Analytic Philosophy Vol. 0 No. 0 xxxx 0000 pp. 1–16

THE PROBLEM OF PEER DEMOTION, REVISITED AND RESOLVED

ENDRE BEGBY
Simon Fraser University

In any domain of inductive reasoning, we must take care to distinguish between (i) which hypothesis my evidence supports, and (ii) the level of confidence I should have in the hypothesis, given my evidence. This distinction can help resolve the problem of peer demotion, a central point of contention in the epistemology of peer disagreement. Peer assessments are inductively formed beliefs which, if they are to be rational, must be continuously updated in light of new evidence. It is true that disagreement does not provide evidence that I am right and you are wrong about the question under consideration. But it need not, in order to lead to peer demotion: instead, significant disagreement can erode my confidence in the initial peer assessment to a point where I can no longer rationally sustain the belief that we are equally reliable. One consequence of this solution is that the line dividing the Equal Weight view from its competitors emerges as an artifact of a mistakenly one-dimensional picture of the epistemic significance of peer disagreement.

I. Peer Promotion, Peer Demotion

Suppose I have a passing acquaintance, Peter, who I am inclined to think is a bit of a dunce. But over the course of conversation one night, Peter surprises me by coming across as well-informed, sensible, and articulate. At the end of the evening, I have a much more positive view of him than previously and would accordingly be more inclined to give some weight to his opinions. Even in a case where we disagree, I might be disposed to think, "well, perhaps he's right and I'm wrong."

Call this an instance of *peer promotion*. There seems to be nothing problematic or even particularly puzzling about peer promotion so described. In fact, it seems to be a perfectly correct response to our most recent interactions: I *should* come to think more highly of Peter's intellectual capacities after that evening.

But now imagine another encounter the following week: Peter is back to pointless sermonizing and alternative-fact-peddling. In light of this new evidence, I start to think that maybe last week was an aberration, and I come to question my judgment of his epistemic character.

Analytic Philosophy Vol. 0 No. 0

From this point on, I would be less inclined to give much credence to his utterances in cases where we disagree.

Call this an instance of *peer demotion*. Unlike peer promotion, peer demotion has received a rather sizable amount of discussion in literature on the epistemology of disagreement. Why, exactly, should peer demotion be so problematic? The answer is typically something like the following: to judge that Peter is your peer is just to judge that he would be as likely as you to "get it right," even in a case where you disagree. And so, once you have judged that Peter is your peer, you cannot, without begging the question, simply take the fact of disagreement itself as evidence that it is Peter, and not you, who is wrong about *p*. And since you can't assume that it is Peter who is wrong, it follows that your disagreement also provides no ground for demoting him from peerhood.

This point is often couched in terms of a principle of "Independence."[1] David Christensen formulates this principle as follows:

> In evaluating the epistemic credentials of another's expressed belief about P, in order to determine how (or whether) to modify my own belief about P, I should do so in a way that doesn't rely on the reasoning behind my initial belief about P. (Christensen 2011: 1–2).

The Independence principle is integral to the position known as the Equal Weight view, that is, the claim that "[w]hen you count an advisor as an epistemic peer, you should give her conclusions the same weight as your own" (Elga 2007: 478).[2] As Frances and Matheson point out (2018: 5.1), the Equal Weight view holds a distinctive prominence in the current literature, in the sense that "[c]ompetitor views of peer disagreements are best understood as a rejection of various aspects of [it]." From a theoretical point of view, what primarily speaks in favor of the Equal Weight view is that it recommends quite simply that one be firmly guided by one's own prior conditional probabilities when updating one's beliefs in light of new information. So if you believe that Peter is your peer, then your prior conditional

[1] For discussion, see for example, Christensen 2007, 2009; Elga 2007; White 2009; Lackey 2010; Kornblith 2010; Comesana 2012; Kelly 2013; Cohen 2013; Lord 2013; Simpson 2013, Moon 2018.

[2] The Equal Weight view is often presented simply as a representative member of a broader family of views known as "conciliationism," and sometimes discussion proceeds as though the two terms were essentially interchangeable. I believe that these tendencies are misguided. Conciliationism is simply too broad to serve as a fruitful anchoring point of philosophical argument. In brief, I view conciliationism as a spectrum of positions that has Equal Weight at one end and "Steadfastness" as a sort of asymptotic limit point on the other end. Virtually all contributions to this discussion are, in this sense, *more or less conciliationist*: very few would hold that peer disagreement could, under no circumstances whatsoever, provide a reason (weak or strong) to rationally reconsider one's commitment to a disputed proposition.

probability is that in a case of disagreement, he is as likely to be right as you are. Peer demotion, then, appears to be just an epistemically irrational refusal to update in light of one's priors, a matter of reverting, without justification, to the assumption that he is most likely to have made the mistake.

