that one is justified to invest in the proposition at the centre of a disagreement (and so *sometimes* alters the doxastic attitude that one is justified in taking towards it).

The next section elaborates on the kind of disagreement that interests us in this paper. Discussion in all following sections is restricted to these disagreements. Section 3 then examines an argument by Thomas Kelly in favour of steadfast views. The argument is separated into three parts. The first part argues that each of the parties within a disagreement is justified in privileging their own view. This argument begs the question, however, both against the other party within the context of the disagreement and against conciliatory views within the context of this paper. In the second part of his argument, Kelly accepts that disagreement gives rise to higher-order evidence about the pre-disagreement evidence for the proposition at the centre of the disagreement, but denies that such higher-order evidence has any bearing on the proposition itself. In the third part of the argument, Kelly argues that even if higher-order evidence bears on the relevant proposition, it is trivialised in any case of disagreement. The second and third parts of Kelly's argument will be shown to rely on a mistaken conception of higher-order evidence.

The theory of higher-order evidence that is presented throughout Section 3 presses us to accept some sort of conciliatory view. Section 4 begins by outlining the flaws of a particular conciliatory view, which Kelly calls the Equal Weight View. The same section then motivates a less worrying sort of conciliatory view, and introduces three objections: that conciliatory views are incoherent, that they demand that we suspend belief on all difficult questions, and that they fall foul of the principle that 'ought' implies 'can'. The first will be dealt with in Section 4.

Responses to the final two will be postponed until after the discussion of justification and Evidentialism that takes place in Section 5.

2. Disagreement

Disagreement as a state and disagreement as an activity.

We can distinguish disagreement as a state from disagreement as an activity.[4] Two agents are in a state of disagreement concerning some proposition p when one believes that p and the other believes that not-p, but only engage in the activity of disagreement when each challenges the other about their belief. If someone implores you to 'stop disagreeing with me!' they most likely demand that you stop disputing the things they say, not that you alter all of your beliefs to align with theirs. If A is in a state of disagreement with B, then B is in a state of disagreement with A. Not so the activity of disagreement, however, for B can refuse to challenge A's belief even though A challenges B's. Disagreement as a state is neither necessary nor sufficient for disagreement as an activity; agents may be in a state of disagreement without engaging in the activity of disagreement by expressing their beliefs, or may dispute beliefs which they in fact hold, whether to play devil's advocate or for the sheer thrill of argument. This paper will therefore focus on states, rather than activities, of disagreement. One way to discover that I am in a state of disagreement is for the other party to disagree with me actively. This is not the only way, however. I may, for example, overhear one agent disagree actively with another, or read their view expressed in a paper, and then realise that we are in a state of disagreement. Henceforth talk

[4] Cappelen and Hawthorn make a similar distinction between agreement as a state and agreement as an activity at Cappelen, H. and Hawthorn, J. 2009. *Relativism and Monadic Truth*. Oxford: Oxford University Press. p. 60.

of disagreement should be taken to refer to the state of disagreement rather than the activity.

Not all cases of disagreement will interest us here because not all cases of disagreement are of even *prima facie* epistemic significance. Suppose that my eight-year-old nephew Andy believes the product of twelve and seven to be eighty-five, while I believe that the product is not eighty-five but eighty-four. Andy's belief does not result from erroneous performance of the calculation, but simply from misremembering his multiplication table, while my belief is simply the result of correctly remembering my multiplication tables. Given the ease with which we can check our beliefs by working through the calculation on paper, it would be rash at this stage to take the disagreement as evidence that I am mistaken, and rasher still to respond by changing my belief or significantly altering the confidence with which I hold it. Suppose that, upon becoming aware of the disagreement, I scratch out the sum on a piece of paper and verify my answer. Andy, who did not work through the calculation with me, retains his original belief but it would if anything be more unreasonable now to take Andy's belief as evidence that could warrant revision of my own, because the residual disagreement can very reasonably be attributed to a divergence of our evidence. The disagreement persists because I have evidence that Andy lacks; I have seen the sum worked through on paper, whereas Andy has not. Were Andy to perform the same calculation, to all intents and purposes bringing his evidence in line with mine,[5] I might reasonably expect his beliefs to fall in line with my own. If I am to

[5] Perhaps Andy's evidence can never be quite the same as mine, even with respect to the calculation, but what matters is that our evidence is similar enough to warrant the same arithmetical conclusions. I performed the calculation under different conditions from Andy (at a different time, under slightly different light levels, and so on) but if these unavoidable differences in conditions force differences in our evidence, they need not be differences that are arithmetically relevant. Of course, if I performed the calculation under perfect light levels while Andy performed the calculation in near total darkness, our disagreement

take any action in response to our disagreement, the obvious course is to go through the calculation with Andy. If Andy's belief then falls in line with mine, then the disagreement dissolves. Suppose, however, that Andy retains his belief as to the product of seven and twelve in the face of this new evidence. Perhaps the divergence of our beliefs can no longer be attributed to a divergence of our evidence, for we have both seen the calculation performed, but it might reasonably be attributed to a divergence in our arithmetical abilities. If Andy cannot check his answer, as I can, by performing the simple calculation on paper in the appropriate conditions, then it seems that he is just not very good at maths, and it would not be reasonable to doubt my belief in response to disagreement with someone so poor at arithmetic. A final response to this disagreement could be to try and teach Andy multiplication, but there might be little reason to exert myself so.

Epistemic peers.

The disagreements that interest us here are those that, unlike the disagreement with Andy, the participants cannot justifiably attribute to a difference of evidence or ability. This paper will follow Thomas Kelly in taking interesting disagreement to be between *epistemic peers*. Two individuals are epistemic peers if and only if they satisfy the following conditions:

> "(i) they are equals with respect to their familiarity with the evidence and arguments which bear on that question, and

might reasonably be attributed to a difference of our evidence. A similar point is made at Kelly, T. 2010. 'Peer Disagreement and Higher-Order Evidence'. *Disagreement*, Feldman and Warfield (eds), Oxford: Oxford University Press. p. 152.

> (ii) they are equals with respect to general epistemic virtues such as intelligence, thoughtfulness, and freedom from bias."[6]

Kelly considers it something of a simplification to regard agents as intelligent or thoughtful simpliciter, and takes an agent's epistemic virtues as relative to some domain of enquiry.[7] Whether or not we can make sense of epistemic virtues abstracted from a particular domain, there are certainly cases in which an agent's epistemic virtues in some domain are more relevant to determining peerage than their epistemic virtues generally. Consider two scientists who disagree about some scientific proposition, one of whom is fanatical about football and totally biased about anything relating to their team. While their bias is limited to football, it looks like the fanatic's peerage with the other scientist depends only on their epistemic virtues in the field of science. The fact that they are a football fanatic can be written off as irrelevant to their scientific disagreement. Matters become more complicated if we allow the scientist's fanaticism to spread to other areas, however, as it looks gradually more like a coincidence that the fanatic has so far avoided fanaticism in science. When the fanatic's fanaticism spreads far beyond football, we have to take it into account in determining whether the two scientists are peers, even if it has not so far affected the fanatic's scientific work. Notice also that one's epistemic virtues, whether general or field-specific, are not constant. They (we might hope) increase with age and education, until their inevitable decline in old age. Lack of sleep, the use of drugs, and much more besides, can affect one's epistemic virtues and therefore peerage.

---

[6] Kelly, T. 2005. 'The Epistemic Significance of Disagreement'. *Oxford Studies in Epistemology*, Vol. 1, Oxford: Oxford University Press. p. 175. Kelly's footnote 11 addresses similar concerns to those raised in my footnote 6.

[7] Kelly. 'The Epistemic Significance of Disagreement'. p. 174, footnote 10.

