

DISAGREEMENT AND DEEP AGNOSTICISM

Eric GILBERTSON

ABSTRACT: One defense of the “steadfast” position in cases of peer disagreement appeals to the idea that it’s rational for you to remain *deeply agnostic* about relevant propositions concerning your peer’s judgment, that is, to assign no credence value at all to such propositions. Thus, according to this view, since you need not assign any value to the proposition that your peer’s judgment is likely to be correct, you need not conciliate, since you can remain deeply agnostic on the question of how the likelihood of your peer’s judgment bears on the likelihood of your own. This paper argues that the case for deep agnosticism as a response to peer disagreement fails. Deep agnosticism (as a general thesis) implies that it is sometimes permissible to withhold judgment about whether there is a *non-zero* chance of a proposition’s being true. However, in cases of disagreement where deep agnosticism is supposed to support the steadfast position, such withholding isn’t rational. This is because of constraints placed on rational credence by objective probability or chance, which ensure that rational credence adequately reflects strength of evidence.

KEYWORDS: epistemic peerhood, disagreement, deep agnosticism, credence

1. Introduction

Consider the following case (from Grundmann¹):

Election. There is a very close presidential race between two competing candidates, X and Y. After the people have voted, two equally renowned election forecasters using different, but equally representative samples come up with conflicting predictions regarding who won. Forecaster A predicts that X has won by a narrow margin. Forecaster B predicts that Y has won by a narrow margin.

In this case it certainly seems as if A, upon recognizing that B has made a different judgment, acquires evidence that defeats her initial justification for the belief that X has won. Thus it seems to be a case in which disagreement with an acknowledged *epistemic peer* (roughly, someone who has equally good evidence and is equally good at reasoning on the basis of such evidence) requires one to reduce confidence in one’s initial judgment, and perhaps even to suspend belief.

In suspending belief on the matter of who won the election, A would give equal weight to the possibility that A’s judgment is incorrect and to the possibility that B’s judgment is incorrect.² Thus *Election* and other similar cases involving peer

¹ Thomas Grundmann, “Why Disagreement-Based Skepticism Cannot Escape the Challenge of Self-Defeat,” *Episteme* (2019): 1–18, <https://doi.org/10.1017/epi.2019.14>.

² Jane Friedman (“Rational Agnosticism and Degrees of Belief,” in *Oxford Studies in Epistemology*,

disagreement are taken by some to support the *equal weight view*: roughly, that in cases of disagreement with epistemic peers you should give equal credence to the possibility that you have made a mistake as to the possibility that your peer has.³ Depending on precisely how the equal weight view is understood, the view known as *conciliationism* may in one sense go further, as it holds that you are rationally required to adjust your credence in the disputed proposition in the direction of your peer in such cases.⁴

But is it really rationally required to suspend judgment, or to adjust your credence, in such cases? These positions have seemed to some to depend on a principle of indifference which is suspect.⁵ That is indeed a legitimate concern. Roger White, however, provides a strong defense of principles of indifference, arguing that, at least in the case of a multiple partitions problem, the apparently absurd results do not depend on them.⁶ So these principles may support the equal

Vol. 4, ed. T. S. Gendler and J. Hawthorne (Oxford: Oxford University Press, 2013), 57-81, 57) argues that there are cases where suspension of belief is clearly epistemically permissible regardless of whether it's rational to assign any credence to the relevant propositions. However, these are cases where the subject is completely ignorant about the matter in question. For example, Friedman suggests that it's permissible to suspend belief about what the price of copper will be in 100 years, or whether the Hill 50 Gold Mine was Australia's most profitable mine between 1955 and 1961. In the sort of case of peer disagreement we are considering, however, one is not completely in the dark. In such circumstances, suspension of belief seems to require giving equal weight.

³ It should be noted that the equal weight view as stated is somewhat imprecise. In particular, as Jehle and Fitelson ("What is the 'Equal weight view'?", *Episteme* 6, 3 (2009): 280–293) have shown, there are a number of different possible precisifications of the view which clearly specify how someone should update her credence upon discovery of disagreement with an epistemic peer. I will however rely on the less precise notion of giving equal weight, as I think that however the notion is precisified it will not affect my argument; the crucial claim that I defend is that whether or not *regarding someone as an epistemic peer* entails *giving equal weight* to the respective judgments, it entails assigning some (perhaps indeterminate) credence value to relevant propositions concerning your and your peer's judgments. Hence, precisely what the rule for updating should be on this view is a further question.

⁴ If it is possible – and rationally permissible – to give equal credence to the possibility that your and your peer's judgments are incorrect without adjusting your credence in the disputed proposition, then conciliationism is not implied by equal weight. But if this is not rationally permissible – perhaps because giving equal weight requires "splitting the difference" between degrees of belief – then conciliationism is implied by equal weight.

⁵ See, for instance, D. Gillies, *Philosophical Theories of Probability* (London: Routledge, 2000); B. van Fraassen, *Laws and Symmetry* (Oxford: Oxford University Press, 1989).

⁶ Roger White, "Evidential Symmetry and Mushy Credence," in *Oxford Studies in Epistemology*, eds. T. S. Gendler and J. Hawthorne (Oxford: Oxford University Press, 2013), 161–186.

weight view or conciliationism, after all. Moreover, Adam Elga presents a “bootstrapping” argument for the equal weight view,⁷ and thus against the view that it’s rationally permissible to remain steadfast in the face of disagreement with epistemic peers⁸—again, suggesting that some form of conciliationism is correct.

The case for the equal weight view and conciliationism has been challenged in another way, however, by Jason Decker, who argues that the arguments for these views overlook the possibility of *deep agnosticism*, by which he means a kind of agnostic credal state which involves a partially-defined (“gappy”) function that fails to assign any value to certain propositions.⁹ Thus Decker defends the rationality of your reacting to peer disagreement with such agnosticism—about the likelihood that your peer’s judgment, rather than your own, is correct—rather than with conciliation or suspension of belief. By withholding judgment on this question, Decker argues, you can remain steadfast in your initial belief: because you take no stand on how likely it is that you peer’s judgment is correct, your recognition of that judgment need not diminish your confidence in your own.

In this paper, I argue that the case for deep agnosticism as a response to peer disagreement fails. First, deep agnosticism (DA)¹⁰ implies that it is sometimes rationally permissible to refrain from assigning any positive credence value at all, whether determinate or not, to a proposition. Given plausible constraints placed on rational credence by objective probability or chance, this implies that it is permissible to withhold judgment about whether there is a *non-zero* chance of the proposition’s being true. However, in those cases of disagreement in which deep agnosticism is supposed to be rational, such withholding isn’t rational. Second, I argue that it is inconsistent to assign a high credence to the proposition that another

⁷ Adam Elga, “Reflection and Disagreement,” *Noûs* (2007): 478–502.

⁸ Decker (“Disagreement, Evidence, and Agnosticism,” *Synthese* 187 (2012): 753–783) labels this view “Millianism.”

⁹ Decker, “Disagreement, Evidence, and Agnosticism;” Decker, “Conciliation and Self-incrimination,” *Erkenntnis* 79 (2014): 1099–1134.