II. Peer Demotion and "Bootstrapping"

Despite its simplicity and apparent intuitive support, the Equal Weight view has struck a number of philosophers as being deeply perplexing. Juxtaposing the putatively problematic peer demotion scenario with the apparently unproblematic peer promotion scenario can help shed new light on why. Peer promotion apparently does *not* require independent evidence.[3] Presumably, a significant part of what I have to go on in coming to believe that Peter is my peer is just that what he says *seems agreeable* to me.[4] So why shouldn't the case of peer demotion be entirely symmetrical? That is, why shouldn't the fact that what he says next week seems *disagreeable* to me constitute a reason to demote him from peerhood?

[3] At least so far as the literature gives us any reason to think. As I will argue below (fn. 13), the theoretical model developed in this paper shows how peer promotion is also subject to a form of the Independence constraint, thereby allowing us to further strengthen the sense that peer demotion and peer promotion are strictly symmetrical forms of updating on evidence from disagreement.

[4] To say that I find these judgments "agreeable" is not to say that they must agree with my own judgments prior to the disclosure of the disagreement: see Section 5 for discussion of a case in which I rationally defer to my peer's judgment precisely by noting the disagreement in question. Note also that the basis for such peer promotion would be undermined if I learned, for instance, that Peter was not in fact expressing his own beliefs, but was merely saying what he thought I would like to hear (or was in any other way "copying my answers"). Does this mean that in order to make epistemic use of Peter's avowed agreement (or disagreement), I would first have to ascertain that he has independently arrived at the beliefs he expresses? The situation here is very similar to the way that defeaters might be treated, for instance, in the epistemology of testimony (see Fricker 1994: 142–145 for an influential account). The hearer's epistemic right to take a speaker's testimony obtains so long as (i) there are no defeaters, and (ii) the hearer displays an appropriate sensitivity to the presence of any would-be defeaters. This sensitivity need not be fool-proof. Nor does one need to display one's sensitivity by first conducting an exhaustive search for defeaters (since this would be a potentially infinite task). We can certainly ask, then, if *all* I knew was that Peter's answers matched mine, would I be (defeasibly) entitled to assume that his answers express his own reasoned, evidence-based view of how things are? In many situations, I believe the answer is clearly yes: while I cannot conclusively rule out that he has copied my answers, it is significantly more likely, on my total evidence, that he has not. (Many thanks to an anonymous reviewer for raising this question.)

This is a worry because peer assessments are beliefs like any other.[5] In particular, peer assessments involve a significant empirical component. And so, if they are to be rational, they must themselves be grounded in evidence. In standard cases, the evidence in question must presumably come, at least in part, from something like an induction over that person's history of agreeing or disagreeing with me on relevant issues. My peer assessment, at any given time, reflects my best current estimate of the evidence, positive and negative, regarding your peerhood. Accordingly, the prohibition on peer demotion essentially amounts to saying that, from this (arbitrarily chosen) point onwards, I can no longer update my inductively formed credence in your expertise in light of new evidence.

But that is clearly not a recipe for epistemic rationality. No good inductive theory will insist that a judgment should only reflect evidence collected up to some arbitrarily chosen time *t*, and that any subsequent evidence must be dismissed as irrelevant. This seems to be what is motivating Thomas Kelly, for instance, when he argues that this stance on peer demotion would render our epistemic situations hostage to highly contingent facts about the temporal order in which the evidence came to us: if you first showed yourself in a favorable light, and only subsequently started spouting nonsense, I would be forced to take a very different stance on your peerhood (and consequently, on those of my own beliefs that you disagree with) than I would if you spouted the nonsense first.[6]

Nor will a good inductive theory require me to take positive evidence from agreement as indicative that you are my peer, but not permit me to take negative evidence from disagreement to indicate the contrary: this could only result in a seriously lop-sided data set. David Enoch seems to be arguing along these lines when he reminds us that "in most cases, a significant part of your evidence as to someone's reliability on some topic is her track record [. . .], that is, how often she—as you believe—got things right" (Enoch 2010: 973). If this is so, then "why on Earth should you not see Adam's belief not-*p* as reason to believe he is less reliable than you otherwise would take him to be?" (ibid.: 979).

These are serious concerns, obviously. Nonetheless, Adam Elga, perhaps the most prominent defender of the Equal Weight view, seems perfectly content to bite the bullet on this point. His defense takes the form of a reductio: any viable strategy for taking disagreement as

[5] It is important to note that the epistemological puzzles of peer disagreement arise in light of A's *belief that* B is a peer, and not simply in virtue of *the fact that* A and B are peers. A's belief that B is a peer might well be false, but still give rise to the same puzzles. Conversely, B might well in fact be A's peer, but no puzzle will arise if A has no evidence to that effect.

[6] See, e.g., Kelly 2010: 164–165; but also Gardiner 2014, for more general discussion.

evidence for peer demotion would lead to the even worse consequence of accepting *Bootstrapping*.[7]

> suppose that you and your friend independently judge the same long series of races. You are then allowed to compare your friend's judgments to your own. (You are given no outside information about the race outcomes.) Suppose for *reductio* that in each case of disagreement, you should be 70% confident that you are correct. It follows that over the course of many disagreements, you should end up extremely confident that you have a better track record than your friend. As a result, you should end up extremely confident that you are a better judge. But that is absurd. (Elga 2007: 486–487).