Although, for simplicity's sake, this paper will primarily consider cases of epistemic peers, the significance of a disagreement depends less on whether the agents involved are actually epistemic peers and more on their evidence as to whether or not they are epistemic peers.[8] If I have excellent evidence that another agent is my epistemic peer, disagreement with them is compelling evidence that I am mistaken, while if I have excellent evidence that another agent is not my epistemic peer, disagreement with them is far less compelling, whatever the facts about our peerage. Determining whether another agent is one's epistemic peer can be a complicated business. In the dispute with Andy, it involved gathering further evidence by performing a calculation on paper. There may be cases of disagreement in which the disagreement itself gives reason to judge that another agent is not one's epistemic peer, for example when an agent's belief appears to be a symptom of mental illness. Such cases are interesting, but are set aside here. We will simply assume that agents are epistemic peers and that their evidence justifies the belief that they are epistemic peers.

Asymmetrical cases of disagreement.

The disagreement with Andy concerned an arithmetical proposition about which there is certainly a fact of the matter. In order to avoid quibbles as to whether the participants in a disagreement are justified in having any belief at all about the truth of the proposition they disagree on, the only cases that interest us here are those in which the proposition has some definite truth value. For the same reason, our primary example will feature agents whose common evidential set in

[8] Elga, A. 2007. 'Reflection and Disagreement'. *Noûs*, Vol. 41. p. 484 and footnote 14.

fact justifies belief in the relevant proposition.[9] This last stipulation introduces an asymmetry that allows us to consider the disagreement from two angles, both from the point of view of an agent who has recognised the full force of the available evidence and from the point of view of an agent who has misjudged it.

One might worry that there is an easy answer to any case of disagreement in which the shared evidence favours some conclusion: whoever has the belief justified by the shared evidence is in the right, so the other participant should simply defer to them. This response begs the question at issue in this paper, however, by assuming that the evidence justifies the same doxastic attitude before and after the disagreement has come to light. One participant has responded appropriately to the initial evidence, but disagreement with an epistemic peer introduces additional evidence such that their total evidence may now justify a different attitude. An awareness of our fallibility is central to the issues discussed in this paper. In some areas it can be very difficult to tell what the available evidence favours, so sincere, intelligent, and reasonable individuals can come to different conclusions. We are imperfectly rational creatures, so in addition to the evidence that bears directly on some proposition, we must also consider evidence to the effect that we have misjudged the evidence.

[9] We can add that the shared evidence does not justified disbelief or suspension of judgement as to the relevant proposition. The assumption that some evidential sets justify a single doxastic attitude is certainly plausible, but this paper stops short of endorsing uniqueness principles, which claim that *any* set of evidence can justify only a single doxastic attitude. More on uniqueness can be found in: Feldman, R. 2007. 'Reasonable Religious Disagreements'. *Philosophers Without Gods: Meditations on Atheism and the Secular Life*, Antony (ed.). Oxford: Oxford University Press. pp. 194-214; Conee, E. 2010. 'Rational Disagreement Defended'. *Disagreement*, Feldman and Warfield (eds), Oxford: Oxford University Press. pp. 69-90; Kelly. 'Peer Disagreement and Higher-Order Evidence'. pp. 119-121.

3. Steadfast views

Given the assumptions of the previous section, one might think that the only justified response to disagreement is to suspend judgement as to the truth of the proposition the disagreement centres on. Suppose that two epistemic peers, A and B, share an evidential set E. A draws the belief that p on the basis of E, while B is equally confident that not-p on the basis of the same evidence. A's conclusion is in fact that justified by the E. As A and B become aware of their disagreement, higher-order evidence about the nature of E is made available; B's belief is evidence for A that E justifies the belief that not-p and A's belief is evidence for B that E justified the belief that p. A's new evidence is misleading, for although we have not stipulated whether p or not-p, we have stipulated that the evidential set E in fact justifies the belief that not-p. Still, misleading evidence is evidence nonetheless[10] and it is no rational defect to take it into account.

For A to retain their belief that p in the face of this new evidence is to dismiss B as mistaken as to the force of the original evidence, but if B is A's epistemic peer then such a judgement is not reasonable. In dismissing B as mistaken, A therefore dismisses B as failing to meet the standards of an epistemic peer. By hypothesis, however, A already considers B as an epistemic peer and a single disagreement is not sufficient grounds to alter this judgement. This suggests that A must either suspend judgement as to the truth of p, or adopt the belief that not-p. If it is appropriate for A, whose belief was in fact justified by the evidence before the disagreement arose, to adopt the belief that not-p, then an equal degree of modesty is appropriate for B, whose belief was not in fact justified by the original evidence. If A were only justified in adopting B's belief,

[10] If this were not so, the notion of weighing the evidence for and against a proposition would make little sense; there would only be evidence for *or* against the proposition, depending on whether the proposition were true or false.

then B would only be justified in adopting A's. If both adopted the doxastic attitude they were justified in holding, disagreement would reoccur and the two would be locked in constant oscillation of belief. It appears, therefore, that the appropriate response is rather for both to suspend judgement

No evidence from disagreement.

Thomas Kelly offers reason to dispute this conclusion. As we have stipulated, A cannot reasonably privilege their view over B's on the grounds of superior evidence or epistemic virtue, as B is A's epistemic peer and A recognises that their evidence supports this fact. Kelly notes, however, that it is consistent to at once recognise another as your equal in some area of ability and yet judge that they have, on some particular occasion, failed to successfully implement this ability. The fact that A and B are of equal ability does not mean that both are correct whenever one is, just as "Two chess players of equal skill do not always play to a draw; sometimes one or the other wins".[11]

From the perspective of either A or B it can seem natural to judge that the other has made a mistake. A believes that p is true and of course if p is true then others who believe that not not-p are incorrect; B believes that not-p and if not-p then others who believe that p are incorrect.[12] To retain their belief that p in the face of disagreement, A need not deny that B is their epistemic peer, only suppose that B has made a mistake in their evaluation of E that A has not; a diagnosis that is a natural consequence of A's belief that p. A parallel line of reasoning is of course available to B. Following this reasoning, both agents can retain their

[11] Kelly. 'The Epistemic Significance of Disagreement'. p. 179.

[12] Assuming, as we must for the disagreement to be at all interesting, that the proposition that p does not shift its truth value relative to its assessor, as is argued with regard to propositions concerning personal taste in MacFarlane, J. 2007. 'Relativism and Disagreement'. *Philosophical Studies*. Vol. 132. pp. 17–31.

respective beliefs, each judging that the other has made a mistake. In the case at hand, A is in fact correct. B is mistaken, but B's mistake is one of having failed to reach the correct conclusion based on the evidence E, not one of failing to take account of additional evidence provided by the disagreement.

It is certainly consistent to judge both that an agent is your epistemic peer and that they have fallen prey to an error that you have not, so the assumption of peerage does not decide the epistemic significance of disagreement, but neither is the question decided simply by the assumption that A begins with the belief that p. This paper has assumed that disagreement with those one takes to be epistemic peers is of at least *prima facie* epistemic significance, yet this argument concludes that A is justified in retaining their belief without even considering whether A's disagreement with an epistemic peer presents new evidence.

