

PEER DISAGREEMENT AND THE LIMITS OF COHERENT ERROR ATTRIBUTION

Nicholas TEBBEN

ABSTRACT: I argue that, in an important range of cases, judging that one disagrees with an epistemic peer requires attributing, either to one's peer or to oneself, a failure of rationality. There are limits, however, to how much irrationality one can coherently attribute, either to oneself or to another. I argue that these limitations on the coherent attribution of rational error put constraints on permissible responses to peer disagreement. In particular, they provide reason to respond to one-off disagreements with a single peer by maintaining one's beliefs, and they provide reason to moderate one's beliefs when faced with repeated disagreement, or disagreement with multiple peers. Finally, I argue that, though peer disagreement is rare, the occasions on which it does occur tend to be especially important, and the kind of response supported here is correspondingly important. In particular, how leading researchers spend their time and effort depends, in part, on how they respond to peer disagreement. And only a response of the kind supported here strikes the right balance between allowing individual researchers to freely pursue what seems to them to be worthwhile projects, and requiring that they pursue those research projects that the community of experts as a whole believes to be likely to yield significant results.

KEYWORDS: disagreement, epistemic peers, epistemic reasons

1. Preliminary Issues

Say that *A* and *B* are *epistemic peers relative to question Q* if and only if all of the evidence that either one has, which bears on the answer to *Q*, is shared by the other, and their ability to reason from this evidence is identically well developed. Understand all future references to “epistemic peers” to be implicitly qualified: it is *relative* peerhood that is my concern.

Disagreement between epistemic peers gives rise to a number of perplexing epistemological questions, both theoretical and practical. My interest here is with questions that are practical, in the sense that a proper answer to them should provide epistemic agents with advice about what to *do*.

The most pressing practical question about peer disagreement might be best put like this:

(The Guidance Question) If one takes oneself to disagree with an epistemic peer, ought one to moderate one's disputed beliefs?

Where moderating one's beliefs can be understood in terms of surrendering them, reducing the confidence with which they are held, or reducing the gap between the credence that one assigns to a proposition and the credence that one assigns to its negation. Construing "belief moderation" in this way allows us to frame a single practical question that can be (and has been) answered from within a variety of epistemic frameworks.

It is worth noting that it is the *appearance* of disagreement, whether veridical or not, that is practically significant, not the fact that one actually is a party to a peer disagreement. For if a disagreement is unrecognized, or if the fact that the parties to it are epistemic peers is unrecognized, the mere fact of the disagreement need not provide either party with a reason to *do* anything.¹ Moreover, if one mistakenly takes oneself to be party to a peer disagreement, one's practical position will be the same as that of someone who recognizes that they are party to a genuine peer disagreement. Whatever reasons tell for or against belief moderation in the latter case should tell for or against belief moderation in the former case as well. I will omit the relevant qualifications below, but it is always with the appearance of peer disagreement (whether veridical or not), and with those who are taken to be epistemic peers (whether they actually are or not), that I am concerned.

We can recognize several kinds of peer disagreement. Call one-off disagreements with a single epistemic peer 'simple disagreements.' These disagreements contrast with those in which one repeatedly disagrees with a single epistemic peer, or in which one disagrees with many different epistemic peers.

Some views take it that all peer disagreements, simple and otherwise, call for identical treatment. Two such views are:

Strong Conformism: Whenever an agent is party to any peer disagreement, that agent ought to moderate his or her disputed beliefs.

and

Strong Non-Conformism: Whenever an agent is party to any peer disagreement, that agent ought not moderate his or her disputed beliefs.

¹ This fact has recently been noted in the literature. See Nathan L. King, "Disagreement: What's the Problem? Or A Good Peer is Hard to Find," *Philosophy and Phenomenological Research* 85 (2012): 252-253, Axel Gelfert, "Who is an Epistemic Peer?" *Logos & Episteme: An International Journal of Epistemology* 2 (2011): 507, and Jennifer Lackey, "A Justificationist View of Disagreement's Epistemic Significance," in *Social Epistemology*, ed. Adrian Haddock, Alan Millar, and Duncan Pritchard (Oxford: Oxford University Press, 2010), 302-303.

There are, however, positions of intermediate strength. Of particular interest will be:

Weak Non-Conformism: Most simple disagreements ought to be met with a refusal to moderate one's disputed beliefs, but one ought to react to most other peer disagreements by moderating one's disputed beliefs.

And, for the sake of completeness, we can mention:

Weak Conformism: One ought to react to most simple disagreements by moderating one's disputed beliefs, but most other peer disagreements ought to be met with a refusal to moderate one's disputed beliefs.²

Obviously these characterizations do not exhaust the space of possible responses to the Guidance Question, and the weak versions admit of degrees of strength themselves.³

Some have recently argued that peer disagreement is relatively uncommon.⁴ They are, I suspect, right about this. Peer disagreement is bound to be uncommon, if for no other reason than that epistemic peers are uncommon. But some of the situations in which peer disagreement *does* arise are of special epistemic significance. Peer disagreement is relatively likely to arise amongst leading researchers in mature branches of inquiry. There are two reasons why this is so. First, in mature branches of inquiry it is usually relatively uncontroversial which findings are significant, and leading researchers should be expected to be familiar with them. Second, in such branches of inquiry, there is a relatively robust consensus among leading researchers about how one ought to go about answering questions in their field.⁵ The former consideration provides reason to believe that the evidence sets possessed by two leading researchers in a mature branch of inquiry, that bear on some question in their field, are relatively likely to be

² Weak conformism is not a promising view, and, as far as I know, no one defends it. I mention it merely for the sake of completeness. It will be ignored below.