He concludes: "Without some antecedent reason to think that you are a better judge, the disagreements between you and your friend are no evidence that she has made most of the mistakes" (ibid.: 487).[8] Such antecedent reason is, virtually by definition, precisely what you would lack in any case where you judge your friend to be an epistemic peer (provided your peer assessment was rationally grounded in the first place).

In effect, we appear to be caught in a dilemma: either accept this highly counterintuitive restriction on peer demotion or violate Independence and allow Bootstrapping. In the remainder of this paper, I will argue that this dilemma is illusory: evidence from peer disagreement can rationally lead to peer demotion without violating the Independence principle and without licensing Bootstrapping.

III. Disentangling Two Hypotheses at Stake in Peer Disagreement

The recognition that peer assessments are grounded in evidence reminds us that we need to distinguish carefully between questions of

(i) which hypothesis my evidence supports (e.g., is Peter my peer?) and
(ii) what confidence (or credence) I should have in the hypothesis, given my evidence

In fact, it is somewhat ironic that this distinction has not been foregrounded much in the epistemology of disagreement, insofar as this literature fully trades on the notion that one's ground-level beliefs—

[7] To forestall confusion, note that the sense of "Bootstrapping" at stake here is that of Vogel (2000) and Cohen (2002), not that of Glymour (1980).

[8] I take it that Elga's requirement of an "antecedent reason" signals an endorsement of the Independence principle.

that is, the judgments about which peers are made to agree or disagree in the standard examples—should be modeled in terms of credences (ranging from 0 to 1), and not in terms of full beliefs. However, when it comes to beliefs about peerhood, most contributors seem happy to work on the assumption that they are essentially an all-or-nothing affair.[9] But we can now see why this is a mistake: if peer assessments are based in evidence, then our confidence should match the strength of the evidence. This is a perfectly general point: in any domain of inductive reasoning, I can be in a situation where, for instance, I have weak evidence for a strong hypothesis. "Fossil found in Uzbekistan suggests tyrannosaurs evolved sophisticated senses before later growth spurt," reads a recent story in *Nature News*.[10] Similarly, I can come to believe that Peter is highly knowledgeable in a particular field, based only on his performance in a Tuesday night pub quiz.

To see where Elga's analysis goes wrong, then, we must first consider the fact that my disagreement with Peter constitutes evidence which simultaneously bears on two distinct hypotheses:

> Hypothesis 1 is simply the issue under dispute, that is, whether p is the case;
>
> Hypothesis 2, by contrast, is the question of whether Peter is my peer.

Let's say I come into the situation believing both that p is the case and that Peter is my peer. Now I gain a new piece of evidence: *Peter disagrees with me about p*. Once the situation is spelled out in these terms, it should be fairly easy to see that this piece of evidence can bear on both hypotheses simultaneously. How, exactly, does it bear on each hypothesis? We will get into more detail shortly, but for now, we can simply note that it bears negatively on each hypothesis. That is, in the kind of contexts that we have been working with, learning that Peter disagrees with me about p should simultaneously lead me to become less confident in p (H1) *and* less confident that Peter is my peer (H2).

If this comes as a surprise, consider the following: what if I had discovered instead that Peter *agreed* with me about p? In that case, I should presumably become *more confident* both that p is correct and

[9] For exceptions, see White 2009; Worsnip 2014. Note also that it is important to distinguish this point from the more common observation that I can believe that a person is my peer to some degree, i.e., I think he is smart and well-informed, but not quite as smart and well-informed as I am. This is orthogonal to the point at stake here: in these cases, I believe that Peter is fully my peer, but I have relatively low confidence in that belief.

[10] Levy 2016.

that Peter is my peer. The latter is simply an application of our peer promotion scenario spelled out above. Peer demotion, then, is simply a structurally identical example, with the evidential significance flipped from confirming to disconfirming. Symmetry considerations require that I be prepared to let disconfirming evidence deflate my confidence in a hypothesis in all of those cases where I am prepared to use confirming evidence to inflate it.[11] (Notice that we haven't said anything yet about *the degree to which* the evidence should increase or reduce one's confidence in the various hypotheses at stake. We will return to the topic shortly.)

An example might help here: say I own a headphone amplifier which is supposedly capable of delivering an output of up to 5 V into 32 Ω. It's a highly regarded model from a reputable manufacturer, and I have no reason to believe that the advertised value is incorrect. Nonetheless, I am curious, and figure it would be good to be completely sure before I invest in a new pair of headphones with highly specific power demands. At a yard sale, I come across an old voltmeter for $1. The seller is in no position to guarantee its working

[11] It might be helpful also to consider how this works in a more formal perspective. Since many of these discussions are pitched in terms of Bayesian conditionalization, we can also adopt that framework. According to Bayes' rule, one's posterior probability (i.e., one's prior probability updated in light of new evidence) should equal the probability of the evidence given the hypothesis divided by the unconditional probability of the evidence, times the prior probability:

