International Journal of Philosophical Studies
Publication details, including instructions for authors and subscription information:
http://www.tandfonline.com/loi/riph20

Discovering Disagreeing Epistemic Peers and Superiors

Bryan Frances [a]
[a] Fordham University, United States
Published online: 10 Feb 2012.

To cite this article: Bryan Frances (2012): Discovering Disagreeing Epistemic Peers and Superiors, International Journal of Philosophical Studies, 20:1, 1-21

To link to this article: http://dx.doi.org/10.1080/09672559.2011.629366

International Journal of Philosophical Studies Vol. 20(1), 1–21

Discovering Disagreeing Epistemic Peers and Superiors

Bryan Frances

Abstract

Suppose you know that someone is your epistemic peer regarding some topic. You admit that you cannot think of any relevant epistemic advantage you have over her when it comes to that topic; you admit that she is just as likely as you to get P's truth-value right. Alternatively, you might know that she is your epistemic superior regarding the topic. And then after learning this about her you find out that she disagrees with you about P. In those situations it appears that the confidence with which one holds one's belief should be significantly reduced. My primary goal in this essay is to present and reflect upon a set of cases of disagreement that have not been discussed in the literature but are vital to consider. I argue that in the new cases one is reasonable in not lowering one's confidence in the belief. Then I articulate and defend an ambitious principle, the Disagreement Principle, meant to answer the question 'Under what conditions am I epistemically blameworthy in retaining my belief with the same level of confidence after I have discovered recognized peers or superiors who disagree with me?'

Keywords: disagreement; epistemic peer; epistemic superior; equal weight view; conciliationism; knowledge

What should you do when you discover that someone firmly disagrees with you on some claim? Suppose you know that someone has seen all your evidence and you have seen all hers. Suppose further that you know that both of you have evaluated that common body of evidence for about the same length of time and with the same care. You also know that she is about as clever, thorough, and open-minded as you are, both generally and with respect to the issues at hand. You know that you have about the same relevant biases. At this point, before you find out her opinion on some claim P relevant to the topic, you fully admit that you cannot think of *any* epistemic advantage you have over her when it comes to the topic in question; you admit that she is just as likely to get P's truth-value right as you are (whether or not you are right about that will not matter). Let us say that under these conditions she is your *recognized epistemic peer* with regard to P (I will relax some of these conditions below). And then after learning all this about her

International Journal of Philosophical Studies
ISSN 0967–2559 print 1466–4542 online © 2012 Taylor & Francis
http://www.tandfonline.com
http://dx.doi.org/10.1080/09672559.2011.629366

you find out that she thinks P is false, whereas you had already concluded to yourself that P is true.

It might seem that in such a situation you should reduce your confidence in your belief in P. After all, two seconds ago you insisted that she was just as likely as you to get the right answer on P's truth-value; and now you have seen that she thinks P is false. Perhaps you should withhold belief entirely. Or, maybe you should split the difference with your disagreeing epistemic peer. If you were inclined to give claim P a probability of .80 and you learned that she was inclined to give it a probability of .20, maybe you should now revise your opinion to .50.

But there are even more compelling cases: I initially believe P but then learn about Smith. I fully admit that she is an outright genius and knows much, much more than I do about the issues relevant to P. I know that she has all my evidence as well as much more evidence. I also know that she is smarter than I am and has thought about and investigated P much more than I have. I know full well that when it comes to the topics germane to P she is not my peer but significantly exceeds me. Prior to finding out her opinion on P, I would have insisted that she is *much more likely* than I am to get P's truth-value right (again, it will not matter to my arguments whether this judgment is correct). Let us say that under these conditions she is my *recognized epistemic superior* with regard to P. Then I find out that she firmly believes ∼P. It may seem obvious that upon such a discovery I should at the very least reduce my confidence level in P by some significant amount, if not withhold belief entirely or move closer to her view than my old view.

In both of the situations described above, it appears that some principle such as the following applies: when one learns of disagreement with one's recognized epistemic peers or, especially, superiors over a particular held belief, the confidence with which one holds that belief should be significantly reduced if not abandoned. Principles like that one are *conciliationist*, as they claim that in certain standard cases of disagreement one must, epistemically, significantly reduce one's confidence in one's belief. Richard Feldman, Adam Elga, and David Christensen have defended theses along this line.[1] I agree with them that some principle like that one has to be true, but I also think we need to work harder on figuring out the correct principle(s).

My primary goal in this essay is to present and reflect upon a set of cases of disagreement that have not been discussed in the literature but are vital to consider when exploring the above issues and evaluating principles like the one in the previous paragraph. In the first two sections I will present and explore the new cases, which generate counterexamples to that principle. For instance, I can be fully rational and justified (in both internalist and externalist senses) in not altering my confidence level in P when I find out that the recognized epistemic superior genius

Smith thinks ∼P; in fact, after encountering her I can know P and know that she is mistaken. Then in section 3 I articulate and defend an ambitious principle, the Disagreement Principle, meant to answer the question 'Under what conditions am I epistemically blameworthy in retaining my belief with the same level of confidence after I have discovered recognized peers or superiors who disagree with me?' In section 4 I will examine more familiar cases of disagreement that seem to support conciliationism; I will argue that some of them are best understood as inconsistent with conciliationism. Finally, I will offer a few short remarks on the question 'How often and for which beliefs are the conditions of the Disagreement Principle *actually* satisfied in our lives, thereby requiring us to lower our confidence levels?'