³ Lackey distinguishes conformist from non-conformist positions. As Lackey defines these terms, however, it would not make sense to speak of "weak" versions of these views. My use of this vocabulary, therefore, represents a slight modification of her use. See Lackey "A Justificationist View," 298-302.

⁴ This is one of the main points of King, "Disagreement."

⁵ Concerning a given mathematical theory, for example, it is uncontroversial what constitutes a valid proof. This is not to deny, of course, that there are competing mathematical theories that do not recognize the proofs of their competitors as valid. Within classical mathematics, for example, it is uncontroversial that indirect proof is valid, even though it is uncontroversial, within intuitionistic mathematics, that it is not.

identical. The second consideration provides reason to believe that their respective capacities to reason from this evidence are relatively likely to be identical.

My purpose here is to argue for weak non-conformism. Weak non-conformism has a number of virtues. First, I will argue that competing views require those who are party to peer disagreements to attribute rational errors to epistemic agents in situations in which they cannot coherently do so. See sections two and three. I will also argue, in section four, that only weak non-conformism strikes the right balance between allowing individual researchers in mature branches of inquiry to freely pursue what seems to them to be worthwhile projects, and requiring that they pursue those research projects that the community of experts as a whole believes to be likely to yield significant results. Strong conformism would too-easily force researchers to abandon promising lines of research, and strong non-conformism would allow them too much license to pursue unpromising lines of research. The time and cognitive resources of leading researchers in mature branches of inquiry have considerable social value, as their work is relatively likely to yield results of practical social significance. Moreover, their work is usually supported by governments, universities or businesses, and so its social significance is also partly a product of the claims that it makes on institutional resources. So support for weak non-conformism comes from two directions: alternative views require parties to peer disagreements to attribute errors to agents in situations in which they cannot coherently do so, and other views misallocate important social resources.

2. Simple Disagreement

2.1. Sources of Disagreement

In *some* cases of simple disagreement, it is obvious that one ought not moderate one's beliefs. These are situations in which one can tell that the other party to the disagreement holds their belief as a result of something other than the evaluation of evidence or the exercise of their rational faculties. If one has a peer who is generally rational, but who, in this instance, holds a belief because of, say, fear or prejudice, and one can tell that this is his/her reason for holding the belief in question, one is under no obligation to moderate one's beliefs.

Now, this is not to say that all disagreements that arise through non-rational causes are without epistemic significance. Those that are without significance are those produced by non-rational causes that one does not recognize as reliably producing true beliefs. Say that Mary can reliably distinguish ravens from crows on sight, but that her ability does not depend on reasoning from the

evidence that she has available, and the mechanism whereby it operates is opaque to her. If Laura knows that Mary has this ability, that Mary disagrees with Laura about whether or not a particular bird is a crow or a raven is, at least plausibly, an epistemically significant fact. And, indeed, it may require Laura to moderate her beliefs. However, even if Mary and Laura are epistemic peers, the fact that this disagreement is of epistemic significance is due not to the fact that they share evidence and reasoning abilities, rather, it is due to the fact that Mary has a particular skill that operates independently of her evidence and reasoning ability. I propose, therefore, to put such cases aside.

Let a *purely epistemic reason* be a reason for holding a belief that concerns only the goal of knowing which propositions are true and which are false. Let a *direct reason* be a reason to hold a belief that concerns only the particular proposition at issue. So, for example, my visual perception that there is a book on the table in front of me is a direct purely epistemic reason to believe that there is a book on the table in front of me. By contrast, the fact that I was offered \$100 to believe that there is a book on the table in front of me is a direct but not purely epistemic reason to believe that there is a book on the table in front of me. Finally, if it is a fact that believing that there is a book on the table in front of me will allow me to come to know many *other* truths (perhaps because I can read about them in the book), then that fact is a purely epistemic, but not direct, reason to believe that there is a book on the table in front of me.

The examples of peer disagreement that are discussed in the literature are largely disagreements in which both parties hold their beliefs for direct and purely epistemic reasons.⁶ Likewise, we can expect the most common peer disagreements – those between leading researchers in mature branches of inquiry – to be such that both parties hold their views for direct and purely epistemic reasons. It is with such cases that I will be concerned. Assume that all of the beliefs discussed below are held for direct and purely epistemic reasons, and all disagreements mentioned are disagreements over beliefs held for such reasons. What Laura ought to do, if she takes Mary's knack for telling ravens and crows apart to be as reliable as her own reasoned discriminations, is an interesting question, but not one that will be addressed here.