$P(H|E) = (P(E|H)/P(E)) \times P(H)$

Here's a useful pair of corollaries:

$P(H|E) > P(H)$ iff $P(E|H) > P(E)$

$P(H|E) < P(H)$ iff $P(E|H) < P(E)$

In other words, whenever the probability of the evidence given the hypothesis is greater than the unconditional probability of the evidence, acquiring this evidence should lead one to increase one's confidence in the hypothesis under consideration. Conversely, whenever the probability of the evidence given the hypothesis is lower than the unconditional probability of the evidence, one should decrease one's confidence in the hypothesis. Why? Because one obtains the posterior probability by multiplying one's prior probability with the quotient of $P(E|H)/P(E)$. Where that quotient is greater than 1, $P(H|E) > P(H)$ simply because we obtain $P(H|E)$ by multiplying $P(H)$ with a factor greater than 1. Conversely, where the quotient is smaller than 1, $P(H|E) < P(H)$ simply because we obtain $P(H|E)$ by multiplying $P(H)$ with a factor smaller than 1.

The question on which everything hangs, then, is simply whether we would be right to suppose that in these cases $P(E|H) < P(E)$. This is a question that can only be settled by recourse to substantive argument, not formal reasoning. But reflection supports this supposition: after all, the probability that a peer will *falsely* believe not-*p* (i.e., believe that not-*p*, when *p* is the case) is lower than the probability that a peer will believe not-*p*, supposing nothing about whether *p* is true. Similarly, the probability that *a peer* will disagree with me is lower than the probability that some random person will disagree with me. So, learning that Peter disagrees with me about *p* should lead me to reduce my confidence both in *p* itself (H1) and in Peter's peerhood (H2). Nonetheless, if my initial confidence in H1 was significantly higher than my initial confidence in H2, I can end up in a position where I continue to believe that *p* (albeit with less confidence than before), while I no longer believe that Peter is my peer.

condition, but I have every reason to think that these things are generally built to last. Once home, I run a test and the voltmeter reads "5 V." Once again, there are two distinct hypotheses in play:

Hypothesis 1: The advertised output is correct.

Hypothesis 2: This voltmeter is reliable.

Coming into the situation, I could be reasonably said to believe both of these hypotheses. Having determined that the voltmeter reads 5 V, I should increase my confidence in both hypotheses. But what if instead it had read, say, 3.2 V? Well, presumably I should reduce my confidence in both. That I should reduce my confidence in H1 follows from the fact that the voltmeter reading raises the evidential probability that the label on the device is wrong. That I should reduce my confidence in H2 follows from the fact that the voltmeter failed to calibrate with what I have reason to believe is a genuine reference value.

Now, consider the following: although I can reasonably be said to believe both H1 and H2, I can, in many contexts, have significantly more confidence in H1 than in H2. For instance, I have more and better evidence bearing on H1 than I have on H2, even though all the evidence that I have bearing on H2 inclines me to believe that it is true. While the 3.2 voltmeter reading also requires me to reduce my confidence in H1, the degree to which it requires me to reduce depends in part on my initial credence in H1 and my initial credence in H2. In light of the considerations offered above, it should be easy to see how my confidence in H2 can "run out" long before my confidence in H1.

This reasoning transposes nicely to the case of peer disagreement. Let's say I have strong initial confidence in p (H1) and comparatively lower confidence in *Peter is my peer* (H2). Peter now manifests his disagreement with me. I should become less confident in p as a result. By how much? Well, that's determined in part by my precise credence in his peerhood going into the situation. But at the same time, I should also become less confident that he is my peer. After all, he disagrees with me about p, and that's just one more piece of evidence to add to inductive base bearing on my assessment of his peerhood. And if my assessment of his peerhood was weakly supported to begin with, this could well be the piece of evidence that cancels out my confidence in Peter's peerhood altogether, leading me, as theorists as fond of saying, to adopt a stance of "suspended belief." My evidence, such as it is, just doesn't support any conviction either way. Accordingly, I should adopt a credence of 0.5 in the hypothesis that Peter is my peer.

In such a case, it should be apparent that Peter has been subjected to peer demotion. I did, up until his disagreement was made manifest, regard him as a peer; now that I am aware of this disagreement, I no longer regard him as a peer. Moreover, it is clear that it is the fact of disagreement which led me to demote him from peerhood. Is there a problem with this?

IV. Meeting the Standard Objections to Peer Demotion

Let us note, first, that it is not a violation of Independence: in the face of Peter's disagreement, I am not using the reasoning that led me originally to believe that *p* as a grounds to conclude that *p* must be true notwithstanding. This can be seen most easily from the fact that I am actually brought to *reduce* my confidence in *p* as a result of Peter's disagreement. This shows that I acknowledge Peter's disagreement as providing new evidence regarding *p*, moreover evidence indicating (however weakly) that my initial judgment was wrong. It is true, of course, that I rely on the reasoning that initially led me to believe that *p* in determining what to think about *Peter is my peer* in light of Peter's disagreement. But there is nothing question-begging about this: *Peter is my peer* is a separate hypothesis with its own proprietary evidence base, distinct from (though partially overlapping with) the evidence base that bears on *p*. My relying on my confidence in *p* (adjusted in the face of the present disagreement) to adjust my credence in Peter's peerhood is not a violation of Independence, but a simple matter of good inductive score-keeping.

