

Conciliation and Peer-Demotion in the Epistemology of Disagreement

Juan Comesaña

1 Introduction

What should your reaction be when you find out that someone that you consider an “epistemic peer” disagrees with you? Two broad approaches to this question have gained support from different philosophers. Precise characterizations of these approaches will be given later, but consider for now the following approximations. First, there is the “conciliatory” approach, according to which the right reaction to a disagreement is to move one’s opinion towards that of one’s peer, in proportion to the degree of trust that one accords to that peer—for instance, if you thought that, in case of disagreement, you are equally likely to be right, then the conciliatory approach would have it that you should meet your epistemic peer halfway.¹ The other, “non-conciliatory” approach, holds that one’s reaction to a disagreement need not be perfectly in line with one’s prior degree of trust in the other party. Notice an important asymmetry between these two approaches: the conciliatory approach

has it that conciliation is the right reaction to *any* disagreement, whereas the non-conciliatory approach has it that there are *some* possible disagreements the correct reaction to which is not to conciliate (or not to the extent mandated by conciliatory views).² This paper examines the dispute between conciliatory and non-conciliatory views by distinguishing and examining possible answers to four different questions. As a first approximation, the questions are the following:³

(I) The Metaphysical Question: *What is it* for a subject *S* to have a certain level of epistemic competence regarding proposition *p* and evidence *e* relative to subject *S**?

(II) The Psychological Question: What is it for a subject *S* to *consider* someone else as having a certain level of epistemic competence regarding proposition *p* and evidence *e* relative to *S*?

(III) The Epistemic Question: Under what conditions is a subject *S* *justified* in considering someone else as having a certain level of epistemic competence regarding proposition *p* and evidence *e* relative to *S*?

(IV) The Main Question: Suppose that *S* is justified in considering another subject *S** as having a certain level of epistemic competence regarding proposition *p* and evidence *e* relative to *S*. Now suppose that *S* finds out that there is a disagreement between *S* and *S** about *p*. *How should that affect S's opinions about p?*⁴

2 The Epistemic Virtues Answer to the Metaphysical and Psychological Questions and a Non-Conciliatory View About the Main Question

Some philosophers have thought that two subjects are epistemic peers just in case, roughly, they are approximately equal when it comes to general epistemic virtues such as intelligence, thoughtfulness, freedom from bias regarding p , etc. This answer to the Metaphysical Question naturally suggests a corresponding answer to the Psychological Question: to consider somebody one's epistemic peer is a matter of *having opinions about that someone*—in particular, opinions about how he fares relative to oneself with respect to those general epistemic virtues.⁵ The following answers to our first two questions can be extracted from this position (the answers are concerned with epistemic peerhood, but epistemic superiority and inferiority can be dealt with in the obvious way):

The epistemic virtues answer to the Metaphysical Question: S and S^* are epistemic peers relative to a proposition p and evidence e if and only if S and S^* both have e and they are roughly equal regarding general epistemic virtues such as intelligence, thoughtfulness, and freedom from bias with respect to p and e .

The epistemic virtues answer to the Psychological Question: S considers S^* his peer regarding proposition p and evidence e if and only if S believes that both S and S^* have e and also believes that S and S^* are roughly equal regarding general epistemic virtues such as intelligence, thoughtfulness, and freedom from bias with respect to p and e .

Suppose that the epistemic virtues answer to the Psychological Question is correct. How does that answer bear on the Main Question? Consider a particular case. Suppose that Tom takes Adam to be his epistemic peer regarding math, in the sense that Tom thinks that he and Adam are approximately equally endowed with intelligence, thoughtfulness, freedom from bias, etc., when it comes to mathematical propositions. Now, on a particular occasion Tom finds out that he and Adam disagree regarding some math question—say, about what is 376 times 215. We are now in a position to ask an instance of the Main Question applied to this case: how should finding out about the disagreement affect Tom’s confidence that his answer is the correct one? It should be obvious that the answer is underdetermined by the information given so far. The mere fact that Tom considers that Adam is just as good as he is regarding all relevant epistemic virtues when it comes to math (or even when it comes to a particular mathematical proposition) doesn’t settle how Tom should react to news that he and Adam have different opinions about what is 376 times 215. What *would* settle that question? Well, information about the specific kind of disagreement is obviously relevant. Suppose, for instance, that Tom’s answer is 80840 and Adam’s is 80480. In that case, it seems plausible to think that Tom should considerably lower his confidence in his answer—he should, for instance, re-check his calculation. But suppose instead that Adam’s answer is 23. In that case, it seems plausible to think that there is no reason for Tom to become less confident in his answer—rather, there is ample reason for him to think that Adam isn’t paying attention, or that he misunderstood the question, or that he is terribly drunk, etc.⁶

So the fact that Tom thinks that Adam is his peer in the epistemic virtues sense *underdetermines* how Tom should react to a disagreement with Adam. How should Tom react to this disagreement depends, in addition, on what the *rest* of his evidence

is like. Therefore, given the epistemic virtues answer to the Psychological Question, a non-conciliatory answer to the Main Question seems obvious: sometimes, at least, you don't need to conciliate with somebody who you justifiably take to be your peer. Whether you do or do not will depend on what the rest of your evidence is like.

