Doubts about philosophy? The alleged challenge from disagreement

Thomas Grundmann, University of Cologne

Abstract: In philosophy, as in many other disciplines and domains of discourse, stable disagreement among peers is a widespread and well-known phenomenon. Our intuitions about paradigm cases, e.g. Christensen's Restaurant Case, suggest that in cases of peer disagreement suspension of judgment is rationally required. This would prima facie suggest ubiquitous suspension of judgment in philosophy. But we are still lacking a deeper theoretical understanding of why and under what conditions suspension is rationally called for. In the first part of this paper, I will focus on this question. After a critical survey of some recent alternative approaches (disagreement as a threat to reliability, as a decision problem, and as providing an undercutting defeater), I will argue that discovering peer disagreement provides one with a *rebutting defeater*, but only if some further non-trivial conditions are satisfied – among them acknowledging one's opponent as an epistemic peer. In the second part of the paper, I will explore the skeptical implications that this account seems to have for philosophy in particular. Here, I will defend two claims. First, skepticism about philosophy is rationally required only if the relevant peerness assumption can be justified. Second, there is no rational basis in philosophy that would support the relevant peerness assumption. If this peerness agnosticism is correct, then we are not forced into skepticism about philosophy and we may rationally retain our philosophical beliefs even in the face of controversy.

As philosophers we are often confronted with colleagues who do not share our views. In the face of controversy we use arguments and present evidence in order to overcome these disputes. Sometimes we succeed in using these rational tools to persuade our opponent that the better reasons are on our side, or that her own argument was invalid or unsound, or that, on balance, her position has the higher theoretical costs or the lower explanatory power. At other times, it is just the other way round. But most of the time the situation in philosophy is very different. After having exchanged arguments and evidence for a long time, the opponents in a dispute

nevertheless have not made any significant progress in overcoming their disagreement. Both sides continue to find their own view more attractive and, overall, more plausible. These irresolvable disagreements about most 'big questions' in the discipline seem to be a highly specific feature of philosophy. In contrast, opposing scientists tend to converge much more often in their views, at least in the long run. In the case of irresolvable disagreement rationality seems to require that both sides suspend judgment. Since such kind of disagreement is widespread in philosophy, skeptical implications for the discipline in general seem unavoidable. Consider the following *simple skeptical argument*:[1]

(1)   Whenever people irresolvably disagree about some issue, suspending judgment on that issue is rationally required.

(2)   Irresolvable disagreement among philosophers is extremely widespread.

Therefore,

(C1)   In philosophy, hardly any judgment can be rationally retained.

On reflection, however, this argument is clearly too simple to succeed, for premise (1) is much too strong. Sometimes I am involved in an irresolvable disagreement with an opponent, but I can still rationally retain my judgment. A flatearther might persistently disagree with me about whether the earth is spherical. Nevertheless, this does not provide me with any reason to lower my confidence in my judgment. Only disagreement with people who are epistemically on a par requires me to suspend judgment. Let me appeal to the common notion of *epistemic peerhood* here.[2] Two epistemic agents, A and B, are epistemic peers with respect to a certain proposition p

---

[1] This is a simplified version of the argument considered by King 2011, 1f.
[2] Compare Christensen 2009, 756..

if and only if (i) A and B have access to the same evidence relevant to p and (ii) A and B are equally competent in assessing the relevant evidence. It seems reasonable that disagreement requires suspending judgment only if we take the opponent to be our epistemic peer prior to the dispute.[3] And this condition is clearly violated in the case of flatearthers. They probably do not use all of the relevant evidence, and they probably also assess this evidence less competently than we do when we evaluate the evidence at hand.[4] In disputes with epistemic superiors or inferiors we can always defer to the view of the relative epistemic authority – who sometimes happens to be ourselves.[5] Furthermore, premise (1) is also too strong for a second reason. Sometimes the very fact that disagreement among peers does not dissolve indicates that there is no real conflict, but a merely verbal disagreement among the interlocutors.[6] In the case of verbal disagreement one side asserts 'p' and the other side asserts 'non-p', but both sides may express true beliefs with their assertions since they use syntactically identical sentences with different meanings. Now, if there is no real conflict among the opponents, then both sides may rationally hold on to their judgment.

Whereas the preceding two caveats are often highlighted in the literature on disagreement, there is a further condition which has received much less attention. Consider a case where you are

---

[3] Here I leave aside the question of whether we might downgrade our opponent on the basis of the prevailing disagreement alone.

[4] For example, flatearthers ignore most of the evidence from space travel and space telescopes (since they believe that all these things are part of a big hoax), they have interpreted experiments incorrectly (e.g., Samuel Rowbotham, who was one of their historical leaders), there is no expert astrologist among them, and most of them are strongly biased by religious background assumptions. Compare Garwood 2007

[5] The relative epistemic authority is the most competent agent in the dispute. She need not be an authority or expert in any absolute sense.

[6] This conjecture was already made by David Hume 1975, 80f: "From this circumstance alone, that a controversy has been long kept on foot, and remains still undecided, we may presume that there is some ambiguity in the expression, and that the disputants affix different ideas to the terms employed in controversy. (...) (I)t were impossible, if men affix the same ideas to their terms, that they could so long form different opinions of the same subject; especially when they communicate their views, and each party turn themselves on all sides, in search of arguments which may give the victory over their antagonists." Verbal dispute is also lively discussed in the current metaphilosophical debate. See e.g. Chalmers, Manley & Wasserman 2009.