Second, my demoting Peter from peerhood is not an instance of my irrationally failing to be guided by my own prior conditional probabilities when updating my beliefs in light of new evidence. It would be a mistake to assume that these priors are correctly specified in terms of the question, how should I adjust my confidence in *p*, given that *Peter* disagrees with me? Rather, they are specified in terms of the question, how should I adjust my confidence in *p*, given that a *peer* disagrees with me? And as we have seen, whether Peter is my peer is a hypothesis under simultaneous consideration. I also have prior conditional probabilities determining how I should update this belief in light of new evidence: what should I believe about Peter's peerhood upon discovering that he disagrees with me about *p*? Again, by parity of reasoning, if I should become more confident in his peerhood in a case where he agrees with me ("peer promotion"), I should become less confident in a case where he disagrees with me ("peer demotion").

Finally, let us consider whether the reasoning that leads to peer demotion must involve a problematic form of Bootstrapping, as Elga alleges. Recall that, according to Elga, permitting peer demotion would have the "absurd" consequence that over the course of

repeated disagreements, I should become increasingly confident that I am a better judge than Peter. But how absurd is this consequence, really? Note first that I can clearly be in a position where, after updating on new evidence, I cease to believe that Peter is my peer without thereby also coming to believe that I am a better judge than Peter. The case we looked at above, where I end up rationally suspending belief on the question of whether Peter is my peer, is just such a case: if I did believe, at that point, that I am a better judge than Peter, then I couldn't be said to have suspended my belief in that hypothesis after all (since I would, in fact, believe its negation). But what might happen subsequently as a result of yet further disagreements is a very different matter. If Peter continues to manifest disagreement with me on matters which I take to be clear as the day, then I probably *should* —even by Elga's lights—gradually come to increase my confidence that I am the better judge.[12] But this is not problematic: after all, at this point, I no longer believe that Peter is my peer.

At this point, it is not really clear how Bootstrapping is supposed to be implicated in any of this. Elga apparently assumes that peer demotion could only arise as a result of you "giv[ing] your own evaluations more weight than those of a friend who you initially count as a peer" (Elga 2007: 488). This would then permit Bootstrapping in the sense that you could illicitly become increasingly confident that "you are a better evaluator than the friend merely by noting cases of disagreement, and taking it that the friend made most of the errors" (ibid.). In the cases I have described, however, you might perfectly well start out by giving your friend's evaluations the same weight as your own when it comes to determining whether p is correct, just as the Equal Weight view requires. But at the same time, you also give your own evaluation whether p is correct some weight (as you must) in determining whether your friend is a peer after all. Now, one might think that peer demotion could only result if one sneakily entitled oneself to first update the peer estimation (in light of one's friend's disagreement whether p) and only subsequently updating on p, in light of the disagreement with a friend whom one no longer regards as a peer (but only because they disagreed with me about p). But as we have seen, this is not true. Peer demotion result could even from the reverse update sequence. That is, even after I have appropriately

[12] Though, of course, I might also be brought to consider competing hypotheses, such as "Peter is just trying to troll me."

reduced my confidence in *p* as a result of noting your disagreement (in light of my previous peer assessment), I can still be sufficiently confident in *p* that I take your disagreement as providing new evidence that you are not my peer.[13]

Here's the upshot: I am in no position to determine the epistemic significance Peter's disagreement regarding *p* until I have set my credence in Peter's peerhood. This credence can only reflect such evidence as I have regarding Peter's peerhood. Moreover, my credence in Peter's peerhood must be continuously updated in light of new evidence. Peter's disagreement with me about *p* is just such a piece of evidence. Given plausible starting points, my updating on that piece of evidence may lead me to no longer believe that Peter is my peer. But it is important to recognize that even in a case where I have fully demoted Peter from peerhood (i.e., suspended belief in the proposition *Peter is my peer*), I should still reduce my credence in *p* as a result of his disagreement. Of course, the degree to which I should reduce my confidence will be significantly less than it would be in a case where I had relatively high credence in his peerhood. But that I should reduce my confidence to *some* degree follows simply from the

[13] This suggests another interesting symmetry between positive and negative cases. Assume that I have never met Peter before, and so have no particular view about his peer status. Now Peter arrives on the scene and voices judgements matching mine. I come to think that Peter seems both smart and well-informed, and so I upgrade my credence in his peerhood. So far, so good. But what if I additionally come to increase my confidence in *p* as a result of the fact that Peter, who after all seems smart and well-informed, agrees with me about *p*? *Now* it seems that something has gone wrong. The sense in which something has gone wrong is, I believe, indicative of the sense in which peer promotion, and not just peer demotion, is subject to Independence (contrary to what we assumed earlier). But the model we have now developed shows how either mode of updating—peer promotion and peer demotion—is quite consistent with this constraint. In both cases, one should update one's confidence in *p* in light of one's *antecedent credence* that Peter is one's peer, not one's credence adjusted in light of Peter's agreement (or disagreement) regarding *p*. So parity is retained: neither case involves a violation of Independence.