¹⁰ It’s important to understand that Decker sometimes uses the term ‘deep agnosticism’ to refer not to this general position (which has nothing essentially to do with disagreement), but instead to a more specific position that licenses credal gaps with respect to specific propositions concerning a peer’s evidence or her judgment. To make matters yet more confusing, Decker also uses the term to refer to the *credal state* one is in with respect to propositions which one’s credence function has assigned no value. In what follows, I’ll frequently rely on context to indicate which sense of the term is intended. However, I’ll use ‘DA’ to refer to the more general position and ‘DA_d’ to refer to the more specific position Decker defends, that is, that it’s rationally permissible to assign no credence value at all to relevant propositions concerning an epistemic peer’s judgment in cases of disagreement. I think it is clear, moreover, in discussion of certain instances of peer disagreement, when I intend to refer to the more specific thesis and when I intend the “credal state” sense.

person is an epistemic peer while assigning no credence at all to the proposition that her evidence, and her ability to handle that evidence, is as good as your own. Since (I argue) the view that deep agnosticism is a rational response to peer disagreement implies that this is rational, it should be rejected.

The failure of deep agnosticism shows that the proponent of the “steadfast” position, according to which it is rationally permissible to stick to one’s initial judgment in cases of peer disagreement (i.e., to refrain from conciliation), must somehow accommodate the fact that in such cases your acknowledgement of an epistemic peer as such requires you to accept that there is a significant (objective) chance that your peer’s judgment is correct, and that this provides you with some reason to think that your own judgment is incorrect.¹¹ The challenge for the proponent of the steadfast position is to show that such an acknowledgement does not in fact require conciliation. If I am right, then this challenge cannot be avoided simply by appealing to the possibility of withholding judgment about the status of epistemic peers’ judgments, as deep agnosticism maintains.

The paper proceeds as follows. First, in §2, I explain Elga’s bootstrapping argument in support of the equal weight view. In §3, I explain Decker’s criticism of Elga and his defense of deep agnosticism (DA_d), and I present my main argument against the view. In §4, I consider whether deep agnosticism may provide a satisfactory solution to the problem of multiple partitions, and I argue that it does not. In particular, I argue that deep agnosticism has certain unacceptable implications in this context. Moreover, I argue that Lewis’s Principal Principle, which holds that a rational agent conforms her credences to known objective chances, provides further support for the claim that it’s rationally impermissible to assign no credence value to the proposition that your epistemic peer is more likely to be correct than you. Finally, in §5, I respond to an objection to my argument’s reliance on the notion of epistemic peerhood.

Before proceeding I should explain more precisely how I am understanding the notion of an epistemic peer. First, I take epistemic peerhood to be *domain-relative*. That is, X is an epistemic peer of Y with respect to a certain issue or subject matter (the morality of abortion, climate change, mathematics, modal logic, etc.). As issues in different domains may be independent of one another, it’s possible to regard someone as an epistemic peer in some cases of disagreement and not in others. I’ll sometimes speak of X’s being an epistemic peer of Y with respect to a certain

¹¹ Of course, it may be that mere recognition of your fallibility on the relevant matter also provides you with such a reason. Whether or not this is so, the point here is that disagreement with an acknowledged epistemic peer provides an independent reason to think that your judgment is incorrect.

proposition p . It would be more accurate to attribute peerhood with respect to some relevant domain of truths which includes $p/\text{not-}p$. But I assume then that epistemic peerhood holds relative to propositions, if it holds relative to domains. So it is not strictly improper to formulate things as I do, and it is convenient to do so. It is also important to understand that epistemic peerhood, as I understand it, doesn't require—although it is compatible with—*identical evidence*; it is sufficient, as far as the disputants' first-order evidence on the matter in dispute goes, that their evidence is *equally good, though not identical*.¹² In the election case, for example, the two forecasters rely on different evidence, but their evidence involves equally representative samples of voters. Finally, epistemic peerhood is understood as involving *equality of higher-order evidence*, in the following sense. Provided X and Y are epistemic peers, the dispute-independent evidence regarding X's and Y's epistemic credentials and trustworthiness on the matter in question no more supports the claim that X has made an error than that Y has. This condition (sometimes referred to as the Independence principle) is important, as it requires that you assess others for epistemic peerhood without relying on your particular reasoning or evidence regarding the disputed matter itself. Reliance on such reasoning or evidence would seem question-begging, as it would allow you to dismiss someone's judgment, and regard them as an epistemic inferior, simply because it conflicts with your own, when you have independent reasons to regard them as being in an equally good position to judge (on the matter in question).¹³

So my understanding of epistemic peerhood is this:

X and Y are *epistemic peers* with respect to a proposition p iff (i) X and Y have the same, or equally good, relevant (first-order) evidence regarding p and (ii) X and Y are equally competent in reasoning based on this evidence and (iii) the higher-order evidence available to X and Y concerning X's and Y's epistemic credentials

¹² Cf. Grundmann, "Why Disagreement-Based Skepticism Cannot Escape the Challenge of Self-Defeat," 3.

¹³ Cf. David Christensen, "Conciliation, Uniqueness, and Rational Toxicity," *Noûs* 50, 3 (2016), fn. 7: "Conciliatory views require an agent to determine the seriousness with which she takes another's disagreement in a way that's *independent* of the agent's own particular reasoning on the disputed matter. The thought behind this is that insofar as disagreement can raise doubts about the cogency of one's reasoning on a certain matter, it would be question-begging to dismiss those doubts in a way that relied on the very reasoning under dispute." It should be noted that not everyone accepts this. See, for instance, Jennifer Lackey, "What Should We Do When We Disagree?," in *Oxford Studies in Epistemology*, ed. T.S. Gendler and J. Hawthorne (Oxford: Oxford University Press, 2010); S. Benjamin, "Questionable Peers and Spinelessness," *Canadian Journal of Philosophy* 45, 4 (2015): 425–444; T. Kelly, "Disagreement and the burdens of judgment," in *The Epistemology of Disagreement: New Essays*, ed. D. Christensen and J. Lackey (Oxford: Oxford University Press, 2013): 121–166.

Eric Gilbertson

and trustworthiness (and specifically, the fact that X and Y believe *p* or *not-p* based on the same, or equally good, evidence) equally supports the proposition that X's judgment regarding *p* is mistaken and the proposition that Y's judgment regarding *p* is mistaken.¹⁴

Finally, let me be clear about my aims. I do not intend to resolve the dispute between proponents of the steadfast position and conciliationists. While my argument, supposing it succeeds, may provide some support for conciliationism, I don't think that such support is very strong. The point of the argument is to show that one particular attempt to defend the steadfast position, which depends on deep agnosticism, fails; and to show that this follows from certain assumptions about what is entailed by regarding someone with whom one disagrees as an epistemic peer. If this is right, it clearly does not mean that the steadfast position is simply indefensible, or that it has any untoward consequences which are independent of deep agnosticism. It may be rationally permissible to remain steadfast in cases of peer disagreement, though not because this is entailed by deep agnosticism. The position that I defend here is that deep agnosticism does not explain why remaining steadfast is coherent and rational, because deep agnosticism is (in these circumstances) irrational.

2. Elga's Bootstrapping Argument for the Equal Weight View

Elga's argument for the equal weight view appeals to the following case:

The horse race. You and an acknowledged epistemic peer are watching a horse race. After the race is run, you're convinced that Cavonnier won and your peer is convinced that Grindstone won.

Elga writes:

When you learn of your friend's opposing judgment, you should think that the two of you are equally likely to be correct. For suppose not – suppose it were reasonable for you to be, say, 70% confident that you are correct. Then you would have gotten some evidence that you are a better judge than your friend, since you would have gotten some evidence that you judged this race correctly, while she misjudged it. But that is absurd.¹⁵

¹⁴ It is also worth noting that it is important that the notion of epistemic peerhood, in particular (iii), be understood in terms of the bearing of higher-order evidence, rather than in terms of *total* evidence, on likelihood of correctness. For if it's understood in terms of (Xs or Ys) total evidence, then the equal weight view would follow trivially from evidentialism. (Thanks to an anonymous referee for helping me see this point.)