1. The Four Cases in Which We Reasonably Don't Lower Our Confidence in Our Belief

Case 1

Suppose I start out knowing or believing (even dispositionally) nothing of what anyone thinks (even dispositionally) about P. At time t_0 I come to believe P is true. Many of our conscious beliefs are held on the basis of evidence; say that my belief in P is based on my evidence E_1. Then I learn via an amazing new poll that with regard to P I had 2,200 agreeing recognized peers, 50 disagreeing recognized peers, and 50 recognized peers who had withheld judgment after much reflection. Pretend that I know that the poll results are representative of my class of peers and I recognize each name in the list of poll participants and somehow know that they are my peers (I will try to omit the impracticality below). Since these were my recognized peers, I know that they have E_1 and E_1 alone when it comes to evidence regarding P. Note the past tense: they *were* my peers. After learning about the poll I now have additional evidence E_2 regarding P – indirect evidence, but evidence all the same. The additional evidence E_2 is what I will call 'the poll results': with regard to P and E_1 I had 2,200 agreeing recognized peers, 50 disagreeing recognized peers, and 50 recognized peers who had withheld judgment after much reflection. Presumably, I am epistemically blameless in retaining my belief in P upon learning the poll results at time t_1, as I have learned that only about 2 per cent of my (previous) recognized peers disagreed with me regarding P. I next encounter several people whom I come to recognize, by time t_2, as my current peers with regard to the issues involved in P. In order for them to be my recognized peers at *this* time, t_2, I need to know that they have seen the same evidence as I have at this time. The poll results are evidence E_2 for P's truth. So by t_2 I must know that each of them knows all about the poll results, where this

means among other things that each of them knows that all of the people polled were their peers with respect to E_1 and P. So we all have E_1 and E_2. At this point I do not know their opinions on P. I would, however, insist at t_2 that they are not only just as generally capable as I am concerning the topics relevant to P but that *they are just as likely as me to get the right answer on P itself*. Then at t_3 I learn that some of these new recognized peers believe $\sim$P.

Is it plausible to think that this discovery at t_3 of a dozen, say, new disagreeing recognized peers should make me *lower my confidence in P by any significant amount*? I do not think so; I am reasonable – in any externalist or internalist sense you please – in not altering my confidence level at t_3. Here is the reasoning process that I might well go through after encountering the new disagreeing peers.

> E_1 seems to me to go strongly in favor of P; and when we are restricted to E_1 96 per cent of my recognized peers agree with me that P is true. Thus, I have truly excellent evidence that E_1 is very strong evidence for P. I have since then encountered some people who are sticking with $\sim$P even though I know that they are aware of E_1 and E_2 but *no* other evidence. I know that E_2 does not support $\sim$P; so I guess they must think that E_1 is fantastic evidence for $\sim$P even though 96 per cent of their recognized peers think just the opposite. The two most obvious explanations for their behavior that pop in my head are these: (a) these new recognized peers have made an error regarding P and E_1 or (b) 96 per cent of my recognized peers radically misjudged P and E_1. I think it is pretty clear that (a) is much more likely than (b).

Through the reasoning process just summarized, I come to think that I have been given no good reason to reduce my confidence in P, and so I make no such reduction. *Even if I have made some mistakes in my above reasoning*, which strikes me as implausible but not outrageous, it is difficult to argue that I am being unreasonable in any interesting sense of 'unreasonable' in sticking with my initial confidence level in P upon going through the above line of thought. Now we move on to consider an analogous case with recognized epistemic superiors instead of peers.

Case 2

I start out knowing or believing nothing of what anyone thinks about P. At t_0 I come to believe P is true; this belief is based on my evidence E_1. Then I learn via a new poll at t_1 that with regard to P I had 2,200 agreeing recognized superiors, 50 disagreeing recognized superiors, and 50

recognized superiors who had withheld judgment after much reflection (I learn nothing of their more specific levels of confidence regarding P, if any there be). Suppose further that I know that the poll results are representative of my superiors. Since these were my recognized superiors, I know that they had E_1 and much more relevant evidence or ability (but of course they did not have knowledge of the poll results, because the results have just been released). Presumably, I am epistemically blameless in retaining my belief in P upon discovery of the poll results, as I have learned that only 2 per cent of the superiors I know about disagree with me regarding P. Perhaps I would not be reasonable if I had good evidence that the 50 disagreeing recognized superiors were vastly superior to the 2,250 other superiors. But let us temporarily set that possibility aside: I have no relevant judgments regarding the relative peer and superiority relations among any of those superiors, as the poll contained no such information. I next encounter several people who I come to recognize, by time t_2, as my superiors with regard to the issues involved in P. In order for them to be my recognized superiors at *this* time, after learning about the poll, I need to know that they have seen at least all the evidence I have (as well as being epistemically superior to me in some way vis-à-vis P at t_1). The poll results are evidence for P's truth that I have. So by t_2 I must know that they know all about the poll results just like I do. All this means is that they know that everyone on the list was *my* recognized superior when I was limited to E_1 (this will strike them as utterly boring information), and of course they know that 2,200 of them agree with P, 50 disagree, and 50 suspend judgment. At this point I do not know their opinions on P. I would, however, insist at this point that they are not only *more* generally capable than I am concerning the topics relevant to P but that *they are more likely than me to get the right answer on P itself*. Then at t_3 I learn that some of these new recognized superiors believe $\sim$P.

Is it plausible to think that this discovery at t_3 of a dozen, say, new disagreeing recognized superiors should make me lower my confidence in P? Well, I would probably be unreasonable upon retaining my confidence level if I had excellent evidence that the dozen new recognized superiors were vastly superior to the 2,300 old recognized superiors (e.g., imagine that I know that one of the new recognized superiors is God), but let us assume that I have no relevant evidence regarding the relative peer and superiority relations among any of the new or old recognized superiors. With such an assumption in place it seems as though I am blameless in retaining my confidence level in P after t_3. Here is the reasoning process that I might well go through after encountering the opinion of my dozen new recognized superiors.

