EPISTEMIC PEERHOOD, LIKELIHOOD, AND EQUAL WEIGHT

Marc Andree WEBER

ABSTRACT: Standardly, epistemic peers regarding a given matter are said to be people of equal competence who share all relevant evidence. Alternatively, one can define epistemic peers regarding a given matter as people who are equally likely to be right about that matter. I argue that a definition in terms of likelihood captures the essence of epistemic peerhood better than the standard definition or any variant of it. What is more, a likelihood definition implies the truth of the central thesis in the debate on peer disagreement, the so-called *Equal Weight View*, according to which we should give the opinions of our peers the same weight we give our own. Adopting a likelihood definition, however, does not end the debate on peer disagreement, because the alleged theoretical alternatives to the Equal Weight View, reinterpreted in the light of a likelihood definition, can in fact be shown to be compatible with this view—though the reinterpreted versions may appear less plausible than the original ones.

KEYWORDS: disagreement, peers, Adam Elga, likelihood, epistemic rationality

1. Introduction

A considerable part of this paper is based on a footnote by Adam Elga. More precisely, it is based on footnote 21 of Elga's seminal paper "Reflection and Disagreement". Since this footnote is so central to my line of argument, I quote it at full length:

> My use of the term 'epistemic peer' is nonstandard. On my usage, you count your friend as an epistemic peer with respect to an about-to-be-judged claim if and only if you think that, conditional [on] the two of you disagreeing about the claim, the two of you are equally likely to be mistaken. On more standard usages, an epistemic peer is defined to be an equal with respect to such factors as 'intelligence, perspicacity, honesty, thoroughness, and other relevant epistemic virtues' (Gutting 1982, 83), 'familiarity with the evidence and arguments which bear on [the relevant] question', and 'general epistemic virtues such as intelligence, thoughtfulness, and freedom from bias' (Kelly 2005). In defense of my use, suppose that you think that conditional on the two of you disagreeing about a claim, your friend is more likely than you to be mistaken. Then however intelligent, perspicacious, honest, thorough, well-informed, and unbiased you may think your friend is, it would seem odd to count her as an epistemic peer

> with respect to that claim, at least on that occasion. You think that on the supposition that there is disagreement, she is more likely to get things wrong.[1]

For my purposes, it does not matter to which passage of Elga's text the footnote is attached. What does matter is that defining epistemic peerhood in terms of *likelihood* reveals the essence of the concept. Elga's exact specification of this definition appears defective, however, and several arguments that rely upon it—among them those for his main thesis, namely that you have to take into account your peers' opinions only if their views are by and large sufficiently similar to your own—are in fact incorrect.

Besides indicating the false consequences Elga draws, however indirectly, from his peerhood definition, I explain some further implications for cases of peer disagreement, not drawn by Elga, that derive from defining peerhood in terms of likelihood. The astonishing result is that the central thesis in the debate on peer disagreement, namely the so-called *Equal Weight View*, according to which we should give the opinions of our peers the same weight we give our own, is easily seen to hold. Furthermore, this is, contrary to appearance, not the end of the debate.

In short: section 2 is a detailed reflection on the issues touched on in the footnote. Sections 3 and 5 are each dedicated to a point at which my account of peerhood departs from Elga's as presented in the footnote. The upshot is, in section 3, that his revised version of the *Equal Weight View* is unmotivated, and, in section 5, that his main thesis is wrong. In section 4, the Equal Weight View is proved, and it is shown that its main alternatives are in fact compatible with it. Finally, in section 6, some loose ends are picked up, and the actual limits of the applicability of the term 'epistemic peer' are pointed out.

2. Two Definitions of Epistemic Peerhood

In the footnote, Elga states, and defends, a definition of epistemic peerhood that he calls 'nonstandard'. The standard definition derives from Gutting, to whom the term 'epistemic peer' is commonly attributed, and is usually quoted in its most concise version from Kelly.[2] Both are mentioned by Elga. According to this standard definition, an epistemic peer is an equal with respect to a certain number

---

[1] Adam Elga, "Reflection and Disagreement," *Noûs* 41 (2007): 499. All brackets except the first pair are in the original.

[2] Gary Gutting, *Religious Belief and Religious Skepticism* (Notre Dame: University of Notre Dame Press, 1982); Thomas Kelly, "The Epistemic Significance of Disagreement," in *Oxford Studies in Epistemology: Volume 1,* eds. Tamar Gendler and John Hawthorne (Oxford: Oxford University Press, 2005).

of factors that influence a person's ability to judge a given matter. The various versions of the standard definition differ slightly as to which factors are included in the list. Following Kelly, we can roughly group the candidates into two categories, labelled, for example, *familiarity with the relevant evidence* and *general epistemic virtues*, or, as I will mostly refer to them in what follows, *well-informedness* and *competence*. Well-informedness may be taken to include, besides knowledge of all relevant facts, an awareness of all relevant lines of argument, enough time and willingness to consider them properly, and access to equipment such as scrap paper or calculators that may help to process the information. Competence may be taken to include, besides intelligence and expertise, virtues such as thoughtfulness, thoroughness, open-mindedness, intellectual courage, ingenuity, and incorruptibility. Arguably, not every potentially relevant factor can correctly be subsumed under one of the two headings. For example, it may seem implausible that well-informedness should embrace sufficient time or willingness. However, the implausibility of such a subsumption would not speak against the basic idea of defining epistemic peerhood by giving a list of factors on which candidates have to be equals.

Given this standard definition (or list definition, as I will occasionally call it), the central question in the debate on peer disagreement is whether it is reasonable to stick to one's belief when one encounters a peer who differs. It seems that the peer, being equally well-informed and competent, might just as easily be as right as oneself. If so, it is hard to justify why one should not revise one's beliefs in favour of an agnostic position when one is faced with a peer disagreement. Yet abandoning, for instance, some political or ideological belief just because some peer fails to share it appears spineless and submissive rather than deliberate and reasonable.

Before we see what Elga's non-standard definition can teach us about this puzzling situation (and what it cannot), let me clarify several notable characteristics of the standard definition, which are only implicitly mentioned by Gutting and Kelly, if at all. First, an epistemic peer with respect to one matter need not be an epistemic peer with respect to another. Our definition does not allow us to take two persons to be peers *simpliciter*; peerhood has to be *relativised to a subject (or a proposition)*. In addition, it should also be *relativised to an occasion (or a time)*, for one's degree of well-informedness concerning a given question as well as one's level of competence might change over time. One might, for example, gain extra evidence, acquire new skills, or forget formerly known

facts.[3] Thus, the very same persons may be peers regarding a given proposition at one time but not at another.

Second, evidence does not include sensations and intuitions. More precisely, it does not include the *phenomenal character* of these sensory or rational seemings; it does not include *what it is like*, for example, to see the Niagara Falls, or to feel the conclusiveness of Gettier's argument. It does include everything that can be communicated, namely the content of a perception or intuition, the fact that one perceives or intuits this content, facts about one's feelings, and so on.[4] The phenomenal character, however, is incommunicable; telling you how the Niagara Falls look does not bring about in you the same feeling you would have if you saw them. Similarly, merely emphasising how convincing Gettier's argument appears does not have the same persuasive effect as simply stating it and letting you judge for yourself. Evidence does not include qualitative experience because it is difficult to see how incommunicable experience might help to establish whether some proposition on which we disagree is true (but see the discussion on the *Extra Weight View* in section 4).

Third, 'being an equal with respect to certain factors' does not mean that a peer must be an equal with respect to *each* of the factors. It is *overall equality* that is required, not equality in every respect. This implies in particular that sameness of evidence, which is often taken to be obligatory for peerhood,[5] is not necessarily required. Although this overall equality specification admittedly makes it more difficult to assess whether two given persons are peers, it is a natural qualification of the concept of peerhood that preserves what is valuable, namely that the

---

[3] See Jonathan Matheson, *The Epistemic Significance of Disagreement* (New York: Palgrave Macmillan, 2015), 24.

[4] In other words, evidence is propositional. See Timothy Williamson, *Knowledge and its Limits* (Oxford: Oxford University Press, 2000), 194–200, for a defence of this not uncontroversial view.

[5] Cf. e.g. the role that 'full disclosure' of evidence and arguments plays in Richard Feldman, "Epistemological Puzzles about Disagreement," in *Epistemology Futures*, ed. Stephen Hetherington (Oxford: Clarendon Press, 2006), or the line of argument in Jennifer Lackey, "Disagreement and Belief Dependence: Why Numbers Matter," in *The Epistemology of Disagreement*, eds. David Christensen and Jennifer Lackey (Oxford: Oxford University Press, 2013), 254, which implicitly depends on the assumption that peerhood implies sameness of evidence. On the other hand, Feldman remarks that a peerhood definition requiring exactly identical evidential possessing would be useless because it could not be met in any real case of disagreement. See Richard Feldman, "Evidence of Evidence is Evidence," in *The Ethics of Belief*, eds. Jonathan Matheson and Rico Vitz (Oxford: Oxford University Press, 2014), 288.

relevant persons are, all things considered, in an equally good epistemic position to evaluate some proposition.

Fourth, note that peers need not be experts: two equally well-informed and competent persons might easily be two ignorant fools. While the most intriguing cases of peer disagreement are clearly those in which the peers have extensive skill and knowledge in the relevant field (for it is in these cases that confidence in our beliefs is most affected), the definition does not rule out cases in which the parties to the disagreement never had any reliable justification for their respective beliefs.

Fifth, it is sometimes argued that there are hardly any peers at all in non-idealised cases of disagreement because equality in both possessing and processing evidence is difficult to establish in real-world scenarios.[6] However, even if this is right, it does not prove the debate on peer disagreement to be pointless. Someone fails to be my epistemic peer by being either my *epistemic superior* or my *epistemic inferior*, and if there is any difficulty in discerning which, they are at most *slightly* superior or inferior. Under the assumption that we should give our peers' opinions the same weight we give our own, it seems plausible that we should give the opinions of those who are only slightly superior only a little more weight than our own and those who are only slightly inferior only a little less. Under the assumption that we should not give our peers' opinions the same weight, however, it seems plausible that we should not give the opinions of those who are only slightly superior or inferior to us almost the same weight. Whatever the insights of the peer disagreement debate may be, they seem to carry over to other, more asymmetric cases of disagreement. Hence the debate has a clear impact even on real-world disagreement.

Finally, some authors who define peerhood in the standard way have factors on their list that I find problematic. According to Elgin, peers have to have the same background assumptions; according to Vorobej, they have to have similar and mutually intelligible manners of reasoning as well as comparably good track records.[7] Roughly, a track record is an account of former successes and failures acquired over the course of many related performances. For example, if we often

---

[6] See e.g. Nathan King, "Disagreement: What's the problem? or a Good Peer Is Hard to Find," *Philosophy and Phenomenological Research* 85 (2011).