⁶ See, for example, Lackey, "A Justificationist View," which discusses disagreements about how to split the bill at a restaurant, King, "Disagreement," Peter van Inwagen, "We're Right, They're Wrong," in *Disagreement*, ed. Richard Feldman and Ted Warfield (Oxford: Oxford University Press, 2010), which address disagreement over the nature of free will, and Earl Conee "Rational Disagreement Defended," in *Disagreement*, ed. Richard Feldman and Ted Warfield (Oxford: Oxford University Press, 2010) which discusses disagreement about mereology.

2.2. Peer Disagreement as Involving Attribution of Rational Error

The first premise in my argument for weak non-conformism is that one must see those beliefs that one holds for direct and purely epistemic reasons as being rationally required of one. This is fairly intuitive, but an argument for it is not hard to construct. Say that Rachel believes that there is a book on the table because she seems to see a book on the table. This is, we said above, a direct and purely epistemic reason for believing that there is a book on the table. Consider what must be the case if Rachel holds this belief, but does *not* consider it to be rationally required of her. She could:

- A. Believe it, but hold that some incompatible belief is rationally required of her.
- B. Believe it, but hold that suspension of belief is rationally required of her.

or

- C. Believe it, while holding that no particular belief is rationally required, though any of several are permitted.

It is clear that A and B are unsatisfactory. If Rachel were to believe that some incompatible belief is rationally required of her – the belief that there is *no* book on the table, say – she would, in an effort to comport with her purely epistemic reasons, surrender her current belief and adopt that one. By hypothesis, this is not what she does. Similar reasoning shows that she must not hold that suspension of belief is rationally required of her. If she believed that suspension of belief was rationally required of her, then, in an effort to comport with her purely epistemic reasons, she would suspend belief. Since, by hypothesis, she does not suspend belief, we can conclude that B is also unsatisfactory.

Consider, then, the final option. If one believes that *P* for direct and purely epistemic reasons, can one take it that no particular belief is rationally required, but that any of several (among which *P* is found) are permitted? No. If any of several are permitted, then one's reasons do not tell in favor of *P* over its rivals. By hypothesis, however, Rachel takes it that her reasons *do* tell in favor of *P*. There are no other alternatives. So, one must take it that those beliefs that one holds for direct and purely epistemic reasons are rationally required of one.⁷

⁷ On a related point, see Roger White, "Epistemic Permissiveness," *Philosophical Perspectives* 19 (2005): 445-459. White argues that epistemic theories that do not require a unique belief of agents, given some evidence, are problematic because they would permit agents to choose

Now, what one is rationally required to believe is a matter of one's evidence and one's ability to reason from that evidence. Therefore, one significant result of the fact that epistemic agents must see the beliefs that they hold for direct and purely epistemic reasons to be rationally mandatory for them is that they must, likewise, see those beliefs as rationally mandatory for their epistemic peers.⁸ Those with different evidence may believe differently, and those with more- or less-well developed reasoning faculties may believe differently, but those with identical evidence sets, and equally well-developed reasoning faculties, are rationally required to share beliefs.

This fact is important. It means that one must see at least one party to any peer disagreement in which one takes part as having suffered a failure of rationality. That is, if one must see one's direct and purely epistemic reasons as rationally mandating that one hold some particular belief, one must see at least one party to the dispute as holding a belief that that person's reasons forbid him or her from holding.

Now, attributing a failure of rationality to someone does not require saying that they are not, in general, rational. Being rational, in the relevant sense, is a matter of having some ability; it is an ability to, perhaps, reason well from one's evidence. And abilities are such that those who genuinely possess them can, in isolated cases, fail to successfully exercise them. For example, if Derek Jeter, a great baseball player, has the ability to hit curve balls, the fact that he swung at, and missed, a curve ball in the game this afternoon does not mean that he does not, in fact, have the ability to hit curve balls. Since he often does hit curve balls, one does better in explaining his failure this afternoon by saying that he suffered from a "performance error": that some unusual feature of the game this afternoon prevented his ability from being exercised successfully. In general, failure to exercise an ability does not show that one does not have the ability in question. So attributing a failure of rationality to someone does not require attributing *irrationality* to that person, it requires attributing only a *rational performance error* to him or her.

That being said, any failure to exercise a skill is some (however weak) *evidence* that the agent does not possess the skill. That Jeter missed a curve ball might be *best* explained by attributing to him a performance error; nevertheless,

randomly between permitted beliefs, but that randomly deciding what to believe is never rationally permissible.

⁸ Conee denies a related claim (see Conee, "Rational Disagreement," 71). He does not require two people to have identical evidence sets in order to count as "epistemic peers," however, so he is actually discussing a different (though related) issue.

that he missed the curve ball on this occasion is *some* reason, though a very weak one, to believe that he does not have the ability to hit curve balls, despite what we may have previously thought.