> E_1 seems to me to go strongly in favor of P; I also know that when they are ignorant of the poll 96 per cent of my recognized superiors agree with me that P is true. Thus, I have excellent evidence that E_1 as well as other expertly available evidence is very strong evidence for P. Now I have encountered a dozen people who are sticking with ∼P even though they are aware of E_1 and E_2 and perhaps some additional evidence. Clearly, E_2 does not support ∼P. The three most obvious explanations for their behavior that pop in my head are these: (a) these dozen new recognized superiors have made an error regarding P, (b) 96 per cent of my recognized superiors radically misjudged P, or (c) these new superiors are far more informed or smart compared to the 2,300 I already know about. I think it is pretty clear that (a) is much more likely than (b) or (c).

Just like in the previous case, through the reasoning process just summarized I come to think that I have no good reason to reduce my confidence in P, and so I make no such reduction. *Even if I have made some mistakes in my reasoning*, which strikes me as implausible but not outrageous, it is difficult to argue that I am being unreasonable in any interesting sense of 'unreasonable' in sticking with my initial confidence level in P upon going through the above line of thought. Thus, even though I have just learned that some people I judge to be my recognized epistemic *superiors* with respect to the topics involved in P firmly disagree with me regarding P's truth, I am reasonable in retaining my belief in P with undiminished confidence (or a marginal change at most).

The numbers as well as the extreme conditions on peerhood are obviously somewhat unrealistic (but the conditions on recognized superiority are less so). The lack of realism is convenient but inessential. For instance, as graduate students in philosophy we often first hear about and study theses without knowing what the profession as a whole thinks of them. Later we learn what other philosophers think; we also come to have views about our abilities compared to those philosophers. For instance, suppose that working with standard philosophy of mind textbooks and articles as a beginning graduate student I come to think that analytic behaviorism is false. I have looked at a great deal of the relevant considerations but have not paid attention to what the profession thinks about behaviorism (this scenario was more realistic a few decades ago when there was not so much communication among philosophers). My belief that analytic behaviorism is false is reasonable based on the epistemically significant considerations available to me, or so let us assume. Then suppose I discover that virtually all professional philosophers of mind over the last few decades have thought that analytic behaviorism is false; I consider these philosophers my peers or superiors

with respect to the philosophy of mind generally and behaviorism specifically. Naturally, I am reasonable in retaining my belief. Later still I find out that there are a couple renegade philosophers of mind Smith and Jones who still believe analytic behaviorism is true. If I think philosophical testimony has some evidential weight I *might* find this a bit puzzling, since I know that Smith and Jones are high-quality philosophers (including being better than me) who have looked at all the considerations I and the other peers and superiors have seen.[2] But even so, I am hardly unreasonable, to any degree, in continuing to stick with my old belief with the very same level of confidence (or, a level of confidence changed a very small amount). There is nothing amiss in my inferring that it is highly probable that they made some odd mistake somewhere, despite their credentials and their superiority over me. After all, their opinion is outweighed by that of hundreds of other professional philosophers who still think that analytic behaviorism is false and do so even after looking at the arguments of Smith and Jones. I have used philosophy as an example, because readers will be familiar with it, but science will work at least as well for those who think philosophical disagreement is relevantly different; indeed, it is even more convincing in the case of science.[3]

Examples involving simple beliefs and simple bodies of evidence can make the same points that cases 1 and 2 make. Here is one:

Case 3

I am part of a study that tests elementary mathematical skills. Two hundred subjects are given a series of simple problems, such as adding up a list of two-digit numbers in just a minute. I know that all but one of the subjects are roughly equal in mathematical ability (indeed, I might be the person who conceived and competently carried out the study). I am one of the subjects and under the same test conditions as everyone else I get the answer '540' for one of the trials. I later see that a total of 192 out of the 200 subjects got '540'. Presumably, I am epistemically blameless in retaining my belief in P (P = 'The total of the numbers is 540') upon such a discovery, as I have learned that only 4 per cent of the peers I know about disagree with me regarding P. Then I find out that my best friend, who also was a subject and who is obviously much better at elementary math than I am, got '530'. In this case, I know that his superiority to me is not measured in extra evidence but in extra relevant cognitive ability. I realize that in order for him to have all my evidence, at *this* time, he has to know about the study results as well (so he knows that 192/200 got '540', 199/200 are my peers, and 199/200 are his inferiors). After learning about the poll he sticks with his opinion,

knowing that he is a mathematical whiz and concluding that the others just made an odd error.

I do not see any good reason for insisting that I must lower by any appreciable amount my confidence in my answer of '540' after encountering my friend, despite the fact that I know full well that he is my superior in mathematics (most relevantly: my superior at doing elementary arithmetic in the head), saw just the same 'evidence' that I did, was apparently completely sober and alert when taking the test, etc. Here is the reasoning process that I might well go through after encountering the opinion of my friend.

> Apparently, he thinks the odds of a simple common mistake are not that bad; I beg to differ. I think he has made two errors: first, he made an error in addition; second, he is either overestimating the odds that so many people (and percentage of people) could make the same error, or he is overestimating his own arithmetic ability. After all, I know that evidence E_1 (the numbers I am given to add up plus the instructions) seemed to me to go very strongly in favor of P; I also know (this is evidence E_2) that when we are restricted to E_1 96 per cent of my recognized peers agree with me that P is true; I also know that the math problem in question was quite simple relative to the abilities of the 200 study participants. Thus, I know that I have truly excellent evidence E_2 that E_1 is very strong evidence for P. Now I have encountered one person who is sticking with $\sim$P even though he is aware of E_1 and E_2 and has superior arithmetic ability. I know that E_2 does not support $\sim$P. The two most obvious explanations of his behavior that pop in my head are these: (a) this recognized superior of mine has made an unlikely error regarding P or (b) 96 per cent of my recognized peers radically misjudged E_1 even though the E_1-P connection is incredibly simple. I think it is pretty clear that (a) is much more likely than (b). Don't get me wrong: both (a) and (b) are very unlikely. The point is that (a) is far more likely than (b). If I knew he was the greatest mathematician of all time and everyone else who took the test including myself was really dim, then I might lower my confidence in (a). Or, if I knew he was God. But those scenarios are silly.