[7] See Catherine Elgin, "Persistent Disagreement," in *Disagreement*, eds. Richard Feldman and Ted Warfield (Oxford: Oxford University Press, 2010), 53; Mark Vorobej, "Distant Peers," *Metaphilosophy* 42 (2011), 711. More precisely, Vorobej distinguishes remote peers, who satisfy Kelly's definition but neither of his two additional constraints, from distant peers, who satisfy Kelly's definition and one of the additional constraints—it does not matter which—and perfect peers, who satisfy all conditions. He then goes on to argue that we should respond to different types of peer differently.

discuss tomorrow's weather and notice that we are almost equally reliable when it comes to guessing whether it will rain, we have a comparably good track record. The reasons why background assumptions, methodological preferences or track records should be excluded in defining peerhood closely connect to the reasons why the standard definition, in whatever variant, is problematic. To these reasons I turn now.

Recall that, according to Elga, "you count your friend as an epistemic peer (...) if and only if you think that (...) the two of you are equally likely to be mistaken."[8] The first thing to note here is that this is not a proper definition. It does not state necessary and sufficient conditions for *being* an epistemic peer, rather it states conditions for *counting* someone as an epistemic peer. So let us instead assume that a peer is defined as *someone who is antecedently equally likely to be mistaken*, and let me note three further details concerning this definition before we proceed.

First, we need a word like 'antecedently' in the definition. Assume, for instance, that I am in a much better epistemic position regarding p than you are, but that the two of us have as yet neither formed nor exchanged any belief about p. Then my *prior* likelihood of being right about p is higher than yours. Assume further that, once we have made up our minds regarding p, we happen to agree that p is true. Then my *posterior* likelihood of being right about p is the same as yours—after all, our beliefs are identical. The term 'antecedently' signals that, in the definition, likelihood is to be understood as prior likelihood.

Second, throughout this text, 'likely to be right' is not to be understood as 'likely to hit the truth' but as 'likely to hit the view best supported by the available evidence' (the same holds *mutatis mutandis* for similar expressions). This means, for instance, that a person who arrives at a true belief due to some misinterpretation of what is in fact a deceptive body of evidence is not right; a person who correctly interprets the misleading evidence and hence arrives at a false belief, on the other hand, is right.

---

[8] Although this definition is still non-standard, some authors have joined Elga in defining peerhood in terms of probability. Moffett, for example, regards a definition such as Elga's as "a very plausible account of the notion of an epistemic peer" (Marc Moffett, "Reasonable Disagreement and Rational Group Inquiry," *Episteme* 7 (2010), 357), Enoch defines a peer as "someone who is, somewhat roughly, antecedently as likely as you are to get things right (on matters of the relevant kind)" (David Enoch, "Not just a Truthometer: Taking Oneself Seriously (but Not too Seriously) in Cases of Peer Disagreement," *Mind* 119 (2010), 956), and White defines peers as equally reliable persons (Roger White, "On Treating Oneself and Others as Thermometers," *Episteme* 6 (2009), 235).

Third, if one talks about likelihood, one should specify what kind of likelihood one means. Elga refers in his definition with 'likely' to the subjective probabilities of the alleged peers. As a result, he defines under which conditions a person considers another person to be his or her peer. (In section 5, we will see where this noteworthy peculiarity leads him.) Another option would be to interpret 'likelihood' as objective probability. However, this kind of probability is ontologically obscure and epistemically difficult to access. A better alternative is to refer to the subjective probability of a neutral observer. Hence, the likelihoods are determined by the relevant credence functions that such an observer would have. Invoking a neutral observer is in fact a parallel to a conventional list definition, because by ascribing or denying peerhood to persons on the basis of such a definition, regardless of whether or not those persons take themselves to be peers, we act like impartial outsiders who aim at judging with maximal neutrality (which of course does not mean that we cannot be wrong).

A peer, I said, is someone who is antecedently equally likely to be mistaken. What reasons could one possibly have to prefer this characterisation in terms of likelihood over a list definition? Elga tells us in the footnote:

> [S]uppose that you think that conditional on the two of you disagreeing about a claim, your friend is more likely than you to be mistaken. Then however intelligent, perspicacious, honest, thorough, well-informed, and unbiased you may think your friend is, it would seem odd to count her as an epistemic peer with respect to that claim, at least on that occasion.

Here, Elga lets us imagine that the two definitions come apart; that, according to his *likelihood* definition, your friend is not your peer but rather your inferior, while, according to some suitable *list* definition, your friend may very well be your peer. Then, he maintains, the result we get from his definition obviously trumps the result we get from the list definition. This seems correct, because anyone who is less likely than you to judge the truth value of some proposition correctly is *ipso facto* not your equal in judging that truth value and should thus not be regarded as your peer. In other words, equal likelihood is necessary for peerhood.

Sufficiency is harder to establish, and is not argued for by Elga (although he takes it for granted). To see a problem with the claim that an equal likelihood of being right is sufficient for peerhood, imagine two people A and B and some highly theoretical proposition p, which A considers to be true and B considers false. While A is an expert on the relevant field, B is merely a layman. B's reason for denying p is her knowledge that C is a well-known expert regarding the matter under consideration, and that C believes $\sim p$. Because of this piece of testimonial

evidence, *B*'s belief is in fact antecedently as likely to be right as *A*'s. Thus we have a scenario according to which the likelihood definition tells us that *A* and *B* are peers regarding *p*, while *A* is clearly more knowledgeable than *B* and hence seems to be epistemically superior.[9]

Compare this scenario with the following variant of Christensen's well-known restaurant case:[10] in a restaurant, we agree to give a 10% tip and split the bill evenly. Then you and I each calculate how much everyone has to pay. You are excellent at doing maths in your head; I, who normally perform poorly at this kind of task, use a calculator. Since we have often done computations in this fashion and compared the results, we know that it is as likely for you to make a mistake as it is for me to enter a wrong number. Are we peers?

The question, asked in this way, is ambiguous. We are clearly not peers regarding mental maths. We are, however, fully peers regarding this specific calculation. This is because whether one finds the correct result depends not only on one's calculating ability; access to useful equipment may also help. It should thus be covered by a good and detailed list definition. And since it is overall equality that matters, and not equality in every respect, a lack of competence can be compensated for by the use of technical means.

It seems prima facie fairly plausible not to consider access to testimonial evidence when assessing peerhood; likewise, though to a minor degree, it might appear reasonable not to take the use of a calculator into account. Concerning other resources, it seems reasonable to a still minor degree to disregard them in assessing peerhood. The crucial point now is that no categorical gap seems to lie between reliance on one kind of resource and reliance on another. Surely, sufficient time to consider the evidence properly should be on our list of factors, for insufficient time affects the respective peer statuses. What non-arbitrary justification could we have to include sufficient time but not access to sufficient scrap paper? Or access to sufficient scrap paper but not to relevant measuring equipment, or electronic means? None, it seems. If so, the apparent specialty of some resources is easily resolved. Even access to expert knowledge is, when you come to think of it, merely a means of obtaining a certain result. (And of course there is nothing special about relying on other human beings rather than

[10] David Christensen, "Epistemology of Disagreement: the Good News," *Philosophical Review* 116 (2007).

machines. Since computer programmes are able to answer astonishingly many questions correctly and serve increasingly often as experts, it would be ad hoc to claim a categorical difference between consulting human experts and consulting computers.) Therefore B, in the example above, should indeed be regarded as A's peer, and 'access to expert knowledge' should be included in a good list definition.

Admittedly, this answer may appear counterintuitive. Keep in mind, however, that B is A's peer only regarding one single proposition, p; with regard to closely related issues, A is probably much more likely to be right. The oddity of counting A and B as peers decreases once we see clearly that we ascribe or deny peerhood only with respect to an extremely narrowly limited subject matter.

(Why not define epistemic peerhood relative to a field of knowledge, rather than a proposition? The answer is that our definition would then be less significant. For if A and B disagree on p, and we know both that A is B's superior on the area of knowledge to which p belongs, and that they are equals with respect to p, then the latter, more specific fact is the decisive one; it defeats the information that we get from the less specific fact. Having said this, I concede that under normal circumstances, it suffices to know whether two people are equals with respect to a certain area of knowledge, because we can quite reliably deduce from this whether or not they are peers with respect to specific propositions in that area.)

In sum: seemingly obvious counterexamples to the sufficiency of equal likelihood of being right for peerhood do not in fact show what they are intended to show; quite the contrary: they help to reveal how peerhood should be understood, and how closely it has to be tied to equal likelihood. If a list definition therefore yields a different result from the one we get from a likelihood definition, so much the worse for the list definition.

But could a thoroughly formulated list definition really yield a different result? For the sake of argument, take a list definition that includes familiarity with the relevant evidence and arguments, sufficient time and willingness to consider the evidence, access to whatever equipment is helpful in processing the evidence, intelligence, expertise, freedom from bias, sobriety, honesty, thoroughness, open-mindedness, intellectual courage, and creativity. Compare this definition to a likelihood definition. How could the two ever come apart? In order to see how, consider

> DAY OF BIRTH. The day of the birth of my first child has finally arrived. My wife has been in labour for hours, and there is still no end in sight. So, on the midwife's advice, I go to a nearby restaurant for a quick meal with my in-laws. I barely eat anything. Never in my life have I been so excited. To calm my nerves, I reach for the bill and calculate what each of us would have to pay if we gave a

> 20% tip and split the total amount evenly. But I find it difficult to concentrate on the computation, and ultimately arrive at a sum quite different from the one my in-laws (who joined me in this game) have figured out. Recalculation proves that I have mixed up the numbers terribly.

There are several points to make here. First, something like 'freedom from extraordinary excitement' is not on our list of factors. Thus, if we are of the opinion that my extreme nervousness on the day of birth and the resulting computational impairment makes me my in-laws' inferior regarding mental arithmetic, the list is incomplete. In support of regarding me as *epistemic inferior* in DAY OF BIRTH, we can adduce that, given my excitement, I am clearly less likely than my in-laws to do the calculation correctly. Hence, either DAY OF BIRTH represents a scenario in which our list definition yields a different result from a likelihood definition like Elga's, or DAY OF BIRTH shows that our list definition is defective. And this defectiveness cannot be easily resolved. Of course, we can simply add 'freedom from extraordinary excitement' to our list and thereby make it immune to the specific counterexample presented by DAY OF BIRTH. But the general problem is that one can easily invent other scenarios that disclose further characteristics that are missing from our list. Philosophers are ready to come up with counterexamples that show the significance of hitherto overlooked attributes. The list is not only not exhaustive as it stands, but cannot be made so as a matter of principle, thanks to the vast variety of potentially relevant properties. A likelihood definition, on the other hand, summarises the effect a 'perfect' list would have. For the only plausible justification for putting further items on the list is that, by putting these items on the list, we let the respective probabilities of the supposed peers' being right converge with each other.