To press Kelly's line further we might point out that disagreement is of little or no *prima facie* epistemic significance when it can justifiably be attributed to a mistake on behalf of the other party, so once A reasons that B is mistaken perhaps the epistemic significance of the disagreement dissolves. The reasoning that A uses to conclude that B is mistaken is worryingly question-begging on two levels, however, begging the question both against B within the context of their disagreement by inferring that B is mistaken from the very belief that B disagrees with, and against conciliatory views within the context of this paper by relying on a belief which, conciliatory views argue, might lack evidential support in light of the disagreement. A made an initial judgement that p on the basis of evidence E. Now, however, A is faced with the evidence of B's belief, that another agent, A's epistemic peer, has judged that not-p on the basis of E. Even though A's initial response to E was justified, there is no guarantee that A's total evidence still justifies this belief after the disagreement has come to light. Perhaps A's enlarged

evidential set now justifies the belief that not-p, or suspension of judgement as to whether or not p. If A is to dismiss B as mistaken in order to establish the epistemic insignificance of their disagreement, A must do so without relying on the belief whose evidential support has been called into question by the disagreement.[13]

Inferring that the other party is mistaken from the mere fact of disagreement also has the unattractive consequence that you cannot regard anyone as an epistemic peer if they disagree with you enough. Suppose that A and B disagree about a number of propositions other than p. Following the reasoning above, A concludes that B is in error regarding all of these propositions and has mistaken the force of the available evidence in each case. As the mistakes mount up, it becomes more and more difficult for A to regard B as their epistemic peer and once they reach a sufficient level, A's evidence will indicate decisively that B is not A's epistemic peer. The second condition of epistemic peerage demands that A and B be equals in terms of epistemic virtues, and the best evidence that A has to this effect is the track records of the two parties. Once B's mistakes mount up high enough, A must conclude that their own track record is far superior to B's. Two chess players of equal skill do not always play to a draw, but if one player wins consistently over a large number of games, it becomes very hard to preserve the assumption that they are of equal skill. Adam Elga considers this an absurd consequence: "Without some antecedent reason to think that [A is] a better judge, the disagreements between [A and B] are no evidence that [B] has made most of the mistakes."[14]

---

[13] Christensen discusses this requirement in Christensen, D. 2011. 'Disagreement, Question-Begging, and Epistemic Self-Criticism'. *Philosophers' Imprint*, Vol. 11, No. 6. pp. 1-22. This paper will assume that A has no independent reason to think that B is mistaken.

[14] Elga. 'Reflection and Disagreement'. p. 487.

Higher-order evidence with no bearing on the relevant proposition.

Kelly's initial argument fails to consider the possibility that the disagreement between A and B provides each with higher-order evidence about the belief justified by E. Kelly presents a further argument, however, in which he considers this additional evidence but argues that whether or not A can regard B's belief as evidence about E (specifically, evidence that E justifies belief that not-p), A should not regard B's belief as evidence about p because A would then be in "the awkward position" of treating B's belief as reason to believe that not-p, despite the fact that B doesn't treat their belief as epistemically significant.[15] To establish this claim, Kelly offers three considerations. Firstly, if B is asked to list their evidence that not-p, the fact that B believes that not-p will not feature among them. Secondly, suppose that B starts to change their mind towards believing that E favours the conclusion that p. B will not (and, presumably, cannot justifiably) treat their earlier belief that not-p as evidence which could restore their confidence that not-p. B's belief that not-p is the result of the evidence that not-p, not itself further evidence to this effect. Thirdly, if B's belief that not-p was evidence that not-p, B could increase their confidence that not-p simply by reflecting on an existing belief that not-p.

Jonathan Matheson notes that we rarely list all our evidence for a proposition when asked, only "a contextually relevant or particularly interesting bit of the evidential story",[16] so the mere fact that B will not include their belief in a list of their evidence that not-p does not show that B does not treat their belief as evidence that not-p. The second and third are more convincing. At the very

[15] Kelly. 'The Epistemic Significance of Disagreement'. p. 187.
[16] Matheson, J. 2009. 'Conciliatory Views of Disagreement and Higher-Order Evidence'. *Episteme*, Vol. 6. p. 272.

least, Kelly's examples show that there are limits to B's application of any evidence that arises from B's belief.

Even if B's belief that not-p is evidence for A that E justifies the belief that not-p, Kelly argues that it would be awkward for A to treat B's belief as evidence that not-p, given how B treats their belief. Note, however, the awkward position that A would be in if they deemed B's belief as evidence that E justifies the belief that not-p but not as evidence that not-p, despite the fact that E is the rest of A's evidence. This position is difficult to make sense of, whereas the asymmetry between how A generally treats B's belief and how B generally treats B's belief is easily explained. B's belief can bear on the question of whether or not p by serving as a check on A's imperfect rational faculties. Because A recognises B as their epistemic peer, the fact that B's belief is inconsistent with A's belief is evidence that A is mistaken, and evidence that A is mistaken is evidence that E supports the conclusion that not-p, given that that A believes that p.[17] B's belief cannot serve as a similar check on B's reasoning because B's belief is always entirely consistent with B's belief, but note that perhaps B's previous belief can serve as such a check on B's present belief if B starts to change their mind as in Kelly's second consideration above.

Evidence from disagreement is defeated in all cases.

Having neutralised the argument that B's belief does not constitute evidence (for A) that not-p, we can now move, as Kelly does, to consider how the confidence that A is justified to invest in p is affected by this new evidence. A believes correctly that E justifies the belief that p, encounters disagreement with B, and so

[17] Christensen, D. 2010. 'Higher-Order Evidence'. *Philosophy and Phenomenological Research*, Vol. 81. pp. 189-190.

acquires higher-order evidence to the effect that the shared evidence E justifies the belief that not-p. Kelly argues that, even granting that B's belief constitutes evidence that not-p, this evidence cannot have any effect on the confidence that A is justified to invest in p, for if B's belief that not-p is evidence that not-p, A's belief that p is evidence that p. After the disagreement has come to light, A and B's extended evidential set E' consists of

1) the original evidence E,

2) the fact that B believes that not-p on the basis of E, and

3) the fact that A believes that p on the basis of E.

There is no reason to think that E' justifies a different level of confidence from E. If we allow that 2) and 3) are equally strong pieces of evidence, as might seem reasonable given that A and B are epistemic peers, then 2) and 3) effectively cancel each other out and the question of the conclusion favoured by E' reduces to the question of the doxastic attitude justified by E.[18] A's attitude is justified, while B's is not. Again, however, B's mistake is simply one of failing to reach the attitude justified by E, not one of failing to take account of the evidence provided by the disagreement. The disagreement has not provided any evidence that could affect the confidence that either A or B is justified to invest in p.

This argument relies on considering 2) and 3) as equivalent and counterbalancing pieces of evidence. Whether 2) and 3) are equivalent does not depend simply on their syntactic form but on the roles that these pieces of evidence play. Despite having the same form, we have already observed that 2) and 3) can play very different roles for A and B. B's belief can play an important

[18] The argument appears in Kelly. 'The Epistemic Significance of Disagreement'. p. 190. Kelly revisits the argument in more detail at Kelly. 'Peer Disagreement and Higher-Order Evidence'. p. 143. In the latter work, Kelly rejects the conclusion of the argument but for different reasons than given in this paper.

evidential role for A that A's own belief cannot, serving as a check on A's imperfect reason. It makes little sense to disregard evidence with a unique role on discovering evidence that bears a superficial syntactic similarity.

2) and 3) are not therefore equivalent in the relevant sense and do not simply cancel each other out. The question for A is not simply that of the doxastic attitude and level of confidence justified by E, but rather that of the doxastic attitude and level of confidence justified by 1)-3), which might well be different to that justified by 1) alone due to the higher-order evidence of B's belief.[19] From the point of view of an observer with no belief as to whether or not p, perhaps it is appropriate to disregard 2) on encountering 3), each of which plays a similar role for the impartial third party; 2) being evidence that not-p and 3) being evidence that p. When one encounters two epistemic peers who disagree as to some proposition, the thing to do is to look at their original evidence and make up your own mind.