Through the reasoning process just summarized, I come to think that I have no good reason to reduce my confidence in P, and so I make no such reduction (or significant reduction). Even if I have made some mistakes along the way, it is difficult to argue that I am being unreasonable in any interesting sense of 'unreasonable' in sticking with my initial confidence level in '540'.

Clearly, the numbers and percentages of peers or superiors make a difference in these three cases. For instance, if you find out that 97 per cent of your 1,000 (total) superiors concerning claim P disagree with you about P, you will probably be certain that you are missing some key evidence regarding P, evidence that goes against P rather strongly. (Alternatively, you'll think you have the same evidence but have committed a performance error). If only 55 per cent of them have that position, or if it was 100 per cent but there were just three recognized superiors, then you will be much less worried. If just 2 per cent of the recognized superiors disagree with you, you probably would not be worried at all. These are psychological generalizations, not epistemic claims, but assuming the topics involved with P are ones in which evidence is dominant in determining expert opinion (e.g., science), appropriate epistemic claims hold as well. The greater the numbers and percentage, of superiors – both are important – who independently and confidently hold that your belief in P is mistaken (coupled with the assumption that the percentage is representative of the whole class of superiors), the stronger the evidence that you have missed some key and undefeated evidence against P.[4] Clearly, the response 'Well, I already know P is true; so I have a defeater for anyone who says P is false' won't be reasonable in many cases (e.g., 99 per cent of one's many recognized superiors firmly believe P is false).

In addition, some of the cases meant to support the Feldman-Christiansen-Elga line are not as straightforward as they appear. But this time there is no crucial temporal element to the case; neither is there any comparison of agreeing with disagreeing peers (or superiors).

Case 4

I am at a restaurant with friends (cf. Christensen 2007). Joe gets the bill and says that with tip the total is $289. We decide to split the bill evenly among the six of us. If I quickly do the math in my head, announce the answer $48, but then Joe says he got the answer $46, I will suspend judgment, as I should. Part of the reason: I know that I did the calculation quickly in my head, and I know that that method is not terribly reliable for me. So it is easy to tell the story in such a way one should suspend judgment. But we can also tell the story in such a way that one should not suspend judgment, or so it seems to me (and I think Christensen would agree). Suppose I actually worked out the calculation on paper, doing the long division neatly, fully, carefully, and to the first couple decimal places. I also know that I am fully sober, not sleep deprived, quite alert, etc. Up until the point at which Joe announced '$46', I would have insisted that Joe is not only just as generally capable as I am when it comes to long division but that *he is just as likely as me to get the right answer this time around* – especially

since I see that Joe is also slowly doing the calculation on a piece of paper and I have every reason to think that he is just as sober and alert as I am.

But after he announces his answer I am not going to suspend judgment, and I do not see anything odd about that. As soon as he announces his answer I will probably have just a few obviously relevant and contending explanations in front of me: (a) despite being perfectly sober and alert, and doing the calculation quite carefully, I made an error; (b) despite seeming to be perfectly sober and alert, I am either drunk or very sluggish; (c) despite my initially thinking that Joe is just as alert and sober as I am, he actually is not, or he has some other cognitive problem that is currently interfering with his ability to do long division. (Naturally, there is also the explanation that he is playing a joke on me, and thus is not really disagreeing with me at all, but let us assume that I have done whatever is needed to rule that out as extremely unlikely.) *In a realistic scenario*, I am going to reasonably think that (c) is much more likely than the others because, and this is the key point, I do not know about his alertness, sobriety, and current cognitive functioning as well as I know of my own. Frankly, given that I feel perfectly fine in every way and I went through the problem with real care, I am going to judge that it is more likely that he is not really feeling up to scratch and as a consequence has made a mistake of some kind.[5] I am not privileging myself in this judgment; I am merely going by the impersonal odds. Such a judgment might be incorrect, but it is hardly unreasonable! After all, it is a *very* simple math problem. When he counters me with 'But I'm fine! I'm perfectly sober and coherent! Really!' I am going to guess that he is really trying to save face, or he is unwittingly fooling himself. I *know* that I am perfectly well; in any realistic case I will not be capable of being as certain that he is perfectly well.

It's crucial to realize that in making this judgment in my favor I'm not treating myself differently from how I'm treating Joe; I'm not engaging in some first-person bias. Suppose the situation had been slightly different: I haven't seen the bill, I have watched Joe and Moe do the calculation with equal care and thoroughness, I know that they are peers when it comes to this type of thing, and I have seen that Joe thinks the answer is '$46' while Moe thinks it's '$48'. If my evidence that Moe is sober and alert is significantly better than my evidence that Joe is sober and alert, then I'm going to conclude that Moe is right and Joe is wrong (or, at least, I'll put significantly more weight on Moe's word than Joe's word). Now if we let Moe = me, we have the original situation described above.

I think a similar result holds for a case that Feldman claims offers clear support for his view.[6] You are in my university office and we both look out on the quad. I seem to see perfectly well a person there with a

blue coat on, and so I tell you that I see such a person. You have a good look at the quad but then say there is no person there with a blue coat on; in fact, there is no one there at all (so we are not just arguing over a borderline case of blueness or coatness or personhood or location). I know that I feel perfectly fine and that the scene I seem to be seeing looks utterly normal. Feldman says, rightly, that after encountering your opinion I should conclude that something strange is going on. But he also says that I should suspend judgment on whether there is someone in the quad with a blue coat on: 'I would not be reasonable in thinking that the problem is in your head, nor would you be reasonable in thinking that the problem is in mine'.[7] This strikes me as highly implausible (which also makes me think that Feldman might have wanted to say something else). Obviously, in any real-life case I will initially think there is no disagreement, and you are pulling my leg. If I can somehow rule that possibility out (it is not child's play to add to the example to make this realistic), then I will be faced with two obvious possibilities akin to (a)–(c) from the restaurant case: (d) I have temporarily lost my mind, even though I feel perfectly fine and everything seems utterly in order visually and with my reasoning skills, or (e) although you appear to me to be functioning perfectly well (or at least as well as me), you are functioning quite badly right now and your reports to the contrary are just plain false. Just as in the restaurant case, (e) will strike me as more likely than (d), and I do not see anything unreasonable in my making that judgment (and neither am I engaging in some first-person bias or implying any outlandish claims about the power of introspection).