involved in a *genuine* disagreement about the proposition p with someone who is, at least from your perspective, your *epistemic peer*. Are you required to suspend judgment? The correct answer depends on the further details of this case. Assume that there is a group of 100 philosophical peers discussing one of the big questions in the discipline. Whereas 98 peers agree with me on p, there is only a single peer disagreeing with me on that proposition. Intuitively, numbers do count in this situation. If the majority takes sides with me, I need not be worried by only one dissenting peer. We can explain this intuition on the basis of the Condorcet Jury Theorem. According to this theorem, the following is true: if the majority of people whose judgments are independent of each other and who each have a more than 50% chance of getting it right agree on a certain proposition in the relevant domain, then the probability of the majority view being true soon approaches 1 when the number of judging peers further increases. So, if I share the majority view in my peer group, it is more likely that I am right than the one dissenting peer. Or, to put it differently, if there is a majority view and a minority view within some group of epistemic peers, then the subgroup subscribing to the majority view is – taken as a group – not a peer to the subgroup subscribing to the minority view. The former group is epistemically superior to the latter group, although their individual members are epistemically on a par. So, when I am confronted with peer disagreement, I should defer to the view of the clear majority, as long as nothing indicates that the individual judgments in that group are dependent on each other. Only if there is no clear majority view, I am rationally required to suspend judgment.

We may conclude that the discovery of irresolvable disagreement requires us to suspend our prior judgment, only if

(i)   the dissenters are epistemic peers,

(ii) the disagreement is genuine, and

(iii) the controversy is symmetric, i.e., the opposing parties are of the same size.

It goes without saying that conditions (i) and (iii) require only approximate equality.

We can now articulate an *improved version of the skeptical argument*:

(1') If there is a *genuine* and roughly *symmetric* disagreement among *epistemic peers* (= GSE-
disagreement), all parties are rationally required to suspend judgment.

(2') In philosophy, GSE-disagreements are widespread and stable.

Therefore,

(C1) In philosophy, hardly any judgment can be rationally retained.

The revised premise (1') seems to fit the intuitive cases much better. Consider, e.g., Christensen's *restaurant case* (Christensen 2007): Some friends go out for dinner at a restaurant. After dinner, they decide to split the bill evenly. Since no calculator is available, two of the friends, Paul and Paula, do some mental math. Usually, both are roughly equally good at that. But in this case they come to slightly different results about the shares. There is no obvious reason why one rather than the other should have made a mistake. They are both neither drunken nor distracted, nor is their performance disturbed in some other way. Intuitively, both are rationally required to withhold their judgment about the shares, as soon as they realize their disagreement. This is just what premise (1') says about the case: Paul and Paula have a genuine dispute (since they cannot both be right about the shares); they are clearly epistemic peers (since they use the same

5

evidence, namely the bill and the number of their friends, and both have proved to be equally good in mental arithmetic in the past); and the opposing parties in this dispute are of equal size (namely one person on each side). In this case, suspension of judgment is mandatory according to (1').

Premise (2') also looks pretty plausible. Here I will take it for granted that at least many disputes in philosophy are genuine and not merely verbal. Whether disagreement in this domain is widespread, stable and symmetric is basically an empirical question. The findings about Philosophy faculty members and PhDs provided by a recent informal survey by the online database PhilPapers suggest that it is.[7] On most of the big questions the disputes are roughly symmetric. When you look carefully at the results you will also find some notable exceptions. Currently, there are some clear majority views in philosophy, for example concerning *idealism* (PhilPapers found only 6.8% adherents of idealism left among the academic philosophers participating in the survey, in contrast to 76.6% external world realists)[8] and *moral cognitivism* (moral cognitivism seems to be the dominant view with 62.7% adherents, in contrast to 18.8% non-cognitivists). But the typical disagreement is, at least roughly, symmetric. This is illustrated by the following sample results from PhilPapers: there are 36.3% adherents of *platonism* versus 40.8% *nomimalists*; 38.1% *empiricists* versus 26% *rationalists*; 43.7% *epistemic externalists* versus 26.1% *epistemic internalists*. These findings suggest that there is no perfect symmetry

---

[7] See PhilPapers Survey
(http://philpapers.org/surveys/results.pl?affil=Philosophy+faculty+or+PhD&areas0=0&areas_max=1&grain=coarse)
, access on 2, 29, 2012.
[8] The sum does not equal 100% since there are also competing non-standard views which I do not mention here.

between the opposing parties. Nevertheless, there is a lot more approximate symmetry than in the sciences.[9]

What about epistemic peerness? It seems save to claim that in irresolvable disputes the parties often *take* each other to be peers, especially if the dispute occurs within a single philosophical tradition.[10] Are they correct in doing so? We can say at least two things to support the assumption of peerness. *First*, if experts in the field are involved in a stable dispute, then they will typically know all (or most of) the relevant arguments, evidence and literature on the topic. From regular conversation, they will also know each other's arguments and evidence. This makes it highly likely that the opponents share roughly the same evidence. *Second*, in real disputes, both parties have the opportunity to experience each other in 'philosophical action.' Over time, they will realize whether the opponent is equally smart and intelligent, whether she comes up with reasonable (though for her opponent not completely persuasive) responses to objections, whether she knows the field equally well, whether she is as dedicated, acute and honest in working hard on the problems. This personal impression of the opponent can be backed by her reputation in the scientific community, at least partly based on her track-record of peer reviewed publications.

So far, the improved skeptical argument appears to be sound. At the same time, it supports a devastating and surely surprising conclusion, namely that doing philosophy does not provide us with rational or justified beliefs about most of its issues. In this paper, I want to do two things. *First,* I will focus on the question *why* GSE-disagreement requires us to suspend judgment, if it

---

[9] Compare for example the debate on whether global warming is human-induced. In a study by Anderegg, W. ; Prall, J.; Harold, J.; Schneider, S. 2010 it was established that 97-98% of the 1.372 surveyed climate researchers agree on the human origin of global climate change. Even more consensus will be found on quantum physics or relativity theory. I do not want to dispute that there are other, more speculative parts of physics (as cosmology or the unification of quantum physics and relativity theory) which are much more controversial among the relevant experts. However, it seems clear that no scientist would claim that beliefs in these parts of physics have been established.