A non-conciliatory answer to the Main Question, then, is a very attractive part of the total package comprised of the epistemic virtues answer to the Psychological Question and a non-conciliatory answer to the Main Question. But the next section will introduce reasons to prefer an alternative answer to the Psychological Question. Once that alternative is in place, moreover, a conciliatory answer to the Main Question becomes natural: namely, the proportional weight view, which is presented in section 4.

3 The Dispositionalist Answer to the Psychological Question—and Goodbye to the Metaphysical Question

The problem with the epistemic virtues answer to the Psychological Question has two sources: first, it is insufficiently relativized; second, it focuses on the wrong kind of phenomena to characterize what it is to consider someone a peer. Take this last concern first. Recall that the epistemic virtues answer to the Psychological Question has it that our taking somebody else to be our epistemic peer is a matter of *our having opinions about that other subject*. For good or bad reasons, however, your explicit opinions about somebody else's competence may not be in line with how you change your mind when confronted with the fact that the subject in question disagrees with you. Take the example from the preceding section: even though Tom thinks that Adam is roughly as good as he is regarding simple arithmetical calculations, his

opinion about the result of a simple calculation is not changed once he learns Adam's preposterous opinion about it. This suggests that Tom doesn't really consider Adam his peer (at least not *with respect to this particular disagreement*, a point to be reconsidered later). In general, there are reasons to accept a broadly functionalist account of belief because somebody's behavior may indeed speak louder than words when it comes to revealing their deep-seated opinions—for instance, dispositions to racist behavior may betray racist beliefs despite explicit (and even sincere) protestations to the contrary. The alternative answer to the Psychological Question defended here extends this idea to the area of relative epistemic competence—and, for this reason, it will be called the “dispositionalist” answer to the Psychological Question. The dispositionalist view has it that whether we take somebody else as a peer regarding a certain proposition is revealed in our *dispositions* to change opinions about that *proposition*, rather than in our *beliefs* about *the other subject*.

To clarify, the dispositionalist answer to the Psychological view is not the claim that a belief that so-and-so is my epistemic peer with respect to proposition *p* just is, or even can be reduced to, a cluster of dispositions to modify my opinions regarding *p* given disagreements with so-and-so. Rather, it is the claim that what it is to consider somebody else a peer should not be understood in terms of our explicit *beliefs* about her, but rather in terms of our dispositions to change our opinions when a disagreement with her arises. The belief and the dispositions are not, of course, unrelated, but no direct reduction is being advocated here.⁷

Relatedly, and unlike the epistemic virtues account, the dispositionalist answer to the Psychological Question does not suggest an answer to the Metaphysical Question. Indeed, if what it is to consider somebody your epistemic peer is determined by your dispositions to change credences, then it is not at all clear that it makes sense to ask

what it is for somebody to *be* your peer. The epistemic virtues conception is naturally understood as first giving an answer to the Metaphysical Question (two subjects are epistemic peers of each other just in case they are roughly equal regarding general epistemic virtues) and then offer an answer to the Psychological Question based on the answer to the Metaphysical Question (a subject considers another one his peer just in case he believes that they are roughly equal regarding general epistemic virtues). The dispositionalist account, by contrast, cannot be understood this way. It is directly and primarily an answer to the Psychological Question, and it is doubtful that an answer to the Metaphysical Question can even be built upon it.⁸

The other root of the problem with the epistemic virtues answer to the Psychological Question is that it is insufficiently relativized. It is true that, just as the questions are relativized to a specific proposition p and evidence e , so too are the corresponding epistemic virtues answers. But the pair of cases considered in the previous section illustrates that these relativizations do not go far enough. For it is tempting to describe the situation as follows: Tom considers Adam his peer with respect to some disagreements, but not with respect to others. Before presenting the dispositionalist answers, then, we should pause to reformulate the questions so as to take into account all the necessary relativizations. So far we have a vague idea of what the needed relativizations are: it is clearly relevant, for instance, whether the disagreeing parties are drunk or sober, attentive or bored, etc. I shall leave a more precise characterization of exactly what to include in the circumstances of the disagreement for later, but we can take it for granted that they include the fact that the parties share their evidence. The reformulated questions, therefore, look like this:

- (I) The Metaphysical Question: *What is it* for a subject S to have a certain level of epistemic competence regarding proposition p and cir-

cumstances C relative to subject S^* ?

(II) The Psychological Question: What is it for a subject S to *consider* someone else as having a certain level of epistemic competence regarding proposition p and circumstances C relative to S ?

(III) The Epistemic Question: Under what conditions is a subject S *justified* in considering someone else as having a certain level of epistemic competence regarding proposition p circumstances C relative to S ?