A second point is that, contrary to such characteristics as ignorance of evidence, shortness of time, lack of intelligence, bias, or drunkenness, an extraordinary state of excitement influences a person's examining abilities only very occasionally to such a degree that that person is thereby considerably less likely to get things right. Moreover, whether extreme excitement may influence a given person's examining abilities at all is highly relative to the specific characteristics of that person. Hence including something like 'freedom from extraordinary excitement' on our list would make the definition too restrictive. However, a more specific description, involving, for example, the fact that a child is about to be born to a parent who tends to get extremely nervous in this kind of situation, would make our first problem more apparent, namely that we would need to add, *per impossibile*, virtually infinitely many more descriptions to our list in order to make it exhaustive.

One could also come up with the idea to include a more general characteristic than freedom from extraordinary excitement on the list, for example the ability to concentrate, and take it to exclude exactly the epistemically significant cases of extreme nervousness. The notorious difficulty with such general characteristics, however, is the vagueness of their entailment conditions. For instance, does the ability to concentrate include sobriety, or absence from test anxiety? It appears that the very same property that enables us to argue that somehow the right cases of extreme nervousness get excluded makes it hard to apply the resulting definition in concrete cases of peerhood assessment. The more general the characteristic is, the less helpful it proves to be.

(In addition, a specific problem regarding the ability to concentrate is that one could suffer from infrequent lacks of concentration in very exceptional situations, while generally having a high level of concentration. So in order to diagnose sufficiently severe lacks of concentration *before* the relevant deliberation processes start, one needs to know much more about the relevant persons than their general ability to concentrate. Equivalent problems will arise for similar characteristics.)

Speaking of unhelpfulness, we should indicate that, unlike a detailed list definition, a likelihood definition provides little guidance for judging whether or not someone is an epistemic peer.[11] We can easily compare the epistemic statuses of given subjects on the basis of the various properties named in a list definition; a far more abstract likelihood condition, by contrast, leaves us alone and unaided with that task. In order to assess someone's likelihood of being right concerning a specific question, we actually need to draw on their familiarity with relevant evidence and arguments, their intelligence, lack of bias, sobriety, and so on. We need the information contained in a detailed variant of the standard definition. Such a variant, however, is almost certain to be wrong, as is shown by various counterexamples, such as DAY OF BIRTH, which are ready at hand. A variant of the standard definition that uses umbrella terms such as 'well-informedness' or 'competence', on the other hand, is both potentially imprecise—there might also be cases in which it differs from a likelihood definition—and comparatively uninformative—it is, for instance, unclear whether sobriety or, for that matter, the ability to concentrate should play a role in deciding whether someone is a peer. It seems that in defining epistemic peerhood, the cost of precision is uninformativeness.

It is important to note, however, that this does not pose a serious problem. A likelihood definition is as precise as could be wished and captures the entire

---

[11] See e.g. Axel Gelfert, "Who Is an Epistemic Peer?," *Logos & Episteme* 2 (2011), 512.

intent of our concept of epistemic peerhood. This is all we need in contexts in which a precise understanding of this concept matters. When it comes to applying the concept, we of course need a detailed list of properties that we can check, but no such list can itself be suitable for a definition. Moreover, in compiling such a list, we have to allow ourselves to be led by the likelihood criterion.

An analogy may be helpful. One can define *water* as the colourless stuff that falls from the skies, fills our lakes and rivers, and flows from our taps. Or one can define it as the substance whose molecular structure is H_2O. While the first definition is much more appropriate for identifying water in everyday contexts, the second is usually taken to be the correct one, the one to which we are to refer in hard cases. Similarly, the likelihood definition reveals to us the essence of peerhood, and is to be preferred in cases of doubt, while list definitions generally allow us to assess peer statuses more directly.[12]

The fact that likelihood considerations should guide us, directly or indirectly, in assessing peerhood is, by the way, the reason why background assumptions, methodological preferences and track records should preferably not be included in a list definition: they do not necessarily help to identify exactly those people who are equally likely to be right. This seems to be easy to see in the case of background assumptions and methodological preferences. After all, if you disagree with an equally well-informed and competent person on some given proposition, and the two of you track down your disagreement to a clash of deeply held assumptions, methodological or whatever, it is hard to see why your assumptions are more likely to set you on the right path than the other person's assumptions. (In fact, the matter is a bit more complex, and related to the difference between *being an epistemic peer* and *counting someone as an epistemic peer*, which I will address in section 5.)

In the case of track records, the claim that including them in the definition does not necessarily improve it may appear more surprising. To see its truth, note first that *having an equally good track record* would itself make a plausible *definiens* in a definition of epistemic peerhood. This is no coincidence: probability is often interpreted as the limit of a series of relative frequencies, and track records

---

[12] The analogy is suggestive but not perfect. For one thing, one could plausibly define *water* not by its actual chemical structure but by its functional role. Additionally, and relatedly, whereas the results of the rivers-and-lakes definition of water are at most slightly different from those of the H_2O definition *in the actual world*, the results of the two definitions differ considerably *in many other possible worlds*, which arguably yields some awkward consequences for the H_2O definition. Both points have no parallels in the case of the peerhood definitions, which makes the case for the likelihood definition on closer consideration far more compelling than the case for the H_2O definition in fact is.

contain such a series. A sufficiently long track record would thus give us the same information as the likelihood criterion. The track record definition, however, faces several problems. The most obvious one has its roots in the fact that the resulting probability is based only on the examination of former failures and successes, whereas in the likelihood definition we are allowed to take into account everything that might be relevant. Thus a track record definition lacks the resources to analyse the various cases of disagreement—among them those cases of enduring disagreement in fields such as philosophy or politics that interest us most—in which it is virtually impossible to get a reliable track record, not to mention a long one. To be sure, one could lower the demand for reliability and view people whose opinions on philosophical or political issues are generally reasonable as having a reliable track record, regardless of whether their beliefs are in fact true. But then surveying a track record would not amount to more than a superficial check of competence and well-informedness.

Even more importantly, the track record definition focuses on the likely effects of essential characteristics such as well-informedness and competence, not on these characteristics themselves. For this reason, it may occasionally produce wrong results: a comparably short track record might be misleading by sheer bad luck; or we might fail to see that former scenarios are not sufficiently similar to the one under consideration. For example, my track record in mental calculation, or even in doing mental calculation in states of excitement, is of no help in evaluating whether I am my in-laws' peer in DAY OF BIRTH because, in this scenario, the likelihood of my being right depends on other factors than my former performances. Moreover, as I explained above, we should take epistemic peerhood to be relative to time; this, too, does not fit well with the idea of deriving the relevant probability from the past. For example, if I begin to work out square roots in my head, my ability to do so correctly might improve rapidly, the result being that, at a certain point in time, my track record up to that point would be considerably worse than my competence. In sum, adding a track record criterion is misleading insofar as it suggests putting too much weight on past performance and too little on the specific conditions of the case under consideration.[13]

---

[13] Lam defines epistemic peerhood as equal reliability, which he measures by comparing degrees of credence regarding relevant propositions to the truth values of those propositions (Barry Lam, "On the Rationality of Belief-Invariance in Light of Peer Disagreement," *Philosophical Review* 120 (2011)). Thus, Lam's definition is in fact a refined version of a track record definition, and hence faces the same difficulties. See also footnote 26 for further remarks on his conception of peerhood.

A corollary of this is that Elga's bootstrapping argument[14] is flawed. Very briefly, this argument is concerned with the case of two people, you and I, say, who disagree about many not too elaborate problems that are sufficiently similar to allow for deriving reliable track records. (Elga's example is judging by eyesight which horse won one of a long series of races.) Suppose that, for each problem, it is reasonable for me to be a bit more than 50% confident that I am in fact more likely to be right than you are. Thus, as we go along, I come to regard my track record as considerably better than yours. As a consequence, I should become exceedingly confident that I am more likely to be right than you are. Yet it seems absurd to base a significantly increased certainty of being more likely to be right merely on the fact that we occasionally disagree. Therefore, so the argument goes, it cannot be reasonable for me to be a bit more than 50% confident that I am in fact more likely to be right.

Whatever other aspects of this argument might be problematic,[15] it surely is not sound if we refrain from assessing other people's peer statuses primarily by their track records. For then it does not follow from the fact that my track record is markedly better than yours that I should become more confident that I am more likely to be right. Whether I am more likely to be right depends first and foremost on essential characteristics such as competence and well-informedness, whose respective degrees must remain the same throughout the whole process if the argument is supposed to make any sense. Therefore, given that in assessing someone's peer status those characteristics are more decisive than track records, my confidence of being more likely to be right should not rise, the momentous absurdity can be avoided, and the bootstrapping argument fails.

There is a third point to be made about DAY OF BIRTH, for the reason that much of what I have said so far about the likelihood definition is not quite right by Elga's lights. Look again at Elga's version of definition: "you count your friend as an epistemic peer with respect to an about-to-be-judged claim if and only if you think that, conditional [on] the two of you disagreeing about the claim, the two of you are equally likely to be mistaken." Here, Elga explicitly relativises epistemic peerhood to an "about-to-be-judged claim"—but he clearly does not relativise it to a specific time or occasion. This latter fact is surprisingly central for understanding Elga's version of the so-called *Equal Weight View*, as I will explain now.

---

[14] Elga, "Reflection," 486–488.

[15] For critical examinations of Elga's bootstrapping argument, see Duncan Pritchard, "Disagreement, Skepticism, and Track-Record Arguments," in *Disagreement and Skepticism*, ed. Diego Machuca (London: Routledge, 2013), and Jonathan Weisberg, "The Bootstrapping Problem," *Philosophy Compass* 7 (2012).

3. Elga and Equal Weight

According to a standard formulation of the *Equal Weight View* (henceforth EW), one should give the opinions of one's epistemic peers the same weight one gives one's own. This is also roughly the understanding of EW with which Elga starts on p. 484 of "Reflection and Disagreement." Four pages later, he starts refining this picture, and ends, on p. 490, with this formulation:

> **Equal weight view** Upon finding out that an advisor disagrees, your probability that you are right should equal your prior conditional probability that you would be right. Prior to what? Prior to your thinking through the disputed issue, and finding out what the advisor thinks of it. Conditional on what? On whatever you have learned about the circumstances of the disagreement.