There are still some complications to note, though: in many contexts (perhaps most), one should probably not be indifferent to the question of whether some arbitrary person is one's peer (to some degree). That is to say, in many domains (e.g., domains that are not highly politically charged, requiring esoteric insight, etc.), one should assume that it is more likely that a random person would get the answer right than that they would get it wrong. Therefore, when a random person turns out to agree with me on *p*, this does in fact increase the evidential probability that *p*. Second, note how things will now look for the next round: at this point, I have an evidentially grounded, positive credence in Peter's peerhood. Peter's continued agreement with me on relevant matters will now constitute stronger evidence that I am right than it did in the previous round. (And similarly, if he were to disagree with me, it would constitute stronger evidence that I am wrong. So again, parity obtains.) Finally, one might worry that an update procedure such as this is a recipe for creating an echo chamber. But note that while echo chambers are potentially pernicious epistemic phenomena (I wouldn't say, necessarily pernicious), this doesn't mean that there is no epistemically rational way that one may end up in one. On this, see, inter alia, Nguyen 2018a, 2018b, Begby 2020. (Many thanks to an anonymous reviewer for pressing me on this point.)

fact that it remains epistemically possible—that is, consistent with my evidence—that Peter is my peer, and on that possibility, his disagreement constitutes evidence (albeit weak evidence) that I am wrong about *p*. This is borne out by analogy with the voltmeter example: even though I use my confidence in the advertised voltage output as a reason to reduce my confidence in the voltmeter, the fact that the voltmeter gave a divergent result is *also* a reason to downgrade my confidence in the advertised output voltage. After all, it does raise the overall evidential probability that the advertised voltage is incorrect.

In sum, we can hold, contra Elga, that disagreements are evidence which can rationally lead me to demote Peter from peerhood without opening the door to Bootstrapping or violating Independence. The question at stake is which particular hypothesis we measure the evidence against. We can agree with Elga that disagreement is *not* evidence that I am right and Peter is wrong in our judgments regarding *p*. But it need not be in order to lead to peer demotion. Instead, it can provide evidence that I have previously overestimated Peter's peerhood.

It's important to note that this is not because I have now come to think that I previously misjudged the evidence regarding Peter's peerhood. My previous peer assessment might well have been perfectly correct in light of the evidence in my possession at the time. Rather, Peter's performance has provided me with *new* evidence, as a result of which I am now required to update my peer assessment. Quite naturally, this update can lead me to decrease my confidence in your peerhood, which is just to say that it can lead to peer demotion. Even though at this point there's a sense in which I am not more confident in my own abilities than I am in my peer's, I nonetheless *am* more confident in *my estimation* of my own abilities than in my peer's, simply because it is based on stronger evidence. Just like past instances of agreement helped build my confidence in the peer assessment, so these instances of disagreement will decrease my confidence in that assessment, maybe to a point where I can no longer functionally sustain the belief that we are equally reliable.

V. Moving Beyond the Impasse

Nonetheless, Elga's Bootstrapping argument has been remarkably successful in framing the debate between the Equal Weight view and its competitors. For instance, Kelly notes briefly that "[t]he possibility of rationally downgrading someone from the status of peer in this way will be especially apparent in cases in which one's initial judgment that the other person is a peer was itself based on relatively insubstantial evidence" (Kelly 2010: 164). This sounds promising. But it is clear from the ensuing discussion that he takes the evidence from disagreement to bear directly on the question of who is right about *p*, thereby

exposing himself to Elga's Bootstrapping argument. In fact, Kelly pretty much acknowledges this: "I hold that [. . .], because there are at least some possible cases in which such bootstrapping clearly is permissible, no view which generally proscribes it can be correct" (ibid.: 165). It is quite clear from context that the cases Kelly has in mind are precisely the cases at stake in his debate with Elga, viz. cases of peer demotion.

David Enoch (2010) is driven to a similar position: what is wrong with the Equal Weight view is that it holds that "the disagreement itself—the mere fact that Adam believes not-*p* when you take *p* to be true—is not legitimate asymmetrical evidence against his reliability" (ibid.: 979). Enoch argues that in cases of disagreement about *p*, we are entitled to take *p* to be true, and to use it "as a premiss" in an argument for peer demotion (ibid.: 980). Moreover, like Kelly, Enoch concedes Elga's main point to the contrary: the Bootstrapping problem is one that we may well have to "learn to live with" in order to get a handle on the problem of peer demotion (ibid.: 991).

We are now in a position to move beyond this impasse. We can agree with Kelly and Enoch that disagreement can provide evidence that can rationally lead to peer demotion. But we can account for this in a way that does not force us to incur the burden of allowing Bootstrapping. The key lies in carefully laying out which hypothesis the evidence is evidence for (or against). Disagreement is not evidence that I am right and you are wrong about *p*. Nonetheless, it is evidence that bears on my estimation of your peerhood: much like previous evidence from agreement led me to upgrade my confidence in your peerhood, the present evidence from disagreement can lead me to downgrade it. This follows naturally from the recognition that peer assessments are a species of inductive generalization, and so must remain continuously open to new evidence.