(IV) The Main Question: Suppose that S is justified in considering another subject S^* as having a certain level of epistemic competence regarding proposition p and circumstances C relative to S . Now suppose that S finds out that there is a disagreement between S and S^* about p under circumstances C . *How should that affect S 's opinions about p ?*

The dispositionalist answer to these questions assumes a fine-grained account of the doxastic attitudes according to which subjects can invest several different degrees of confidence (“credences”) in propositions, rather than a coarse-grained account according to which there are three (or two, depending on how you count) attitudes: belief, disbelief and suspension of judgment. Within that framework, you and another subject count as disagreeing about p just in case your credence in p is some number m and the other subject’s credence in p is some *different* number n . Consider now the class of all possible disagreements about a specific proposition p between you and a certain subject S^* . This class will include cases where your degree of belief in p is .9 and S^* ’s is .1, cases where these credences are reversed, etc. Think now about all of your dispositions to assign degrees of confidence in p given all the possible

disagreements between you and S^* . The following is then the dispositionalist answer to the Psychological Question:

Dispositionalist answer to the Psychological Question: S considers S^* his peer with respect to proposition p and evidence e under circumstances C if and only if, when faced with a disagreement with S^* regarding proposition p under circumstances C , S is disposed to revise his credence in p so that his new credence is the mean of S 's and S^* 's pre-disagreement credences.⁹

A noteworthy feature of this answer to the Psychological Question is that you may consider a subject your peer regarding one circumstance of disagreement but not another. Remember that we are postponing the question of what exactly should count as part of the circumstances C (that question is examined in detail in sections 5 and 6), but none of the issues to be discussed later are prejudiced by noting that the other subject's credence in p and other propositions are allowed to be part of C . That has the consequence that you may consider a subject a peer with respect to p and a certain range of credences over p , but not with respect to other ranges of credences. So, for instance, it may well be that although Tom considers Adam his peer regarding the disagreement that arises when Tom thinks that $376 \times 215 = 80840$ and Adam thinks that the result is 80480, Tom doesn't consider Adam his peer regarding the disagreement that arises when Adam thinks that the result is 23, or when Tom can see that Adam is very hungry (for Tom may know that when Adam is hungry his generally reliable judgment is affected). Notice also that C can include, for instance, information regarding how the subject reacted to your disagreement, or even about the fact that a disagreement has taken place. There is no specific restriction governing how

the dispositions that characterize your judgments of relative epistemic competence should evolve over time.¹⁰

4 The Proportional Weight View

The Main Question can now be reconsidered in light of the dispositionalist answer to the Psychological Question. Consider a disagreement between you and a peer regarding a proposition p . Your credence in p is 0.9 and your friend's 0.1. That you consider your friend your peer about p means (according to the dispositionalist answer to the Psychological Question) that you have a bunch of dispositions to modify your credence in p given different possible disagreements between you and your friend. In particular, it means that you have the disposition to assign a credence of 0.5 to p given that your pre-disagreement credence is 0.9 and your friend's is 0.1 (and given that the disagreement takes place under certain circumstances). Now, given that (as is assumed in asking the Main Question) you are *justified* in having that disposition, how should you change your opinion about p once you find out that her degree of belief in p is 0.1 (under the circumstances mentioned)? When formulated this way, a natural answer seems to be that your credence should be 0.5—that is, your credence in p after finding out about a certain disagreement should match the credence that you were justifiedly disposed to have given that disagreement. This is precisely what the proportional weight view claims. More generally, friends of the proportional weight view put forward the following principle:

Proportional weight answer to the Main Question: Subjects must react to a disagreement with respect to a proposition p in circumstances C by modifying their credences in p in accordance to their (justified)

pre-disagreement dispositions.

What can be said in favor of the proportional weight view? At least this much is true: failing to abide by its prescription constitutes a failure of coherence, because if you violate it you fail to form a belief that you have a justified disposition to form. More specifically, the following principle seems to enjoy considerable plausibility and to offer a reason for believing in the proportional weight view:

Dispositional justification: if S is justified in being disposed to assign credence n to proposition p under circumstances C and S is in circumstances C , then S is justified in assigning credence n to p .¹¹

Notice that the principle of dispositional justification is plausible only if the circumstances C are assumed to satisfy the maximality condition that there are no other epistemically relevant circumstances C^* such that C is a proper part of C^* and S is in C^* . Otherwise, S may well be justified in being disposed to assign credence n to p under circumstances C but not justified in assigning credence n to p under circumstances C^* —if C^* contains evidence that is relevant to S 's assessment of p . Moreover, even if it is assumed that the circumstances in the principle of dispositional justification satisfy this maximality condition, that principle can be used to justify the proportional weight view only if it is also assumed that the circumstances in the formulation of that view also satisfy that condition. In other words, to have a good argument from the principle of dispositional justification to the proportional weight view it must be assumed that the circumstances C mentioned in the formulation of the proportional weight view include everything epistemologically relevant that is in S 's possession at the time of disagreement.

So, *should* it be assumed that C includes all the evidence available to S at the time of disagreement? There are in the literature two views about this. According to some philosophers (who are usually associated with the proportional weight view and in general with conciliatory views) part of the evidence that S has must be ignored or “bracketed” when considering what to include in C . According to other philosophers (usually associated with non-conciliatory views), this would lead to a violation of an extremely plausible principle of total evidence. The next two sections address this controversy.