Now, the significance of this fact is that each time one attributes a performance error to an agent, the case becomes stronger that the agent has not been suffering from performance errors, but that he or she does not, initial appearances to the contrary, possess the ability in question at all. So there may be no trouble in explaining Jeter's failure to hit curve balls on one or two occasions by attributing performance errors to him. But if he *repeatedly* misses curve balls – if he misses, say, twenty in a row – those repeated failures are an indication that our attributions of performance errors to him were mistaken, and that we ought, instead, to have determined that he does not have the ability to hit curve balls. So it is with the attribution of rational performance errors. One can charitably explain a single instance of failed reasoning by attributing to the agent a rational performance error, but with each additional attribution, further attributions become less tenable. Having attributed, say, twenty rational performance errors to an agent (over a short period of time, or without interruption), one is forced to conclude that one's previous attributions of rational performance error were mistaken, and that the subject is not, despite one's original belief to the contrary, generally rational.

These considerations tell in favor of some kind of non-conformism. Conformism, of any kind, requires that one (often or always) respond to simple peer disagreements by moderating one's beliefs. Doing so requires attributing a failure of rationality to oneself: it requires saying that the belief that one holds, or the credence-level that one assigns to it, is, in fact, prohibited by one's evidence. Now, this failure of rationality can be, indeed, must be, thought to be a mere rational performance error. But each attribution of a rational performance error to oneself makes any additional such attributions less reasonable than those that came before. Just as Jeter's repeated failure to hit curve balls is some evidence that he is *unable* to hit curve balls, so one's own failure of rationality is some evidence that one is not generally rational.

The problem with conformism, and so the root of the case for non-conformism, is that attributions of rational error to oneself are, in a certain way and to some degree, *self-undermining*. In short, the problem is that attributions of rational error to oneself give one evidence that one is not, in general, rational, and if one has (sufficiently strong) evidence that one is not, in general, rational, one ought not to trust one's own judgments. In particular, one ought not trust *that*

judgment. That is, one ought not trust one's own judgment that one ought not trust one's own judgments.

Look at the situation in the following way. Prior to encountering an epistemic peer who disagrees with you, you have some confidence in your own rational faculties. If you respond to the disagreement by attributing a rational performance error to yourself, you thereby admit that there is some reason to believe that this confidence was misplaced. Now, the determination that you committed the rational error was a product of your rational faculties. So the net effect of this line of reasoning is that your rational faculties provide you with some reason to believe that your rational faculties are unreliable. That is, they provide you with a reason not to trust the reasons that they provide for you, *including this one*. And this is incoherent.

So one must not, on pain of incoherence, attribute too many rational performance errors to oneself. But, because simple disagreements are not repeated, there is no similar problem with attributing rational performance errors to those with whom one disagrees. Since one's peers, by definition, enjoy no advantage in evidence or reasoning ability, the balance of reasons usually supports attributing the error to the other party and refusing to moderate one's beliefs. Notice, however, that weak non-conformism leaves open the possibility that there are some cases of simple disagreement, limited in number, in which moderating one's beliefs may be the thing to do. If one attributes only a small number of rational performance errors to oneself, the evidence that they provide to the agent that she is generally unreliable may be quite weak. And so, given that she has attributed only a small number of rational performance errors to herself, she can say that she is, typically, reliable, but that she happened to make a mistake in some particular case.⁹

Now, a complete account of the epistemic significance of peer disagreement should be able to sort those cases in which one may attribute the error to oneself

⁹ Compare this to Egan and Elga's view. They say: "News of unreliability can come by way of the very faculty whose reliability is called into question. The news need not completely undermine itself, since the reasonable response can be to become confident that the faculty is unreliable, but has happened to deliver correct news in this instance." (Andy Egan and Adam Elga, "I Can't Believe I'm Stupid," *Philosophical Perspectives* 19 (2005): 82.) If the faculty at issue, however, is the one whereby *this* judgment (that the faculty is unreliable but has, in this instance, provided correct information) is made or evaluated, then the course of action recommended by Egan and Elga is incoherent. For in that case, one should not trust that this judgment was correctly made. As far as I can see, in order for Egan and Elga's view to work, there would need to be "nested" faculties: for each (or, more reasonably, for some) there would be another faculty that checks to see if it provided the right verdict in each case.

from those in which one must not, but my project is not so ambitious.¹⁰ To support weak non-conformism, it suffices to show that most simple peer disagreements ought not be met with belief moderation, but that most non-simple peer disagreements ought to be. I take it that the case has been made for the former claim. The latter claim will be vindicated below.

3. Complex Disagreement

3.1. Repeated Disagreement with a Single Epistemic Peer

Elga has an argument that was designed to show that parties to a peer disagreement are not permitted to prefer¹¹ their own views, to those of the other parties to the dispute. Call this his “no-bootstrapping” argument. Say that Adam takes Katie to be his epistemic peer. He finds, however, that they repeatedly disagree with each other. If Adam is, on each occasion, permitted to prefer his own view to Katie’s, he can inductively infer that he is more reliable (on the matters under dispute) than is Katie. The inference is simple: he was right and Katie was wrong the first time they disagreed, as well as the second time, and the third time, and so on. Therefore, he will probably be right in the future, and Katie will probably be wrong. But it is obvious that this inference is no good: the mere fact that he prefers his own views to Katie’s does not give him grounds for