If I were to take the truth of a contentious belief as a premiss in an argument purporting to show that I am right and Peter is wrong about *p*, I would indeed be violating Independence and opening the door to Bootstrapping. But, contra Kelly and Enoch, this is not what is happening here, as can be seen from the fact that I fully allow that Peter's disagreement constitutes evidence that I am wrong. Instead, I take the truth of the contentious belief (or more broadly, my confidence in my own power of judgment) as a premiss in an argument purporting to show that Peter is not my peer. But this is a different matter altogether.

One might worry that this proposal will sanction a blanket apology for favoring oneself in cases of disagreement (along the lines of the "extra weight view" discussed for instance in Elga 2007, Kelly 2010, and Enoch 2010). But this worry would be ill-founded. I referred above to the fact that in "typical cases" of peer disagreement (i.e., the sorts of cases that are usually under consideration in the literature) I will, as a matter of fact, have stronger evidence for my own reliability

than for my peer's. But this is clearly a contingent matter: in other cases, the values might well be reversed. Suppose my friend is a solid amateur chess player maintaining a stable rating around 1800 ELO, with many years' experience playing in tournaments, in clubs and in online fora. For my own part, I've been playing quite a bit of online blitz lately, and I'm very pleased with my own development, having recently achieved a rating of 1850. If pressed, I would say that my analytical skills are probably as good as my friend's: after all, we are similarly rated, and that's essentially all the evidence I have to go on. Now we each sit down and strategically assess White's winning chances in a series of complex positions, say with an eye toward placing bets on an ongoing tournament. Before we disclose our analyses to each other, I might think as follows: if we agree about all or most of these cases, that should strengthen my confidence in the hypothesis that my analytical abilities are as strong as his. But what if we disagree? Surely the fact that my friend has maintained a stable rating around 1800 over a number of years and across multiple platforms is better evidence of his abilities than my having recently achieved a similar rating on a particular online chess forum is of mine. And so I would do well to defer to his judgment on White's winning chances. But as in the above cases, this is not because the disagreement itself provides first-order evidence that his analyses are better than mine; rather, it provides evidence that rationally erodes my confidence in the hypothesis that my abilities are as sharp as his. Accordingly, there is nothing in my proposal which entails that I have a presumptive epistemic right to favor my own judgments in the face of disagreement simply because they are mine.

VI. Concluding Remarks

Where does this leave us? As we have seen, the Equal Weight view retains its standing as something like a default position in the epistemology of disagreement, at least in the sense that rival accounts must be framed as departures from the Equal Weight norm. The problem of peer demotion has long been thought to constitute a central point of contention—and demarcation—in these disputes. Drawing on general principles of inductive reasoning, this paper has sought to disentangle importantly different dimensions of evidential significance arising from peer disagreement. The resulting argument shows that evidence from peer disagreement can rationally lead to peer demotion without violating the Independence principle and thereby incurring the burden of accepting Bootstrapping.

In closing, it is reasonable to ask whether the Equal Weight view could not be amended to accommodate this insight. So amended, it would amount to the view that one should give equal weight to the judgment of one's dissenting peer, *proportional to one's confidence that*

they are in fact one's peer. The strict lessons espoused by Elga would now hold only in the limiting case where I not only believe that Peter is my peer, but where my rational confidence in Peter's reliability also exactly matches my confidence in my own.[14] In all other cases, we should refrain from saying how a case of peer disagreement should impact one's confidence in p until we have tallied *all* of the evidence that bears on the situation, both one's first-order evidence bearing directly on p, and the higher-order evidence bearing on one's estimates of the reliability of the parties involved.

On the face of it, this doesn't appear like an unattractive position. So what could be the problem? In addition to forcing Equal Weighters to reconsider their conclusions in a wide range of examples that have helped give shape to the debate (such as the Horserace case above, or the Restaurant bill case), it would now become increasingly difficult to say how the resulting position differs in any interesting way for instance from the "Total Evidence view," which Thomas Kelly (2010) has espoused as a direct *competitor* to the Equal Weight view. The idea that the Equal Weight view and its competitors carve out importantly different orientations on the problem is placed under serious strain: instead, that whole dialectic emerges as an artifact of a mistakenly one-dimensional picture, apparently endorsed on all sides, of the epistemic significance of peer disagreement.[15]

Acknowledgments

In preparing this paper, I have benefitted from discussions with Holly Andersen, David Christensen, Arianna Falbo, Nic Fillion, Matt Maxwell, Mike Perry, and Somayeh Tohidi. Thanks also to an anonymous reviewer of this journal. Funding for this project was provided by SSHRC Insight Development Grant 430 2016 0460.