My thesis thus far is about the reasonableness of a *reaction* to new information; I have not said anything about how reasonable my *initial* belief was, before I knew about what anyone thought about P. Even so, it is clear that that belief, formed in isolation, can be reasonable (it might amount to knowledge, even high-grade knowledge). Furthermore, it is clear that retaining one's confidence level upon learning about the (highly confirming) poll results in cases 1, 2, and 3 is reasonable. Thus, I *start out* by being reasonable in coming to believe P, I *retain* that reasonableness after learning about the poll or study results (in cases 1–3), and, if my arguments above are correct, I *retain* that reasonableness after encountering the new disagreeing recognized peers or superiors even if I don't diminish my confidence level to any interesting extent. Indeed, I can retain knowledge of P.

Matters are more interesting if my initial belief in P is significantly irrational or unreasonable. Suppose that in case 1 (poll of peers) I start out (before learning the poll results) with confidence level 0.8 in P when my evidence supports a level of just 0.6 (and grant the further assumptions that are needed to make this first assumption make sense). Suppose further that the poll results are depressing: 2,200 of my 2,300 recognized epistemic

peers concerning P have a confidence level of 0.9 (based on the very same evidence). So we have all made mistakes. Now things get interesting: what level of confidence in P should I have after hearing the poll results?

The question admits of many answers. On the one hand, I know I have confidence level 0.8 and I have just found out that 96 per cent of my 2,300 peers have level 0.9 based on the very same evidence. Naturally, this should – epistemically should – suggest to me that the proper confidence level is 0.9, or at least closer to 0.9 than 0.8. So, I should increase my confidence level in P.

Then again, the peers have made a large error regarding P. The evidence supplied by the poll is highly misleading. In another epistemic sense of 'should', I should not let it influence my opinion. So perhaps I should – in this alternative epistemic sense of 'should' – stick with 0.8.

And of course in yet another epistemic sense of 'should' I should decrease my confidence level, as my initial evidence supported P to level 0.6 only. So perhaps I should move to 0.6. In any case, I think these considerations are enough to show that the question 'What confidence level should – epistemically should – I have after learning of the poll results?' admits of highly divergent answers depending on what is being asked for.

2. The Epistemology of Other Relevant Beliefs in the Four Cases

Although I do not need it for my arguments thus far, I think that the disagreeing recognized peers (and superiors) can be reasonable in their $\sim$P belief *before* finding out about the poll. Using the very same evidence base E_1 two people come to opposite yet reasonable beliefs. I do not see much reason to doubt this assumption (although it is less reasonable in the arithmetic cases). I guess I can see the plausibility in the 'maximality' thesis that given a single body of evidence and proposition, there is just one *maximally* reasonable degree of confidence to have in that proposition based on just that evidence. I do not believe that thesis myself, in part because I think one highly reasonable person might be generally epistemically quite cautious while another highly reasonable person is generally epistemically risky, and this difference might generate different 'maximally' reasonable doxastic attitudes based on the same evidence (but I am not sure what 'maximally reasonable' comes to here, so I am not really sure about this). There is more on that matter below. But even if the maximality thesis is true, I think the bar for reasonable belief is lower than that for maximally reasonable belief. And, so, my assumption regarding the reasonableness of the disagreeing peers and superiors is true even if the maximality thesis is true too.

It is a bit more difficult to defend the reasonableness of the dozen disagreeing peers *after* they have come to know about the poll/study results in case 1 (in which the poll considered peers, not superiors). Consider one

of these folks, Peter the peer. He just learned that among his peers concerning the topics relevant to P, 2,200 of them disagree with him, 50 agree, and 50 suspend judgment. And he is sticking to his guns, continuing to believe that P is false. Can Peter be reasonable in retaining his belief? Keep in mind that he cannot, for instance, think that he has some stunning new evidence that others have not seen: since these 2,200 people who disagree with him are his *recognized* peers, he knows that they have seen *all* his evidence and were just as likely as him to get P's truth-value right.

Despite all that, I think he can be rational in retaining his belief. He could just say to himself 'Well, I guess this is one of those many, many cases in which people have really screwed up, despite their intelligence and efforts'. Clearly, these cases of widespread error are extremely common, and Peter can know that fact. Whether this reaction of his is rational depends in part on the height of the bar for rationality, and I think the fuzziness of 'rational' precludes anything close to a definite height. But even if he isn't being *maximally* rational in sticking with his belief in ∼P (I'm setting that question aside), *as long as he doesn't think too hard about the situation* (which I think makes it harder to be rational while sticking to one's guns), I don't suppose there is anything significantly irrational in his shrugging his shoulders and concluding that this is one of those times when the clear majority has just made a mistake – despite his also saying that these people are his peers.

It is interesting that this issue – my disagreeing epistemic peer may not be (entirely) reasonable in retaining her belief in ∼P after learning about the poll – does not apply to the superior case, case 2. Suppose that Sandy is one of the dozen superiors of mine who have learned about the poll results and are sticking with their belief in ∼P. She has learned that among *my* superiors concerning the topics relevant to P, 2,200 of them disagree with Sandy, 50 agree, and 50 suspend judgment. Sandy is perfectly free to think that she is the superior of the 2,200; she might even know this! More to the point, I can conclude (after hearing that she knows about the poll and is sticking with her belief in ∼P) that she is just a bit misinformed about her relative abilities or information relevant to P. This is consistent with my insisting that she is my superior on the relevant issues.