¹⁰ It may be possible to combine the present view with other views currently in the literature, to yield a complete account of the significance of peer disagreement. For example, Lackey holds that belief moderation is required in inverse proportion to the degree of justified confidence with which one’s disputed belief is held (see Lackey, “A Justificationist View”). One might accept this, but use the fact that one ought not, in most cases, moderate one’s beliefs in light of simple disagreements, to fix a lowest bound of justified confidence, above which belief moderation is not permitted. That is, one could order one’s beliefs by the degree of justified confidence with which they are held, and take the line that separates those that may be moderated given a simple disagreement, from those that may not, to be given by the fact that, if one moderates all of those beliefs below the line, as well as at least one above it, one will be able to draw the incoherent conclusion that one’s rational faculties are unreliable. I do not intend to endorse this strategy here; only to note that weak non-conformism is consistent with some answers to the Guidance Question that may already be found in the literature, and to point to the possibility that it might be usefully combined with them.

¹¹ Elga cashes out this notion of “preference” in terms of the relative weights that are to be given to the credence levels that the various parties to the dispute assign to their beliefs, which are then averaged. But this is optional. For example, one could understand “preferring one’s own views” in terms of refusing to suspend judgment when faced with peer disagreement, and the argument will work just as well.

inferring that he is more reliable than she is.¹² So, in such cases, one must not be permitted to prefer one's own views to those of one's peers.

Now, we must be careful not to draw a stronger conclusion from Elga's argument than it warrants. Notice that Adam has the opportunity to inductively infer that he is more reliable than Katie only if they *repeatedly* disagree with each other. And so if the no-bootstrapping argument tells against conformist views, it only tells against strong versions of conformism. It does not tell against views that permit agents to stick with their beliefs only when they encounter simple disagreements, or those that permit them to only infrequently prefer their own views when faced with repeated disagreements. Simple disagreements are a one-time affair, and so they provide no firm inductive base from which to argue that those with whom one disagrees are at an epistemic disadvantage. So, if Elga's argument is sound, and if the argument of the last section, which shows that one usually has good reason to refrain from moderating one's beliefs when faced with simple disagreement, is any good, then we have an argument for weak non-conformism. The view supported is non-conformist because my argument from the last section indicates that simple disagreements should usually be met with a refusal to moderate one's beliefs, but it is a weak version of non-conformism, because Elga's argument indicates that repeated disagreements call for moderating one's beliefs.

Unfortunately, the situation is not as simple as this suggests. I will argue that Elga is right about repeated disagreements: they usually ought to be met with belief moderation. But his no-bootstrapping argument likely applies to fewer cases of repeated peer disagreement than it may initially appear to, and so other considerations must be found to fill the gap. Consider again the case of Adam and Katie. Adam initially judges that Katie is his epistemic peer. However, whenever they disagree Adam holds that his view is correct and that Katie is in the wrong. I will argue that this may or may not give Adam grounds for inductively inferring that he is more reliable than Katie, but that, even when it does not, Adam still has reason to moderate his beliefs.

Consider the following.¹³ There are computer programs that have been designed to produce "academic papers" appropriate for one discipline or another,

¹² See Adam Elga, "Reflection and Disagreement," *Noûs* 43 (2007): 486-487, for Elga's presentation of the argument.

¹³ This thought experiment was inspired by the Sokal Hoax. Alan Sokal, a physicist, wrote a paper that was a "farrago of deliberately concocted solecisms, howlers and non-sequiturs" (according to Paul Boghossian, "What the Sokal Hoax Ought to Teach Us: The Pernicious Consequences and Internal Contradictions of 'Postmodernist' Relativism," *Times Literary Supplement*, December 13, 1996, 14), and submitted it to the journal *Social Text*, where it was

simply by stringing together jargon, and citing authorities in the field. Imagine that one of these papers has been submitted to a journal, and that Paul has been asked to review it. Paul fails to notice that it is without intellectual merit, and recommends that the journal publish it. Prior to reviewing the paper, Paul had a certain level of confidence in his ability to reason well concerning matters in his field. When it comes to light that the paper is of no intellectual merit, how ought Paul to respond? It would seem that the proper response is for Paul to reduce his confidence in his own competence in the field. If he could reason at all well, about matters in his field, he should be able to differentiate papers that deserve to be published from those that were produced by a computer program that produces imitations of academic papers.

Return now to Adam, who believes that Katie is his epistemic peer, but who repeatedly disagrees with her. Each time they disagree, Adam refuses to moderate his beliefs, and maintains that Katie has made a mistake. Prior to disagreeing with her, Adam had some degree of confidence in his own reasoning ability, and he was just as confident in Katie's reasoning ability (this is part of what it is to take her to be his epistemic peer). Now, after many disagreements, Adam inductively concludes that he is more reliable than Katie. That is, he inductively concludes that, despite his original opinion to the contrary, Katie is not really his epistemic peer. What ought this conclusion do to Adam's confidence in his own reasoning abilities? I would like to suggest that it should reduce it. Just as Paul's inability to sort meritorious papers from those that are hoaxes should reduce his confidence in his own abilities in his field, so Adam's inability to sort those who are his epistemic peers from those who are not, should reduce his confidence (though to a smaller degree) in his.