5 Two Versions of the Proportional Weight View

The question, then, is: do the circumstances C in the dispositionalist answer to the Psychological Question and the proportional weight answer to the Main Question include all the evidence that S has to go on in making up his mind about p ? Of particular interest is the question whether the reasoning that led the parties to their respective views about p should be included in C or not. Notice that this is a version of the Epistemic Question, for what is being asked is whether the pre-disagreement reasoning can be appealed to when determining peerhood—that is, what is being asked is whether a difference in the pre-disagreement reasoning can *justify* a difference in the dispositions that determine peerhood.

Some philosophers sometimes talk as if the original reasoning should be *ignored* or *bracketed* when determining whether you are justified in counting somebody else as a peer or not. Here, for instance, is Elga:¹²

Note that the circumstances of a disagreement should not include a detailed specification of the chain of reasoning that led you to your con-

clusion. For if they did, then making the relevant conditional probability judgment would involve thinking through the disputed issue—and hence would not be prior to your doing so. (...)

In general, circumstances of disagreement should be individuated just coarsely enough so that the relevant conditional probability judgment is genuinely prior to your reasoning about the disputed issue. (This coarseness constraint is what makes the equal weight view nontrivial. *For otherwise—if the view simply required that one’s new opinion should equal one’s prior opinion, conditional on all of one’s new information—the view would be tantamount to the requirement that one conditionalize on one’s new information.*)¹³

In a similar vein, Christensen (2007) argues that “I should assess explanations for the disagreement in a way that’s independent of my reasoning on the matter under dispute” (p. 16). And, more recently, he defends the following principle:

Independence: In evaluating the epistemic credentials of another’s expressed belief about P, in order to determine how (or whether) to modify my own belief about P, I should do so in a way that doesn’t rely on the reasoning behind my initial belief about P.¹⁴

For some other philosophers, on the other hand, the nature of the reasoning in question is crucial. They believe that your degree of belief in p should take into account not only the disagreement and the circumstances under which it takes place, but also the reasoning that led from the evidence to the disagreement. Here is Kelly:

[I]f you and I have arrived at our opinions in response to a substantial body of evidence, and your opinion is a reasonable response to that evidence while mine is not, then you are not required to give equal weight to my opinion and to your own.¹⁵

[W]hat it is reasonable to believe 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 believe as they do.¹⁶

I propose that we interpret the disagreement between Elga and Christensen, on the one hand, and Kelly on the other as a disagreement over the answer to the Epistemic Question.¹⁷ Under this interpretation, Kelly would hold, while Elga and Christensen would reject, the following thesis:

Relevance of original reasoning: The circumstances *C* relative to which a subject *S* is justified in considering somebody else *S'* as having a certain level of epistemic competence regarding proposition *p* may include the original reasoning that led *S* and *S'* to their respective opinions regarding *p*.

One clarification regarding the principle is needed. “Reasoning” will not always be the best word for what I mean to be referring here—namely, that thing that philosophers like Elga and Christensen think should be “bracketed off” when determining peerhood or when determining how to react with a disagreement with a peer. Unfortunately, there is no perfect terminological choice. One option (sometimes used in the literature) is “pre-disagreement evidence,” but this is also not a good name, for Elga and Christensen would of course think that it is perfectly legitimate to appeal

to *vast amounts* of the evidence amassed before the disagreement took place (for instance, one's general understanding of the world, or of how people think, etc.) in determining how to react to that disagreement or even who to count as a peer. In addition, Elga and Christensen often seem to want to make a distinction between the evidence and what one does with the evidence (as shown in the quotes from them just above), and argue that it is this second thing that should be bracketed off. In any case, I will argue in the next section that one should *always* take into (proper) account *all* of one's evidence, and so that, whatever exactly Elga and Christensen mean to bracket off, they are wrong.¹⁸

As a result of their different attitudes towards that principle, Elga and Christensen also disagree with Kelly over the proper interpretation of the proportional weight view. Although all three of them would agree with the letter of the proportional weight view (one should react to a disagreement under circumstances *C* in accordance with one's justified pre-disagreement dispositions to so react under those circumstances), they would mean different things by it because they include different things under the circumstances of disagreement. Expressed in a way which brings this difference to the fore, Elga and Christensen would accept the following different versions of the proportional weight view:

Proportional weight - EC version: Subjects must react to a disagreement with respect to a proposition *p* in circumstances *C* by modifying their credences in *p* in accordance to their (justified) pre-disagreement dispositions—*where C doesn't include the original reasoning.*

Proportional weight - K version: Subjects must react to a disagreement with respect to a proposition *p* in circumstances *C* by modifying

their credences in p in accordance to their (justified) pre-disagreement dispositions—where C includes the original reasoning.