[10] Compare, e.g., van Inwagen 1996, 138; Feldman 2006, 217, 219; Kelly 2005, 167f; Goldberg 2009, 106.

does. So far, this has only been suggested by intuitive cases, like the restaurant case. In addition, I want to gain a systematic understanding of the sources of this requirement from a theoretical perspective. In the literature, three different models have been suggested to explain why GSE-disagreement has the epistemic effect of neutralizing the rationality of one's own judgment. According to the *reliability argument*, the diversity of judgment threatens the reliability of the relevant methods. According to the *decision problem argument*, peer disagreement puts us in a position in which we cannot rationally decide what to believe. Finally, according to the *undercutting defeater argument*, disagreement creates a certain kind of defeater which undermines the rationality of our own judgment. My assessment of these different models is guided by one crucial adequacy condition, namely that a promising explanation of the epistemic effects of disagreement has to be neutral with respect to the metaepistemological positions of internalism and externalism. The reason is simply that internalists as well as externalists both feel the pull of the challenge from disagreement. So, the challenge should not crucially depend on either of these views. Related to this issue, there is the question of how we have to understand this disagreement when claiming that it neutralizes justification. Is it *the fact* of GSE-disagreement (factive reading), *the belief* that GSE-disagreement occurs (doxastic reading) or *the justified belief* that GSE-disagreement occurs (justificationist reading) that does the epistemic work here?[11] Concerning the general question of which of the above models one should prefer, I will argue that neither of them works. I will then present my own systematic understanding of what is epistemically going on in a case of GSE-disagreement. According to my view, in facing GSE-disagreement we instead acquire a *rebutting defeater*. This calls for a doxastic reading of GSE-disagreements. I will continue by arguing that, nevertheless, the skeptical argument can only succeed if we have good reason to believe that there is (widespread) GSE-disagreement in

---

[11] Enoch 2010, p. 970f, is very clear on this issue.

philosophy. In the ***second part*** of this paper, I will defend the view that the skeptical argument is unsound, since we have no good reason to believe that there is (widespread) GSE-disagreement in philosophy. If this is correct, philosophers need not be worried by the challenge from disagreement after all.

<p style="text-align:center">I</p>

According to the **reliability argument**, there is in fact so much diversity of opinion in philosophy that the methods used by philosophers cannot be (or are highly unlikely to be) reliable. For the reliabilist this by itself can establish that philosophical beliefs are never *prima facie* justified. But even the internalist can use a similar argument. She can argue that as soon as we *realize* that philosophical methods are (very likely) unreliable (e.g., by considering the reliability argument) we acquire an undercutting defeater which *ultima facie* removes all justification from our philosophical beliefs. For the internalist it is not the *fact* that reliability is lacking, but the agent's *belief* that reliability most likely is lacking that does all the epistemic work. The reliability argument thus passes the above defended adequacy condition of metaepistemological neutrality.[12] So far, so good.

But let's have a closer look at the argument itself. As I said, the argument starts from the observation of widespread and stable diversity in philosophy and concludes that it is very unlikely that there is a philosophical method that is reliable. By which steps does it get there?

---

[12] In Goldberg 2009 the reliabilist argument runs under the label 'unreliable methods argument,' whereas the internalist argument is called 'undercutting defeater argument.'

Goldberg (2009), one of the major proponents of this argument,[13] describes the situation in philosophy as follows:

> "Some progress can be made (...) by noting the diversity of incompatible opinions on contested matters of philosophy. (...) If there are not all that many methods available, then it is likely that each of the methods available will be shared by a great many people. Under these conditions it would be nothing short of incredible if it turned out that, while (taken as a group) the belief-forming methods used in philosophy do not produce a preponderance of true belief, nevertheless there is a distinctively philosophical method that is reliable." (Goldberg 2009, p. 106f)

Goldberg seems to have something like the following argument in mind here:

(R1) There is a widespread disagreement in philosophy.

(R2) There are not all that many philosophical methods available.

(R3) If (R1) and (R2) are true, then it is highly unlikely that there is a philosophical method that is reliable.

Therefore,

(CR) It is highly unlikely that there is a philosophical method that is reliable.

Before analyzing the crucial premise (R3), some conceptual distinctions and a few comments on the other premises are required. Let us first consider the concept of *reliability*. Reliability is basically a gradual notion. When applied to belief-forming methods the degree of their reliability is determined by the truth-ratio of the methods' output-beliefs. In the above argument, reliability

---

[13] For a similar, but more simple version of this argument compare Brenan 2010, 1: "Dissensus shows that philosophical methods are unreliable instruments of truth."

is used as an absolute notion. Consequently, a method is either reliable or unreliable. In order to define the absolute notion on the basis of the gradual one we need a certain threshold. Where this threshold exactly lies in the case of belief-producing methods is, at least to a certain extent, either indeterminate or determined by contextual factors. For reasons of simplicity, let us just stipulate that a method must produce more than 70% true beliefs in order to count as reliable in the absolute sense. One further conceptual clarification is needed here. We may attribute a truth-ratio to different things. We can talk about the total truth-ratio of a *discipline*. It is determined by the ratio of true to false beliefs produced in the discipline (e.g. philosophy) in general. Apart from this, we can also talk about the truth-ratio of an *epistemic agent*. It is determined by the number of true and false beliefs that the agent acquires. Finally, we can talk about the truth-ratio of *methods*.