Even if we were to grant, against the advice of this paper, that 2) and 3) cancel each other out, consider the situation in which A encounters disagreement, not only with B, but also with C, D, E, and so on for as many peers as we fancy. As long as each of these beliefs provides evidence that not-p, there will be some number such that, if A is in a state of disagreement with that number of agents, the resulting higher-order evidence will be sufficient to shift the doxastic attitude

[19] And perhaps because of the evidence of A's own belief. Although the evidence of B's belief plays a different role from A's belief, and a role more relevant to the questions this paper is concerned with, this paper does not intend to claim that one's own belief is never evidence for oneself. Jessica Brown has suggested to me in correspondence that such a view of evidence could lead us to mischaracterise the nature of two-against-one disputes. Suppose that A and B both believe that p and encounter disagreement with C. Unless we allow that A's belief provides evidence for A, we mischaracterise the two-against-one disagreement by equating A's evidence with the evidence A would have if A observed one-against-one disagreement between B and C: the fact that B believes that p and the fact that C believes that not-p (along with whatever evidence A has that bears directly on the relevant proposition).

justified by A's total evidence. Kelly recognises this point in a later paper, arguing that "At some point, when the number of peers grows large enough, the higher-order psychological evidence will swamp the first-order evidence into virtual insignificance."[20] So Kelly's argument, even when granted more assumptions than this paper has argued is reasonable, does not establish the truth of any steadfast view. When A is faced with enough disagreement, the higher-order evidence can affect the confidence A is justified to invest in the proposition at the centre of the disagreement. Exactly what level of confidence or doxastic attitude is justified by the evidence will depend on further details of the case, largely the extent to which E justifies the conclusion that p. If it only weakly justifies the belief that p, the higher-order evidence from B's belief might well shift the balance far enough that A's total evidence now justified the conclusion that not-p.

4. Conciliatory views

If disagreement presents us with evidence that bears on the subject of the disagreement, and this evidence is not immediately trivialised by counterbalancing evidence, then the level of confidence that one is justified to invest in a proposition can be affected when disagreement comes to light. We might be tempted further towards some conciliatory view according to which the level of confidence that one is justified to invest in a proposition is *always* affected when disagreement comes to light. Whether higher-order evidence affects the doxastic attitude one is justified in taking towards the proposition depends on the rest of one's evidence, but perhaps it will always affects the level of confidence that one is justified in investing in the proposition, even if only by a small amount.

[20] Kelly. 'Peer Disagreement and Higher-Order Evidence'. p. 144.

This section will first discuss, and dismiss, a conciliatory view according to which one must always 'split the difference' in a disagreement with an epistemic peer. It will then move to consider Kelly's argument in favour of a conciliatory view. Finally, this section will consider three objections to conciliatory views: that conciliatory views are incoherent, that they demand suspension of judgement on difficult questions, and that they fall foul of the principle that 'ought' implies 'can'. The first will be dealt with in this section, but responses to the final two objections will be offered at the end of Section 5, in light of that section's discussion of justification and Evidentialism.

Problems with 'splitting the difference'.

Kelly discusses a conciliatory view which he calls the Equal Weight View, according to which "it would be unreasonable for [A or B] to retain his or her original opinion … [A or B] should give equal weight to his or her opinion and to the opinion of the other in arriving at a revised view."[21] Giving 'equal weight' to the opinion of others does not mean taking higher-order evidence such as 2) and 3) as equivalent and counterbalancing, for as Kelly argued above, all evidence arising from the disagreement itself would be rendered impotent and the question of the level of confidence justified after the disagreement would reduce to the question of the level of confidence justified before the disagreement. On such an understanding of 'equal weight', it is therefore hard to see how the Equal Weight View could deem it unreasonable for A and B to retain their original opinions; at the very least it seems that A would be justified in retaining their original opinion, given that it was justified by the pre-disagreement evidence.

[21] Kelly. 'Peer Disagreement and Higher-Order Evidence'. p. 112.

The thought here is rather that A and B should *split the difference* by adopting the level of confidence in p that is intermediate between A's original level of confidence and B's. As A and B are equally confident that p and not-p, respectively, the justified response for both is suspension of judgement. Kelly points out, however, that this view is inadequate because it ignores the asymmetry in cases such as that of A and B.[22] Only A's belief was justified by the original evidence E, but equally extensive revisions are now demanded of each party regardless. The level of confidence that A and B are justified to invest in p after the disagreement has come to light depends only on their prior levels of confidence; the original evidence E drops out as entirely irrelevant. We have not specified the degree to which E supports the belief that p, but whether the original evidence supported belief that p weakly or strongly has no affect on the level of confidence that p that is justified post-disagreement, according to the Equal Weight View. While the theory of higher-order evidence outlined so far accepts that disagreement can affect the doxastic attitude that is justified by one's evidence, this is in virtue of shifting the balance of one's total evidence, not in virtue of completely supplanting whatever evidence was available previously.

Christensen notes a response to this criticism, which is to specify that the Equal Weight View only advises agents "what the proper response is to one particular kind of evidence … the evidence provided by their peer's disagreement."[23] Splitting the difference is the appropriate response to this evidence, but given that agents will normally have additional evidence, the question remains what they should do taking all of their evidence into account. While such a move certainly lessens the worrying consequences of the Equal Weight View, it does so at the cost of making it far less interesting. This new view

[22] Kelly. 'Peer Disagreement and Higher-Order Evidence'. p. 123.
[23] 'Disagreement, Question-Begging, and Epistemic Self-Criticism'. p. 4.

amounts to the very plausible claim that, in a case of peer disagreement where no other evidence is available, both parties must make equally extensive revisions to secure a justified level of confidence. Suppose that A and B have forgotten their original evidence.[24] Here the only available evidence is the higher-order evidence that at some point A deemed the original evidence (whatever it was) to justify belief in p, while A deemed the original evidence to justify belief in not-p, and this limited evidence points to suspension of judgement as the only justified state. This version of the Equal Weight View is, at best, only part of the answer we set out to find, which was how disagreement evidence interacts with one's original evidence to affect the level of confidence that one is justified to invest in the proposition under disagreement.

Kelly's argument for a conciliatory view.

Splitting the difference is a radical general policy and errs by ignoring a potentially enormous set of evidence. Kelly argues, however, that while it does not typically demand that one split the difference, disagreement will typically exert some epistemic pressure to move one's level of confidence some way towards that of the other party.[25] The argument moves as follows. Suppose that you encounter disagreement with a large number of epistemic peers, in fact, with all of them.[26] You are not justified in retaining your belief with the same confidence as before the disagreement arose. But "if you are rationally required to be less confident after *all* your peers have disagreed with you, then it would seem that you are also required to be at least somewhat less confident after even *one* of your peers

[24] Kelly. 'Peer Disagreement and Higher-Order Evidence'. p. 124.
[25] Kelly. 'Peer Disagreement and Higher-Order Evidence'. pp. 136-137.
[26] We can also assume that your peers have arrived at their beliefs independently from each other and that they are numerous. For more on the importance of independence see Kelly. 'Peer Disagreement and Higher-Order Evidence'. pp. 146-149.

disagrees with you."[27] Suppose you encounter disagreement with all of your peers, not all at once, but one at a time. If you are justified in giving "zero weight" to the opinion of the first peer and so retaining the same confidence before and after the disagreement comes to light, then you are justified in responding in the same way to the second peer, and so on, in which case you are justified in retaining the same confidence even after you have been exposed to disagreement with all of your peers, contrary to what we just said.

This consequence is not unavoidable after allowing that one can retain one's belief with undiminished confidence after registering disagreement with a single peer. One could recognise that the peer's belief constitutes evidence, without considering the evidence sufficient to warrant revision of one's confidence. Because one takes peer disagreement as evidence, however, disagreement will warrant revision of confidence as it mounts up. This response is only as plausible as the claim that one can add evidence that p to one's evidential set without that increasing, even by a tiny amount, the level of confidence that one can justifiably invest in p. Although Kelly's example does not suffice to conclusively establish a conciliatory view, there is plausibility in the idea that disagreement evidence will always affect the level of confidence that one is justified to invest in the proposition at the centre of the disagreement, even if only by a small amount. This section will close by introducing three objections to conciliatory views.