¹⁵ Elga, "Reflection and Disagreement," 486.

Elga's thought is that if we consider a series of such events, where in each it is thought to be reasonable for you to be more confident that you got it right than that your peer did, this leads to the conclusion that you get inductive evidence that you're a better judge, thus bootstrapping your way to expertise.¹⁶ Since it is absurd to think you could do this, it follows that it's not rational, after all, to be more confident in your own judgment than your peer's. Supposing, for instance, that you were more confident in your own judgment because you judged that you have a better track record than your peer, you would thereby take the fact that you disagree as evidence that s/he's made more mistakes. That would violate Independence. Without some antecedent reason to think you're a better judge, the fact that you and your peer disagree isn't evidence that s/he's made more mistakes.

So, Elga concludes, you should instead give equal weight to both judgments and think that you're no more likely to be right than she is. And the same goes for other disagreements with epistemic peers. Any case in which you and someone you take to be an epistemic peer disagree on a claim based on the same evidence is one in which you have just as much reason to think you have erred as that she has. For if it were reasonable for you to give more weight to the proposition that your peer has erred in such a case, then it would be possible for you to gain knowledge, based on a series of such disagreements, that you are a better judge than your peer. Since this isn't possible, it's only reasonable for you to give equal weight.

3. Decker's Criticism of Elga and Defense of DA: Avoiding the Bootstrapping Problem without Giving Equal Weight

Let's suppose that Elga is right that it isn't possible to bootstrap your way to expertise by sticking to your guns in the face of peer disagreement. Decker agrees. He does not agree, however, that giving equal weight (i.e. judging that your and your peer's judgments are equally probable) is required in order to avoid the problem. In his view, you can stick to your initial judgment and take yourself to be justified in doing so without being committed to thinking that you're *more likely to be correct* than your peer; you can simply withhold judgment about this. Decker claims that it's a

¹⁶ It may be useful to compare the "easy knowledge objection" to reliabilism: According to reliabilism, it's possible to acquire knowledge e.g. that a thing is red (by using your reliable faculty of color vision) even if you don't know that the method by which you arrived at such knowledge is reliable. But, according to reliabilism, in coming to know that x is red you thereby come to know that your color vision has operated correctly. So, by acquiring such knowledge on many occasions, you come to know that your color vision is reliable. But, the objection goes, it is absurd to think you can acquire knowledge of the reliability of your color vision in this way. (Cf. Elga, "Reflection and Disagreement," 488.)

mistake to think that to remain steadfast is to give any *extra* weight to your judgment, beyond what's involved in simply making the judgment. Consequently, even if there's a series of similar cases, and you confidently stick to your judgment throughout, you need not take this as inductive evidence that you're a better judge than your peer.

As I indicated earlier, Decker's argument rests on a position he calls *deep agnosticism* (DA). According to it, some propositions are such that our credence functions assign no value (or value range) at all. In the present case, the thought is that you might have no view at all on how to assess the judgments of you and your epistemic peers. Consequently, your credence function might assign no value to the proposition that your peer's judgment is as likely, or more likely, to be correct as yours. (Decker writes, "You might well feel that one would need a kind of external God's-eye perspective of the situation in order to make such an assessment, and, lacking that perspective, you might find yourself *deeply agnostic* on who has done a better job at the task of judging."¹⁷) In Elga's view, sticking to your guns in these cases entails taking such a stand: if you stick with your first-round judgment in the face of peer disagreement, you thereby take your judgment to be more likely to be correct than your peer's. This involves a second-order judgment. In Decker's view, sticking to your guns need not involve any such second-order judgment, as you can simply withhold judgment about who's more likely to be correct. Moreover, if you withhold such judgment, this does not mean you are required to give *equal* weight to you and your peer's being correct; for you can withhold judgment on whether they are *equally* likely to be correct, too.

So the issue is whether deep agnosticism provides an adequate defense of the steadfast position. Is it indeed rational to be deeply agnostic about relevant propositions concerning the judgments of your epistemic peers? First, I propose that

¹⁷ Decker, "Disagreement in Philosophy," in *The Palgrave Handbook of Philosophical Methods*, ed. Chris Daly (MacMillan, 2015), 151. Decker presents another case in which you and a peer are official judges at a diving competition: You've been given the option of scoring the other judges on *their* performances, in addition to scoring the divers on theirs. Moreover, you have the opportunity to revise your scores for the divers on the basis of your knowledge of the other judges' scores. You might, for instance, after finding out your peer's score differs from your own, split the difference and revise your score accordingly. Indeed this might seem like the only rational response to your peer's score. Decker argues otherwise:

When one sticks with one's first-round judgment in the face of peer disagreement, one is not somehow getting evidence (or taking oneself to have somehow gotten evidence) that one is a better judge than one's peer. One isn't taking a stand on that issue one way or another. And that's fine. Contrary to the impression that is sometimes given in the literature on disagreement, it's okay to be deeply agnostic on propositions. ("Disagreement in Philosophy," 151.)

we grant that it may be reasonable to be deeply agnostic about *some* propositions regarding you and your epistemic peers.¹⁸ It's important to see that even supposing this is the case, it doesn't follow that you can take someone to be your epistemic peer, in circumstances in which the relevant higher-order evidence no more supports the claim that you've made an error than that she has, while at the same time withholding judgment about *whether her judgment is reasonable*. In fact it is plausible that to take someone to be an epistemic peer in these circumstances *is* to take a stand on this question.¹⁹ And it is even more plausible that in order for you to regard someone as an epistemic peer you must assign *some positive credence value* to the proposition that her judgment is reasonable. Moreover, this apparently gives you some reason to think that your peer's judgment is *correct*. For in judging that your peer's judgment (that *p*) is reasonable, you apparently acknowledge that your peer has some reason (given her evidence and her starting points) to think that *p* is true. Thus you appear to be committed to thinking that there is *some* reason to think *p* is true.²⁰ But, if you take a stand on *that*, then you can't rationally refrain from judgment about how her judgment bears on the likelihood that yours is correct,

¹⁸ I don't know whether it *is* reasonable, but I'll grant that it is, since I'm only interested in whether DA can provide sound defense of the steadfast position, not whether deep agnosticism is *ever* rational. For what it's worth, one proposition to which it seems it's rationally permissible for your credence function to assign no value (or no value range), is the proposition that your peer is deeply agnostic about some proposition *p*.

¹⁹ More precisely, it is plausible that regarding someone as an epistemic peer requires taking such a stand *provided you take your own judgment to be reasonable*. I am assuming that in asking what our response to peer disagreement should be we are interested only in such cases. (Thanks to an anonymous reviewer for pointing out the need to clarify this point.)

²⁰ It might be argued that (rightly) regarding another as an epistemic peer does not entail that you have some reason to think your peer's judgment is correct, at least if some form of permissivism is true. Suppose, for example, two ancient Greek philosophers, who both thought the world is finite in duration, disagreed about whether it is more than 10,000 years old. Because they lacked significant evidence that would bear on the issue, both views appear reasonable, and each is reasonable to think the other's view is, from his point of view, reasonable. Is this a reason for either to think that the other's view is at all likely to be correct? It might seem it is not. (Thanks to an anonymous reviewer for this point and for offering this example to illustrate it.) But while I grant that neither individual has a reason to think that the other's judgment is *more likely than not* to be correct, I claim that it is irrational for either individual to regard the other's view as reasonable while assigning *no credence value* to the proposition that it is correct. To do so would be to withhold judgment not only on whether the view is correct but on whether it is reasonable. In committing to the view that the other's view is reasonable, each individual rejects the claim that there's no reason at all to think the other's judgment is correct – even if (from their point of view) there may be no such *weighty* reason, or any reason that is particularly worth considering in certain contexts.

since confidence in her judgment implies (some degree of) doubt about your own.²¹ Taking a stand on the epistemic standing of your peer requires taking a stand on this question.