So, the practical effect of reducing Adam's confidence in his own reasoning abilities should be to reduce his confidence that he has judged correctly that he is right, and Katie is wrong, about the matters under dispute. After a small number of disagreements, his opinion of Katie's abilities ought to decline only a small amount, and so he can conclude that Katie is *almost* his epistemic peer. In this situation it is not difficult to see how he could have thought that she really is his peer, and so his confidence in his own abilities need be lowered only a small amount. But after they have disagreed many times, he must conclude that her

published. Sokal's purpose in perpetrating the hoax was to expose abuses of science in some branches of the humanities, but I think it can also be useful to reflect on the epistemic position of those who agreed to publish his paper. See Alan Sokal, "Transgressing the Boundaries: Toward a Transformative Hermeneutics of Quantum Gravity," *Social Text* 46/47 (1996): 217-252.

reasoning abilities are substantially impaired. And an explanation of how he could have mistaken someone whose reasoning abilities are substantially impaired for an epistemic peer requires reducing his confidence in his own abilities a great deal. Taken far enough, this is incoherent, because it would require him to judge that his earlier judgments (that he was right and Katie was wrong) were the product of an unreliable reasoning faculty.

For the sake of clarity, it may be useful to assign values to the confidence levels that Adam assigns to his own reliability, and that of Katie. All numbers, however, are for the purposes of illustration only. Say that, prior to disagreeing with Katie, Adam believes that *P* is the case, and he is 80% confident in his ability to judge correctly in such matters. He finds that Katie, whom he takes to be his epistemic peer, disagrees with him. Adam decides not to moderate his beliefs, and so concludes that Katie is wrong. He can explain Katie's false belief by ascribing to her a rational performance error. But they also disagree about *Q* and *R*, say. Again, he refuses to moderate his beliefs, with the result that he has the basis for an inductive argument that he is more reliable on these matters than is Katie. But now he finds that he was previously unable to sort those who are his epistemic peers from those who are almost his epistemic peers, which calls into question his ability to reason well about the matters under dispute. Say that being unable to sort peers from almost-peers ought to reduce his confidence in his own reasoning abilities by 10%.

So after disagreeing with Katie about *P*, *Q* and *R*, Adam is 72% confident in his ability to accurately judge such things. Now say that he finds that Katie also disagrees with him about *S*. Again, he refuses to moderate his beliefs. After *this* disagreement, he must say that he was unable to sort those who are his peers from those who are quite far from being his peers. And if he is unable to do so, his confidence in his own abilities ought to drop by another 10% (call this Adam's 'epistemic discount rate', and remember that numbers in this example are for illustrative purposes only). So his confidence in his own abilities falls to 64.8%. And so on, each time they disagree. As his opinion of Katie drops, so his opinion of himself should drop. Once his opinion of himself has fallen far enough, he should no longer be confident that he was correct when he maintained his belief in *P* against Katie's opposition. Adam might be able to maintain his obduracy if he and Katie disagree only a few times, but as the number of times that they disagree

increases, it becomes more and more rational for Adam to moderate his beliefs, *whatever* he thinks about Katie's reasoning abilities.¹⁴

Adam can claim that he is more reliable than Katie only if he finds (what he takes to be) many occasions on which he has arrived at a true belief whereas Katie arrived at a false one. But given that he previously took Katie to be his epistemic peer, as these cases proliferate, his confidence that they *are* cases in which he arrived at a true belief ought to fall. Whether or not there are any cases of repeated peer disagreements in which one can bootstrap one's way into an argument that one is more reliable than one's peers depends on what one's epistemic discount rate is.¹⁵ If the discount rate is low, then Elga's no-bootstrapping argument shows that one ought to moderate one's beliefs when one repeatedly disagrees with an epistemic peer. If the discount rate is high, one cannot inductively infer that one is more reliable than those with whom one disagrees, but then the considerations surveyed here provide independent support for the claim that, when one repeatedly disagrees with an epistemic peer, one usually ought to moderate one's beliefs. Either way, repeated peer disagreement usually calls for belief moderation.¹⁶

3.2. Disagreement with Many Epistemic Peers

Lackey takes one of the principle weaknesses of non-conformism to be that it does not recognize that when many of your epistemic peers hold their views independently, and yet disagree with you, you have reason to moderate your beliefs.¹⁷ Only strong non-conformism, however, is required to say that one need not moderate one's beliefs in this situation, weak non-conformism has room to say

¹⁴ We should note that Adam can avoid concluding, on this basis, that he is generally unreliable, by moderating his beliefs before his opinion of his own abilities falls below the threshold beyond which that conclusion is rationally mandatory.

¹⁵ If the epistemic discount rate is very high, then preferring one's own opinions to those of one's peers will never permit one to argue that one's own judgments are more reliable than those of one's peers. Imagine, for example, that the discount rate is 100%. Then as soon as Adam has enough evidence to judge that Katie is not his peer, he ought to lose all confidence in his own reasoning abilities. Now, of course no one's discount rate is 100%, but it need not be that high in order to undermine Adam's confidence in his own beliefs.

¹⁶ As any peer disagreement is uncommon, we should expect repeated peer disagreement to be extremely uncommon. If it is, then moderating one's beliefs when one encounters it should not force one to draw the incoherent conclusion warned against in the last section.