The dispute between philosophers allied with Elga and Christensen on the one hand and those allied with Kelly on the other is usually put this way: friends of Elga and Christensen argue that we shouldn't take into account what I am calling the original reasoning *when determining what the rational reaction is in case of a disagreement with a peer*, and friends of Kelly disagree. Thus understood, the dispute is between conciliatory and non-conciliatory views, for if we are barred from appealing to the original reasoning when deciding how to react to a disagreement with a peer the natural option is to conciliate, whereas appeal to that reasoning opens the door to non-conciliation in at least some cases. As the issue is understood here, however, the dispute is not about how to react to disagreement with a peer (the proportional weight view gives the answer to that), but rather about what evidence can influence who we are justified in considering a peer and who we are not. Thus understood, it is a dispute within the conciliatory family. Proponents of both the EC and the K version of the proportional weight view agree that the answer to the Main Question is that you should always conciliate, but they disagree about their answer to the Epistemic Question, and so they may disagree about the range of cases that the Main Question is concerned with. This may seem like no difference at all: according to one way of viewing the dispute they disagree about whether it is always rational to conciliate in a disagreement with a peer, whereas according to the other way of seeing the dispute they agree that it is always rational to conciliate in those cases, but disagree as to when we are justified in taking somebody else to be a peer. The two ways of viewing the disputes are indeed equivalent as far as the answer to the Main Question is concerned, for they give the same answer to that question. But it doesn't follow

that it doesn't matter which way the dispute is viewed. It matters (in part) because if the dispute is viewed as an inter-conciliatory one, then the Epistemic Question is brought to the fore, whereas it is usually ignored (sometimes purposefully so) in the recent literature on disagreement. The treatment of the dispute between the EC and the K versions of the proportional weight view in the next section (and, relatedly, the treatment of the Epistemic Question) will have a somewhat deflationary flavor, for it will be argued that whether the original reasoning can make a difference as to whether we are justified in considering somebody else a peer or not is something that will vary from case to case, and that there is no disagreement-specific principle that will correctly classify those cases. Even so, the very fact that the answer to the Epistemic Question should take this deflationary form is something that is easily missed if the question is never confronted to begin with.

6 Demotion from peerhood

One prominent argument in favor of (in effect) the EC version of the proportional weight view has to do with *illegitimate demotions from peerhood*:¹⁹

Illegitimate demotions from peerhood argument

1. If it is correct to allow the original reasoning to be part of C , then illegitimate demotions from peerhood are justified.
 2. Illegitimate demotions from peerhood are not justified.
- Therefore,
3. It is not correct to allow the original reasoning to be part of C .

The following example will illustrate why many philosophers believe 1 to be true.²⁰ Suppose that you and a friend are at the Kentucky Derby, sitting next to each other. As the Derby is about to finish, Pioneer of the Nile and I Want Revenge are head to head. After the race is over, but before the official result is announced, you tell your friend that you are pretty sure that Pioneer of the Nile won, and your friend tells you that she is pretty sure that I Want Revenge won. It would be absurd for you to now reason in the following way: “I used to consider her my peer when it came to simple perceptual judgments such as who won a horse race; but *this time she got it wrong* (since she thinks that I Want Revenge won, whereas the truth is that Pioneer of the Nile won), so I have a reason to demote her from peerhood; therefore, even though before this particular disagreement her opinion would have had considerable weight in what I should believe regarding the outcome of the race, the very fact that she got it wrong this time shows that her opinion should have less weight now.” The absurdity of your reasoning can be dramatically displayed by imagining that you keep on disagreeing with your friend about close races and, on the basis of the fact that (as you see things) she keeps getting it wrong, you come to the conclusion that her opinion about such things is almost worthless.²¹

Your demoting your friend from peerhood in that way is clearly illegitimate and, therefore, unjustified (premise 2 of the argument may be taken as trivially true). What goes wrong in your argument that she ought to be demoted, though? Proponents of the illegitimate demotions from peerhood argument claim that the diagnosis of your mistake lies in noticing that you are using the original reasoning²² (that it looks like Pioneer of the Nile won) in order to decide the worth of your friend’s disagreeing opinion. In judging the worth of your friend’s opinion, they will argue, you should bracket the fact that (as you see things) Pioneer of the Nile won—otherwise, you are

bound to find your friend's opinion worthless. So, according to this line of thought, what goes wrong in your argument for demoting your friend from peerhood is that you allow the original reasoning to have a say regarding the worth of your friend's disagreeing opinion—and so this shows that premise 1 of the argument is right.

But consider a modification of the case. Suppose that, instead of being pretty sure that I Want Revenge won, your friend is pretty sure that Hold Me Back won. Part of the evidence available to you is that even though it was a close race between Pioneer of the Nile and I Want Revenge, you can clearly see that Hold Me Back is still hundreds of meters from the finish line when the first two horses cross it. You are perfectly justified, then, in discounting your friend's opinion that Hold Me Back won—and you are therefore perfectly justified in demoting her from peerhood (although, as pointed out below, it is not strictly speaking right to say that what you are doing is demoting her from peerhood). But notice that your discounting your friend's opinion is based on your original reasoning—if that reasoning were not allowed to be part of the circumstances of disagreement, like proponents of the illegitimate demotions from peerhood argument would have it, then this perfectly legitimate demotion from peerhood would count as illegitimate. In addition to illegitimate demotions from peerhood, there are *legitimate* demotions from peerhood, and these legitimate demotions are indeed properly based on the original reasoning. Therefore, what explains what goes wrong in the case where you illegitimately demote your friend from peerhood cannot be that you are using the original reasoning—premise 1 is false. These considerations can be summarized in the following argument in favor of the unrestricted version of the proportional weight view:

Legitimate “demotions from peerhood” argument

1. If it is not correct to allow the original reasoning to be part of *C*, then legitimate “demotions from peerhood” are unjustified.
 2. Legitimate “demotions from peerhood” are justified.
- Therefore,
3. It is correct to allow the original reasoning to be part of *C*.