Are conciliatory views incoherent?

Disagreement knows few bounds, if any, and the epistemology of disagreement is no exception; just witness the disagreement between advocates of conciliatory

[27] Kelly. 'Peer Disagreement and Higher-Order Evidence'. pp. 137-138.

and steadfast theories. Suppose that I am justified in my acceptance of a conciliatory view. I then meet an epistemic peer who advocates a steadfast view and, in accordance with my conciliatory view, reduce my confidence in my conciliatory view. Adam Elga argues that any unrestricted conciliatory view is therefore incoherent, as it calls for its own rejection.[28] Elga opts instead for a "*partially conciliatory*"[29] view, according to which the confidence that one is justified to invest in a proposition is reduced in all cases of disagreement, except for disagreements about the correct theory of disagreement.

First note that this is not the perhaps more recognisable problem of self-defeat on which the assumption that the theory is true entails that it is false. Elga's charge is not that conciliatory views are logically contradictory, as he gives no reason to think that conciliatory views logically entail their falsity. The charge is that they are incoherent in some other way. This incoherence does not consist in conciliatory views calling for their rejection in the face of any disagreement as to their truth, for as defined by both Elga and this paper, conciliatory views need only claim that the level of confidence that one is justified in investing in a proposition is reduced by the recognition of disagreement with an epistemic peer, not that one must abandon one's belief in the face of such disagreement.[30,31] We can imagine possible cases, however, in which disagreement reduces the

[28] Elga, A. 'How to Disagree about How to Disagree'. *Disagreement*, Feldman and Warfield (eds), Oxford: Oxford University Press. p. 179.
[29] Elga. 'How to Disagree about How to Disagree'. p. 184.
[30] Christensen marks this point by distinguishing "between principles which automatically self-undermine, and principles which do so only potentially – that is, they self-undermine only under particular evidential circumstances" and noting that conciliatory views "belong to the latter category". Christensen. 'Disagreement as Evidence: The Epistemology of Controversy'. p. 762.
[31] Elga. 'How to Disagree about How to Disagree'. p. 175. Elga talks about *respected advisors* rather than epistemic peers. Elga also talks about how one *should* alter one's confidence in response to disagreement, rather than how disagreement affects the confidence that one is justified in. Presumably Elga intends this 'should' to be distinctively epistemic so that one should always adopt the level of confidence that is epistemically justified, even if there are, for example, moral or prudential reasons to adopt a different level of confidence. Different sorts of justification will be discussed further in Section 5.

confidence one is justified in investing in one's conciliatory view to such an extent that it must be abandoned (when, for example, one encounters disagreement with a vast number of epistemic peers) and Elga will argue that this possibility renders conciliatory views incoherent.

The argument continues by way of an analogy. Just as theories of disagreement offer advice on how to respond to cases of disagreement, consumer advice magazines offer advice on what products to buy. Suppose that *Consumer Reports* magazine tells consumers to buy only toaster X. *Consumer Reports* magazine also rates consumer advice magazines and rates another magazine, *Smart Shopper*, higher than *Consumer Reports*, advising its readers to follow the consumer advice found in *Smart Shopper*. *Smart Shopper* advises readers to buy only toaster Y. *Consumer Reports* therefore advises consumers a) "Buy only toaster X" and b) "Buy only toaster Y", but "it is impossible to follow both pieces of advice."[32] Just as *Consumer Reports* lands itself in incoherence by recommending two incompatible responses to the same situation, any theory of disagreement will land itself in incoherence if it recommends two incompatible responses to the same disagreement, but this is exactly what "[a theory of disagreement] does if it ever recommends a competing [theory] over itself."[33] Suppose that theory M advises you to enter belief state X in response to a particular disagreement, theory N advises you to enter belief state Y in response

[32] Elga. 'How to Disagree about How to Disagree'. p. 180. On one interpretation of the advice, both a) and b) can be followed if the consumer doesn't buy any toasters. I can follow the advice *only throw punches on a Wednesday* without ever throwing a punch, but it would be odd to describe me as following both this advice and the advice *only throw punches on a Tuesday* on the basis that I never punch. The problem, if there is one, can be avoided by interpreting a) as *buy a toaster and buy only toaster X* and b) as *buy a toaster and buy only toaster Y.*

[33] Elga. 'How to Disagree about How to Disagree'. p. 181. Elga abstracts from talk about theories of disagreement, to talk about inductive methods. As everything said of inductive methods is supposed to hold of theories of disagreement, we can safely ignore this complication here.

to the same disagreement, and that belief state X is incompatible with belief state Y in that no one can be in both states simultaneously. If theory M also advises you to abandon theory M in favour of theory N in response to the disagreement, then theory M effectively advises you to enter both belief state X and belief state Y. As "it is impossible to follow both pieces of advice … [theory M] gives incoherent advice about how to respond" to the disagreement. [34]

Consider now the possible case in which one's conciliatory view calls for its rejection. This case is importantly different from the case of theory M. While theory M recommends the acceptance of two belief states, X and Y, in response to the same situation, a conciliatory view need only recommend one response to the relevant disagreement: the rejection of one's conciliatory view. If theory M recommended only that one should abandon theory M and accept theory N, it would not recommend two incompatible responses. Similarly, if the only toaster-buying advice *Consumer Reports* offered were to follow the toaster-buying advice found in *Smart Shopper*, it too would avoid incoherence.

Aside from a particular recommendation as to how to respond to the present case of disagreement, any conciliatory view will make a general recommendation that applies to possible future cases of disagreement. Perhaps Elga considers it incoherent to recommend, at some time t_1 (the time at which you encounter disagreement with a steadfast peer), both that you lower your confidence in response to possible future case of disagreement at time t_3 later than t_1 (as per the general recommendation of a conciliatory view) and to recommend that at t_1 you abandon your belief that you should lower your confidence in response to possible future case of disagreement at time t_3 (as per the recommendation that you abandon your conciliatory view). There is no

[34] Elga. 'How to Disagree about How to Disagree'. p. 182.

incoherence here, however, because one could carry out both these recommendations. You could very well abandon your belief that you must lower your confidence in response to future cases of disagreement, and yet still lower your confidence in future cases of disagreement; you could even do so without any peculiarity if at some t_2 between t_1 and t_3 you were to regain confidence in your conciliatory view. Christensen gestures to one of the ways in which it might well be reasonable to regain your confidence in a conciliatory view between t_1 and t_3.[35] After rejecting one's conciliatory view, one might come to believe that it was over-conciliatory to respond to the disagreement at t_1 by abandoning one's view, in which case one might readopt it.

We might still feel slightly uneasy about a view recommending its own rejection in any circumstance, but it need not amount to contradiction or to incoherence in Elga's sense of making two recommendations that cannot both be followed. Christensen notes that many plausible views of disagreement, such as the view that "If I have thought casually about P for 10 minutes, and have decided it is correct, and then find out that many people, most of them smarter and more familiar with the relevant evidence and arguments than I, have thought long and hard about P, and have independently but unanimously decided that P is false, I should become less confident in P",[36] will recommend their own rejection in certain circumstances. Bear in mind that to recommend that one abandon one's conciliatory view is not to say that the view is false. Recommending their own dismissal in certain extreme cases (and depending on the precise conciliatory view, it may only be in very extreme cases) is in keeping with the modesty and openness to new evidence that conciliatory views champion. One can be justified

[35] Christensen. 'Disagreement as Evidence: The Epistemology of Controversy'. p. 762.
[36] Christensen. 'Disagreement as Evidence: The Epistemology of Controversy'. p. 763.

in rejecting a view that is in fact true, so long as this response is justified by one's evidence.