¹⁷ Jennifer Lackey. "What Should We Do When We Disagree?" in *Oxford Studies in Epistemology: Volume 3*, eds. Tamar Szabó Gendler and John Hawthorne (Oxford: Oxford University Press, 2010), 278.

that when many of your epistemic peers disagree with you, you ought to moderate your beliefs.

Weak non-conform, unlike the strong version, takes it that peer disagreement is of *some* epistemic significance. One reason that it is significant is that disagreement calls for explanation. When it is widespread it can be more easily explained by attributing to oneself a rational performance error than by attributing such an error to many other people. Consider: there must be some mechanism whereby a performance error can be explained. If Jeter fails to hit a curve ball, we can say that he suffered from a mere performance error, but in order to do so we must suppose that there is some fact that caused him to fail to successfully employ his (typically reliable) ability to hit curve balls. Likewise, if a rational performance error is to be attributed to someone, we must suppose that there is some particular (and unusual) reason that that person was unable to successfully employ his or her rational faculties.

But if the dispute is with many (independent) epistemic peers, it is more plausible that an explanation of a performance error in the single individual can be found, than it is that an explanation of many performance errors in those who disagree with him/her can be found. Say that I am speaking to a group that I take to consist of my epistemic peers. If they all disagree with me about some claim that I make, it is more likely that I am, say, tired, and so have, in this case, reasoned poorly, than that everyone in my audience is tired, absent outside information to that effect. So in a situation in which a single agent disagrees with many epistemic peers once, the agent has good reason¹⁸ to attribute the performance error to him/herself,¹⁹ and so moderate his/her views in light of the view held by the other parties to the dispute.²⁰

¹⁸ Although this reason can be outweighed if, for example, one has good evidence that the other parties to the dispute are unusually tired.

¹⁹ As noted above, this cannot be done too many times. But as a practical matter (and it is with practical matters that I am concerned), this need not trouble us. Even simple peer disagreement is relatively uncommon. Disagreements with groups of epistemic peers should, then, be expected to be extremely uncommon. So it is unlikely that disagreements with groups of epistemic peers would push one to attribute outright irrationality to oneself.

²⁰ Assuming, of course, that they do, or at least, appear to, share a view. If one disagrees with many epistemic peers, *who disagree with each other*, the case for retaining one's prior belief gets stronger. In such a case, rational performance errors must be attributed to many people, no matter what happens. The more thoroughly the community of epistemic peers is fragmented, the more reasonable it is to attribute the performance errors to others and not to oneself.

4. The Social Advantages of Weak Non-Conformism

That it is incoherent to believe that your own rational faculties are unreliable provides the rationale for some version of non-conformism. Elga's no-bootstrapping argument, as well as the argument of section 3, shows that it is a weak version of this view that we must opt for. I would now like to argue that there is substantial utility, in particular, substantial *social* utility, in weak non-conformism. And this utility is not shared by strong conformism, *nor* by strong non-conformism.

We said above that we can define a notion of a *relative* epistemic peer. Relative epistemic peers are probably rare, but it is reasonable to believe that they occupy some socially significant epistemic roles. King argues that novices to a discipline are fairly likely to be relative epistemic peers,²¹ but of special interest here is the fact that leading researchers in some fields are also fairly likely to be relative epistemic peers. Fields that are most likely to produce relative epistemic peers are those in which there is widespread agreement on how to answer questions in that field, or within a particular sub-branch of it, as well as widespread familiarity with the results of other researchers whose work bears on whatever question might be at issue.²² These features are characteristic of "mature" branches of inquiry.²³

Now, strong conformism has some undesirable social consequences in such fields. It would have the leading scholars in mature branches of inquiry moderate their beliefs whenever their epistemic peers (other leading scholars) disagree with them. This gives individual scholars with idiosyncratic views too much power to direct the research of their colleagues. Imagine that Alice and Beth are leading researchers in a mature branch of inquiry. Alice, along with most of her colleagues, believe that *P*. Beth disagrees, and has been unable to win any converts to her view. If strong conformism is true, then Alice should respond to this disagreement by either suspending her belief on the question of whether or not *P*, reducing her confidence in *P*, or reducing the gap between the level of credence that she assigns to *P* and that which she assigns to not-*P*. Reacting in any of these ways would stifle research. Moderating her beliefs would force Alice to abandon,

²¹ King, "Disagreement," 255.

²² These are some of the features of what Kuhn calls 'normal science.' They are, according to Kuhn, characteristic of most periods of research in well-developed branches of inquiry. (See Thomas Kuhn, *The Structure of Scientific Revolutions* (Chicago: University of Chicago Press, 1962), chapters 3-4.)

²³ So we should expect to find more epistemic peers about questions concerning, for example, mathematics, and fewer concerning questions about morality.

or make it more likely that she would abandon, a promising line of research, simply because one colleague disagrees with her. The practical problem with strong conformism is that it assigns too much weight to the views of idiosyncratic individuals, and those individuals can then close off promising avenues of research for others.