In cases 1–4 after I learn of the disagreement am I reneging on my judgment, made just before finding out that the peers/superiors believe ∼P, that they were my recognized peers/superiors? Before I found out their ∼P belief I judged them to be just as or more likely than me to get P's truth-value right. After I have learned their view on P and yet have retained my belief in P, I might – depending on what kind of epistemic values I have – adopt a more cautious view. With regard to case 1 (poll of peers), I may well feel pressure to renege my judgment. After all, I will have just learned that the dozen new peers continue to believe ∼P on evidential basis E_1 even though they know that about 96 per cent of

their peers think E_1 shows that P is true. This might strike me as so strange that I will be inclined to think that something is wrong with their judgment on this matter – although I will have no good reason to think them generally inferior to me on the topics at hand.

But I might be the type of person who puts very little weight on 'They are my peers and disagree with me'. This type of fact doesn't 'move me' significantly in a psychological sense. If so, then I might even *empathize* with the dozen new peers who are sticking with $\sim$P even when hearing the poll results that are so unfavorable to them. I won't 'withdraw' my initial view that the dozen new people are my peers, as I think, correctly, that I would do the same thing if I were in their situation.

In case 2 (poll of superiors) it may cross my mind that it is possible that the new superiors think that they are superior to the polled superiors. Although I have concluded that such a thought is unlikely to be true, I might well imagine how one of the new superiors might be unfortunate enough to believe it. So I retain my belief that the dozen new superiors are indeed my superiors. Case 3 is similar to case 2 on this matter.

Case 4 is different from cases 1–3 since in the former I will retract my initial view that my restaurant friend Joe was just as likely as I was to get the right answer to 'How much do we owe?' If Joe and I have both done the calculation very carefully, as described in the second variant of the case, and he still is getting a dollar amount different from the one I calculated, and I adopt reaction (c) (viz. 'despite my initially thinking that Joe is just as alert and sober as I am, he actually is not, or he has some other cognitive problem that is currently interfering with his ability to do long division'), then I am saying to myself that Joe was not in as good a position as I was to do the calculation. The same holds for Feldman's case of seeing a person with a blue coat on the quad.

3. The Disagreement Principle

Now I would like to switch gears. Here are two of the most interesting questions regarding the epistemology of disagreement:

Q1 Under what conditions am I epistemically blameworthy in retaining my belief with the same level of confidence after I have discovered recognized peers or superiors who disagree with me?

Q2 How often and for which beliefs are those conditions actually satisfied in our lives, thereby requiring us to lower our confidence levels?

A short, probably correct, but relatively uninformative answer to Q1 is 'When the learning of the fact of disagreement presents me with really good undefeated evidence that my belief is false'. A little more informative: I will be epistemically blameworthy in retaining my belief if via the discovery of disagreement I learn of *new and impressive* undefeated evidence that P is false. We can use cases 1–4 to suggest an answer to Q1 that is more interesting and informative than those two. Near the end of the essay I will make a few remarks on Q2.

In case 2 (poll of superiors), if the poll results had gone the other way, with 96 per cent of the polled superiors *dis*agreeing with me, then upon learning the poll results I would have to change my view in order to maintain reasonableness. Roughly put, when informed of the poll results in this case I will, if I have some good sense, justifiably think *there is an excellent chance I am missing something*. A little more precisely, the idea in the superiors case is this: I know that since these folks are my superiors, it is very probable that they collectively have not merely all my evidence regarding P but additional evidence; and I also know that since they disagree with me, the additional evidence (that I do not have) includes some evidence against P – evidence that must be pretty strong since it is strong enough to convince almost all of my superiors to reject P. Thus, I know that there probably is some powerful evidence out there against P, evidence that I have yet to take into account. Not only that: since they have all my evidence, I probably do not have any defeater corresponding to the evidence they have that I lack. So, I should significantly reduce my confidence in P.

On occasion, though, what I learn from the fact of disagreement is not necessarily that I *lack* some bit of evidence. Instead, my problem might be that I have failed to sufficiently *appreciate* some evidence I already have. Perhaps I have failed to 'put together' or 'properly digest' various pieces of evidence to see where they lead or amount to. This is often what happens in the peer case when I discover that 2,200/2,300 of my recognized peers disagree with me: since they have exactly the same evidence I have and yet disagree with me so thoroughly, I am probably not digesting that evidence properly.

However, it is even more complicated than that: discovering disagreeing recognized peers or superiors sometimes does not suggest that you have missed out or failed to revealingly combine some bits of evidence. In the restaurant case in which I do the calculation quickly (and would be unreasonable to stick with my initial answer of '$48'), Joe's disagreeing with me does not suggest I have missed out on or failed to combine some pieces of *evidence*, at least how I am inclined to understand 'evidence'. Instead, it suggests I made an error in reasoning with the common body of evidence. When one fails to appreciate evidence one possesses, this can be a mistake of omission; when one makes an error

in reasoning with some evidence, this can be a mistake of commission; that is the distinction I am after here (never mind the aptness of the choice of vocabulary). Thus, when I discover disagreeing recognized peers or superiors I obtain evidence that *either* I have missed out or failed to appreciate some evidence *or* I have not done so but have made an error in reasoning (such as a calculation).

But there is a fourth possibility, an intriguing one. When I find out that the recognized peers or superiors disagree with me, this might give me evidence that, although I have not made a reasoning error or missed or failed to appreciate some evidence, my 'starting points' might be faulty. Here is a case based on that idea.

Case 5

Peter the peer, Sandy the superior, and I are on a jury. Peter and Sandy say we should vote that the butler is not guilty; I think we should vote that he is guilty. I may have good reason to think that my disagreement with Peter and Sandy comes not from evidence or reasoning or differing beliefs about reaching verdicts but different epistemic starting points – where the latter could be extremely general beliefs but might be general epistemic propensities instead. Peter and Sandy are epistemically extremely cautious, being very keen on 'avoid falsehoods and going down blind alleys'; whereas I am keen on 'gather truths, knowledge, understanding, promising methods, and fertile concepts'. They naturally live their epistemic lives one way; I live mine another. After many hours of fruitful discussion I know that we have really seen the same evidence, have processed it just as well as each other, and have made no error in reasoning. It is just that the evidence is being plugged into epistemic systems that operate with different ground rules so to speak.