My suggestion, then, is that the proponent of the equal weight view should deny that it is coherent to hold both (i) that it's rational for you to withhold judgment on who is a better judge²² in this case (and on whether both judges are equally good), *and* (ii) that it is *not* rational to withhold judgment on *whether the other disputant is an epistemic peer*. Hence, if it is rationally required to regard the other disputant as an epistemic peer²³ this constrains rational credence in related propositions about her, and your, epistemic standing, including the proposition that your peer's judgment is more likely to be correct. In particular, it would be

²¹ Here it may seem I am assuming Uniqueness—i.e., that given one's total evidence, there is a unique rational doxastic attitude that one can take to any proposition. Since uniqueness is controversial, this assumption would be unwarranted. But my claim that taking a stand on the likelihood that your peer's judgment is correct rationally requires taking a stand on the likelihood that yours is incorrect does not imply uniqueness. I do not assume that you and your peer cannot be equally *rational* in this case; I do however assume that your *judgments* can't both be *correct*. Hence, to the extent that you judge that your peer's judgment is likely to be correct, you are rationally required to reduce confidence in the correctness (though not necessarily the rationality) of your own judgment. (I return to the issue of uniqueness and epistemic peerhood in §5.)

²² As an anonymous referee notes, the question of *who's a better judge* (in this particular case, or in general) is not the same as the question of *who's more likely to be correct*. It's possible, then, to be agnostic about one of these and not the other (i.e., taking a stand on one doesn't require taking a stand on the other). However, if 'better judge' is not understood in terms of track record, or somehow in terms of the likelihood of being correct (either in general, or in the particular matter at hand), then it seems to me it must be understood in terms of reasoning capability; and in that case, given the understanding of epistemic peerhood I'm relying on, I think it would be straightforwardly inconsistent to regard a disputant as an epistemic peer while withholding judgment on the question of whether s/he's a better (or worse, or equal) judge. So unless there is some other notion of what makes one judge *better* than another, which doesn't rely on considerations of track record or reasoning capability, I claim it's not rational to be agnostic on either question, provided the other is regarded as an epistemic peer.

²³ It might be argued that in many of the cases of disagreement under consideration you may remain steadfast by being deeply agnostic about *whether the other disputant is an epistemic peer*. I don't deny that this may *sometimes* be permissible and that it would allow you to remain steadfast. But I think that the typical sort of case which is of interest to philosophers is one where you have good (perhaps conclusive) reasons to take someone to be an epistemic peer – that is, it's one where your evidence doesn't permit withholding judgment about this. Moreover, Decker never suggests it is permissible to be agnostic about *whether one's peer is an epistemic peer*. He is interested rather in deep agnosticism with respect to certain propositions which presuppose epistemic peerhood, and so understanding of which actually requires you take the other to be an epistemic peer.

inconsistent for you to remain steadfast while adjusting your credences in such propositions.

Let me now present my argument more precisely and clarify my assumptions.²⁴

- (P1) If you (outright) believe that p , then you assign a credence value of 0 to the proposition that there's no chance that p is true (thus you assign credence 1 to the proposition that there's some chance p is true). (To put it another way, believing that p requires that it is doxastically necessary (for the believer) that there's some chance that p is true.)²⁵
- (P2) If DA is correct, then in cases of peer disagreement regarding p , it is rational for you to withhold judgment on whether there's a zero chance that your peer's judgment is correct. (DA implies that it's rational to assign *no value* to the proposition that your peer's judgment is correct, and to the proposition that it's more, or less, likely to be correct than yours is.)
- (P3) If it's rational for you to withhold judgment on whether there's a zero chance that your peer's judgment is correct, then – since in such cases you take your disputant to have equally good evidence, and to be equally capable of reasoning on the basis of such evidence – it is rational for you to withhold judgment on whether there's a zero chance that your own judgment is

²⁴ I should note that I'm focusing on cases of disagreement that involve *outright* belief, rather than degree of belief (supposing the latter notion makes sense). While this is not the only kind of case of peer disagreement that may be of interest, I think it is typical. My argument doesn't show that deep agnosticism about judgments of an acknowledged epistemic peer is incoherent in cases where one has less than outright belief, I think it strongly suggests that it's incoherent even in such cases. Regardless, I take it that establishing incoherence in cases of outright belief is significant.

²⁵ Dylan Dodd ("Belief and Certainty," *Synthese* 194 (2017): 4597-4621) defends a credence-doxastic possibility link: If S has a positive credence in p , then p is doxastically possible for S . My claim that if S has positive credence in p sufficient for outright belief that p , then the proposition that there's no chance that p is not doxastically possible for S . This is a kind of credence-doxastic possibility-chance link. (Dodd defines doxastic possibility as follows ("Belief and Certainty," 4604): p is doxastically possible for S iff S 's doxastic state leaves it open that p is true—iff S 's point of view or perspective is consistent with p .) He argues that if you believe that p then p is doxastically necessary for you. I take no stand on this, since I do not wish to commit to the view that is arguably implied by it (given further assumption that S 's having a positive credence in p implies that p is doxastically possible for s), that believing that p requires having credence 1 in p . (Recent proponents of this view include Roger Clarke, "Belief is Credence 1 (in Context)," *Philosophers' Imprint* 13, 11 (2013): 1-18, Dodd, "Belief and Certainty," and Dominik Kaus, "Credence as Doxastic Tendency," *Synthese* (2018), <https://doi.org/10.1007/s11229-018-01938-4>.) My argument is compatible with the Threshold (Lockean) View that outright believing a proposition p consists in assigning a sufficiently high credence, and with the (less popular) view that outright believing p requires assigning credence 1 to p .

correct.

- (P4) If DA_d is correct, then in cases of peer disagreement regarding p , it is rational for you to withhold judgment on whether there's a zero chance that your own judgment is correct.
- (C) Since it is not rational for you to withhold judgment on whether there is a zero chance that your own judgment is correct while remaining steadfast (as it is not doxastically possible for you that there is a zero chance that your judgment is correct, in such circumstances), DA_d is incorrect.

Let me offer a few additional remarks in defense of the argument. First, regarding P1, it might seem there are clear counterexamples to the claim that believing that p requires assigning credence 0 to the proposition that there's no chance that p . Suppose, for instance, that I carry out a formal logic proof and conclude on this basis that the argument in question is formally valid. I acknowledge that I might have made a mistake and even that the argument might not in fact be valid. It seems, then, that – assuming that this is a domain of necessary truths – I must acknowledge that there's a possibility that there is in fact *no chance* that the argument is valid. That is, it seems that this is a doxastic possibility for me – something my beliefs leave open. However, I think that as long as I outright believe that the argument in question *is* valid, my belief does in fact rule out the possibility that there's no chance of its being valid.²⁶ First, note that this claim is much weaker than the claim that I believe that p only if p is doxastically necessary for me. Since the latter claim is plausible,²⁷ we appear to have good reason to accept the weaker claim. In addition, I maintain that your having credence 1 in the proposition that there's some chance that p is compatible with acknowledgement of your fallibility on this matter. So your having credence 0 in the proposition that there's no chance that the argument is valid is also consistent with your acknowledgement of your fallibility regarding this proposition.