The reason for the refinement is the insight that the actual circumstances of the disagreement can influence your or your peers' probability of being right on a given question. To give an example, Elga lets us assume that the weather gets extremely hot, and that you know that your friend has severe problems concentrating in such circumstances. Not so you: your mental abilities are usually not affected by excessive weather conditions. If the two of you then evaluate some proposition, it may be the case that, although both of you are equally likely to be right under normal circumstances, you are more likely to be right on this particular occasion. In other words, your prior conditional probability that you will be right is not 50% but, for example, 80%. This probability is *prior* in the sense that it has to be calculated independent of what you and your friend think about the proposition that is to be evaluated; it reflects what you should[16] have said about your likelihood of being right under specific conditions before you made up your mind, before the actual disagreement arose, and before you knew whether these specific conditions would actually occur. This probability is *conditional* in the sense that the specific conditions under which the disagreement actually occurs are to be taken into account.

To continue with the hot weather example, assume that you should have said prior to evaluating the claim under consideration, prior to the occurrence of a disagreement and prior to what you later learn about the specific circumstances of the disagreement that, in the event of hot weather, you should be 80% confident of being right in evaluating a proposition of a specific kind. Then Elga's version of

---

[16] Here and at a similar location below, the Elga of "Reflection and Disagreement" would have written 'would' instead of 'should'. Enoch argues—similarly to my argument in section 5—that Elga's descriptive understanding does not capture the epistemic force we are confronted with here, and suggests a normative revision (Enoch, "Truthometer," 970–972). Enoch even reports that Elga, in conversation with him, agreed to the revision.

EW tells us that you should indeed be 80% confident when a disagreement actually arises and the weather actually gets very hot.

What has all this to do with DAY OF BIRTH? The answer is, of course, that this scenario is quite similar to the hot weather case in that my lack of concentration is limited to a special occasion: while I am normally my in-laws' equal in doing mental maths, I am their inferior on the particular day on which the birth takes place. My impairment is temporary, not permanent. As his reasons for reformulating EW suggest, Elga prefers not to deal with the possibility of such temporary impairments by relativising the definition of epistemic peerhood to specific times or occasions; he rather incorporates into EW a mechanism that prevents us from viewing our peers (in Elga's time-invariant sense) as equally likely to be right in case the circumstances of the disagreement are epistemically unfortunate either for us or for our peers.[17] Hence Elga would not regard DAY OF BIRTH as a scenario that shows the advantages of the likelihood definition over the standard one; he would regard it as a scenario that indicates that the standard formulation of EW needs refining.

(Whether or not we take DAY OF BIRTH to be appropriate as a base for our arguments in favour of the likelihood definition is inessential for the dialectical force of these arguments. This is because we can easily invent an alternative scenario in which a long-lasting or even permanent change precludes me from being your peer. Assume, for example, that I have to take pills for severe depression, and that a side effect of these pills is that they reduce my attention span significantly. Since this makes it harder for me to follow lengthy lines of argument, something like 'length of attention span' should be on a standard definition's list. And again, if that is already included or entailed, other examples in which hitherto unconsidered *and enduring* characteristics play a central role are easy to find.)

The alternative to refining EW is relativising peerhood not only to an "about-to-be-judged claim" but also to a time. If we do this, we can deal with scenarios such as the hot weather case while keeping EW in its original form. And this is preferable for several reasons, which I will list in the next paragraph. Before

---

[17] I think that this interpretation is well supported by Elga's considerations regarding the best understanding of EW. There is, however, a passage in footnote 21 that speaks against this interpretation. In defending his definition, Elga writes: "[S]uppose that you think that (…) your friend is more likely than you to be mistaken. Then however intelligent, perspicacious, honest, thorough, well-informed, and unbiased you may think your friend is, it would seem odd to count her as an epistemic peer with respect to that claim, *at least on that occasion*" (my italics). Here, Elga indeed relativises epistemic peerhood to a specific occasion. So his use of the concept seems not to be perfectly consistent in this respect.

this, I should mention that Elga gives a second argument for his version of EW, namely that it enables us to deal with peer disagreements in which we find the opinions of our supposed peers obviously irrational. The most famous example of such a disagreement is the extreme restaurant case:[18] as in the normal restaurant case, you and I each calculate how much everyone has to pay. But this time the result you get is virtually impossible; it is, for example, ten times greater than the full amount stated on the bill. Do I still have to give your belief that yours is the right result the same weight as my belief that mine is correct (after all, you are my peer)? Elga observes that if we had known beforehand that we would find our peer's answer absurd, we would not have given her opinion equal weight (at least not if the case is asymmetric in that our peer does not find our answer absurd). Therefore, according to his version of EW—and his descriptive formulation: recall the 'would or should' problem addressed in footnote 16—I need not regard the probability that you are right as equal to the probability that I am right in such a case of extreme disagreement.[19] In section 6, I show how we can solve this problem without revising EW.

Turning now to the reasons for preferring relativising peerhood to times over refining EW, we have to note, first, that the former proposal, in particular the formulation of EW, is considerably simpler. Second, it is somewhat arbitrary, in Elga's account, what cognitive or mental shortcomings are to be ascribed to the special circumstances of the disagreement, and what discredits someone as a peer. For example, how long-lasting must an effect of, say, pills for depression be in order to concern peerhood? What if the pills reduce my attention span for only three days (a week, a month)? There is no corresponding arbitrariness in the likelihood account, because, due to the relativisation to points in time, everything that concerns the specific circumstances of the disagreement is automatically taken into account when peerhood is ascribed or denied. Third, recall Elga's reason for favouring the likelihood definition over the standard definition: it does not primarily matter how well-informed and competent other people are concerning some proposition p; if they are less likely to be right concerning p than we are, it appears odd to regard them as peers concerning p. The same argument can be used to show why a time-relative likelihood definition is better than a time-invariant one: it does not primarily matter whether other people are usually equally likely to be right concerning p; if they are less likely to be right concerning p than we are on a specific occasion, it appears odd to regard them as peers concerning p *on that occasion*. Finally, our original reasons for preferring a

---

[18] Christensen, "Good News," 199.
[19] See Elga, "Reflection," 490–491.

time-relative definition still hold: one's familiarity with the relevant evidence as well as one's level of competence might change over time (and sometimes they change pretty quickly). In order to account for this, it is good to have a time-relative definition.

In this section and the previous one, I have discussed different ways of defining epistemic peerhood. Elga, in his footnote 21, states his usage of 'epistemic peer', then states the standard usage, and finally argues that his is preferable. In much more detail, I have explained the standard usage and why it is lacking (and, in many but not all respects, my arguments here are merely a specification of Elga's short remark). I have also examined Elga's definition, but have left the discussion of the discrepancy between *being an epistemic peer* and *counting someone as an epistemic peer* for section 5. This being said, the result of my discussion is that we should define epistemic peerhood and EW as follows:

> **Definition 1.** $P_1, ..., P_n, n \in \mathbb{N}$, are **epistemic peers** regarding a proposition p and a time t if and only if $P_1, ..., P_n$ are antecedently equally likely to be right when evaluating p at t.

> **Definition 2.** The **Equal Weight View** holds that one should give the opinions of one's epistemic peers the same weight one gives one's own.

4. Why the Equal Weight View Is True—and Why This Is Not the End of the Debate

Elga takes care to emphasise that the disagreement itself does not count as evidence for whether or not the parties to the disagreement are peers. Hence we are *not* allowed to argue in the following way: "I believe that *p*, and I believe that you are my peer. Upon finding out that you believe ~*p*, I have two options: either I can revise my belief that *p* and become agnostic about that matter; or I can revise my belief that you are my peer on the basis of your poor judgement concerning *p*. Which alternative is better depends on the specific proposition under consideration and the depth of my respective beliefs."

Elga claims that the second option—revising the belief that some people are our peers—is open to us only if we agree *prior* to the disclosure of a potential disagreement that we will not regard them as our peers if the disagreement actually arises. This is perfectly plausible. For assume the opposite: that I first claim that if some specific person disagrees with me, I will regard her opinion as likely to be right as I regard mine, and that I then, when the disagreement actually arises, nonchalantly downgrade her reliability. Such behaviour could hardly be

deemed internally consistent.[20] So either we should be prepared right from the beginning to take an occurring disagreement as a reason for not counting people as peers, and consequently deny from the outset that there could be any peers at all, or else we should not normally take the disagreement itself to be evidence for or against peerhood. Elga, for instance, allows for counting the disagreement itself as evidence only in cases in which the disagreement is rather peculiar, as it is, for example, in some variants of the extreme restaurant case. (In section 6 I will say more about taking the disagreement itself as evidence.)

According to our likelihood definition of epistemic peerhood, peers regarding some proposition p are those who are antecedently equally likely to be right about p. Observe now that any reason we might have for degrading someone's opinion is *ipso facto* a reason not to count him or her as a peer. For let us consider those reasons one by one. If a list definition were to be preferred, it would perhaps be possible to find some hidden factor that we had overlooked while compiling the list. In such a case, peers in the sense of that list definition would not be equally likely to be right about p. Yet, as I explained, we should abandon list definitions. If, secondly, the specific circumstances of the disagreement were not taken to affect the likelihoods of being right of the people involved, we would sometimes be permitted to downgrade a person's belief irrespective of his or her likelihood of being right. Yet, by relativising our definition to time, we make sure that the specific circumstances of the disagreement are taken to affect the relevant likelihoods. Thirdly, we would also be permitted to downgrade people's beliefs but not their antecedent likelihoods if the disagreement itself were relevant not only in exceptional situations like the extreme restaurant case but also quite generally. Yet, as I argued in the preceding paragraph, this condition contradicts our assumption that we consider a case of *peer* disagreement. Since the antecedents of all these conditionals have thus been ruled out, it appears that every conceivable reason for discounting someone else's opinion concerning a given question is also a reason to regard him or her as less than equally likely to be right. But if this is so, how could one *not* give the opinion of someone who is equally likely to be right as much weight as one gives one's own? Given that the peers aim at believing about p what is best supported by the evidence at hand,[21] they are therefore compelled to give their respective beliefs

---

[20] Proponents of the Right Reasons View and the Total Evidence View, both outlined below, are likely to disagree. As my discussion of these views will show, however, they should be presented in a way that is compatible to this claim.

[21] This is a variant of evidentialism, the view that all reasons to believe are evidential reasons, in contrast to, for example, pragmatic reasons. Evidentialism is widely acknowledged (see e.g.

the same weight. In other words, EW follows, almost trivially, from our definition of epistemic peerhood.