It is not clear (to me anyway) what I should actually *do* at this point. Am I supposed to switch to their cautious attitude? Is that even possible for me to control? And what would it even mean to suspend judgment on which epistemic attitude is best? Many other examples generate the same questions. I will not investigate those questions here, as interesting as they are. My encounter with Peter and Sandy has given me some evidence, albeit nowhere near conclusive, that some of my starting points might be faulty (we cannot say they are false, as they might not be beliefs).

Summing up the last few paragraphs, when I learn about disagreeing peers or superiors I often (depending on the details) obtain sociological evidence that (a) I have missed some crucial evidence, (b) I have failed to appreciate some evidence I already possess, (c) I have made an error in reasoning, or (d) I have faulty starting points. No doubt, I have missed some other, perhaps equally important, possibilities, but let us

start with those four. The sociological evidence in question might be good enough to require me to change my level of confidence in P; then again, it may not. It depends on the correct answer to question Q1. Perhaps this relatively precise and ambitious (i.e., almost certainly false but useful at this stage of investigation) principle is enlightening as a response to Q1:

> **The Disagreement Principle**: I should withhold my belief (or at least significantly reduce my confidence level) upon learning of peers or superiors who disagree with me if and only if that discovery gives me evidence E_X that is (a) significant for me, (b) new for me, and that (c) strongly suggests that either (1) I made an error in reasoning, (2) one of my starting points is faulty, or (3) there is some other evidence E_Y that (3a) I do not have or have not adequately appreciated, (3b) goes against P, and (3c) goes quite strongly against P.

Now I need to explain the key parts of the principle.

What I mean by 'E_X is significant for me' in part (a) is this: E_X is greater than or *comparable in force* to the body of evidence I already have regarding P. So E_X is not piddly compared to the evidence I already have. Here are a couple examples I hope make the notion I have in mind in (a) a bit clearer.

In 2011 I travel around the world and in my many, many separate encounters with swans in many, many different environments I end up discovering a total of 1,000 swans each of which is white. This is pretty strong evidence that all swans are white! But just because it is *strong* evidence does not mean it is *significant* evidence. If I have been doing this traveling for 50 years, 1961–2011, and in that half-century have similarly encountered 400,000 swans each of which was white, then the 2011 evidence is not 'comparable in force to the body of evidence I already have'. The 2011 evidence, although strong, does not appreciably change the epistemic status of my 'All swans are white' belief.

Similarly, just because you encounter very strong evidence against your belief does not mean that that evidence is significant for you. You know full well that you did not murder the maid, but you can appreciate the strength of the evidence gathered by the prosecutor who is not aware that the clever butler is trying to frame you. Thus, strong evidence need not be significant for person S, since significance-for-S is relative to the body of evidence already had by S. Perhaps the strength of a piece or body of evidence is its significance relative to another body of evidence; I will not try to sort out the relation between evidence strength and evidence significance.

Now I would like to show how the Disagreement Principle is able to account for some of our key judgments regarding cases 1–4, thereby demonstrating that it has some explanatory power.

Condition (a) from the Disagreement Principle – that the new evidence E_X is significant for me – does not hold in cases 1–3, which is *why* one need not adjust one's confidence level in those scenarios despite the fact that conditions (b) and (c) are satisfied in those three cases. In case 1 for instance (poll of peers), when at t_3 I discover that a dozen of my new peers disagree with me, although the information is new and strongly suggests I have gone wrong somewhere, this new evidence is not significant for me because I already am aware that 2,200/2,300 of my recognized peers agree with me based on E_1, and the additional evidence E_2 (the poll results) is in my favor as well.

Recall that in case 2 (poll of superiors) we noted that I probably would have been unreasonable to retain my belief in P upon first hearing about the poll results if I had thought that the 50 disagreeing superiors were vastly superior (regarding the issue at hand) to the 2,250 other superiors. Happily, the Disagreement Principle accounts for the difference: if I learn that with regard to P I have 2,200 agreeing superiors, 50 disagreeing superiors, and 50 suspending superiors, and I know or firmly believe that the 50 disagreeing ones are *vastly superior* to the remaining 2,250 when it comes to the issues at hand, then I have encountered *significant* new evidence (the poll results plus the knowledge that the 50 disagreeing ones are the superiors of the other 2,250) that strongly suggests (but does not prove) that there is some evidence (that the 50 super-superiors have) that I do not have or have not adequately appreciated, that the other superiors do not have or have not sufficiently appreciated, that goes against P, and that is quite strong (strong because it convinced those 50 super-superiors that P is false). A similar point holds for the variant of arithmetic skills case 3 in which my friend who got '530' is known by me to be an infallible god who never jokes around (never mind how I know this) or I know that the other 199 study participants (which includes me) are completely unreliable at math while my friend is a genius at math.

The Disagreement Principle gives the right results when applied to both versions of the restaurant case 4 as well. In the scenario in which I quickly did the calculation in my head, Joe's disagreeing with me makes (a) true, but when I do the calculation with real care and feel perfectly fine, then (a) is false. In the first scenario (a) is true because Joe's disagreement is being compared to the small amount of evidence I have for my belief that the portion of the bill is \$48: a quick calculation in my head. In the second scenario (a) is false because a similar but stronger bit of contrary evidence – Joe's getting the answer '\$46' and insisting and appearing to be sober and alert – is being compared to the large and strong evidence I have for my belief that the portion is \$48: my

careful calculation, my reliability, my normal behavior, plus my feeling perfectly fine.