Although Elga does not consider this possibility, we can describe a case that is more similar to Elga's example of theory M above. This is the case in which disagreement reduces the confidence one is justified to invest in one's conciliatory view to such an extent that one must abandon it in favour of a steadfast view. In this case, the conciliatory view recommends both conciliation and steadfastness in future cases of disagreement, and one certainly cannot respond to a disagreement in both ways. As described in both this and Elga's paper, however, conciliatory views need never demand that one accept a steadfast view, no matter how extreme the disagreement. Even in the most radical disagreements, the furthest that a conciliatory view must go is to recommend disbelief in itself, that is, belief in its negation. Steadfast views, which claim that disagreement *never* decreases the confidence one is justified in investing in the proposition at the centre of the disagreement, are more radical than the negations of conciliatory views, which would claim only that *some* cases of disagreement do not decrease the confidence one is justified in investing in the proposition at the centre of the disagreement.

Conciliatory views demand suspension of judgement on difficult issues.

Disagreement abounds among epistemic peers in a variety of tricky areas, such as philosophy, politics, and ethics. If the level of confidence in which we are justified is affected by every disagreement with an epistemic peer, then perhaps we must withhold judgement on many of these questions. For any belief we form on a contentious issue, the disagreement of peers should lead us to reduce it drastically. If we were to reduce it to the level of disbelief, however, then we

would face disagreement once more. The only stable position is therefore suspension of judgement.[37] This might not seem problematic in itself. By hypothesis such questions are tricky, and there must be many tricky questions that we are only justified in suspending judgement on. When one considers the breadth of disagreement in areas like ethics, however, the problem appears more pressing. Without some level of ethical and political belief we cannot function within society, and one might have serious moral concerns about too much suspension of judgement on such questions

We will return to this question at the end of the following section, where it will be argued that this consequence is not as damning as it might seem. Even if belief in some propositions is not epistemically justified, it may be justified for moral or prudential reasons.

'Ought' implies 'can'.

If any case of disagreement reduces the level of confidence that one is justified to invest in the proposition under disagreement, some weak cases of disagreement (such as that in which the shared evidence is stacked in favour of the conclusion that p, A recognises the force of the evidence while A's peer fails to, and in which A's peer is not particularly confident in their belief) must reduce it by a very small amount, because one will be justified in relying on one's belief in largely the same situations before and after the disagreement has come to light. One might therefore argue against conciliatory views along the lines of the mantra that 'ought' implies 'can'. As one ought to have a justified level of belief, any justified level of belief is a level that one can actually hold. If conciliatory views entail that disagreement can shift the justified level of confidence by smaller amounts than

[37] Elga. 'Reflection and Disagreement'. p. 484.

we can shift our own level of belief, then such views are false. The next section will look at just what we mean by 'justification' in this context, and it will be argued that, in the relevant sense, justification may in fact be beyond one's reach.

5. Evidentialism

Evidentialism is "the thesis that a person is justified in believing a proposition iff the person's evidence on balance supports that proposition."[38, 39] Evidentialism is certainly appealing. To deny Evidentialism we must claim either that one can be justified in believing a proposition that one's evidence does not on balance

[38] Feldman, R. 2009. 'Evidentialism, Higher-Order Evidence, and Disagreement'. *Episteme*, Vol. 6. p. 294.

[39] This paper leaves open various questions about the nature of evidence, which allows for a variety of responses to possible criticisms of Evidentialism. If you and your brain-in-a-vat counterpart (whose experiences, past and present, are identical to your own) have identical evidence, then Evidentialism entails that your beliefs are justified iff your counterpart's beliefs are justified. This paper takes no issue with this conclusion but notes that it could be resisted by a somewhat externalist view of evidence according to which identical phenomenology is not sufficient for identical evidence. For example, the brain-in-a-vat's evidence might differ from yours simply because it differs in the mechanism that produces it. Your evidence comes largely from reliable mechanisms such as perception, while the brain-in-a-vat's evidence comes from the unreliable machination of scientists, and these different mechanisms could be held to produce different evidence, despite producing identical phenomenology.

Suppose that two agents form the same belief, one on the basis of reliable evidence, the other on the basis of some unreliable evidence. Both then forget the original evidence for their beliefs. It might be argued that both now have the same evidence for their beliefs (the mere fact that they remember forming them), in which case Evidentialism entails that one belief is justified if and only if the other is. Intuitively, however, the belief originally based on reliable evidence is justified, whereas the other is not. This conclusion could be preserved by externalist views of evidence, according to which the two agents have different evidence even after their memory lapses: one remembers forming a belief that was in fact based on reliable evidence, while the other remembers forming a belief that was *not* in fact based on reliable evidence. Difference in reliability of mechanism entails a difference of evidence, despite identical phenomenology.

Another response is available, which is to argue that the nature of the agents' evidence depends on the agents' track records. If an agent knows that they have an excellent record of forming beliefs only on the basis of reliable evidence, then simply remembering forming a belief is excellent evidence that the belief is correct, whether or not it was originally based on reliable evidence If, on the other hand, an agent knows that they have a long record of forming beliefs on the basis of unreliable evidence, then simply remembering forming a belief is not good evidence that the belief is correct. As the agents may have different evidence, Evidentialism allows that one might have a justified belief while the other does not.

support, or that we can lack justification for belief in a proposition even though our evidence on balance supports the proposition. Of course, A might, in a sense, be justified in believing that not-p, even though A's evidence supports the proposition that p, if B puts a gun to A's head in an attempt to resolve the disagreement, but this kind of justification is not the distinctively epistemic justification that Evidentialism is concerned to characterise.[40] There are various reasons for which one's belief might be deemed 'justified' (prudential or moral reasons, for example) but so far we have been discussing a particular kind of justification (which this paper terms 'epistemic justification') and Evidentialism captures the reliance of this kind of justification on one's evidence. We can add that a person is justified in disbelieving a proposition if and only if the person's evidence on balance supports the negation of the proposition, that a person is justified in suspending judgement about a proposition if and only if the person's evidence on balance supports neither the proposition or its negation, and that the confidence with which one is justified in believing or disbelieving a proposition depends on the degree to which one's evidence on balance supports the proposition or its negation.[41] Further questions remain of course; what it is, for

[40]Catherine Elgin and Alvin Goldman, among others, argue that belief is not voluntary, in which case it is not clear that A could really come to believe that not-p in response to B's threat. Were it possible, however, this belief would be justified in some sense. The question of voluntary belief will not be addressed in this paper. Elgin, C. 'Persistent Disagreement'. *Disagreement*, Feldman and Warfield (eds), Oxford: Oxford University Press. p. 60. Goldman, A. 1978. 'Epistemics: The Regulative Theory of Cognition'. *Journal of Philosophy*, Vol. 75. p. 515.

[41] Feldman and Conee label this principle of epistemic justification *EJ* and present it as follows: "Doxastic attitude *D* towards proposition *p* is epistemically justified for *S* at *t* if and only if having *D* towards *p* fits the evidence *S* has at *t*" and with regard to levels of confidence the pair say that "EJ is compatible with the existence of varying strengths of belief and disbelief. If there is such variation, then the greater the preponderance of evidence, the stronger the doxastic attitude that fits the evidence." Feldman, R. and Conee, E. 2004. 'Evidentialism'. *Evidentialism*, Feldman and Conee (eds), Oxford: Oxford University Press. p. 83. I adopt a different expression of Evidentialism in order to avoid the less than clear notion of *fitting the evidence.*

In the Afterword to their paper, Feldman and Conee discuss the account of Evidentialism given here but dismiss it so as to leave open the possibility that suspension

example, for evidence to support a proposition. Such questions are interesting but will not be addressed here. Evidential support is taken as basic.