Now, premise (R1) should be understood in such a way that it implies that the truth-ratio of philosophy as a discipline is below 70%. Is that plausible? Consider the case of *extreme diversity*. In that case true and false beliefs would be distributed equally on every matter of philosophy. Philosophy taken as a whole would then produce only 50% true beliefs.[14] It seems clear that this is not generally the case in philosophy. Not everything is controversial (even opponents sometimes share their views on positions that they dismiss) and not always is there perfect symmetry in philosophical disputes. But it seems not unreasonable to believe that philosophy in general is characterized by *strong diversity*, namely a truth-ratio below 70%. If this is true, then

---

[14] In fact this picture is the result of a strong simplification. Not all philosophical assertions are either true or false. Some sentences may be vague or meaningless. Moreover, there may sometimes be two or more conflicting positions that are all wrong since they are contrary rather than contradictory. For example, in a controversy about the color of a particular car that was involved in an accident, some people may believe that the car was red, others believe that it was yellow and then there are still others who believe that it was black. In fact, all sides might be wrong if the car was white. Moreover, false positive beliefs may often be counterbalanced by true negative beliefs ('The car was not yellow,' 'The car was not red,' 'The car was not black.') In this paper I put aside these kind of complications. (Thanks to Sandy Goldberg for reminding me of this point.)

we can conclude that philosophy, as a discipline, is not reliable. For the sake of the argument, let us also take for granted that there are not all that many philosophical methods available, i.e., that premise (R2) is true.

Premise (R3) now claims that if (R1) and (R2) are true, then it is highly unlikely that there is a philosophical method that is reliable. The above discussion of (R1) made it clear, that the truth of (R1) implies that philosophy (as a whole) is not reliable. Is it true that if philosophy (as a whole) is not reliable and there are not that many philosophical methods, then the existence of a reliable philosophical method is not very likely? Let us consider some cases to get clear about this point. First, suppose that there is only one philosophical method that is applied by all philosophers to all philosophical matters. In this case, it is obvious that the truth-ratio of the discipline can simply be projected onto the truth-ratio of the method. Since philosophy is unreliable the one philosophical method is unreliable, too. Second, consider a case in which there are two or more different methods used in philosophy in such a way that every philosopher uses only one method all the time. Can we proceed as before in projecting the truth-ratio of the discipline onto the truth-ratio of the agents (and then further onto the truth-ratio of the methods)? It seems not. Take the following analogy. You know about a certain basketball team that its strike rate is about 50%. Every second throw by the team scores. If this is the only thing you know about the team, you certainly cannot claim that it is highly likely that every player of the team has a strike rate of 50%. The poor performance of some players might be compensated by the extremely strong performance of other players in the same team. Even if there were only two players in the team, then one of them might always fail, whereas the other might be a perfect player. Things change when one receives the additional information that all players in the team are peers, that is, all are

equally competent players. With this additional information in your hands you can project the strike rate of the team onto the strike rate of every player.

We can use this insight to improve the original reliability argument by adding the further premise (R2a) that most philosophers are epistemic peers.[15] We then get the following revised version of the argument:[16]

(R1) There is a widespread disagreement in philosophy.

(R2) There are not all that many philosophical methods available.

(R2a) Most philosophers are epistemic peers.

(R3*) If (R1), (R2) and (R2a) are true, then it is highly unlikely that there is a philosophical method that is reliable.

Therefore,

(CR) It is highly unlikely that there is a philosophical method that is reliable.

(R3*) looks much more plausible than (R3). If philosophy as a discipline is unreliable, and most philosophers perform equally well (or badly), then it is highly likely that philosophers are in general unreliable. If we now add the further assumption that (nearly) every philosopher uses one and the same method all the time, then we really get to the conclusion that it is highly unlikely any philosophical method is reliable.[17]

---

[15] It seems as if Goldberg has something like this additional premise in mind, when he talks about the peerness of philosophers on p. 106. Nevertheless, he does not include this additional premise in the explicit argument.
[16] Thanks to Sandy Goldberg and Jens Kipper for reminding me that a plausible version of this argument requires some qualifications.
[17] In the final version of the argument, this further assumption should explicitly be added as a premise (R2b). (R3*) should then be modified to (R3**) such that it reads: If (R1), (R2), (R2a) and (R2b) are true, then it is highly unlikely that there is a philosophical method that is reliable.

This looks like a fair reconstruction of the reliability argument. But even under this charitable reading the argument is not a good one. The assumption that nearly every philosopher almost always uses one and the same method is clearly unrealistic. Philosophers use different methods on different occasions which all might have a very different truth-ratio.[18] That blocks any straightforward projection of the truth-ratio of philosophers onto the truth-ratio of the methods they use. But wouldn't the use of different methods by a single philosopher imply a proliferation of philosophical methods, which seems to be in tension with the claim (R2) that there are not all that many philosophical methods available? No. Even if single philosophers use a number of different methods, different philosophers might use the same method on different occasions.

Here is, finally, the case of a *philosopher's toy world* which is meant to illustrate by a very simple example that even under the highly idealized assumptions of (i) extreme diversity in philosophy (with a truth-ratio of about 50%), (ii) epistemic peerness of all philosophers and (iii) a highly restricted number of available methods, it does not follow that the existence of reliable methods is completely unlikely. In our *philosopher's toy world* the number of philosophers is restricted to two, the numbers of available methods is also restricted to two and there are only six philosophical problems waiting for their solution.

Philosophers's toy world

---

[18] One might think that there are only very few philosophical methods available, at least if one restricts one's attention to current analytic philosophy: the evidence seems to be provided almost exclusively by thought experimentation and philosopher's appeal to intuitions, and the assessment of this evidence seems to be based on reflective equilibrium. That's all. But this would be much too restrictive. Apart from the objection that methodological naturalism is a live option for philosophers (one might also think that introspection and appeals to common sense provide additional input), one has to concede that there are many different ways of using intuitions and of implementing the abstract idea of an reflective equilibrium which constitute different fine-grained methods. (Thanks to Jens Kipper for raising this worry.)