If the legitimate “demotions from peerhood” argument is sound and (therefore) the illegitimate demotions from peerhood argument is not, what *does* go wrong with your reasoning when you illegitimately demote your friend from peerhood? The answer is that, although there is nothing wrong with the mere fact that you are taking into consideration the original reasoning, there is indeed something wrong with *the way* in which you are taking it into consideration when illegitimately demoting your friend from peerhood.

Think about your state of mind after you judge the race but before you talk to your friend about it. As you see things, it was a very close race between Pioneer of the Nile (who won) and I Want Revenge, but Hold Me Back was a clear loser. From your point of view, then, there is something clearly wrong with someone who thinks that Hold Me Back won, but it is understandable how someone can think that I Want Revenge won. Indeed, your state of mind is such that you should be ready to revise your opinion about who won the race if new evidence that I Want Revenge won comes in, but not if evidence that Hold Me Back won comes in. It’s not that you should be completely impervious to evidence that Hold Me Back won—if enough evidence starts accumulating that this is indeed what happened (if, for instance, everyone around you thinks this and the official announcement agrees with them) then you should start to think that you seriously misperceived the situation,

and should come to acquiesce in the opinion that Hold Me Back won. But it should be much harder for you to accept that you were wrong in thinking that Hold Me Back lost than it should be for you to accept that you were wrong in thinking that I Want Revenge lost. Your original reasoning doesn't just dictate what you should think about who won the race, it also dictates how easy it should be for you to change your opinion about such things given different kinds of new evidence. In general, evidence plays two different roles in shaping the opinions of rational subjects: it determines what opinions they hold, and it also determines how *resilient* those opinions are to new evidence.²³ So, what goes wrong when you illegitimately demote your friend from peerhood is that you are ignoring this second role of the evidence. Your fault is not that you are taking into account the original reasoning, but that you are doing so in the wrong way.

A toy model of the situation might help here. There are six propositions that you are interested in: that Pioneer of the Nile won the race (p), that I Want Revenge won (q), that Hold Me Back won (r), that your friend thinks that p (s), that she thinks that q (t), and that she thinks that r (u). Because you think that only one horse can win, and because you also think that your friend knows this, you take seriously only nine possibilities. Your evidence after you judge the race but before you talk to your friend determines the credence that you invest in each of these possibilities. Now, given what we said about your evidence, something not too far-off from the credence distribution in Table 1 would be justified by that evidence. In particular, notice that, given that the credence that you invest in a proposition is just the sum of the credences that you invest in the possibilities in which that proposition is true, your credence in p is 0.9, your credence in q is 0.09 and your credence in r is 0.01.

Table 1: Your credence distribution after judging the race but before talking to your friend.

p	q	r	s	t	u	Cr
T	F	F	T	F	F	0.8
T	F	F	F	T	F	0.08
T	F	F	F	F	T	0.02
F	T	F	T	F	F	0.007
F	T	F	F	T	F	0.08
F	T	F	F	F	T	0.003
F	F	T	T	F	F	0.004
F	F	T	F	T	F	0.004
F	F	T	F	F	T	0.002

Table 2: Your credence distribution after your learn that t .

p	q	r	Cr
T	F	F	0.49
F	T	F	0.49
F	F	T	0.01

How should learning that a certain proposition is true affect your credence distribution? One natural answer to this question is that you should “zero-out” all the possibilities where the proposition is false and then re-distribute the leftover credence over the remaining possibilities in proportion to their original credence.²⁴ Tables 2 and 3 show, respectively, how your credence distribution would look after learning that your friend thinks that I Want Revenge Won and that she thinks that Hold Me Back won (propositions with credence of either zero or one have been eliminated from the tables). Notice that, if you follow this procedure, your credence in the proposition that Pioneer of the Nile won after learning that your friend thinks that I Want Revenge won will decrease to less than 0.5, whereas after learning that she thinks that Hold Me Back won will only decrease to 0.8.