Does that speak against the definition? According to the several variants of the standard definition, and even according to Elga's understanding of the likelihood definition, there is room for regarding people as peers without giving their opinions equal weight. This may happen because of some characteristic of a specific scenario that is either not on our list of factors or time-relative, but which nevertheless reduces our peers' actual probability of being right. If my arguments so far are correct, however, it is arbitrary and confusing to disregard that characteristic, once it is identified, in evaluating someone's peer status. The most precise definition of peerhood is indeed not neutral with respect to the truth of EW.

The bad news (if it is bad news) is that this is not the end of the debate. It is merely a shift of focus: instead of discussing EW, we should discuss more carefully what conditions people in fact have to satisfy in order to be peers. The reason is that all well-known alternative theories to EW can be reformulated in a way that makes them compatible with EW.

Take the *Extra Weight View*. According to this view, one should give one's own opinion more weight than one gives the opinions of one's epistemic peers ('peers' understood in the sense of the standard definition).[22] An extreme version of this view is the *Steadfast View*, according to which one should give the opinions of one's epistemic peers (again understood in the sense of the standard definition) no weight at all. Both the Extra Weight View and the Steadfast View come in several variants because of the various reasons one could have for

---

Nishi Shah, "A New Argument for Evidentialism," *The Philosophical Quarterly* 56 (2006) and Jonathan Way, "Two Arguments for Evidentialism," *The Philosophical Quarterly* 66 (2016)) but not universally held (see e.g. Andrew Reisner, "The Possibility of Pragmatic Reasons for Belief and the Wrong Kind of Reasons Problem," *Philosophical Studies* 145 (2009)). We can, however, safely presuppose it here, because otherwise the debate on peer disagreement would not get off the ground. For assume that you are justified to believe *p* just because, for instance, believing *p* makes you significantly happier than not believing *p*. Obviously, this kind of justification, which is not based on evidence, is not normally undermined by learning that some peers disbelieve *p*, for disagreement can at most indicate that one has misevaluated the evidence, not that one has misjudged one's feelings.

[22] Variants of the Extra Weight View are defended e.g. in Gideon Rosen, "Nominalism, Naturalism, Epistemic Relativism," *Noûs* 35(s15) (2001), Peter van Inwagen, "We're Right. They're Wrong," in *Disagreement,* eds. Richard Feldman and Ted Warfield (Oxford: Oxford University Press, 2010), and Michael Huemer, "Epistemological Egoism and Agent-Centered Norms," in *Evidentialism and Its Discontents*, ed. Trent Dougherty (Oxford: Oxford University Press 2011).

disregarding one's peers' opinions partially or completely. The most common variant is to adopt an agent-centred point of view and argue that *having* a certain piece of evidence is epistemically more significant than *knowing* that someone else has that piece of evidence (or a similar one). As a consequence, my intuition that *p* counts for more than my knowing that you have the intuition that ~*p*.

Now, there are two ways to make sense of this. First, one could argue that it is possible that my intuition counts for more than my knowing of your contradictory intuition even if this does not make me more likely to be right. This seems to be Wedgwood's view; it comes down to denying what I took for granted in my deduction of EW, namely that we aim at believing what is most likely to be true.[23] The second way, however, is closer to the point: that my intuition counts for more just means that I am more likely to be right. In other words, even if we are equally competent and well-informed and thus are peers in terms of the standard definition (or its most elementary variant), we are not, according to this understanding, peers in terms of the likelihood definition. Therefore it does not contradict EW (and neither does its extreme variant, the Steadfast View). More generally, those adherents of the Extra Weight View who accept that we should believe what is most likely to be true do not deny that we should give the opinions of those who are equally likely to be right the same weight we give our own. They merely claim that there are, apart from characteristics such as competence and well-informedness, other factors, more closely related to one's individual perspective, that influence—and in fact increase—one's likelihood of being right considerably. In addition, they may hold that your disadvantage of not having my

---

[23] In a nutshell, Wedgwood's argument is as follows (see Ralph Wedgwood, *The Nature of Normativity,* (Oxford: Oxford University Press, 2009), 257–263): in order to avoid scepticism, one needs to trust at least some of the relevant intuitions that are had by reasonable people. The set of intuitions that are to be trusted without any scrutiny should not be too small—otherwise it would be insufficient as a base for non-sceptical views—but, far more importantly, it should not be too large either, because the larger the set is, the more prone to error are the theories that are built upon the intuitions it contains. Thus any *non-arbitrary* way (other ways would of course be ad hoc) of making the set smaller, up to a minimum size, is welcome. One such way, according to Wedgwood, is to trust one's own intuitions to a significantly higher degree than those of others. It is highly controversial whether this is indeed *non-arbitrary*. But even if it were, the primary motivation for the greater trust in one's own intuitions is avoidance of scepticism. Nothing in Wedgwood's line of reasoning suggests that I am presently more likely to hold a true belief about a specific proposition if I stick to my belief just because it is supported by my intuitions and undermined only by those of others. And although anti-sceptical presuppositions might be of significant epistemic value in the long run, they are not generally compatible with the aim of believing what is best supported by the available evidence and hence most likely to be true.

point of view can be compensated for by more competence and well-informedness on your side. In this case, we may be peers, in the sense of being equally likely to be right. Opponents of the Extra Weight View, on the other hand, deny that the fact that a specific assessment is one's own increases one's likelihood of being right, and hence needs (negative) compensation. Therefore, we can revise the Extra Weight View as follows:

> **Definition 3.** The **Extra Weight View** holds that the fact that a specific assessment is one's own increases one's likelihood of being right.

This way of stating the Extra Weight View preserves what is at stake between its advocates and adversaries and adjusts it at the same time to our new conception of epistemic peerhood.[24]

Arguably, the revised version is far less plausible than the original. For consider a disagreement between two equally competent and well-informed people: saying that each of them is justified in giving his or her own opinion more weight than another's does at least not appear contradictory; saying that each has an increased likelihood of being right, on the other hand, is hardly consistent (for instance, their likelihoods of being right cannot add up to more than 100%, but exactly this could happen if the Extra Weight View were true and if it applied symmetrically to both parties to a disagreement). This means that the only reading of definition 3 that appears defendable is as follows: Suppose A and B are equally competent and well-informed people who disagree on p. Then the Extra Weight View holds that, from A's perspective, A is more likely to be right about p than B, because the fact that a specific assessment of p is A's increases A's likelihood of being right. From B's perspective, on the other hand, B is more likely to be right about p than A, because the fact that a specific assessment of p is B's increases B's likelihood of being right. So far, so good. However, we are not so much interested in A's or B's perspective, but rather in the viewpoint of a neutral observer, who

---

[24] Another view that falls under my definition 3 is presented in Enoch, "Truthometer." Enoch argues that the disagreement itself has to count as evidence because of the asymmetry between p (a proposition that I believe prior to the disclosure of the disagreement on the basis of my original evidence) and my supposed peer's belief that $\sim p$. Whereas comparing the piece of evidence that *I believe that p* with the piece of evidence that you believe that $\sim p$ should cause me to become agnostic, given that we are peers, comparing the piece of evidence that p (a proposition I take to be true) with the piece of evidence that you believe that $\sim p$ need not necessarily cause me to become agnostic, according to Enoch. Although Enoch admits that the consequences of his view are precisely the same as those of an appropriate version of the Extra Weight View, he does not take his view to be a variant of the Extra Weight View. The reason is that he defines the Extra Weight View slightly more narrowly than I did.

wants to know whether A is more likely to be right about p than B or vice versa. As the answer to this cannot be 'both', there is no consistent neutral reading of definition 3. This is exactly what speaks against the Extra Weight View, in whatever definition; an impartial observer cannot ascribe likelihoods of being right in accordance with this view.[25] As this embarrassing point comes out clearer in my definition of the Extra Weight View, the task of defending this view might become noticeably harder after the shift of focus that I advocate. As we will see, similar considerations apply for the other well-known alternatives.

One of them is known as the *Right Reasons View*. According to this view, the rational thing to do in a peer disagreement is to stick to one's opinion if one responded rightly to the original evidence and to revise one's opinion if one responded wrongly to the original evidence.[26] As a consequence, it does not matter whether any disagreement, with peers or non-peers, gets disclosed after one first formed one's belief, because what one should do depends solely on whether or not one responded correctly to the original evidence, and not on how other people evaluated this evidence. Should we therefore say that the Right Reasons View allows us to stick to our beliefs—given that we indeed hit the ones best supported by the evidence—, even if our peers favour other positions? No. For recall that our peers are those who are antecedently, i.e. before the disagreement gets disclosed, as likely to be right as we are. Then proponents of the Right Reasons View, who

---

[25] I am grateful to an anonymous reviewer for getting me to clarify this.

[26] Although hardly anyone defends a full-fledged version of this view, elements of it can be found in the works of several authors. Three examples: Lackey's *Justificationist View*, according to which the prior degree of justification that one has for a particular belief is crucial for the epistemic force of disagreement about that belief, rests on an externalist notion of justification (Lackey, "Belief-Dependence," 320). In other words, Lackey implies that one should be justified *for the right reasons*. Secondly, and even less obviously, van Wietmarschen discusses an understanding of the arguments for EW in terms of *evidential support*, a notion which is commonly regarded as being closely related to an externalist conception of justification (Han van Wietmarschen, "Peer Disagreement, Evidence, and Well-Groundedness," *Philosophical Review* 122 (2013)). He points out that, according to such an understanding, EW fails. Thirdly, Lam, who analyses peerhood in terms of *reliability*, considers two measures of reliability, one of which (calibration) relies on a ratio of true propositions to total propositions, the other (Brier Scoring) on closeness to the truth (Lam, "Belief-Invariance"). In both cases, the reliance on truth rather than on something like (internalistically) justified response to evidence results in an externalist track record account of peerhood (cf. footnote 13). Right reasons thus play a role not in evaluating the specific disagreement under debate but in evaluating those surrounding disagreements on which the peerhood assessment is based. Given such an account, it is unsurprising that, as Lam shows, belief revision is not always called for in (non-extreme) cases of peer disagreement.

hold that, in case a disagreement occurs, one party may be right in sticking to their belief and the other wrong, should also hold that, in such a case, one party was antecedently more likely to be right than the other. To be sure, we could not have known which one, and perhaps still cannot know; but this corresponds well to the Right Reasons View, since, according to that view, we cannot know which party responded correctly to the original evidence and thus should keep hold of their opinion. In other words, since the notion of rationality that underlies the Right Reasons View is externalist, we should adopt an externalist understanding of likelihood as well. What does such an understanding look like?