Perhaps the most obvious reason why someone might resist the idea of assigning credence 1 in such cases is that this is seen as incompatible with acknowledgement of fallibility. However, my acknowledgement of my own fallibility regarding p does not require that it is rationally permissible for me to outright believe that p while also believing that it's possible that there's no chance that p . For, as Rosenkranz²⁸ has convincingly argued, acknowledgement of fallibility

²⁶ Cf. Dodd ("Belief and Certainty," 4609): "[I]t is essential to representations that if a representation assertorically represents the world as being a certain way, it doesn't leave open the possibility that the world is otherwise."

²⁷ For a defense of this claim, see Dodd, "Belief and Certainty."

²⁸ Sven Rosenkranz, "Fallibility and Trust," *Noûs* 49, 3 (2015): 616-641.

regarding p can be understood in terms of a *lack* of relevant higher-order beliefs about one's immunity to error, together with a readiness to revise one's belief regarding p in light of future evidence; and such attitudes are perfectly consistent with assigning credence 1 to p .

Finally, on P3, the thought is that you have no relevant evidence which could warrant different attitudes toward your peer's judgment and your own. Hence, if it's permissible for you to withhold judgment about some proposition concerning the likelihood of your peer's judgment, then it must also be permissible for you to withhold judgment about a parallel proposition concerning the likelihood of your own judgment. The only possible exception to this, as I see it, would be that you are required to refrain from withholding in your own case simply because it is yours. But – whether or not this idea is tenable – it is not something that DA_d sanctions. (Recall that Decker explicitly denies that deep agnosticism involves giving such extra weight to your own judgment.)²⁹

I conclude that deep agnosticism about your epistemic peer's judgment is not a rational response to your disagreement, as long as you remain steadfast. Deep agnosticism commits you to withholding credence in the proposition that your own judgment has a non-zero chance of being correct – which you cannot do if you stick to your guns.

In the next section, I will turn to a different set of considerations which may seem to support deep agnosticism in cases of peer disagreement. In particular, deep

²⁹ Levy (“The Surprising Truth about Disagreement,” *Acta Analytica* (2020), <https://doi.org/10.1007/s12136-020-00437-x>) argues that different kinds of disagreement carry different kinds of information, and some (viz., “partisan” disagreements on which left and right seem intractably opposed, e.g., over the morality of abortion, or gun control) do not carry any information about likelihood of error; hence, they don't provide any pressure to conciliate. If this is right, then it may seem to undermine P3. While I do not have the space here discuss the details of Levy's account, or to provide an adequate response to it, I'll offer two brief responses. First, supposing Levy is right about this, there are still plenty of other cases of non-partisan disagreement which do seem to carry information about likelihood. Second, I think that provided partisan disagreements are understood as *genuine* disagreements between epistemic peers – i.e., as cases where at most one of the disputants is correct – we must see them as carrying some information about the likelihood of correctness of the judgments involved. On Levy's account, partisan disagreement provides reason for further investigation of the foundations from which one has reasoned en route to the judgment in question. He argues that such disagreements may raise concerns about irrelevant influences on our judgments, rather than concerns about likelihood of error. But I think these two concerns cannot be independent of one another in the way Levy's account requires. Concerns about irrelevant influences involve concerns about correctness of the associated judgments; hence information about likelihood of the former involves information about the likelihood of the latter.

Eric Gilbertson

agnosticism may seem to provide a solution to the multiple partitions problem, by explaining why the principle of indifference is false. I'll argue that one important lesson that emerges from this discussion is that deep agnosticism (with respect to relevant propositions) does not adequately reflect the strength of your evidence in these cases. While this alone does not show that such credal states aren't justified (arguably, this conclusion would be warranted only assuming the truth of evidentialism), I argue that the further assumption that credence should be guided by objective chance does provide grounds for rejecting *DA_d*.

4. Mystery Square and Multiple Partitions

The *principle of indifference* states that in cases where one's total evidence no more supports one proposition than another (using ' \approx ' to express this relation, and letting $P(\cdot)$ be a *rational* subject's credence function) rational credence in the propositions should be the same:

$$(POI) p \approx q \rightarrow P(p) = P(q).$$

The *multiple partitions problem* for POI is that, in cases where a proposition p is a member of two partitions of different size, the principle gives inconsistent answers to the question of what your credence in p should be. Take the following example.

Mystery Square. A mystery square is known only to be no more than *two* feet wide. Apart from this constraint, you have no relevant information concerning its dimensions. What is your credence that it is less than *one* foot wide?

It seems you have no more reason to suppose the square is less than 1 foot wide than that it's more than 1 foot wide. Now consider the *area* of the square. Do you have any more reason to suppose it's less than 1 square foot, than that it's between 1 and 2 square feet, or between 2 and 3, or 3 and 4? It seems you don't.

This is a problem for POI, because it seems to entail that we assign credence values evenly over the two partitions,³⁰ and this leads to contradiction: Since $A_1 \approx A_2 \approx A_3 \approx A_4$,³¹ we should assign equal credence value to each one, i.e. 1/4. And

³⁰ White ("Evidential Symmetry and Mushy Credence," in *Oxford Studies in Epistemology*, ed. T.S. Gendler and J. Hawthorne (Oxford University Press, 2010), 161-186) suggests a modified version of POI, which states that in cases of evidential symmetry we are to apportion our credence evenly, i.e.: (*POI**: If $\{p_1, p_2, \dots, p_n\}$ is a partition of your knowledge such that $p_1 \approx p_2 \approx \dots \approx p_n$, then for all i $P(p_i) = 1/n$). White takes *POI** to be a corollary of *POI*; Decker argues however that *POI** only follows from *POI* given a principle he labels "Full Spreading," which states that mutually exclusive and jointly exhaustive propositions p_1, p_2, \dots, p_n , one's credence function, C , should be such that $C(p_1 \vee p_2 \vee \dots \vee p_n) = 1$.

³¹ A_1 is the proposition, 1 sq. ft. < *area of mystery square* \leq 2 sq. ft., A_2 is 2 sq. ft. < *area of mystery*

since $L1 \approx L2$,³² we should assign $1/2$ to each of these. But $L2$ iff $A2 \vee A3 \vee A4$; so $L2 \approx (A2 \vee A3 \vee A4)$. White argues however that the problem is really not with POI, since the absurd conclusion follows from other assumptions without the aid of indifference principles. Specifically, he argues that the problem arises from principles of *transitivity* (If $p \approx q$ and $q \approx r$, then $p \approx r$) and *equivalence* (If p and q are known to be equivalent, then $p \approx q$).³³ Decker criticizes certain of White's assumptions and argues the principles of indifference really are the source of the problem.³⁴ I won't try to resolve this dispute here; instead I will focus on the question whether deep agnosticism avoids the problem, as Decker argues.

Decker argues that indifference principles, such as POI, are false, hence they cause no trouble for reasonable peer disagreement. (They would if they were true, because they would demand that you update your credences in such cases in a way that eventually brings you into agreement with their peers.³⁵) They are false, in his view, because they're incompatible with deep agnosticism.

For instance, Decker argues that in the case of Mystery Square you have good reason to believe $A1 \vee A2 \vee A3 \vee A4$ (your credence function is such that $C(A1 \vee A2 \vee A3 \vee A4) = 1$), even though you have *no* reason at all to believe any of the disjuncts (your credence function is simply undefined for each of these). Consequently, it's permissible for you to be deeply agnostic about $A1$, $A2$, etc. $C(A2 \vee A3 \vee A4)$ and $C(A2)$ are also undefined, in Decker's view. (You have *no reason at all*, according to Decker, to think that $A2 \vee A3 \vee A4$, or $A2$, is true.)