Table 3: Your credence distribution after you learn that u .

p	q	r	Cr
T	F	F	0.8
F	T	F	0.12
F	F	T	0.08

This toy model need not be taken very seriously, but it is revealing that, even on the basis of this toy model, it is easy to see why dismissing your friend’s opinion that I Want Revenge won would be illegitimate and also why not giving too much weight to her opinion that Hold Me Back won would be justified. The toy model shows that your original reasoning dictates *both* what would count as legitimate and illegitimate demotions from peerhood—and, therefore, it confirms the conclusion that premise 1 in the illegitimate demotions from peerhood argument is false.²⁵

It is time to correct a possibly misleading way of putting things. To say that, when your friend thinks that Hold Me Back won, it would be legitimate for you to demote her from peerhood is not quite the right way to state what you would be doing. For remember that, according to the dispositionalist answer to the Psychological Question, for you to consider your friend a peer with respect to a certain circumstance of disagreement is for you to have the disposition to “split the difference” with her if a disagreement were to occur in that circumstance. But the circumstance includes what it is that your friend thinks—whether she thinks that I Want Revenge or Hold Me Back won. And, as the toy model helpfully shows, *you never considered your friend a peer with respect to this latter circumstance* (that explains the scare quotes in “the legitimate ‘demotions from peerhood’ argument”, which, as it turns out, is not an argument based on demotions from peerhood). So it is after all true that it is never legitimate to demote someone from peerhood on the basis of the original reasoning.

But this is not because, as the proponent of the EC version of the proportional weight view would have it, it is not legitimate to consider the original reasoning when assessing your friend's testimony. On the contrary, it is because the friend of the K version of the proportional weight view is right that you should always (properly) take the original reasoning into account, and whether you count someone as a peer or not is determined by the dispositions justified by your original reasoning (together with all the rest of your evidence). If the diagnosis of the illegitimate and legitimate demotions from peerhood (and "demotions from peerhood") argued for in this section is correct, then it is the K version that is the correct interpretation of the proportional weight view.

To recap: it can easily be seen in the Kentucky Derby case that your evidence may very well justify you in considering someone who believes that I Want Revenge won a peer, but not someone who thinks that Hold Me Back won. The same sort of situation could arise in more interesting cases of disagreement. What features of the evidence determine the level of resilience that your credences should have? It is with respect to *this* question that the only available answer is a deflationary one: what determines the resilience of your credences in cases of disagreement is the same thing that determines the resilience of your credences in other cases, namely, the nature of your evidence. Rational thinkers respond to their evidence, including evidence that they have a disagreement, by making whatever doxastic changes are required by that evidence. Everyone will accept this. But some philosophers seem to think that there are *special principles* about the rational reaction to evidence when the evidence includes the fact that there is a disagreement. This is what I am denying. We should respond to evidence of disagreement the same way we respond to any old evidence, and there is no topic-specific principle about how to deal with evidence

of disagreement, just as there is no topic specific principle about how to deal with evidence regarding penguins.

7 Conclusion

Recall the four questions about disagreement:

(I) The Metaphysical Question: *What is it* for a subject S to have a certain level of epistemic competence regarding proposition p and circumstances C relative to subject S^* ?

(II) The Psychological Question: What is it for a subject S to *consider* someone else as having a certain level of epistemic competence regarding proposition p and circumstances C relative to S ?

(III) The Epistemic Question: Under what conditions is a subject S *justified* in considering someone else as having a certain level of epistemic competence regarding proposition p circumstances C relative to S ?

(IV) The Main Question: Suppose that S is justified in considering another subject S^* as having a certain level of epistemic competence regarding proposition p and circumstances C relative to S . Now suppose that S finds out that there is a disagreement between S and S^* about p under circumstances C . *How should that affect S 's opinions about p ?*

If the arguments of this paper are right, the Metaphysical Question should be ignored, the Psychological Question should be given a dispositionalist answer and the Main Question should be given a conciliatory answer (in particular, the proportional weight view). But what I take to be the main conclusion of this paper is that the

Epistemic Question should take center stage (or, perhaps, be recognized as being what is deep down at stake between the different positions on disagreement in the literature). I have argued that a partial answer to the epistemic question is given by the adoption of the K version of the proportional weight view, where the original reasoning is allowed to play a role in the determination of who we are justified in considering as peers. As to how to complete the answer to the epistemic question, I have also suggested that the contemporary literature on disagreement is misleading in this respect. For that literature proceeds as if there were specific principles dealing with evidence of disagreement. The time has come, however, to seriously consider the option that there are no such disagreement-specific principles.²⁶

Notes

¹Elga (2010) calls positions of this kind “conciliatory.”

²Authors who advance conciliatory views include Feldman (2007), Elga (2007), and Christensen (2007). Authors who advance non-conciliatory views include van Inwagen (1999), Kelly (2005), Kelly (2010), and Wedgwood (2007).

³The reason why this is only a first approximation to the questions that should interest us will become apparent in section 3.

⁴My (II) is Elga’s 1 and my (IV) a version of Elga’s 2 modified to include explicit reference to the fact that the subject is justified in considering his friend an epistemic peer—see Elga (2007), p. 483.

⁵Kelly (2005) proposes the epistemic virtues answer to the Psychological Question, and credits Gutting (1982) with the basic idea. See also Feldman (2007).

⁶The example is adapted from Christensen (2007).