Assume that I toss a coin. What is the probability of its coming up heads? Here is a surprising answer: either the world is such that the coin will come up heads. Then this event will definitely occur; hence, its probability is 1. Or the world is such that the coin will come up tails. Then *this* event will definitely occur, and the probability of the coin coming up heads is 0.[27] In other words, in a world in which each proposition is either true or false, the probability that a proposition is true is always either 0 or 1. A notion of probability that allows for values between 0 and 1 for future events is grounded in our limited knowledge about how things will turn out, not in the way they will actually turn out. In this sense, it is internalist. If I say "The probability that it will rain tomorrow is 30%," I do not mean that it is not completely certain, given today's worldwide weather conditions, whether or not it will rain tomorrow; what I mean is that in light of all the evidence I have at hand and can evaluate properly, my rational degree of credence is 30%. An externalist notion of probability cannot serve the purpose of reporting such a rational degree of credence between 0 and 1, caused by insufficient information. In the same way an externalist notion of rationality cannot account for the reasonableness of believing, on the basis of misleading evidence (namely one's peer's misevaluation of the original evidence), what is in fact not supported by the original evidence. This analogicity of the two externalist understandings carries over to the respective views so that we arrive at this definition:

> **Definition 4.** The **Right Reasons View** holds that two persons are peers regarding some proposition p if and only if they will either both correctly evaluate the

---

[27] In what follows, I ignore theories of objective probability, according to which the world may be such that the probability of the coin coming up heads is between 0 and 1, because taking those theories into account would merely complicate matters, but not cause essential revisions. The reason is that invoking objective probabilities only has consequences for the values that externalist probabilites can take, but not for the difference between externalist and internalist probabilities. It is this difference on which my reasoning here relies.

> evidence regarding p, or if they will both incorrectly evaluate the evidence regarding p, i.e. if their respective likelihoods of being right are either both 1 or both 0 (whereby likelihood is to be understood in an externalist way).

Like the original version, this reformulated one claims that the actual occurrence of a disagreement is of no epistemic significance; one should revise one's view if one is wrong or has probability 0 of being right, but not if someone else thinks differently. In sum, the Right Reasons View can be made compatible with EW by claiming that, in order to be peers, two persons have to have the same externalist probability of being right.

Relying on externalist probabilities, however, is a very strange thing to do. We are not used to using them, and there is a good reason for that. Talk of probabilities is meaningful insofar as it enables us to distinguish degrees of uncertainty. It does not matter for this aim whether we are uncertain because there is some definite fact of the matter that we do not (yet) know, or because there is no fact of the matter at all (or not yet). To use probability talk to indicate whether some fact holds true is a confusing and needlessly complicated way of speaking. Since the revised variant of the Right Reasons View forces us to interpret the notion of likelihood, as it figures in the peerhood definition, in an externalist fashion, it forces us to adopt that confusing and overly complicated way of speaking.

A third view often seen as an alternative to EW is the *Total Evidence View*, which is most prominently defended by Kelly.[28] According to this view, "what is reasonable to believe [in a peer disagreement scenario] depends on both the original, first-order evidence as well as on the higher-order evidence that is afforded by the fact that one's peers [understood in the sense of the standard definition] believe as they do."[29] The first-order evidence regarding some non-doxastic proposition *p* comprises all evidence regarding *p* except evidence regarding what others believe about *p*. Second-order evidence regarding *p* is evidence regarding what others believe about *p*. Third-order evidence regarding *p* is evidence regarding what others believe about what others believe about *p*. And so on. Higher-order evidence is evidence that is not first-order. The idea behind the Total Evidence View is that, although the disclosure of a disagreement gives us higher-order evidence that supports suspension of judgement, acquiring this higher-order evidence gives us no reason to completely disregard the first-order evidence on the basis of which we formed our original belief; what is reasonable

---

[28] In Thomas Kelly, "Peer Disagreement and Higher-Order Evidence," in *Disagreement*, eds. Richard Feldman and Ted Warfield (Oxford: Oxford University Press, 2010).

[29] Kelly, "Peer Disagreement," 142.

to believe depends on *both* sorts of evidence. As a result, it may—but need not—happen that one should not suspend judgement in the case of a peer disagreement because the first-order evidence clearly supports one position over the other.[30] There is thus no golden rule on what to do in the case of a disagreement with someone who is equally competent and well-informed. The Total Evidence View allows for splitting the difference if the first-order evidence is sufficiently inconclusive, but otherwise it prescribes sticking to one's belief. So it would crucially depend on the quality of the first-order evidence whether one should accept some equally competent and well-informed person as someone who is equally likely to be right in case of a disagreement. If the first-order evidence appears strong, one should rather trust one's own assessment of it than another person's opinion, and since one knows this in advance, before a disagreement arises, one should, according to the Total Evidence View, refrain from considering other people as one's peers from the outset. Hence we can reformulate the Total Evidence View as follows:

> **Definition 5.** The **Total Evidence View** holds that whether two persons are peers regarding some proposition p depends, among other factors such as their competence or well-informedness, on the first-order evidence they have for p.

There may be cases in which the first-order evidence plays virtually no role at all in deciding who is a peer, and there may be those in which the quality of the first-order evidence makes it more or less impossible that there are any peers at all. This version of the Total Evidence View is compatible with EW.

As with the other alternatives, however, the revised version might appear harder to believe than the original. That the first-order evidence adds some information, over and above what the higher-order evidence tells us, to *what it is reasonable to believe* seems prima facie more plausible than that it adds some information, over and above what a person's competence, well-informedness etc. tell us, to *how likely that person is to be right*. For recall that people's likelihoods of being right on a certain matter should be determined *before* they make up their minds on that matter. However, claiming that the first-order evidence influences these likelihoods in effect means that their likelihoods depend on the positions they will take: if they take a position that is clearly more strongly supported by the first-order evidence, their likelihood will increase; if they take a position that is clearly less strongly supported, their likelihood will decrease. Hence the revised

---

[30] See Thomas Kelly, "Disagreement and the Burdens of Judgment," in *The Epistemology of Disagreement*, eds. David Christensen and Jennifer Lackey (Oxford: Oxford University Press, 2013), 50–51.

version of the Total Evidence View seems to be incompatible with our preferred understanding of the likelihood definition.

One might suspect that not only alternative theories of EW are in need of reformulation owing to our revision of the peerhood definition, but also EW itself. I seriously doubt, however, that there is any plausible way of formulating a version of EW that is not entailed by the likelihood definition without failing to capture EW's original intent. Among the implausible ways, one is the way of exclusion. According to this way, EW is defined as the thesis that the Extra Weight View, the Right Reasons View and the Total Evidence View are all wrong. This way is implausible not only because it is unlikely that those three views are the only meaningful rivals one could imagine, but also because it completely leaves in the dark what EW actually says and why we should have any interest in it. Another implausible way is that of refocusing. Here, EW states that, under normal conditions, it is not too hard to find disagreeing interlocutors who are equally likely to be right. (More often than not, one could add, it suffices if they are equally competent and well-informed.) The idea here stems from the observation that, in all three rival theories that I have discussed, peerhood depends on some special condition (point of view; actually correct response to evidence; nature of first-order evidence) that is quite unlike those factors listed in standard definitions. This arguably makes peerhood harder to come by. However, this way is in fact no better than the former exactly because it merely spells out a common strand in the three rival theories and does not add any substantial content. Moreover, an adherent of the true intent of EW should have no problems in allowing as a possibility that disagreeing peers are hard to find. Yet another implausible way of reformulating EW could be called the way of ignorance. According to this way, EW is defined as the thesis that we should give the opinions of those who are equally competent and well-informed the same weight we give our own. Here, there is simply no mention of epistemic peers in the likelihood sense. This way is implausible because our insights from the previous section are ignored and, as a result, a deficient notion of epistemic peerhood is still, though implicitly, in use. In sum, several conceivable attempts to adjust EW to our new understanding of peerhood prove to be inadequate. This is not surprising: what is essential in EW is that it connects the concept of epistemic peerhood with the prescription to give the opinions of those who fall under this concept equal weight. Since epistemic peerhood should be understood in terms of likelihood, it is virtually impossible to conceive of a plausible reformulation of EW whose truth does not follow almost immediately from our likelihood definition.

In the preceding paragraphs, I have not taken pains to present *conclusive* arguments against the several alternative theories to EW that I have discussed (nor will I do that in what follows); the most I have done is to indicate that they may appear less plausible after revision. My aim here has chiefly been to explore whether the most prominent alternatives to EW can be reformulated in a way that is *compatible* with the likelihood definition of epistemic peerhood and, *a fortiori*, with EW itself. The result is that they can indeed be appropriately reformulated; although they are arguably less plausible after revision, they cannot be ruled out as a consequence of adopting a likelihood definition. What we gain from our new framework, then, is that we see old theories in a new light, and that this helps us to understand their respective entanglements better. What is more, the point of view presented here is not merely some new one; since it is, if I am right, the one that arises from the preferable definition of epistemic peerhood, this point of view is the most suitable.

I have not taken pains to present conclusive arguments against the alternatives to EW that I have discussed; but I will, in the next section, argue— conclusively, I hope—against a marginalisation of EW that Elga, who accepts EW (or his refined version of it), puts forward in "Reflection and Disagreement." In doing so, I shall elaborate further on what understanding of the likelihood definition is appropriate.

5. Clusters of Controversy

If what I have said so far is correct, we should not discuss whether EW is true— for it clearly is—but rather what circumstances, apart from well-known ones such as lack of competence or information, might prevent someone from being equally likely to be right. Does disagreement on some very fundamental principles, for instance, suffice for not counting someone as a peer regarding matters stemming from those principles? Elga thinks it does; more particularly, he holds that in a 'cluster of controversy,' that is an extended field of related issues on which disagreement prevails, one is unable to determine whether the people one disagrees with are one's peers. To see what this means, consider

> FOUR-DIMENSIONALISM. I am of the opinion that four-dimensionalism is an extremely useful and well-founded philosophical theory. Our metaphysical worldview gets so much more elegant and straightforward once we adopt it! Unsurprisingly, I find highly plausible many theses that fit well with my adherence to four-dimensionalism: that persons just are mereological sums of person stages; that identity can be contingent; that eternalism is to be preferred over presentism, counterpart theory over accounts of transworld identity, semantic approaches to vagueness over ontic or epistemic ones; and so on. You,

on the other hand, find hardly any of this plausible. You argue that we essentially conceive the world as consisting of enduring material objects (among them, most notably, ourselves), and that the structure of our thought about the world cannot be adequately captured by a metaphysical framework so alien to experience and common sense. It is thus not astonishing that although we are both (as we are happy to admit to each other) very able and well-read philosophers, the two of us favour completely different theories in many areas of modern philosophy.