But if these credence functions are undefined, then two things follow: (i) it's not the case that $C(A2 \vee A3 \vee A4) > C(A2)$ and (ii) $C(A2 \vee A3 \vee A4)$ and $C(A2)$ are not definitely *non-zero*. Like White, I take (i) to be intuitively false, given $(A2 \vee A3 \vee A4)$ is logically weaker than $A2$.³⁶ If $C(A2 \vee A3 \vee A4)$ and $C(A2)$ are not definitely *non-zero*, then, if deep agnosticism with respect to these propositions is rational, this means it's rational to withhold judgment on the question of whether $A1$ ($A2$, etc.) is possible—i.e., on whether there's some objective chance of its being true. But it isn't rational to withhold judgment on that.³⁷

square ≤ 3 sq. ft., etc.

³² $L1$ is the proposition, $0 \text{ ft.} < \text{width of mystery square} \leq 1 \text{ ft.}$; $L2$, the proposition, $1 \text{ ft.} < \text{width of mystery square} \leq 2 \text{ ft.}$

³³ See White, "Evidential Symmetry and Mushy Credence," 165-6.

³⁴ Decker, "Disagreement, Evidence, and Agnosticism," 757-763.

³⁵ Decker, "Disagreement, Evidence, and Agnosticism," 756-7.

³⁶ Decker argues the asymmetry here is merely logical and not epistemic, and he argues (761) that the fact that $(A2 \vee A3 \vee A4)$ is a "safer bet" than $A2$ alone can be explained without positing any epistemic asymmetry. I'm not convinced, but for now I'll set this aside and focus on (ii).

³⁷ Decker compares the undefined credence function to the case of division by zero, noting that it

Now, as Decker correctly points out, the inference from *we must not assign a credence value of zero* to *we must assign a non-zero credence value* is a *non-sequitur*, since we could very well assign *no credence value whatsoever* to the proposition.³⁸ The argument I have just given does not rely on this inference, however. To see this, notice that the inference from *we need/must not assign a (precise) credence value* to *we need/must not assign any credence value at all* is also a *non-sequitur*. In particular, even though I may reasonably refuse to assign any precise credence value to A1, A2, A3, and A4 (in *this* sense, I may be agnostic), this alone does not entitle me to assign *no* non-zero value at all, even one that is highly indeterminate (or, to use Elga's term, "mushy"³⁹). In particular, if I know that there is some objective chance that A1 (A2, etc.) is true, then not only should I refrain from judging that it is definitely not true, I should acknowledge that there is a real

follows from its being undefined that it's not the case that 7/0 is equal to 0, 7, or Julius Caesar. ("Disagreement, Evidence, and Agnosticism," 766.) Clearly, it would be a mistake to infer from the fact that I ought not identify any of these as the value of 7/0, that I ought to assign some *other* value. Similarly, he suggests, it would be a mistake to infer simply from the fact that I ought not have a credence value of 0 or 1 for *p* that I ought to have some other credence value for *p*. But there is an important difference between these cases. In the case of division by zero, the fact that we ought not to assign 0 or 7 as the value of 7/0 is explained by, and grounded in, the fact that division by zero is undefined. In the case of Mystery Square, the fact that we ought not to assign 0 or 1 as credence values for A1 (etc.), is, it seems, explained by the fact that one's evidence provides some reason to think A1 is true (because one's evidence clearly supports the disjunction A1 \vee A2...), and also some reason to think it's false (because it provides some reason to think of each of A2-A4 that it is true). (The nature of the grounding relation is of course controversial, but I take it to be a metaphysical, explanatory relation, such that if a truth *p* is grounded in certain other truths, then these other truths explain (in a sense which is independent of our interests) why *p* is true. So, for example, the fact that Either Trump is human or he isn't, is grounded in the fact that he's human. And the fact that water has certain macroscopic properties is grounded in its chemical composition.)

I agree, then, that having a high credence value for the proposition that you shouldn't have a credence value of 0 or 1 for *p* does not by itself require that you have some credence value for *p*. (Decker, "Disagreement, Evidence, and Agnosticism," 766) But I claim it does require this in cases where high credence value for this higher-order proposition is explained by your evidence for *p*. Cases of disagreement with an acknowledged peer are, I claim, of this sort.

³⁸ "Disagreement, Evidence, and Agnosticism," 760.

³⁹ On the idea of indeterminate credence value, see e.g. Bas van Fraassen, *Laws and Symmetry* (Oxford: Oxford University Press, 1989), Peter Walley, *Statistical Reasoning with Imprecise Probabilities* (London: Chapman and Hall, 1991), James Joyce, "How Probabilities Reflect Evidence," *Philosophical Perspectives* 19 (2005), 153-78, and Scott Sturgeon, "Reason and the Grain of Belief," *Noûs* 32 (2008), 231-46.

possibility that it *is* true. The trouble with deep agnosticism is that it involves withholding judgment as to whether there is such an objective chance.

It is important to ask whether simply refraining from assigning a credence value of zero, without also assigning a positive value, adequately reflects the strength of your evidence in these cases. As we've just seen, it does not. In the horse race case, for instance, if you recognize your peer as an epistemic peer, then you must at least recognize the possibility that she got it right, whereas you got it wrong. So you must give some non-zero credence to the proposition that she got it right. It's not acceptable, then, to give no credence at all to this proposition.⁴⁰

It seems appropriate at this point to raise the question, under what conditions in general is deep agnosticism about a proposition p supposed to be rational? On this matter, Decker has only a couple of brief suggestions. First, he suggests that DA is warranted for propositions that one is unable to represent in thought (e.g., a child may for this reason be deeply agnostic about Gödel's incompleteness theorem). This is clearly irrelevant to the cases of peer disagreement we are concerned with, however. There's no question of your being able to grasp the relevant thoughts concerning your peer's judgment and its comparative likelihood of being correct. Decker also suggests that deep agnosticism may be warranted in some cases of propositions one hasn't yet considered. But he acknowledges that there are of course many propositions one hasn't considered toward which one nevertheless has implicit attitudes. So lack of explicit belief doesn't imply lack of credence value in a proposition. Moreover, the sorts of unconsidered propositions Decker has in mind are ones for which one seems to have no relevant evidence one way or the other (e.g., the proposition that the current U.S. President is at this very moment in the Oval Office). Again, the cases of peer disagreement that concern us are not like this. In these cases, you have evidence that bears on relevant propositions regarding your peer's judgment; it's not as though the information available to you is simply silent on the matter.

Moreover, even if we suppose that in general failure to consider a proposition is sufficient grounds for deep agnosticism (a position that bears some similarity to certain contextualist views of knowledge), this would not provide an adequate

⁴⁰ I think there is a danger here of confusing probabilities and credences. As White notes ("Evidential Symmetry," 162-3), indifference principles should be understood as normative constraints on what your credence may be, not as principles for determining what the objective probabilities or chances are. "You can't get probabilities out of ignorance"—this is granted. Principles of indifference don't say otherwise. So, from the fact that you may have no grounds at all for saying what the objective probability of, e.g., Mystery Square's having 1m sides is, this does not mean that your *credence* in the proposition should be undefined, or even that it should be indeterminate. Agnosticism about these objective probabilities doesn't entail deep agnosticism.