⁷Thanks here to discussion with (suppressed for review). That the belief and the dispositions are not unrelated can be seen by considering cases where they come into conflict. In cases where an explicit belief that she is my peer is in conflict with my dispositions to change my mind about the disputed issue, we think that I should either give up my belief or change my dispositions.

⁸One could try: two subjects are epistemic peers of each other just in case they are both justified in taking each other to be peers. This, however, would entail a kind of infallibility regarding justified judgments of epistemic peerhood.

⁹Within a Bayesian framework where credences are represented by functions that satisfy the probability calculus and change in belief is modeled in terms of conditionalization on newly acquired evidence, relative epistemic competence can be defined in terms of conditional credences in p given all the different possible disagreements.

¹⁰These considerations show that the notion of relative epistemic competence considered here doesn't behave like the notion of "respect" introduced in Lehrer (1976): No "constancy condition" is imposed on your dispositions. To say that there are no *specific* restrictions on how you should change your dispositions doesn't mean, of course, that there cannot be *general* restrictions (not having to do with the fact that we are dealing with a case of disagreement) on how those dispositions can and should evolve.

¹¹Notice that I am not defining what it means for S to be justified in being disposed to assign certain credences in terms of what S would do, and so the principle is not subject to potential problems having to do with the "conditional fallacy." Notice also that whether in those circumstances S would be justified in assigning any *other* credence to p depends on whether a suitable "uniqueness principle" is true—for discussion, see White (2005).

¹²Elga is thinking within a Bayesian framework, where instead of talking about the dispositions to change credences he talks about the conditional credences of the subject. See note 9.

¹³Elga (2007), p. 490, my emphasis.

¹⁴Christensen (2011), pp. 1-2.

¹⁵Kelly (2010), p. 25.

¹⁶Kelly (2010), p. 32.

¹⁷I am not suggesting that this is how the authors themselves have understood their disagreement, for the distinction between the four questions is mine (although it is inspired by what these authors have written). What I *am* suggesting is that this is an illuminating way of seeing the dispute.

¹⁸Thanks here to an anonymous referee for the *American Philosophical Quarterly*.

¹⁹Elga (2007), Christensen (2007) and White (2005) all offer versions of the argument.

²⁰The kind of example is from White (2005).

²¹I am assuming that, for some reason, you never become aware of the official outcome of the races.

²²Here is one of the cases where “reasoning” is not quite the right word.

²³See Skyrms (1977), cited in Egan and Elga (2005).

²⁴This is equivalent to saying that you should re-normalize your credence distribution after zeroing-out the relevant possibilities. In turn, this is equivalent to the traditional Bayesian account of learning, according to which you should replace your original credence distribution with the result of its conditionalization on the proposition learned.

²⁵Notice also that this diagnosis doesn’t depend on the fact that this is a perceptual case. The points made apply just as plausibly to any other case of disagreement.

²⁶Many thanks to David Christensen, Adam Elga, Tom Kelly, Carolina Sartorio and an anonymous referee for this journal for many helpful conversations and suggestions on how to improve previous versions of this paper.

Bibliography

Christensen, D. (2007), “Epistemology of Disagreement: the Good News,” *Philosophical Review* 116(2), pp. 187–217.

Christensen, D. (2011), “Disagreement, Question-Begging and Epistemic Self-Criticism,” *Philosopher’s Imprint* 11(6), pp. 1–22.

Egan, A. and Elga, A. (2005), “I Can’t Believe I’m Stupid,” *Philosophical Perspectives* 19, pp. 77–94.

Elga, A. (2007), “Reflection and Disagreement,” *Noûs* 41(3), pp. 478–502.

Elga, A. (2010), “How to disagree about how to disagree,” in Feldman and Warfield (2010).

Feldman, R. (2007), “Reasonable Religious Disagreements,” in *Philosophers Without God: Meditations on Atheism and the Secular Life*, Oxford University Press.

- Feldman, R. and Warfield, T. (eds.) (2010), *Disagreement*, Oxford University Press.
- Gutting, G. (1982), *Religious Belief and Religious Skepticism*, University of Notre Dame Press.
- Kelly, T. (2005), "The Epistemic Significance of Disagreement," in J. Hawthorne and T. Gendler (eds.), *Oxford Studies in Epistemology*, Oxford University Press, volume 1.
- Kelly, T. (2010), "Peer Disagreement and Higher Order Evidence," in Feldman and Warfield (2010).
- Lehrer, K. (1976), "When Rational Disagreement is Impossible," *Noûs* 10(3), pp. 327–332.
- Skyrms, B. (1977), "Resiliency, Propensities, and Causal Necessity," *Journal of Philosophy* 74(11), pp. 704–713.
- van Inwagen, P. (1999), "Is it Wrong, Everywhere, Always, and for Anyone, to Believe Anything on Insufficient Evidence?" in E. Stump and M. Murray (eds.), *Philosophy of Religion: The Big Questions*, Wiley-Blackwell, pp. 273–285.
- Wedgwood, R. (2007), *The Nature of Normativity*, Oxford University Press.
- White, R. (2005), "Epistemic Permissiveness," *Philosophical Perspectives* 19(1), pp. 445–459.