[14] To be fair, I think such situations exist, even though they are most likely to arise in cases where I don't think either of us should be particularly confident in our abilities. For example, consider Greebles (Gauthier & Tarr 1997): Greebles are a class of novel stimuli used in psychological experiments to help chart the development of perceptual expertise. Coming into a Greeble experiment for the first time, I probably should have *exactly* the same confidence in your ability to determine whether a particular Greeble is a Plok or a Glip as I have in my own. So if you say Plok and I say Glip, I probably should think that your answer is as likely to be right as mine.

[15] Note, for instance, how Kelly appears to hold that the import of higher-order evidence from peer disagreement is limited to that of increasing the probability that my own initial judgment regarding p was flawed. Certainly, that is an important dimension of the significance of higher-order evidence. But another dimension, in many cases as significant or more, is its bearing on my confidence in my initial assessment that you are my peer. Plausibly, Kelly's inability to see his way past the Bootstrapping problem (see Section 5 above) is precisely a result of his assumption that the significance of higher-order evidence is limited to that one dimension.

References

Begby, Endre. "Evidential Preemption," *Philosophy and Phenomenological Research* (2020): early view.

Christensen, David. "Epistemology of Disagreement: The Good News," *Philosophical Review* 116 (2007): 187–217.

Christensen, David. "Disagreement as Evidence: The Epistemology of Controversy," *Philosophy Compass* 4 (2009): 756–767.

Christensen, David. "Disagreement, Question-Begging, and Epistemic Self-Criticism," *Philosophers' Imprint* 11 (2011): 1–22.

Cohen, Stewart. "Basic Knowledge and the Problem of Easy Knowledge," *Philosophy and Phenomenological Research* 65 (2002): 309–329.

Cohen, Stewart. "A Defense of the (Almost) Equal Weight View." *The Epistemology of Disagreement*, eds. David Christensen and Jennifer Lackey. Oxford University Press, 2013.

Comesana, Juan. "Conciliation and Peer-Demotion in the Epistemology of Disagreement," *American Philosophical Quarterly* 49 (2012): 237–251.

Elga, Adam. "Reflection and Disagreement," *Noûs* 41 (2007): 478–502.

Enoch, David. "Not Just a Truthometer: Taking Oneself Seriously (but not Too Seriously) in Cases of Peer Disagreement," *Mind* 119 (2010): 953–997.

Frances, Bryan, and Jonathan Matheson. "Disagreement," *Stanford Encyclopedia of Philosophy* (2018).

Fricker, Elizabeth. "Against Gullibility." *Knowing from Words*, eds. B. K. Matilal and A. Chakrabarti. Dordrecht: Kluwer Academic Publishers, 1994.

Gardiner, Georgi. "The Commutativity of Evidence: A Problem for Conciliatory Views of Peer Disagreement," *Episteme* 11 (2014): 83–95.

Gauthier, Isabel, and Michael J. Tarr. "Becoming a 'Greeble' Expert: Exploring Mechanisms for Face Recognition," *Vision Research* 32 (1997): 1673–1682.

Glymour, Clark. *Theory and Evidence*. Princeton, NJ: Princeton University Press, 1980.

Kelly, Thomas. "Peer Disagreement and Higher Order Evidence". *Disagreement*, eds. Richard Feldman and Ted Warfield. Oxford: Oxford University Press, 2010.

Kelly, Thomas. "Disagreement and the Burdens of Judgment." *The Epistemology of Disagreement*, eds. David Christensen and Jennifer Lackey. Oxford: Oxford University Press, 2013.

Kornblith, Hilary. "Belief in the Face of Controversy". *Disagreement*, eds. Richard Feldman and Ted Warfield. Oxford: Oxford University Press, 2010.

Lackey, Jennifer. "What Should We Do When We Disagree?," *Oxford Studies in Epistemology* 3 (2010): 274–293.

Levy, Anton. "Horse-Sized Dinosaur Sheds Light on T-Rex's Origins". *Nature News*, 14 March 2016.

Lord, Errol. "From Independence to Conciliationism: An Obituary," *Australasian Journal of Philosophy* 92 (2013): 365–377.

Moon, Andrew. "Independence and New Ways to Remain Steadfast in the Face of Disagreement," *Episteme* 15 (2018): 65–79.

Nguyen, C. Thi. "Cognitive Islands and Runaway Echo Chambers: Problems for Epistemic Dependence on Experts," *Synthese* (2018a): https://doi.org/10.1007/s11229-018-1692-0

Nguyen, C. Thi. "Echo Chambers and Epistemic Bubbles." *Episteme* FirstView, 2018b.

Simpson, Robert Mark. "Epistemic Peerhood and the Epistemology of Disagreement," *Philosophical Studies* 164 (2013): 561–577.

Vogel, Jonathan. "Reliabilism Leveled," *Journal of Philosophy* 97 (2000): 602–623.

White, Roger. "On Treating Oneself and Others as Thermometers," *Episteme* 6 (2009): 233–250.

Worsnip, Alex. "Disagreement about Disagreement? What Disagreement about Disagreement?," *Philosophers' Imprint* 14 (2014): 1–20.