Strong non-conformism suffers from a complementary problem. Strong non-conformism never requires epistemic peers to moderate their beliefs in the face of disagreement. Imagine again a leading scholar in a mature field of inquiry who has idiosyncratic views. As far as we, collectively, can tell, pursuing those views is not a promising way to arrive at true beliefs. As far as we, collectively, can tell, the idiosyncratic scholar is wasting his or her time, and thereby misallocating important epistemic resources. Strong non-conformism, however, would allow him or her to pursue his/her idiosyncratic research program. Ideally, our leading scholars would pursue all and only promising avenues of research. Strong conformism is objectionable because it closes off promising avenues of research; strong non-conformism is objectionable because it allows researchers to pursue unpromising ones.

Weak non-conformism, by contrast, allows us to avoid both problems. A few cases of simple disagreement will not compel a leading scholar to abandon his or her research project. That is, as long as the leading researchers in the field largely regard a project as promising, weak non-conformism will not prevent a researcher from pursuing it. But frequently repeated disagreement, or disagreement with many (independent) epistemic peers, will compel a researcher to abandon or re-direct the course of his or her research. And it *should*. In such situations we no longer have reason to trust that the research that the scholar is pursuing is a wise use of epistemic resources (most notably, his or her time and effort). Perhaps research could once have been conducted by a solitary individual. But that is no longer true. Most of our significant research is now done in an institutional setting, whether supported by universities, grants, or private businesses, and so the allocation of the time and effort of researchers is now of considerable social and institutional significance. One advantage of weak non-conformism is that it works to efficiently employ these resources.

5. Conclusions

There are limits to how much rational error one can attribute to an agent, whether oneself or another. This fact tells in favor of a weak non-conformist response to the Guidance Question. I have argued that when faced with a simple disagreement over a belief held for direct and purely epistemic reasons, one

usually ought to refuse to moderate one's beliefs. Repeated disagreements with a single epistemic peer, or a single disagreement with multiple epistemic peers, over a belief held for direct and purely epistemic reasons, I have argued, should be approached differently. In such cases, one usually has good reason to moderate one's beliefs. If these arguments are any good, together they support weak non-conformism.

Additional support for weak non-conformism is found in the fact that plausible alternative answers to the Guidance Question encourage the misallocation of social epistemic resources. Research communities ought not to have the direction of their work changed by the influence of lone dissenting members. One weakness of strong conformism is that it requires members of the larger research community to moderate their beliefs when faced with a single member that disagrees with the group consensus. Strong non-conformism faces a complementary problem. If there is a lone member of a research community that pursues a program that his or her peers see as misguided, s/he ought to let that fact shape his/her research. But strong non-conformism allows idiosyncratic researchers to ignore the opinions of their peers. Weak non-conformism, on the other hand, yields the correct result in both cases. According to weak non-conformism, simple disagreements typically do not require those who are party to them to moderate their beliefs. So researchers whose work represents the group consensus need not change their views in light of a single disagreement with a single peer. Moreover, those who disagree with a group of epistemic peers typically have a reason to moderate their beliefs, so the lone dissenter typically has reason to bring his or her views more closely in line with those that are taken, by his or her peers, to be on the right track.

Notice, finally, that weak non-conformism also leaves room for genuinely innovative research to make an impact in the research community. Peer disagreements arise only when the parties to the disagreement have identical sets of evidence and identically well-developed reasoning abilities. Innovative research can make an impact in the community when it *improves* either our evidence or our reasoning abilities. Newton deserved recognition by the community of physicists not because he was a dissenting peer and they were obligated to moderate their beliefs in light of his disagreement with them, but because he provided a *better* way to think about physics. That he was *not* their peer was what mattered. Weak non-conformism does not allow those who *merely* have novel views from misdirecting research, but it does not prevent innovators who are the epistemic *superiors* of other leading researchers (whether because they have better

evidence, or because they can reason better from their evidence) from directing the course of research.

Much has been made, recently, of the fact that peer disagreement is rare. Indeed, one gets the impression that this observation is supposed to serve as evidence that it is less important than has typically been thought. But the importance of peer disagreement is not a function merely of its frequency, but it is also determined by the importance of the individual disagreements that we *do* encounter. Peer disagreement is unusual, if only because there are few epistemic peers. But I think that we ought not infer, from that fact, that it is unimportant. In fact, those relative peer disagreements that do occur are most likely to occur among leading researchers in mature branches of inquiry, and how those disagreements are resolved is of significantly more importance than, say, how disagreements over how to split a restaurant bill are resolved. How leading researchers approach the practical questions raised by peer disagreement determines, in part, how efficiently resources spent on research and development will be employed. The Guidance Question is a pressing question, not because of the frequency with which it is asked, but because of what turns on the answer given, in the few cases in which it must be answered. The limits on the coherent attribution of rational error suggest that a weak non-conformist answer is appropriate, and giving some such answer also helps ensure that the resources that will be deployed, in one way or another, depending on how we answer the question, are not wasted.²⁴

²⁴ An early draft of this paper was presented as a part of the 2012 Prometheus Seminar Series at The Johns Hopkins University. I would like to thank everyone who attended that seminar for the helpful feedback that they provided.