

Peer Disagreement and the Rational Course of Action

Recently there has been a question of what the rational course of action is when there is a disagreement with an epistemic peer. For the purpose of this paper I will take an epistemic peer to be someone who (1) has equal familiarity with the evidence and arguments relevant to the question and (2) who is equal “with respect to general epistemic virtues such as intelligence, thoughtfulness, and freedom from bias” (Kelly(a) p175). If two people who have about the same likelihood of judging a given set of evidence correctly, viz. arriving at the verdict that the evidence actually supports, and they each come up with different verdicts, then what should they do when faced with the disagreement? Feldman, Elga, and others argue that when faced with a disagreement with an epistemic peer we should give equal weight to each of the peers; thereby splitting the difference. On the other hand, detractors from this view, such as Kelly, say that this "equal weight view" (EWV) has unwanted consequences and therefore must be abandoned. In this paper I'm going to defend the EWV by showing how it should be applied to cases where a limited reading of the view might produce counter intuitive results.

In the first section I will present the motivation for the equal weight view. Next, I'll review why Kelly says that, if you accept the EWV, then you should accept the Uniqueness Thesis. Finally, in the third, fourth, and fifth sections I'll look at the cases where Kelly thinks the EWV produces the wrong result and respond to his objections.

§1 The Equal Weight View, Simply Stated

Kelly presents the EWV thus:

EWV: In cases of peer disagreement, one should give equal weight to the opinion of a

peer and to one's own opinion. (Kelly(b) p2)

We can see how this works in a couple of cases. First, suppose you and an epistemic peer are sitting on a jury together. You both have reviewed the same evidence, but you think the correct verdict is guilty and your peer thinks that the correct verdict is innocent. The EWV says that rationality requires that you give as much weight to your peer's opinion as to your own. Effectively this means that you should suspend judgment on the issue. Given that your peer is just as likely to judge the evidence correctly as you are, there is no reason to suppose that either view is more likely to be correct. Instead, both parties should suspend judgment.

According to Kelly, "certain kinds of cases can make [the EWV] seem almost trivial or obviously true" (Kelly(b) p3). Take a perceptual disagreement case as an example, suppose you and an equally well sighted individual are both watching the finish of a horse race. Independently you determine that Horse A won and your peer determined that Horse B won. Intuitively it seems that, upon finding out that your peer disagrees with you, you should suspend judgment about which horse won. Proponents of the EWV say that we should arrive at the same verdict for epistemic disagreements.

Applying this to a Bayesian view of rationality, we get the "split the difference" principle. For instance, suppose you and a peer both review evidence (E) that it's raining outside. Suppose you then have a .1 degree of confidence that it's raining and your peer has a .7 degree of confidence that it's raining, then, according to the equal weight view, rationality requires that you both split the difference and revise your confidence to a .4 degree of confidence that it's raining.

§2 The Equal Weight View and the Uniqueness Thesis

Kelly's first step to objecting to the EWV hinges on the need to accept the uniqueness thesis (UT) which he states as follows:

UT: For a given body of evidence and a given proposition, there is some one level of confidence that it is uniquely rational to have in that proposition given that evidence.

(Kelly(b) p8-9)

Kelly says this is a problem for proponents of the EWV in this case that he presents:

How things stand with me:

At time t_0 , my total evidence with respect to some hypothesis H consists of E . My credence for H stands at .7. Given evidence E , this credence is perfectly reasonable. Moreover, if I was slightly less confident that H is true, I would also be perfectly reasonable. Indeed, I recognize that this is so: if I met someone who shared my evidence but was slightly less confident that H was true, I would not consider that person unreasonable for believing as she does.

How things stand with you:

At time t_0 , your total evidence with respect to H is also E . Your credence for H is slightly lower than .7. Given evidence E , this credence is perfectly reasonable. Moreover, you recognize that, if your credence was slightly higher (say, .7), you would still be perfectly reasonable. If you met someone who shared your evidence but was slightly more confident that H was true, you would not consider that person unreasonable for believing as she does.

At time t_1 , we meet and compare notes. How, if at all, should we revise our opinions? (Kelly(b) p. 7-8)

Kelly rightly notes that if we accept both the EWV and the UT, then we have to say that both people need to adjust their confidence that E supports H to the same confidence level, viz. split the difference. Furthermore, Kelly says that the proponent of the EWV should accept the UT. For if this case is possible, then it seems ridiculous for the two peers to split the difference when they both already hold rational beliefs about H given E . Holding the UT gets around this by saying that both peers could not hold differing rational beliefs on the same evidence in the first place. So, according to Kelly, the proponent of the EWV should also

accept the UT.

§3 Washing Away the Evidence

Kelly's first objection is that the EWV ignores what the body of evidence actually supports. That is, if (1) you have confidence of .2 that E supports H and (2) your peer has confidence of .8 that E supports H and (3) E actually supports H to .8, then your peer should not revise her confidence down as much as you should up given the new (psychological) evidence. He says that adjusting to .5 as the EWV would have you do is not rational for either of you since the original evidence, E, actually supports .8. Kelly says that the EWV gives the wrong verdict in this case because the EWV allows the original evidence (E) to be washed out by the psychological evidence, i.e. that you have confidence of .2 and your peer has confidence of .8.