Eric Gilbertson

defense of the steadfast position. This is because we are clearly interested not only in cases of peer disagreement where the disputants do not consider relevant propositions about their peers, but also cases where they do explicitly consider such propositions. That is, we want to know how you should respond to disagreement with an acknowledged peer, provided you clearly understand all of the relevant facts concerning their peerhood and their handling of the evidence at hand.

The upshot is that although we may have some examples in which deep agnosticism is plausible, this does not support DA_d. The features of the former which seem to explain the rationality of deep agnosticism are absent in the case of peer disagreement.

Rational Credence and the Principal Principle

I have argued that deep agnosticism with respect to A1-A4 and L1-L2 in the case of Mystery Square does not adequately reflect the strength of your evidence. I suggest that we may find further support for this claim in Lewis's Principal Principle:⁴¹

$$(PP) C(A \mid ch(A) = x) = x.$$

PP states that a rational agent conforms her credences to known objective chances. It is based on the idea that we should use frequency data to guide our subjective probabilities. For example, if you know that the coin that you're about to toss is a fair one and so has a 50% chance of landing heads, then you should give equal credence to its landing heads and its landing tails. My argument rests on a closely related idea, namely, that rational credence is constrained by objective chance, in the following way: if you know that there's *some* objective chance that p , then you should assign some (perhaps indeterminate) positive credence value to p . (Note that an even stronger principle seems warranted, namely that if, *for all you know*, there's some objective chance that p , then you should assign some positive credence value to p . However, I won't assume that this principle holds.) While this principle is not strictly entailed by PP, I think it is clearly suggested by the underlying normative principle that we ought to be guided by objective chance in forming credences.

If you acknowledge someone as an epistemic peer with respect to some issue about which the two of you disagree, you thereby acknowledge that there is some objective chance that s/he got it right (and thus that you are mistaken). You also acknowledge that there is some objective chance that s/he handled the available evidence better than you did. So, given the principle suggested, you are rationally required to assign some positive credence value to these propositions.

⁴¹ David Lewis, "A Subjectivist's Guide to Objective Chance," in *Philosophical Papers, Vol. II* (Oxford: Oxford University Press, 1986).

It might be objected that recognition of epistemic peerhood (and in particular, of the possibility that your peer's judgment is correct) does not require taking a stand on such objective chances, even supposing that credence should be guided by objective chance in the way I've suggested. After all, it may seem that if you have *no reason* at all to think that your evidence (or your handling of it) is superior (inferior) to your peer's, then you should be deeply agnostic about the *likelihood* of your peer's judgment being correct. To see why this is a mistake, let's consider one more example due to Feldman:⁴²

[There are] two suspects in a criminal case, Lefty and Righty. Suppose now that there are two detectives investigating the case, one who has the evidence [against] Lefty and one who has the evidence incriminating Righty. They each justifiably believe in their man's guilt. And then each finds out that the other detective has evidence incriminating the other suspect. If things are on a par, then suspension of judgment is called for. If one detective has no reason at all to think that the other's evidence is inferior to hers, yet she continues to believe that Lefty is guilty, she would be unjustified. She is giving special status to her own evidence with no reason to do so, and this is an epistemic error, a failure to treat like cases alike.

Feldman is claiming that we have reason to accept the equal weight view, as we have no reason to give greater weight to our own evidence in such cases.⁴³ Decker argues however that it is a mistake to think that if one of the detectives sticks to her judgment she is indeed giving special weight to it:

This is a *non sequitur*. If she has *no reason at all* to think that her evidence is inferior to the other's, it would be acceptable for her to be deeply agnostic on the matter of who has better evidence. The situation is better described as her simply continuing to give her evidence status *as* evidence. She needn't be giving it privileged status in the sense of assigning a higher credence value to the proposition that her evidence is superior to the other detective's than to the proposition that the other detective's evidence is at least as good as hers. *She need not be placing any credence at all on those propositions....*⁴⁴ (my emphasis)

So Decker's claim is that A has *no reason* to think her evidence is inferior (or superior) to B's, then it is acceptable—perhaps even required—for A to be deeply agnostic on whether *A's evidence is at least as good as B's*. But *this* is a *non sequitur!*

⁴² Richard Feldman, "Reasonable Religious Disagreements," in *Philosophers Without Gods*, ed. L. Antony (Oxford: Oxford University Press, 2007), 208.

⁴³ It should be noted that although Feldman is explicitly concerned only with the quality of evidence in this case, I take him to be supposing that neither detective has a reason to think that he *handled the evidence* any better than the other. (I take this to be implicit in his statement that "things are on a par.")

⁴⁴ "Disagreement, Evidence, and Agnosticism," 770.

For A's having no such reason is clearly compatible with A's having some reason to think that B's evidence is *equally* good as A's. This, together with the argument given in §3, implies that A's having no reason to think her evidence is inferior to B's is compatible then with A's having *some* reason to think that B's, rather than A's, *judgment* is correct. Moreover, we should note that, supposing A has no reason to think his evidence is *inferior* to B's, this is compatible with his having some reason to think it is at least possible that B's evidence (and, moreover, B's handling of that evidence) is superior, i.e., that there is some objective chance this is in fact the case. And A does have *some* reason to think all of these things *are* indeed the case, provided she recognizes B as an epistemic peer. So, A should give some positive (and perhaps indeterminate) credence to these propositions.

We may grant Decker's claim that A need not assign higher credence to the proposition that her evidence is superior to B's than to the proposition that it is inferior. We may also grant that she need not assign higher credence to the proposition that her evidence is *equal* to B's than to the proposition that they are unequal (although failure to assign such credence does not seem compatible with a clear recognition of epistemic peerhood). It doesn't follow from any of this that it's rational for her to assign *no credence at all* to any of these propositions. In particular, if A recognizes B *as an epistemic peer*, then she should give high credence to the proposition that B's evidence is at least as good as hers—or, at the very least, she should assign a *non-zero* value to this proposition.⁴⁵

I have argued that it is irrational to be deeply agnostic about relevant peer propositions in Feldman's case. That is, it is rationally required that A assign some positive credence to the proposition that B's judgment is correct, and to the proposition that B's evidence is superior to A's, provided A takes B to be an epistemic peer. Does this imply that it would be irrational for A to stick to his original judgment in these circumstances? I don't claim that it does. And I will not take a stand on the issue of conciliationism vs. steadfastness here. My claim is just that it would be irrational for A to be deeply agnostic on such propositions *while at the same time regarding B as an epistemic peer*. Whether A must *conciliate* provided A

⁴⁵ Decker is clear that he only insists that it is rationally *permissible* to refrain from assigning a non-zero value in this case, not that this is rationally *required*. ("Disagreement, Evidence, and Agnosticism," fn. 23.) Is this satisfactory? Arguably, it is not. Provided you have good reason to take someone as an epistemic peer, this gives you a conclusive reason to assign a non-zero value to the proposition that their evidence is as good as yours. That is, it *rationally requires* you to accept that this is a genuine possibility. Thus you cannot simply withhold credence in this proposition. Compare Mystery Square: Since you assign a value of 1 to $L1 \vee L2 \vee L3 \vee L4$, and you know that there is some chance that $L1$ ($L2$, etc.), rather than $L2 \vee L3 \vee L4$, you shouldn't simply refrain from assigning any value to $L1$. Rather, you should assign a positive value to $L1$.

assigns positive credence to these propositions, or whether there is some basis for steadfastness other than DA_d , is a question I will leave open.