Philosophical problems	Philosopher 1	Philosopher 2	Facts
p?	M_1: p	M_2: not-p	p
q?	M_1: q	M_2: not-q	q
r?	M_1: r	M_2: not-r	r
s?	M_2: s	M_1: not-s	not-s
t?	M_2: t	M_1: not-t	not-t
u?	M_2: u	M_1: not-u	not-u

Here is a short instruction how to read this table. Take for example line 1. It says that with respect to the philosophical question whether p is true philosopher 1 forms the belief that p on the basis of method 1, whereas philosopher 2 comes to a contradictory result on the basis of his use of method 2. In fact, p is true such that method 1 leads to a true belief here, whereas method 2 leads to a false belief. You will easily realize that in the philosopher's toy world there is a perfectly reliable method, namely method M_1, which is used by both philosophers on different occasions. The existence of this perfectly reliable method is not at all in tension with the assumption that in the toy world there holds extreme diversity in philosophy, all philosophers are peers and there are not all that many methods used. I do not see why this observation should not generalize to more complex scenarios. Therefore, I conclude that general diversity and peerness in philosophy cannot establish the claim that (very likely) there are no reliable methods available in philosophy. This claim cannot be established unless one adds the completely implausible assumption that nearly every philosopher uses almost always the same method. Even if philosophy were correctly characterized by extreme diversity, this would not rule out the possibility that at least some

philosophical methods are reliable. The reliability argument, therefore, seems not to be the argument we were looking for.[19]

According to a second approach, fully disclosed disagreement poses a certain kind of **decision problem** for the participants in the dispute. In such a situation, no one can rationally decide whom to trust. The problem is the following: when I am confronted with a (symmetric) peer disagreement it seems equally reasonable for me to trust either of the two conflicting parties, since as epistemic peers both share their evidence and are equally competent in assessing it. In that case, I lack any reason to prefer one of them. Such a preference would be completely arbitrary or partial. The judgment of both parties should receive equal weight. Therefore, suspending judgment is rationally required here. Even if I happen to be one of the conflicting parties, the problem seems to persist. My own judgment should not receive extra weight simply because it is mine. Or so the argument goes. In the face of peer disagreement about most of the big philosophical questions I should, according to this argument, suspend judgment on most of the substantial issues in philosophy.

In order to evaluate this argument it seems advisable to distinguish between two different perspectives on a dispute. On the one hand, we may be *external observers* of a disagreement

---

[19] In his 2009 paper, Goldberg introduces a second argument, which he calls 'The Undercutting Defeater Argument.' According to this argument, it suffices to establish that there are not that many reliable methods available in philosophy to create an undercutting defeater: when I use one philosophical method knowing that many philosophical methods are unreliable and that there is no reason to believe that my method is superior to the others, then I am no longer justified in believing that the method which I am currently using is a reliable one. By now, Goldberg thinks that this argument is superior to the reliability argument. For more on that issue see Goldberg forthcoming. – Here is not the place for a comprehensive reply to this argument. But prima facie I have two worries about it. First: it is not fully clear to me whether the information that the majority of philosophical methods is unreliable really is sufficient to provide an undercutting defeater, rather than making suspension of judgment about the reliability of my own method rational. Second: the argument relies on the ambitious assumption of extreme diversity or even ultra-extreme diversity (with a truth-ratio less than 50%) in philosophy. Now, even if it turned out that disagreement in philosophy isn't as widespread as Goldberg assumes, but is restricted to the big philosophical questions, then we should still be able to explain the challenge from disagreement in these more restricted cases. However, Goldberg's new argument couldn't provide us with such an explanation since in general philosophical methods are still reliable, then.

among peers, and on the other hand, we may be involved as *participants* in a peer disagreement.[20] Laypersons are typically in the position of external observers when they are confronted with a dispute among experts. Or consider the case of someone who wants to get information about the outside temperature. When she is then confronted with conflicting readings of two thermometers, then she observes a 'disagreement' among instruments. Now, assume that, as an external observer, you know nothing about the experts or the instruments at hand but that they are usually equally reliable. Further, assume that the disagreement among your sources is equally distributed, such that you cannot defer to any majority view here. In this case, you do have a decision problem, since you do not know whom to trust. Consider, in contrast, the case of active participation in a disagreement. Let us assume that you and your respective peer share all the evidence. Given this evidence, you find the proposition p more plausible, but your peer prefers not-p instead. In such a situation, there is no decision problem that is comparable to the previous problem of an external observer. For, your own perspective on the evidence suggests to you that p is more plausible than not-p. Hence, you are already committed to p prior to the dispute with your opponent, so it is no longer an open question for you whether p or not-p. In other words, since you have already decided to accept p over not-p prior to the dispute, you are not facing a decision between accepting p or not-p anymore. These observations show that there may be a decision problem in the vicinity of peer disagreement, namely the problem that external observes face when they have to decide whom of the opposing parties they should trust. But this problem is orthogonal to the problem at issue here, namely to explain why peer disagreement is a challenge for those who are actively participating in such a dispute.

---

[20] Similarly, Brenan 2010 distinguishes between outsider and insider skepticism.

Many proponents of the challenge from disagreement assume that one's recognition of being involved in a GSE-disagreement gives rise to a *defeater*, more precisely to an **undercutting defeater** for one's antecedently prima facie justified beliefs. Before looking more closely at this argument, let us briefly rehearse how epistemic defeaters basically work.[21] Defeaters are pieces of information acquired by the epistemic agent which remove prima facie justification from the agent's relevant beliefs. When the agent acquires beliefs or evidence that suggest that one of her previously held beliefs is false or based on an epistemically objectionable source, then this removes any antecedent justification that the agent may have had for that belief. If you acquire information which suggests the falsity of one of your previous beliefs, then this constitutes a *rebutting defeater*. Consider the case of a Müller-Lyer illusion where you justifiably believe that the two arrows are of unequal length because one looks longer than the other (and you haven't encountered this kind of illusion before). As soon as you measure the lines, you will realize that the arrows are, in contrast to your visual impression, in fact equally long. This further information provides you with a rebutting defeater for your previous belief. If, on the other hand, you acquire information that suggests that something was epistemically objectionable about the source of one of your other beliefs, this will constitute an *undercutting defeater*. Consider a situation in which you have the impression of a red spot on a black wall and in which, on this basis, you acquire the justified belief that there is a red spot on the wall. Now assume further that you receive the information that the wall is illuminated by red light. Under this condition, even a white spot on the black wall would appear redly. Relying on something's looking red when illuminated by red light is, therefore, not an epistemically acceptable source for believing that it actually is red. As soon as you discover that you are in such a situation you will thus acquire an undercutting defeater.