With Evidentialism in the background, this paper has questioned the effect that disagreement has on one's total evidence, and so on the level of confidence and doxastic attitudes that one is justified in taking towards propositions. Section 3 considered Kelly's arguments that disagreement cannot affect the balance of one's total evidence and so cannot affect the level of confidence in which one is justified. Section 4 then looked at various forms of conciliatory views. This section will examine objections to both Evidentialism's necessity and sufficiency claims, and then proceed to disarm the objections to conciliatory views from the end of the previous section.

The necessity claim.

Evidentialism claims that one is justified in believing a proposition only if one's evidence on balance supports the proposition; evidence which on balance supports some proposition is necessary for justified belief in that proposition. Alvin Goldman has argued that we should seek principles to "regulate or guide our

---

of judgement is the appropriate response to evidence that just barely supports the conclusion that p over the conclusion that not-p. Feldman and Conee. 'Evidentialism'. p. 102. This is a difficult issue and one made more complicated by the fact that there is no clear division between belief and suspension of judgement. One can suspend judgment as to the truth of a proposition while still tending more towards its truth than falsity. Suspension of judgement to the highest degree occurs when an agent has equal confidence in the truth and falsity of a proposition, and becomes weaker as the agent becomes more confident one way or the other, eventually, but at no clear point, becoming a weak form of belief or disbelief. One way to respond to Feldman and Conee's worry is to agree that it is appropriate to suspend judgement when one's evidence just barely supports the conclusion that p, not because belief is not epistemically justified, for it has been stipulated that belief is the attitude supported by the evidence, but for reasons of prudence. We must be aware that we are simply not incisive enough to distinguish very small differences in evidential support, so when the evidence just barely supports some conclusion, prudence demands that we suspend judgement.

intellectual activities."[42] Evidentialism's necessity claim is such a principle, telling us that epistemic justification depends only on evidence, rather than, say, on faith or dice rolls. One might worry that this is of limited value as a practical guide, however, due to the imperfect nature of the human capacity for weighing evidence. A's evidence justifies some level of confidence that p. If A could evaluate the evidence perfectly (and desired only to form the belief justified by the evidence) this is the level of confidence that A would adopt. While it may be interesting to describe the epistemic operation of an agent with a gods'-eye view of the available evidence, mere humans complete with human fallibility and imperfect cognitive processes cannot be held to such a high standard. Not only is it possible to mistake the force of one's evidence, but it is possible to make such a mistake while acting reasonably. Some areas are very difficult to judge, so reasonable agents can reach contrary conclusions on these topics without irrationality or blame.

If we could be held responsible for all of the flaws in our cognitive processes, then perhaps it would be fair to deem beliefs unjustified even when they stem from these flaws, but one's ability to evaluate the evidence correctly is largely the result of natural ability, which no one can really be held responsible for. Considerations such as these lead some to argue that the level of confidence A develops will be justified, as long as A earnestly and reasonably attempts to evaluate the evidence. This view is taken by Hilary Kornblith when he argues that "having justified beliefs is simply doing the best one can in the light of the innate endowment one starts from".[43] If this is so, then it appears that the necessity claim made by Evidentialism is false. One can be justified in one's doxastic

[42] Goldman. 'Epistemics: The Regulative Theory of Cognition'. p. 509.
[43] Kornblith, H. 1983. 'Justified Belief and Epistemically Responsible Action'. *The Philosophical Review*, Vol. 92. p. 46.

attitude even if one's evidence does not on balance support that attitude, as long as one has done one's best, so having evidence that on balance supports some proposition is not therefore necessary for justified belief in the proposition. Rather than enquiring into the doxastic attitude that is justified given A's evidence and how this is affected by disagreement, perhaps we should rather be asking what it is reasonable for A to believe is the doxastic attitude that is justified given A's evidence and how this is affected by the disagreement. As one's evidence can favour some conclusion while it is reasonable to believe that it favours another, these questions cannot be given the same answer.

This worry about Evidentialism's necessity claim stems from the perception that a level of confidence may be epistemically justified, according to Evidentialism, that it is impossible to attain given one's cognitive imperfections, and therefore that there is a gap between epistemic justification and blameworthiness. One may be unable to satisfy the standards for justification according to Evidentialism and yet still avoid blame if one has done the best one can with the capacities one has. Remember, however, that there are many and various reasons for which we might deem a belief 'justified' and our concern here is with only one of these, which we have termed 'epistemic justification'. Kornblith seems to have in mind what this paper will term *deontological* justification. One's belief is justified in this sense if and only if one has "fulfilled one's epistemic duties in forming or continuing to hold that belief"[44] and Kornblith takes one's epistemic duty to be fulfilled when one's beliefs are "responsibly arrived at";[45] when one has done one's best with one's innate capacities. One might reasonably be blamed for holding a belief even if that belief is epistemically justified. If one should always strive to think well of people, then perhaps one

[44] Plantinga, A. 1993. *Warrant: The Current Debate*. Oxford: Oxford University Press. p. 14.
[45] Kornblith. 'Justified Belief and Epistemically Responsible Action'. p. 45.

might reasonably be blamed for an uncharitable opinion of someone, even if it is supported by one's evidence. Similarly, one could hold an epistemically unjustified attitude and yet be immune to blame because one has done the best one can.

We have to distinguish the kind of epistemic merit or defect that comes from having an epistemically justified or unjustified belief from praiseworthiness or blameworthiness, so the fact that one can avoid blame despite failing to meet the standards for epistemic justification laid out by the principle of Evidentialism does not establish that the principle is false. In this way the standards of epistemic justification are akin to the standards of artistic excellence.[46] There are those who can create artistically excellent work but try as I might, doing my best with my natural capacities, I simply cannot satisfy these standards. I am in no way blameworthy for failing to satisfy standards of artistic merit that are beyond the reach of my natural abilities, but that does not mean that we should lower our standards to such a level that I can satisfy them. Similarly, another might create a work of genuine artistic excellence but be blameworthy for another reason; perhaps for abandoning their family to create it.

Merit and praiseworthiness are not always separate. Artistic merit is rarely (if ever) the product *simply* of natural ability. Even the naturally artistic must work to hone their skill, and working to perfect one's talent is a natural reason for praise. We might blame the natural artist who is too lazy to work at perfecting their craft, but as long as they do *something* worthwhile with their time we don't tend to blame someone for neglecting one of their talents.

[46] The analogy is suggested in Feldman and Conee. 'Evidentialism'. p. 87.

The sufficiency claim.

Evidentialism claims that one is justified in believing a proposition if one's evidence on balance supports the proposition, that is, that evidence which on balance supports some proposition is sufficient for justified belief. This claim can be doubted on the grounds that one can lack justification for believing that p, despite one's evidence on balance supporting that proposition, if one believes that p for another reason. Suppose for example that A's evidence supports the conclusion that p, but that A has not recognised this fact. Instead of believing that p on the basis of their evidence, A believes that p because of a dice roll, or simply because A plucked the belief out of the air, neither of which are justifiable practices.

The worry about the sufficiency stems from a perceived incongruity between epistemic justification and epistemic defect. One may satisfy the standards for justification according to Evidentialism and yet one's belief still be defective in that it does not result from an appropriate sensitivity to one's evidence. The response to this worry is to reject the idea that one's belief is immune from all epistemic defect if that belief is justified in the distinctive sense characterised by Evidentialism. Epistemic justification is not the only epistemic virtue, although it is an important one, and a belief may be defective for various other reasons, as Goldman recognises when he notes that "a cognizer's beliefs can be justified no matter how good or bad his evidence-gathering and hypothesis-generating practices."[47]

[47] Goldman. 'Epistemics: The Regulative Theory of Cognition'. p. 521.