4. Conciliationism

Conciliationism is the view that in cases in which you learn of disagreeing recognized peers or superiors, you should significantly reduce your level of confidence in your belief. That characterization is hopelessly uninformative because it ignores all sorts of relevant details. No one endorses it in that bald form. Feldman, Elga, and Christensen are conciliationists of one sort or other.

Thus far I have been arguing that many cases of disagreement are contrary to the spirit of conciliationism. However, when people argue for conciliationism they use cases for which, I suspect, they implicitly assume that prior to the discovery of a disagreeing peer or superior the protagonist *had little or no idea what other peers or superiors thought about P*. So let's consider cases that have that feature, which means setting aside cases 1–4. In the cases conciliationists focus on, in which before the discovery of a disagreeing peer or superior one was ignorant of the opinions of any peers or superiors, is the discovery of a disagreeing peer or superior enough to make one blameworthy if one does not reduce one's confidence level at least a smidgen – as Feldman, Elga, and Christensen suggest?

I think it depends on the details. If I consciously know that the belief is one that *many* of my peers have taken a position on, but I have no idea what percentage will agree with me – and notice that this is a very common occurrence – then before discovering a recognized peer who disagrees with me I can consciously know full well that, human nature being what it is, the odds are extremely high that at least some of the peers will disagree with me. Actually running into one of them will then not be at all surprising. I *already* consciously knew there were such people, so encountering one will mean nothing. In this situation it seems that I am reasonable in sticking to my guns and not lowering my confidence in P (unless I have some reason to think that the peer I've run into is indicative of the great majority of my peers, which of course is further information).

On the other hand, if before learning of any disagreement I was convinced not only of P's truth but that 99 per cent of my peers will agree with me, and then the first recognized peer I asked about P, Joe, disagreed with me on P, I would then have two obvious and competing explanations to choose from: (a) I guess I was *very mistaken* about what the evidence regarding P shows (as it is so unlikely that the first peer I meet would be in that one per cent) or (b) I guess I was *slightly mistaken* when I thought that Joe was my peer (naturally, there are other possible explanations, but they probably won't occur to me in a real-life situation). I could then judge that (b) is significantly more likely than

(a). And then I will fail to have any reason for significant adjustment in my confidence in P. Whether this judgment would be reasonable depends on how confident I am that Joe is my peer – more exactly, how confident I am that he is my peer compared to how confident I am that the evidence proves P. If I am as confident that Joe is my peer as I am confident that the evidence proves P, then I will be puzzled and no longer know what to think of either Joe's status as my peer or the quality of support P gets from the evidence we have. In such a case I probably am unreasonable if I stick to my belief in P with the same amount of confidence. If I am much more confident that he is my peer than I am confident that the evidence proves P, then once again perhaps I should significantly reduce my confidence in P. So the case can be filled in so that it is favorable to conciliationism. But if I am much less confident that he is my peer than I am confident that the evidence proves P, then I have no reason to withhold judgment on P. Now the case is unfavorable to conciliationism. I don't take it to be a counterexample, as conciliationism comes in several forms and in some cases it is ambiguous among several precisifications, but we have seen that the case can be made into a challenge for conciliationism.[8]

Fordham University, United States

Notes

1 David Christensen, 'Epistemology of Disagreement: the Good News', *Philosophical Review*, 116 (2007), pp. 187–217; Adam Elga, 'Reflection and Disagreement', *Noûs*, 41 (2007), pp. 478–502; Richard Feldman, 'Reasonable Religious Disagreements', in Louise Antony (ed.) *Philosophers Without Gods: Meditations on Atheism and the Secular Life* (OUP, 2007), pp. 194–214; Richard Feldman, 'Epistemological Puzzles About Disagreement', in Stephen Hetherington (ed.) *Epistemology Futures* (OUP, 2006), pp. 216–36. Contrary theses are argued for in Thomas Kelly, 'Peer Disagreement and Higher Order Evidence', in Richard Feldman and Ted Warfield (eds) *Disagreement* (OUP, 2010), pp. 111–74; and Thomas Kelly, 'The Epistemic Significance of Disagreement', in John Hawthorne and Tamar Gendler Szabo (eds) *Oxford Studies in Epistemology*, 1, (OUP, 2006), pp. 167–96.

2 Then again, I might be the type of philosopher who doesn't put much, if any, evidential weight on philosophical testimony. The same result follows.

3 One might think the science case is stronger than the philosophy case, if one has doubts about the connection between 'a large number and percentage of philosophers who have expertly studied P think P is true' and 'those philosophers have discovered some strong evidence for P'. I discuss the epistemology of philosophical agreement and disagreement in my essay 'Philosophy Sabotages Knowledge', which is forthcoming in an OUP volume of new papers on disagreement edited by David Christensen and Jennifer Lackey.

4 However, although that shows that numbers count, I doubt whether there are precise numerical principles at work here – principles such as 'Given that you have degree of belief A in proposition B, and then you discover that fraction

C of your representative number D of recognized epistemic superiors regarding B believe ∼B (and you have no idea regarding their degrees of belief), you should lower your degree of belief in B by such-and-such amount'. There are familiar soritical arguments that there simply must be such precise principles, but only believers in sharp cutoffs (e.g., epistemicists) accept those arguments. The rest of us assume that our concepts (e.g., of belief, of justification) have no such precision. I would like to ignore the precision issues here.

5 I'm not saying that my main evidence for my own mathematical confidence in this case is introspection. My evidence in favor of my current competence comes from my knowledge of my past mathematical competence, my knowledge that at this dinner I haven't been drinking (or my knowledge that I have not had enough alcohol to impair my mathematical ability), my knowledge that I haven't been doing anything odd (such as spilling my drink or falling off my chair or eliciting odd comments or looks from others), *and* my introspective knowledge that I feel fine in every way.

6 Feldman, 'Epistemological Puzzles'.

7 Feldman, 'Epistemological Puzzles'.

8 Thanks to the referees for their comments.