---

[21] For a more comprehensive account of defeaters see Grundmann 2010.

Do defeaters have to be beliefs that are themselves justified in order to do their epistemic work? This view might initially look quite attractive: if unjustified beliefs cannot confer justification, then unjustified beliefs cannot remove justification, either. But intuitive cases help to see that this cannot be the correct view. Here is one such case: John is a reliable perceiver of his environment. An evil neuroscientist manipulates his brain in such a way that John acquires the belief that all of his perceptual experiences are induced by an evil scientist. This belief is false and unjustified at the same time. Nevertheless, as soon as John has acquired this belief, all his perceptual beliefs are defeated and thus no longer justified. If this is correct, then defeaters need not possess any positive epistemic status in order to remove justification. Call these kinds of defeaters *doxastic defeaters.* Accordingly, if S has a belief that q that suggests that her belief that p is false or based on an unreliable source, then the justification of S's belief that p is overridden even if her belief that q is in fact false or does not have any positive epistemic status.[22] Besides, there are also defeaters which need not even be believed to do their work. Consider the case of a male detective who has finally gathered all the evidence to realize that the female suspect is not as innocent as she claims. Since he has fallen in love with her, he still doesn't believe that she is guilty. But since evidence to the contrary is available to him, he ought to have acquired a belief that suggests the falsity of his belief that the suspect is innocent. This fact overrides the detective's formerly justified belief in her innocence. We can call these kinds of defeaters *normative defeaters.* Accordingly, if S – in light of the evidence that is available to her – ought to acquire a belief that q that suggests that her belief that p is false or based on an unreliable source, then the justification of her belief that p is overridden.[23] In general, defeaters are either doxastic or normative.

---

[22] See Grundmann 2010, Lackey 2008, p. 45.
[23] See Grundmann 2010, Lackey 2008, p. 46.

If the challenge from disagreement can be explained in such a way that the recognition of disagreement constitutes a defeater of some kind, this would have the following advantage: it is common ground among internalists and externalists in epistemology alike that defeaters can remove justification. Hence, this explanation would be as metaepistemologically neutral as it should be.

As I said before, the challenge of disagreement is often explained within the framework of undercutting defeaters. We thus get something like the following version of the improved skeptical argument:

(1'') If we either believe or epistemically should believe that we are in a GSE-disagreement about p, then we acquire an *undercutting defeater* for believing that p.

(2'') In philosophy, we cannot escape from either believing or being epistemically obliged to believe that we are in GSE-disagreement about almost everything.[24]

Therefore,

(C2) In philosophy, we cannot escape from acquiring undercutting defeaters about almost everything and hence lack ultima facie justification for philosophical beliefs about almost everything.

---

[24] One might ask why premise (2'') makes a modal claim rather than an actual one. In other words, why isn't the claim that in philosophy defeaters are *in fact* ubiquitous sufficient to argue for skepticism about philosophy? The quick answer is that skepticism about justification is itself a modal claim, namely, that in nearby worlds we cannot achieve justified beliefs. In order to argue for a modal claim we need at least one modal premise. (Thanks to Joachim Horvath for raising this issue.)

Why should we think that (1'') is true, i.e., why should we think that the discovery of disagreement gives rise to an undercutting defeater? One such explanation is offered by Matheson 2009:

> "after the discovery of the disagreement, one becomes justified in suspending judgment as to what one's first-order evidence supports. One has found out that an epistemic peer disagrees as to what the first-order evidence supports, and this defeats whatever support one had for believing that the first-order evidence supports what one believed. (...) (I)t seems that such a discovery provides one with an *undercutting defeater* for one's first-order evidence by way of justifying one in suspending judgment toward higher-order propositions concerning what one's first-order evidence supports." (277)

Here is a more transparent reconstruction of Matheson's argument:

(M1)   Discovering that one disagrees with an epistemic peer about p removes one's justification for believing that one's own evidence justifies p (instead of what one's peer believes, i.e., not-p).

(M2)   If one has no justification for believing that one's evidence justifies p (instead of the contrary belief of one's peer), then this undercuts one's justification for p.

Therefore,

(MC)   Discovering that one disagrees with an epistemic peer about p undercuts one's justification for p.

(M1) should be best understood in such a way that peer disagreement is evidence for the fact that our second-order claim that our first-order evidence justifies our first-order belief p is unjustified.

In more technical terms: when we discover peer disagreement about p we acquire a *normative undercutting defeater* of our second-order belief. When we cannot decide whether our common body of evidence justifies my belief or my peer's belief, then the second-order belief that I am the one who is justified is no longer justified. The problem with (M1) is that it presupposes the truth of *uniqueness*, i.e. the claim that one and the same body of evidence cannot justify conflicting judgments for different agents. If a given body of evidence can justify two conflicting judgments, then we need not decide which one is justified. Both could be equally justified. But then the discovery of disagreement does not create an undercutting defeater for our second-order belief. Is uniqueness generally true or are there counterexamples to it? As it seems to me, we can quite easily construct an abstract counterexample to uniqueness. Let us assume that my peer and me have exactly the same body of evidence E at our disposal. To this body of evidence we apply two slightly different methods of balancing the evidence, M_1 and M_2. These methods both possess an equally good track record. The number of cases in which they lead from true evidence to true output beliefs is roughly the same in both cases.[25] But in the current situation, when applied to E, they lead to conflicting results. What is wrong with claiming that in such a case both results are equally justified? Since uniqueness is at least dubious, we have to conclude that the truth of (M1) is also doubtful.