However, it's not clear how Kelly expects the fact of the matter, that E supports H to .8, to play a role in what the peers should do. For they are both just as likely to have judged the evidence correctly and each one's confidence in their level of confidence, that is, their second order confidence that E supports H, is equal. If Kelly is asking you to reevaluate the evidence when you find out about the disagreement then that won't seem to add anything since you already know that E gives you confidence of .2 in H. So, if Kelly wants to say that we should discount misleading evidence because it's misleading, he'll need to give an account of when or how we know one piece of psychological evidence is misleading and the other not. That is what evidence do we have for, for example, our judgment being clearer than the peer? On the other hand, if Kelly is saying that we should have the confidence in H that the evidence supports, then that seems trivially true, but not particularly helpful. For we want to know what we should do when we don't know what the evidence actually supports and we have a disagreement. Trivially, if we knew what the evidence actually supported, we wouldn't have a disagreement.

What this seems to come down to is a confusion about types of rationality. Cohen talks about at least two types of rationality: ideal rationality and intersubjective rationality (Cohen). I won't go into the details of

this view, but it seems like given E there is an ideally rational thing to do in the objective sense, but when the psychological evidence is introduced there is an intersubjective rational thing to do, viz. a thing to do that is rational in the sense that it's not obviously wrong (Cohen). Specifically, Kelly says,

If one were to ask which one of us should revise his or her view at this point, the answer is clear and uncontroversial: while it is reasonable for you to retain your current level of confidence, I should significantly reduce mine, since, ex hypothesi, this is what a correct appreciation of my evidence would lead me to do. (Kelly(b), p11)

But, if we take intersubjective rationality to be an acceptable way of determining what's rational, then it seems like Kelly is just assuming that the person being asked the question is not an epistemic peer. In fact, he seems to suggest the question is asked to someone who knows to what extent E actually supports H. It's not clear how this is relevant to either of the peers, unless Kelly just means that the peers should believe E (or E plus psychological evidence) supports H to the degree that E (or E plus psychological evidence) actually supports H, viz. that they should be rational in an ideal sense, but nobody would deny that there is a sense in which they are rational if they believe what E, or E plus the psychological evidence, actually supports.

§4 The Interpersonal and Intrapersonal Objection

Kelly's next objection aims to draw out the "strangeness" of the EWV in that the EWV has us disregard the original evidence in the face of psychological evidence. He says that the point is more clear in the intrapersonal case. For instance, if you suddenly realize that you have two opposing beliefs, then "it might be perfectly reasonable to resolve the conflict by dropping one of the two beliefs and retaining the other" (Kelly(b) p13). Why does Kelly say it's reasonable? Because

"[a] normative view about how it is reasonable to resolve inconsistencies among one's beliefs which completely abstracts away from facts about which beliefs are better supported by one's evidence, and which would have one treat one's prior beliefs on a par, regardless of how well- or ill-supported they

are by one's total evidence, would not be an attractive one." (Kelly(b) p13)

However, I don't think Kelly wants to say we should drop a belief without reviewing the relevant evidence in light of the inconsistency. It seems like, upon noticing that we have two opposing beliefs, the EWV gets it right. We should suspend judgment until we review the evidence and make some choice about what to believe. If that's right then the "strangeness" that Kelly talks about is not there after all.

§5 Litmus Paper vs An Epistemic Peer

Kelly then wants to "reflect on a much more general question: in what circumstances does it make sense for me to treat the fact that someone else believes as she does as evidence for the truth of that which she believes?" (Kelly(b) p19). His answer is that it is those circumstances where her belief reliably reflects what the evidence supports. He gives an example of a weather forecaster with an 80% accuracy rate. Your belief that the forecaster will accurately predict the weather is just a matter of how consistent the correlation between the predictions and the actual outcome turn out to be. Kelly then goes on to say this isn't much different from your belief that a litmus paper will turn red given a certain amount of acidic content of the liquid you dip it into. If this is right, then why, when there is a peer disagreement, does the EWV give so much weight to the psychological evidence? According to Kelly, the two types of evidence should count for exactly the same.

To further draw out the intuition Kelly imagines building a robot that predicts the weather based on some of the same evidence you use to predict the weather. Suppose your peer accurately predicts the weather 70% of the time and suppose the robot accurately predicts the weather 70% of the time on some of the same evidence. Why not treat the robot as a peer in this case? In addition, suppose both you and your peer ignore the robot when considering the evidence. That is, you both notice it, but consciously do not considerate as part of your reasoning. The EWV requires that when you and your peer become aware of your disagreement that you split the difference. However, Kelly wants to know why you weren't required to give the robot the

same weight as you did your peer. But this isn't an argument against the EWV. For the EWV makes no claim about what is rational before the peer disagreement. Kelly might still want to ask why the robot isn't just treated as a peer. It seems like the answer to that is pretty clear, it's just that the robot isn't an epistemic peer. So, since the EWV doesn't make any claims about rationality as it pertains to non-psychological evidence, there's no reason for it to address the different types of evidence in its treatment of disagreements.

§Conclusion

I've shown that the EWV, understood as a principle of rationality for peer disagreements, does not succumb to Kelly's objections. However, there are some important aspects that must be considered. First, the EWV is only meant to be a principle concerning peer disagreements and considering analogies to intrapersonal cases or differences between types of evidence should be closely scrutinized to insure that the analogy holds. In addition, it seems like the EWV requires the acceptance of at least two types of rationality, namely ideal rationality and intersubjective rationality. The responses to Kelly's criticisms in this paper relied on drawing out these two points to defend the EWV.

Bibliography

1. Thomas Kelly (2005) "The Epistemic Significance of Disagreement"
2. Thomas Kelly (forthcoming) "Peer Disagreement and Higher Order Evidence"
3. Stewart Cohen (Spring 2009) presented at ASU Epistemology Seminar