FOUR-DIMENSIONALISM confronts us with a cluster of controversy about broadly metaphysical issues. If you and I are peers in this scenario, we must, according to EW, suspend judgement about whether four-dimensionalism is true. Of course, we could deny that we are peers by adopting one of the alleged rivals to EW discussed in the previous section. Elga, however, rejects these views. He argues instead that the mere fact that our disagreement is not isolated but concerns a whole bunch of related claims makes it impossible for me to count you as a peer. This is the case because, in order to judge whether you are my peer, I normally consider how reliable your beliefs about related matters are. For example, if I am about to judge whether you are my peer concerning some multiplication problem, my evidence is how good you generally are when it comes to calculating. In FOUR-DIMENSIONALISM, however, I cannot tell how reliable your beliefs about related matters are, as our controversy extends to related matters as well. (In fact, I think it goes as deep as to the question of what demands a good philosophical theory should satisfy.) To suppose that your beliefs are misled because they differ from mine would beg the question. Neither can I suppose, according to Elga, that you are as competent a metaphysician as I am, and that your beliefs in this area are, for that reason, as reliable as mine, because I simply lack the resources to judge whether you are as competent as me. How could I assess your metaphysical competence if your idea of how closely our metaphysical theories should resemble the way we actually conceive of the world differs so much from mine? Therefore, Elga denies that, in cases like FOUR-DIMENSIONALISM, the parties to the disagreement think that they are peers.[31]

(In fact, the last step of the argument is a bit too fast. What is lacking is the premise that mere ignorance about other people's peer statuses does not justify regarding them as peers; a positive doxastic attitude is needed. If one believes neither that some other people are epistemically superior nor that they are epistemically inferior, one would not thereby have a reason to count them as

---

[31] See Elga, "Reflection," 492–497. Elga's own example concerns abortion, not four-dimensionalism. All further arguments in this text apply equally well to both cases.

equally likely to be right. This appears most plausible if there is no conceivable way of comparing two people's respective peer statuses. In this case, we do not just suffer from a removable lack of knowledge; we rather face a scenario in which the question of whether those people are peers becomes meaningless. Perhaps Elga should be interpreted as believing that clusters of controversy do constitute such an unbridgeable gulf for peerhood ascriptions.)[32]

It should be emphasised that Elga, in presenting his line of reasoning, consistently uses expressions like "someone counts another one as a peer," "someone thinks another one is equally likely to get things right," or "someone has a certain opinion about another one's abilities." He says neither anything like "someone *is* another one's peer" nor anything like "someone *should* count another one as a peer." Throughout the whole argument, questions about whether or not certain people are peers are asked, or answered, only from the point of view of these people. There is no 'view from above,' no third-person perspective; neither is there a normative dimension involved. This is quite surprising, because, for one thing, the objective, third-person point of view appears to be considerably more relevant than the first-person one. After all, the parties to a disagreement are usually much more prone to error in judging the other parties' abilities than an impartial observer who specifically concentrates on peerhood issues. For another thing, we cannot avoid bringing in a normative dimension sooner or later because our aim in epistemology is not to describe what kinds of beliefs people actually tend to hold under specific circumstances but to explore what kinds of beliefs they should hold under such circumstances. What interests us is not under what circumstances people generally happen to count others as peers but under what circumstances they are justified in doing so.

So, in short, we can ask peerhood questions in three ways: first-person, third-person, and normative. How are these three ways connected? What are Elga's reasons for asking only the first-person questions? And is he right in doing so?

Recall Elga's definition of peerhood in the footnote: "On my usage, you count your friend as an epistemic peer with respect to an about-to-be-judged

---

[32] Most philosophers seem to accept the premise quite generally anyway. See e.g. Enoch, "Truthometer," 956. An exception is Vulich, whose argument, however, is based on a reformulation of EW, according to which it is not suspension of judgement that is called for in the case of a peer disagreement but reconsideration (Richard Vulich, "Peer-Hood," *Logos & Episteme* 2 (2011)). A more conclusive (and, at second glance, equally pertinent) objection to the premise is King's argument to the point that lacks of clarity concerning peer statuses raise similar epistemic problems as cases of peer disagreement (King, "Disagreement," 267–269).

claim if and only if you think that, conditional [on] the two of you disagreeing about the claim, the two of you are equally likely to be mistaken." Elga defines *epistemic peer* only for first-person, descriptive contexts.[33] His reasons for this are given in another footnote (number 14), in which he points out why he describes a peer disagreement problem in a certain way:

> Note that in setting up the problem, the initial assumption is that you *count* your friend as your epistemic peer. That contrasts with some presentations, in which the initial assumption is that your friend *is* your epistemic peer. The former assumption is appropriate, however. For example, one sometimes is reasonable in thinking wise advisors to be foolish. Evidence, after all, can be misleading. In such cases, one is reasonable in being guided by one's *assessments* of the advisor's ability, even if those assessments are in fact incorrect.[34]

Here, Elga draws our attention to cases in which one's judgement that another one is one's peer is wrong because of, for instance, misleading evidence. In these cases, he claims, we are nevertheless right in considering the other one a peer because our judgement is all we can rely on. One cannot be guided by how things are but only by how one takes things to be. For that reason, it does not matter whether peerhood actually holds; reasons for belief revision can arise only if one *thinks* that peerhood holds.[35]

There are at least two points to make here. First, while it is certainly right that we should give equal weight to the opinions of those whom we have reason to count as our peers, regardless of whether they actually are, it is not clear why this fact presents a reason not to define epistemic peerhood in the way I did. With my definition in place, we can equally well set up peer disagreement scenarios in which we assume that the parties to the disagreement count themselves as peers. (Elga does not explicitly deny this, but refrains nevertheless from giving a full-fledged definition of peerhood.)

The second point is that Elga's considerations do not speak against a normative account, that is against saying that we should give equal weight to the opinions of those whom we *should* count as our peers (where the rational normativity signified by *should* is internalist). In fact, Elga himself brings in a normative element when writing that "one is *reasonable* in being guided by one's

---

[33] The wording in the definition is of course second-person. What matters here, however, is that the ascriber is not abstracted away, and that we thus cannot extract necessary and sufficient conditions for being a peer *simpliciter*.

[34] Elga, "Reflection," 499. All italics are in the original.

[35] See Nicholas Tebben, "Peer Disagreement and the Limits of Coherent Error Attribution," *Logos & Episteme* 4 (2013), 179–180, and King, "Disagreement," 262.

assessments." This is a natural move because beliefs obviously differ with regard to their reasonableness. And the belief that some people are one's peers is reasonable precisely if one is justified in counting them as peers. So the idea that the crucial question, from an epistemologist's point of view, is whether one should count someone as a peer (and not whether one would do so) might be regarded as a natural specification of what is meant here by "counting someone as a peer."

Perhaps Elga had something like this in mind. Perhaps he thought that it goes without saying that one normally counts people as peers if and only if one should do so. Perhaps he chose his specific wording only to emphasise that it is always from a first-person point of view that peerhood is ascribed or withheld, and not to emphasise that a normative understanding is out of place. Be that as it may; the effect of his consistent avoidance of normative vocabulary and his constant adherence to first-person language as far down as to the definition of peerhood is that his argumentation to the point that people involved in a cluster of controversy cannot be peers seems more convincing than it otherwise would. The reason is that in asking the normative question, we often feel forced to reflect more deeply on the matter under consideration, and may in this indirect way eventually come to take on what is in fact a third-person perspective. To be sure, in doing so, we can never acquire the point of view of an omniscient observer; but we can, up to a certain extent, abstract from what we actually believe. In other words, we may widen the first-person perspective, which we cannot cast off in all real cases of disagreement, by reminding ourselves that we can tackle problems from different ends, and that only trying to do so might reveal us a sufficiently broad view. Let me give an example.

Concerning FOUR-DIMENSIONALISM, I have argued on Elga's behalf that I lack the resources to judge whether you are my peer because assessing the quality of your beliefs on related matters is out of question for me for the simple reason that your beliefs on related matters are part of the same cluster of controversy. The consequence, according to Elga, is that I have no determinate opinion about whether you are my peer. This may seem a reasonable description of what one would think in this scenario. But is it a reasonable description of what one *should* think here? If the question is asked that way, other considerations may arise: am I really justified in ignoring your opinion just because our beliefs differ so extensively? How could it be that I am rationally required to hold metaphysical views so completely different from those that you are rationally required to hold (recall that the scenario is symmetric)? Is it really impossible to assess, however imperfectly, the quality of your reasoning in comparison to mine? In order to approach these questions, imagine someone who is neutral on the whole cluster of

issues surrounding four-dimensionalism. Such a person would not care about how extended our controversy is because her ignorance of the entire cluster makes it impossible from the outset for her to assess how reliable our beliefs on related matters are. In trying to assess our peer status, however, she would naturally draw on less specific pieces of evidence such as general intelligence, lucidity of reasoning, or professional reputation. (In a way, she would pay less attention to track records and more attention to the characteristics that usually cause such records. This observation fits well with the criticism of the track record criterion at the end of section 2. Recall also that my critique of the list definition is perfectly compatible with the fact that we should usually check people's peerhood statuses with the help of a detailed list of properties) If the result is that, to the best of her knowledge, the two of us are on the whole equally good metaphysicians, then she is justified in regarding us as peers and, consequently, in becoming agnostic about whether four-dimensionalism is true. We, as parties to the disagreement, should do the same: even if we have to set aside many of our deeply held metaphysical beliefs, there is enough left to enable us to assess, however roughly, whether or not we are equally competent and well-informed, and hence whether or not we are peers. It seems odd that the improvability of this assessment should be the crucial reason for me to stick to my original belief that four-dimensionalism is well-founded (and for you to stick to your original belief that it is not).

In sum, once we realise that, from the viewpoint of people who do not have to set aside many deeply held beliefs, the comprehensiveness of our disagreement need not prevent assessment of our peer status, we should join them in forming a belief about our peer status from a broader point of view. Therefore, according to my line of argument, and *pace* Elga, we cannot generally deny a peer status to people with whom we disagree extensively. A corollary of this is that a list of factors by means of which we assess epistemic peerhood must not include background assumptions or methodological preferences.

6. The Limits of Rationality

So far, I have argued that a thorough understanding of the concept of epistemic peerhood results in a likelihood definition, which in turn leads to immediate acceptance of EW, and that the mere fact that a disagreement is widespread and deep-rooted does not rule out the fact that the parties to the disagreement are peers. Yet how widespread and deep-rooted could a peer disagreement maximally be? Are there any limits? Or, to put the question differently, are there any non-

trivial circumstances that generally allow us not to regard equally well-informed and competent people as peers?