But let us grant that (M1) is true and that we are thus not justified in believing that our evidence justifies our own first-order belief. Does this really undercut the justification of our first-order belief? Here is a reason to think that this is not the case. As pointed out above, undercutting defeaters suggest that the source of the relevant first-order belief was unreliable or untrustworthy. In contrast, Matheson just observes that when we discover peer disagreement we should not

---

[25] This does not require that in the past both always produced the same output beliefs. It may just be that the cases in which they produced false output beliefs on the basis of true evidence cancel each other out.

believe that our first-order belief is justified. But there is a crucial difference between positively *believing* that one's source is *unreliable* and merely *not being permitted to believe* that it is *reliable*. I do not see how the latter can satisfy the structural characteristics of an undercutting defeater. There is a desperate move that Matheson could make to save a slightly different version of his argument. He could replace the controversial premise (M2) by the following premise (M2*).

> (M2*) One is justified in one's first-order belief that p only if one is justified in believing that one's own evidence justifies that p.

In the resulting argument, the work of undercutting defeaters would be restricted to (M1). Now, the problem is that (M2*) looks like an internalist JJ-principle.[26] Using such a principle would obviously conflict with the adequacy condition of metaepistemological neutrality.

There is another recent attempt by Carey (2011) to defend the view that discovering peer disagreement generates an undercutting defeater. Here is what Carey says:

> "This, it seems, is why disagreement with a peer is so epistemically important: it provides an undercutting defeater for my evidence. I think it is precisely this kind of defeat story that makes plausible strong conciliatory views (...), as, since disagreement is with someone I have good reason to think is a peer, my evidence will always support that my reliability *given that we disagree* is 50%, which naturally supports suspension of judgment." (378)

---

[26] According to the JJ-principle, a particular first-order belief is justified only if the agent is justified in believing that his first-order belief is justified. Of course, unqualified versions of this principle would automatically generate a regress of justification.

Carey is certainly right in claiming that in the situation of peer disagreement the probability of my first-order belief being true is at most 50%.[27] The same holds for my peer. But does this really generate an undercutting defeater? Undercutting defeaters indicate that the belief-forming processes underlying the first-order beliefs are either generally unreliable or are executed under epistemically bad conditions that make the performance of a generally reliable process unreliable under the relevant circumstances. So, when someone has used the reading of tealeaves to make some prediction about the future and then is informed that her method is generally unreliable, she acquires an undercutting defeater. But the same is true when someone bases her beliefs about her surroundings on vision, but then learns that the lighting conditions were bad or that she was under the influence of a drug with hallucinatory side-effects. In these later cases, undercutting defeaters also explain why the performance of the process was unreliable. In contrast, discovering peer disagreement does neither indicate that nor explain why the relevant processes are unreliable or executed under epistemically unfavorable conditions. It rather indicates that one side must have made a mistake and that it is unclear which side it was. This surely speaks against the perfect reliability of the involved processes, but it does not speak against a degree of reliability that is far above the relevant threshold.[28] It seems perfectly conceivable that two reliable processes of equal (but not perfect) quality lead to conflicting results about a certain proposition p. But why can't one just say that the reliability of a process *under the condition of peer disagreement* is relevant to the epistemic evaluation of its performance here? Understood this way, the degree of the process' reliability seems clearly too low to count as reliable in the absolute sense. There are at least two problems with this view. First, when we attribute reliability

---

[27] Why at most? If the peers hold conflicting beliefs that are not plainly contradictory, both beliefs might be false.
[28] You might get the impression here that by talking about the reliability of the relevant processes I have already violated the adequacy condition of metaepistemological neutrality. But notice that even internalists accept that if *from one's own perspective* the relevant reliability is low, then prima facie justification will be removed.

to a process not in general, but under relevant conditions of its execution, then we restrict ourselves to those conditions which are *causally operative* in bringing about the judgment (e.g. taking drugs or obtaining of bad lightning conditions). In contrast, the fact of peer disagreement makes no causal contribution to bring about the conflicting judgments – it rather is the result of their being brought about. The disagreement condition therefore seems irrelevant to the assessment of the reliability of the respective first-order beliefs. Second, this view would have an extremely counterintuitive consequence, namely that rebutting defeaters turn into undercutting defeaters. Consider a case in which you have the prima facie justified visual belief that a particular stick is bent and a prima facie equally justified tactile belief that the stick is not bent. It seems natural to say that the tactile belief constitutes a rebutting defeater for the visual belief here. But according to the view at hand, we should rather say that *relative to the conflict between the epistemically equally trustworthy sources vision and touch* the probability of our judgment is at most 50% and hence we acquire an undercutting defeater for our prior belief.