5. Objection and Reply

It might be objected that my argument depends on a conflation of, or perhaps an equivocation between, two different understandings of epistemic peerhood, *rationality peerhood* and *accuracy peerhood*. These notions may be defined as follows:⁴⁶

X and Y are *rationality peers* with respect to a certain issue iff X and Y's opinions on that issue are equally likely (given relevant higher-order evidence) to be rational (i.e. X and Y are equally likely to have reasoned correctly about the disputed matter).

X and Y are *accuracy peers* with respect to a certain issue iff X and Y's opinion on that issue are equally likely (given relevant higher-order evidence) to be accurate.

The worry is that because my criticisms of deep agnosticism seem to depend on the assumption that epistemic peers are *rationality peers*, they also depend on the (controversial) assumption of uniqueness, i.e.:

UNIQUENESS: Given one's total evidence, there's a unique rational doxastic attitude that one can take to a proposition.

This is because I suggest (in §3) that any reason to think that your peer has reasoned correctly on the basis of the available evidence is a reason to doubt that *you* have done so. However, if uniqueness is false, then this is not so. For in that case, from the fact that your peer's judgment is *rational*, it does not follow that yours is not also (fully) rational. And if it is false that any reason to think your peer has reasoned correctly is a reason to doubt that you have done so, then disagreement with your peer doesn't rationally require you to conciliate or to reduce your confidence in your judgment.

I have two replies to this objection. First, the notion of epistemic peerhood as I understand it is not to be identified with *either* the notion of *rationality peerhood* or the notion of *accuracy peerhood*. It is distinct from both, and consequently, my argument cannot be said to conflate these, or to equivocate between them. This is because epistemic peerhood as I understand it implies equality (i.e., equal quality) of relevant first-order evidence, and equality of evidence in my view is not implied by either accuracy peerhood or rationality peerhood. For instance, suppose that I know that my friend is in "Fake Barn Country." In that case, I shouldn't regard him as an

⁴⁶ Cf. Christensen, "Conciliation, Uniqueness, and Rational Toxicity."

accuracy peer, though I may still acknowledge that he's a rationality peer (who is equally likely to have reasoned correctly on the basis of the evidence available to him). Should I take him to be an epistemic peer, then? While I acknowledge that some evidential internalists will probably say 'yes,' I maintain that my peer's evidence is inferior to mine. While we both use the same process of forming perceptual beliefs, we base our respective perceptual judgments on different perceived facts: his judgment, *that's a barn*, is based on his perception of a barn facade, whereas mine is based on my perception of a (real) barn.⁴⁷ So, although we count as rationality peers, we do not count as epistemic peers. For X to be an epistemic peer of Y requires not just that X be equally likely to have reasoned correctly, but also that X have equally good evidence as Y. Accuracy peerhood is also insufficient for epistemic peerhood, since X and Y may be equally likely to have determined the truth about some subject matter on the basis of unequal or incommensurable evidence.

Because rationality peerhood is not sufficient for epistemic peerhood, even supposing disagreement with rational peers doesn't require conciliation it doesn't follow that disagreement with *epistemic* peers doesn't require conciliation. (If rational peerhood is compatible with permissiveness, it doesn't follow that epistemic peerhood is, too.)

My second response to the objection is more significant. It is that my case against deep agnosticism does not depend on uniqueness even if epistemic peerhood is understood as *rationality* peerhood. This is because I do not need to maintain that in the cases I've considered, any reason for being confident in your peer's judgment is a reason for doubting your own. All I need is the assumption that in these cases regarding the other person as an epistemic peer is clearly warranted and that this requires assigning positive credence value to the proposition that their judgment is correct. The problem with being deeply agnostic in these cases is that it requires not only that you assign no value to the proposition that your peer has reasoned as well as you, but that you assign no value to the proposition that their judgment is correct.⁴⁸ While it is hard to see how you could do that while remaining steadfast, my argument does not depend on either conciliationism or the equal weight view.⁴⁹

⁴⁷ I don't claim that my peer's judgment isn't *justified* on the basis of the evidence he has. But I do claim that the quality of his evidence is less than that of mine.

⁴⁸ The permissivist worry that I raised earlier (fn. 19) may arise again here. At this point, I have nothing further to say in response and will simply emphasize how modest the claim in question really is: what is ruled out is that you are rationally permitted to assign no credence value at all, hence what's required is only that you assign *some* value, not that you assign a value that is sufficient for belief, or for belief that it is more likely than not to be true.

⁴⁹ It's also worth noting that the motivations for conciliationism and equal weight do not clearly

Granted, I do claim that in Elga's horse race case, for example, confidence in your peer's judgment implies (some degree of) doubt about your own. But here I am concerned only with the correctness of the conflicting judgments, only one of which *can* be correct. The point is that although your recognition of the other as an epistemic peer and of the fact that she's reached a different judgment of who won the race need not diminish your confidence that you've handled the evidence just as well as she has, it *should* diminish your confidence that your judgment of who won is *correct*. In this respect, at least, peer disagreement requires conciliation. Whether or not disagreement in the horse race case casts doubt on the *rationality* of your judgment, it casts doubt on its *accuracy*.⁵⁰

It might be thought that my notion of epistemic peerhood may be understood simply as *accuracy + rationality* peerhood (call this *AR peerhood*). This may seem to be suggested by my talk of epistemic peers being equally likely to "get it right," as well as equally likely to have handled the evidence correctly. But, while I am open to the possibility that AR peerhood is at least extensionally equivalent to my notion of epistemic peerhood—i.e., X is an epistemic peer of Y iff X is an AR peer of Y—I do not think that I am committed to this, and I do not think that AR peerhood and epistemic peerhood are *a priori* equivalent. For epistemic peerhood is, as I've explained, to be understood in terms of *evidence*, and the relationships between evidence and truth, and between evidence and rationality, are a highly contested matter. So I will refrain from claiming that epistemic peerhood is or is not extensionally equivalent to AR peerhood.

To sum up, deep agnosticism is consistent with the plausible view that it's rationally permissible to stick to your judgment in cases of peer disagreement (thus, to be more confident in your judgment than your peer's). Moreover, it implies (again, not without plausibility) that in doing so you need not take a stand—that is, you may remain deeply agnostic—on the comparative likelihood of your judgment's being true versus your peer's. But deep agnosticism about such propositions is

depend on uniqueness. Authors who argue that they do not include Christensen, "Conciliation, Uniqueness, and Rational Toxicity," Elga, "Reflection and Disagreement," Kornblith, *Disagreement* (Oxford: Oxford University Press, 2010), Frances, "The Reflective Epistemic Renegade," *Philosophy and Phenomenological Research* 81, 2 (2010): 419-63, and Lee, "Conciliationism Without Uniqueness," *Grazer Philosophische Studien* 88, 1 (2013): 161-88. Thus even if my argument against deep agnosticism were to rely on the assumption that either equal weight or conciliationism is correct, this would arguably not commit me to uniqueness.

⁵⁰ Cf. Christensen, "Conciliation, Uniqueness, and Rational Toxicity." Christensen argues that the pressure to conciliate generated by accuracy-based evaluation of peers is less threatened by permissive accounts of rationality than the pressure to conciliate generated by rationality-based evaluation.

Eric Gilbertson

incompatible with a recognition of your peer *as an epistemic peer*, in such circumstances. This recognition requires you to take a stand—i.e., to assign some credence value—on the question of comparative likelihood of correctness. Moreover, although deep agnosticism may avoid the multiple partitions problem in cases like Mystery Square, the way in which it does this is unsatisfactory, since it is not consistent with an adequate view of the strength of your evidence in such cases.