Kornblith discusses an example in which Jones, a headstrong young physicist, presents a paper so that he can hear the praise of his colleagues.[48] After the cogent paper has been presented, all but one of the audience have evidence that supports belief in Jones's thesis, but the lone dissenter soon presents his objection which, to any impartial observer, totally destroys Jones's thesis. Jones is so unable to tolerate criticism that he blocks his mind and fails to hear the objection, or even notice that an objection has been raised. Jones's evidence still supports belief in his thesis. For the rest of the room, however, disbelief is the only justified response to the evidence, for they have heard a devastating objection to which no one, Jones included, has a response. As Jones has acquired no new evidence from the objection, Evidentialism entails that Jones's belief remains justified. Jones's belief is certainly defective, however, having resulted from an unwillingness to criticise his work (albeit an impressively effective one).

If it is possible for Jones to block his mind so perfectly that his evidence is totally unaffected by his colleague's objection, then indeed Jones's belief remains justified by his evidence. Epistemic justification is not the only epistemic virtue, however, so this paper can agree that Jones's belief is not only defective, but epistemically defective, despite being epistemically justified.[49] In particular, Jones's belief is defective because, although the belief is justified, it is not retained *because* it is justified; Jones retains the belief because he is unable to consider new evidence that may harm his high opinion of himself, certainly a bad evidence-gathering and hypothesis-generating practice of the kind Goldman draws our attention to.

---

[48] Kornblith. 'Justified Belief and Epistemically Responsible Action'. p. 36. From here on the example departs in various ways from Kornblith's original.

[49] Jones's belief is *propositionally rational*, but fails to be *doxastically rational*, in the language of Christensen. 'Disagreement, Question-Begging, and Epistemic Self-Criticism'. p. 4.

Do conciliatory views demand suspension of judgement on difficult issues?

Do conciliatory views entail that the only justified doxastic attitude to take towards difficult issues is suspension of judgement? Certainly this will be the case for many difficult questions, but this in itself is not particular problematic; difficult questions are difficult after all, and there are certainly many issues that we are in no evidential position to justifiably believe or disbelieve. Problems arise, however, when we recognise that we must take a stand on various difficult issues in order to function. Moral and political issues are tricky but cannot totally be avoided. Even attempting to abstain on moral issues can be a moral issue. This section has shown, however, that even this need not be seen as problematic. There are various kinds of justification, and moral and political beliefs might be prudentially or morally justified, even if they are not justified epistemically. Of course, the same cannot be said for all difficult issues. Evidentialism therefore encourages us to accept that we are not justified in holding beliefs about many contentious issues, philosophical issues of no practical significance, for example.

'Ought' implies 'can'.

Section 4 concluded with the objection to conciliatory views that, as one ought to have the level of confidence in propositions that is justified and 'ought' implies 'can', conciliatory views must be false because we cannot always adjust our levels of confidence precisely enough to the level that is justified after certain disagreements. Now that we have further explained the notion of justification at play in this paper, however, we can see that this objection fails. Whatever the status of the principle that 'ought' implies 'can' with regard to moral justification,

it cannot be applied to epistemic justification. There are many standards that not everyone can meet all the time, whether the standards for artistic merit, the standards for receiving a First Class degree, or the standards for epistemically justified belief.

6. Conclusion

This paper has considered how disagreement affects one's evidence and therefore the confidence that one is justified to invest in a proposition at the centre of a disagreement. The justification in question is characteristically epistemic and depends entirely on one's evidence, as articulated by the principle of Evidentialism. There are various other important features of belief, both epistemic and non-epistemic, many of which we label 'justification'. Failure to distinguish the many and varied senses in which a doxastic attitude can be justified led to the objections to conciliatory views discussed at the end of the previous section, as well as the objections to Evidentialism from earlier in the section. A doxastic attitude may be justified in the epistemic sense of being justified by one's evidence, and yet be defective in other ways, both epistemic and non-epistemic; similarly a doxastic attitude may be epistemically unjustified and yet still be meritorious and praiseworthy.

For this reason the question of how one should respond to disagreement is a difficult one. In general we feel that agents should try to believe in accordance with their evidence, but sometimes it might well be appropriate for various reasons (ethical, prudential, aesthetic) to hold epistemically unjustified beliefs. Consider a case in which an agent underestimates the force of their evidence. This case allows us to see more clearly the complications inherent in the question of how an agent *should* respond to disagreement, for while it might seem obvious

that agents should enter those doxastic attitudes justified by their evidence, the agent who has mistaken their evidence will be acting somewhat chaotically if they do. This is not to say, however, that there are not various situations in which we can quite reasonably be interested in epistemic justification above all else. In such a situation disagreement could very well have an important influence on the doxastic attitude that one should hold.

This paper has not addressed the question of how to weigh various kinds of justification against each other, or which takes precedence in different situation, but limited itself to asking how disagreement affects one's evidence and therefore how disagreement affects the beliefs justified by one's evidence. Steadfast views were discussed and rejected in Section 3. Section 4 argued for the plausibility of conciliatory views and introduced three objections. One was dealt with in that section, and two further objections were dismissed in Section 5.

Bibliography

Cappelen, H. and Hawthorn, J. 2009. *Relativism and Monadic Truth*. Oxford: Oxford University Press.

Christensen, D. 2007. 'Epistemology of Disagreement: The Good News'. *Philosophical Review*, Vol. 116. pp. 187-217.

Christensen, D. 2009. 'Disagreement as Evidence: The Epistemology of Controversy'. *Philosophy Compass*, Vol. 4. pp. 756-767.

Christensen, D. 2010. 'Higher-Order Evidence'. *Philosophy and Phenomenological Research*, Vol. 81. p. 185-215.

Christensen, D. 2011. 'Disagreement, Question-Begging, and Epistemic Self-Criticism'. *Philosophers' Imprint*, Vol. 11, No. 6. pp. 1-22.

Conee, E. 2010. 'Rational Disagreement Defended' *Disagreement*, Feldman and Warfield (eds), Oxford: Oxford University Press. pp. 69-90.

Elga, A. 2010. 'How to Disagree about How to Disagree'. *Disagreement*, Feldman and Warfield (eds), Oxford: Oxford University Press. pp. 175-186.

Elga, A. 2007. 'Reflection and Disagreement'. *Noûs*, Vol. 41. pp. 478-502.

Elgin, C. 2010. 'Persistent Disagreement'. *Disagreement*, Feldman and Warfield (eds), Oxford: Oxford University Press. pp. 53-68.

Feldman, R. and Conee, E. 2004. 'Evidentialism'. *Evidentialism*, Feldman and Conee (eds), Oxford: Oxford University Press. pp. 83-107.

Feldman, R. 2007. 'Reasonable Religious Disagreements'. *Philosophers Without Gods: Meditations on Atheism and the Secular Life*, Antony (ed.). Oxford: Oxford University Press. pp. 194-214.

Feldman, R. 2009. 'Evidentialism, Higher-Order Evidence, and Disagreement'. *Episteme*, Vol. 6. pp. 294-312.

Goldman, A. 1978. 'Epistemics: The Regulative Theory of Cognition'. *Journal of Philosophy*, Vol. 75. pp. 509-523.

Kelly, T. 2005. 'The Epistemic Significance of Disagreement'. *Oxford Studies in Epistemology*, Vol. 1, Oxford: Oxford University Press. pp. 167-196.

Kelly, T. 2010. 'Peer Disagreement and Higher-Order Evidence'. *Disagreement*, Feldman and Warfield (eds), Oxford: Oxford University Press. pp. 111-174.

Kornblith, H. 1983. 'Justified Belief and Epistemically Responsible Action'. *The Philosophical Review*, Vol. 92. pp. 33-48.

MacFarlane, J. 2007. 'Relativism and Disagreement'. *Philosophical Studies*. Vol. 132. pp. 17–31.

Matheson, J. 2009. 'Conciliatory Views of Disagreement and Higher-Order Evidence'. *Episteme*, Vol. 6. pp. 269-279.

Plantinga, A. 1993. *Warrant: The Current Debate*. Oxford: Oxford University Press.