The lesson we can learn from Elga's line of reasoning is that we need some common ground in order to compare people with respect to their probability of being right on a certain matter; once we set aside too much, we are at a loss concerning whether they are, for example, equally competent. Contrary to what Elga thinks, clusters of controversy are not sufficient for rendering us so uncertain, for we can still rely upon comparatively unspecific but sufficiently informative characteristics such as intelligence, thoroughness, or freedom from bias. As we have seen, we are in fact even required to rely upon these factors because, first, what interests us is not whether we would count someone as equally likely to be right but whether we should do so, and, second, this obligation cannot be fulfilled by ignoring evidence such as reliable information about those characteristics. We would encounter serious problems, however, if we were not able to acquire such reliable information. Then we would indeed have no idea whether we should count someone as our peer. To see how this can happen, consider

> THE AWKWARD COLLEAGUE. I remember that there is a meeting tomorrow. Unfortunately, I have forgotten who will take part. So I ask a colleague of mine, Jane, whom I know to be a reliable and in fact pretty smart person. She says that, apart from me and her, only Simon and Sue will be there, so in total we will be five people. "Wait a moment," I say, "you and I are two people, and Simon and Sue are also two people; that makes four, so four is the number of people who will be there." Jane shakes her head in disbelief and asserts, "No, two and two equals five."

The situation here is supposed to be symmetric: Jane is as baffled as I am when she notices that we differ about what two and two adds up to. So I cannot just refer to the profoundness of my belief that 2+2=4, or the strength of my justification for it, in order to discredit her as a peer.[36] But is there a way of establishing the opposite, namely that she is indeed as likely as I am to be right? (A consequence thereof would be that I have to suspend judgement on whether 2+2=4.)

In assessing whether Jane is my peer in elementary maths, I cannot draw on how reliable her beliefs about other simple computations are. Even if the results she arrives at are frequently different from mine, I am not allowed, according to the argument in section 5, to conclude that her beliefs are erroneous solely from

---

[36] Taking symmetry seriously also rules out Christensen's idea that common-sense checking resolves the problem (Christensen, "Good News," 199–201). This point is also made by Tomas Bogardus, "A Vindication of the Equal-Weight View," *Episteme* 6 (2009), 329.

the fact that we are tangled in a cluster of controversy. If, on the other hand, our disagreement is restricted to very few cases (or perhaps even to the one concerning 2+2), this is not by itself a proof of the reliability of her beliefs because track records—and nothing other do I compile when I compare her results with those that I regard as true—are, as we have seen, only probable but not indefeasible effects of what is actually essential, namely such characteristics as competence and well-informedness. Yet it is exactly her competence that I have reason to doubt when faced with her apparently insane belief that 2+2=5. This single belief is, given all I know about addition, so irrational that it undermines any argument that deduces from her being almost always correct that she is competent in adding numbers. Consider, for instance, the question of whether it is rational for someone who believes that 2+2=5 also to believe that 3+3=6. I do not think that our understanding of rationality allows us either an affirmative or a negative answer to this question; we simply have no idea, and cannot have any idea, what way of adding numbers would be rational under the presumption that 2+2=5. Therefore it does not matter, under this presumption, how good our track record is; any track record is made insignificant by a disagreement on such a basic level.

These considerations suggest that, in extreme cases of disagreement, the disagreement itself can indeed count as evidence for whether or not the parties to the disagreement are peers. Elga would agree, as I mentioned at the beginning of section 4. His argument is that if we agree prior to the disclosure of the disagreement that we will not count someone as a peer in case we find her opinion utterly insane, then actually finding her opinion utterly insane would justify us in regarding her as an epistemic inferior. However, as Elga points out, if the situation is symmetric, that is if our supposed peer has exactly the same doxastic attitude towards our belief that we have towards hers, EW again requires both of us to suspend judgement.[37] This last claim, among other things, is wrong. Before focusing on the errors in Elga's argument, however, let's first see what the correct account looks like.

Although we should not normally use our discussion partner's belief to conclude anything about her peer status, extreme cases such as THE AWKWARD COLLEAGUE are an exception because, in such cases, her belief functions as a defeater of whatever evidence we may have regarding her peer status. In THE AWKWARD COLLEAGUE, Jane seemed to me a pretty smart person, but whatever reasons I had for this belief are undermined by the fact that she thinks that 2+2=5. Could one really call someone smart who thinks that 2+2=5? I do not know; or,

---

[37] Elga, "Reflection," 491.

rather, I think that our very understanding of smartness, or intelligence, or competence, is simply not apt to answer that question. Learning that Jane thinks that 2+2=5 makes it impossible for me to compare her respective degrees of smartness, intelligence, and competence to mine. In order to do so without just begging the question and assuming that I am right about 2+2, I would have to set aside beliefs that are so fundamental that there is not enough conceptual knowledge left to say what smartness, intelligence, competence, etc. actually are. If some disagreeing interlocutor's thinking differs too radically from ours, we cannot apply our qualitative concepts to it without thereby taking sides.[38]

In scenarios such as THE AWKWARD COLLEAGUE, there is therefore no point in comparing my probability of being right with Jane's. Our respective peer statuses are *incommensurable*. Thus, I cannot just regard Jane as an epistemic inferior; I rather cannot say anything about whether she is my superior, inferior, or equal, and that is not because I lack some relevant information. There could be no information that would help in this case. As a consequence, EW cannot be applied in cases like THE AWKWARD COLLEAGUE, because EW tells us only how much weight we should give our peers' opinions, but not how much weight we should give to the opinions of people whose peer statuses are incommensurable to ours. Therefore, nothing forces us to revise our beliefs in extreme cases of disagreement. (Since there is obviously no sharp line that divides cases like THE AWKWARD COLLEAGUE from those in which we are not left stranded in assessing competence, we cannot always tell whether EW should be applied.)

Admittedly, this might appear unsatisfying.[39] The reason is that we can conceive of a generalisation of EW—call it EW*—which holds that we should also give equal weight to the opinions of those whose peer statuses are incommensurable to ours. Then my disagreement with Jane would indeed force me to suspend judgement on whether 2+2=4. In favour of EW*, one could argue that we have no reason to take our actual conception of rationality to be right. The consequence would be relativism about rationality, according to which a judgement p has to be understood as elliptical for a judgement of the form p *relative to the conception of rationality R*—in the same way as Einstein's relativistic conception of mass entails that judgements like *x has mass M* have to be interpreted as *x has mass M relative to spatio-temporal framework S,* or moral

---

[38] Analogous points might be made for other list definition factors, e.g. well-informedness. It appears, however, that uncertainty concerning the application of our terms for those factors goes hand in hand with uncertainty concerning the application of evaluative terms for characteristics such as competence or rationality, so it is the latter we should primarily focus on.
[39] I am grateful to an anonymous reviewer for pressing me on this point.

relativism entails that judgements like *It is morally wrong to φ* have to be interpreted as judgements like *It is morally wrong to φ relative to moral framework M.*[40] Against EW*, one has to adduce that relativism about rationality entails what could be called epistemic antirealism about truth, namely that there is no proposition of which we could know that it is objectively true (and not just true relative to some specific conception of rationality). The reason is that every judgement is contestable by someone whose underlying conception of rationality is sufficiently different, so that, according to EW*, disagreement would force us to suspend judgement on whether it is true. (It is of little help to point out that we quite rarely, if at all, encounter extreme cases of disagreement, and that we are therefore almost never forced to suspend judgement on propositions of whose truth we are highly confident. For one thing, it would be odd if our justification to believe such propositions as *2+2=4* depended on the contingent fact of whether or not we have met someone who honestly denies their truth; and for another thing, it has been argued that merely possible disagreements are of epistemic significance as well,[41] so that mere contestability might already suffice for suspension of judgement.) Hence, there either are no propositions that are objectively true, or, if there are any, we will never be able to establish their truth. In short, EW* is on the one hand more demanding as EW, since it requires a non-trivial further assumption, and comes on the other hand with considerable theoretical costs regarding the nature and epistemology of true propositions. I therefore ignore EW* in what follows, and continue to discuss EW, which, as we have seen, cannot be applied to cases like THE AWKWARD COLLEAGUE, in which the protagonists' peer statuses are incommensurable.

Regarding Elga's way of dealing with extreme cases of disagreement, this means that he is wrong both in thinking that EW requires us to suspend judgement if the extreme case is symmetric (like THE AWKWARD COLLEAGUE) and in thinking that we should regard our discussion partner as epistemically inferior if she does not find our belief as irrational as we find hers. Since in both symmetric and asymmetric cases we cannot reasonably compare the other one's peer status with ours, EW allows us to disregard his or her opinion in both kinds of cases—though in each case not on the basis of epistemic inferiority.

---

[40] See Gilbert Harman, "Moral Relativism," in *Moral Relativism and Moral Objectivity*, Gilbert Harman and Judith Thomson (Oxford: Blackwell, 1996), for a comparison between relativism about mass and relativism about morals (1–3, 13, 18–19, 41) as well as for an outline of how moral relativists can explain actual disagreements (32–44). This outline would by and large carry over to the case of relativism about rationality.

[41] See e.g. Marc Andree Weber, "Armchair Disagreement," *Metaphilosophy* 48 (2017).

Recall from section 3 that Elga gives two reasons for revising EW. One is that we have to take into account the specific circumstances of the disagreement (e.g. hot weather). I argued that this should rather be done by relativising peerhood to time. The other is that we should have a plausible way of dealing with extreme cases of disagreement. We now see, however, that the potential occurrence of such cases does not put us under any pressure to refine EW in the way Elga suggests: since the parties to an extreme disagreement cannot reasonably be regarded as peers, EW, in the sense of definition 2, simply does not apply. Thus Elga's refined version of EW remains unmotivated.

7. Conclusion

The two main insights in this paper are, first, and anticipated in a footnote of Elga's, that a likelihood definition—preferably my definition 1—captures the essence of epistemic peerhood better than any kind of list definition and should therefore be adopted; and, second, that this adoption necessitates a shift of focus in the debate on peer disagreement: the alleged theoretical alternatives to EW, which are in fact compatible with it, should not be taken to entail that we should give our peers' beliefs less than equal weight, but only that even very competent and well-informed people may easily fail to be our peers. Revised accordingly, however, these alternatives might appear harder to believe.

Further insights in this paper include criticisms of several of Elga's views that are based on or related to his understanding of epistemic peerhood (the bootstrapping argument; his version of EW; the thesis that EW does not hold in clusters of controversy; his way of dealing with extreme cases of disagreement). They also include the idea that in extreme cases of disagreement—but only in those—we lack the common ground needed to compare people's likelihoods of being right. This incommensurability, however, sets only somewhat inextensive limits to our ability to ascribe peerhood and apply EW. As you will never in your life encounter a person like my colleague Jane, you should always give your peers' opinions equal weight.