On the one hand, it has become clear that none of the extant explanations of the challenge of disagreement really works – neither the argument from reliability nor the diagnosis of an alleged decision problem nor even the claim that the discovery of peer disagreement constitutes something like an undercutting defeater. On the other hand, it also seems intuitively clear that the discovery of peer disagreement removes our previous justification in some way or other. So, what is the correct explanation of this epistemic phenomenon? According to my own view, the discovery of peer disagreement indeed gives rise to a defeater. But the relevant defeater is of the *rebutting* rather than of the undercutting kind. This explanation also has the advantage of being metaepistemologically neutral, since, as I said before, both internalists and externalists accept that defeaters can remove prima facie justification. The basic idea behind the **rebutting defeater**

approach is the following.[29] Assume that I am prima facie justified in believing that p on the basis of my evidence. When I then discover that a peer disagrees with me about p, I come to believe that someone whom I take to be equally competent on this matter asserts 'non-p.' On this basis I come to believe that there is some testimonial evidence against p. It is highly controversial in the current debate about testimony whether, in general, one needs some information about the testifier's trustworthiness or reliability in order to acquire justified testimonial beliefs on the basis of her testimony.[30] But it seems clear that when you already have some prima facie justified belief about a certain matter, you should respect testimonial evidence to the contrary only to the extent to which you take the testifier to be reliable on the matter at hand. Given that this condition is satisfied, then my belief that my opponent's assertion of 'non-p' is equally reliable than my own belief removes all my prior justification for p. It does so by relying on the information that some equally reliable testimony speaks for the falsity of p. But this just looks like a rebutting defeater, then. From my perspective the truth and the falsity of p seem equally supported by my evidence in that situation and, therefore, all prior justification is neutralized.

The crucial point in the previous argument is the following: when we discover peer disagreement we acquire a defeater only if the peerness assumption makes it plausible that our opponent is equally reliable in getting at the truth. If epistemic peerness were completely dissociated from reliability, as it may be on some internalist readings of evidence and competence, then the discovery of peer disagreement would not support my belief that there is something that speaks against the truth of my prior belief. Hence, we no longer would acquire a rebutting defeater. Notice that the same is true when we understand the challenge from disagreement as providing us with an undercutting defeater. According to this view, I come to believe that I am mistaken about

---

[29] Thanks to Joachim Horvath for prompting the following clarifications of my view.
[30] Reductionists claim that one needs such additional information, whereas anti-reductionists deny it.

the controversial proposition with a probability of 50% when I discover disagreement with a peer. But why should we think that peer disagreement implies a probability of being mistaken of 50%, unless we assume that peers are equally reliable? Therefore, I conclude that epistemic peerness has to be related to reliability in order to play any role in defeating justification. This is true for externalist and internalist views alike. Therefore, we can revise the original definition of epistemic peerness in such a way that two agents are epistemic peers with respect to proposition p if and only if they are equally reliable in their judgments concerning p. For the reason given above this revision does not conflict with the adequacy condition of metaepsitemological neutrality. We are now in a position to formulate the rebutting defeater argument more precisely:

(G1)   I am prima facie justified in believing that p. (assumption)

(G2)   I believe that there is someone who is my epistemic peer about p and who expresses her belief that non-p. (based on the discovery of peer disagreement)

(G3)   Epistemic peers about p are (roughly) equally reliable about p. (by the definition of 'epistemic peer')

(G4)   What I believe about my peer speaks as strongly for the truth of non-p as my prior evidence speaks for the truth of p. (from G1, G2, G3)

(G5)   If (G4) is true, then the prima facie justification of my belief that p is neutralized by a doxastic rebutting defeater. (from the definition of doxastic defeaters)

(G6)   Ultima facie, I am no longer justified in believing that p. (from G4, G5)

In this argument, the discovery of peer disagreement is understood as providing us with a rebutting defeater of the doxastic kind. So we only need to *believe* that we are disagreeing with a peer in order to acquire a defeater. This gives us the final version of the skeptical argument:

(1''') If we believe that we are in a GSE-disagreement, then we acquire a (doxastic) rebutting defeater for believing that p.

(2''') In philosophy, we cannot escape from believing that we are in GSE-disagreement about almost everything.

Therefore,

(C3) In philosophy, we cannot escape from acquiring rebutting defeaters to almost everything and hence we lack ultima facie justification for almost all our philosophical beliefs.

In this first part of my paper, I have made a systematic suggestion how we should understand the challenge from peer disagreement in philosophy. If this is basically on the right track, this challenge relies on the epistemic role of defeaters, or, more precisely, on the role of rebutting defeaters. In the second part of this paper, I will consider possible replies to this challenge.

II

You might think that there simply cannot be a convincing reply to the final version of the skeptical argument. In some sense it seems invulnerable. Since one finds disputes among people who *take* each other as peers almost everywhere in philosophy, it seems unavoidable to acquire rebutting defeaters to almost every philosophical belief (unless one philosophizes in complete isolation). As long as *beliefs about peerness* are sufficient to generate doxastic defeaters, you

might think, it is simply irrelevant whether these beliefs are also justified. Surprisingly, much of the debate has focused on this later question. People tend to think that suspension of judgment is rationally required only if you have *good reason* to believe that your opponent is your peer. Consequently, the defenders of philosophical knowledge try to argue that the peerness assumption is not justified, whereas the skeptics insist that the peerness assumption is reasonable. One might wonder whether this whole debate rests on a mistake, namely that defeaters actually have to be justified in order to do their epistemic work As we have seen above, this assumption seems incorrect. Defeaters need not themselves be justified in order to remove prima facie justification. Purely doxastic defeaters are already sufficient for that.

Now, many philosophers think that suspension of judgment is rationally required in cases of peer disagreement only if the opponents are justified in taking each other as peers. Here are some representative voices on the matter:

> "conciliation is required *only to the extent* that the dispute-neutral assessment provides one with a *strong positive reason* for thinking the other person likely to have gotten things right." (Christensen 2009, 760, my italics)[31]

> "It is uncontroversial among both skeptics and non-skeptics that peer disagreement is a threat to rational belief *only if* the rational subjects have *reason* to think that it obtains. (...) (R)easonable belief that a peer disagreement obtains is doing the epistemic work." King 2011, 13f)[32]

---

[31] See for a similar commitment also Christensen 2011, 15f.
[32] Notice that King himself dismisses this view in his 2011 paper.

"The cleaner question, then, is that of how we should respond to disagreement with those

we *justifiably* take to be our peers." (Enoch 2010, 972)[33]

But if there are purely doxastic defeaters, then justified defeaters are not actually required. So, the

above claims seem much too strong. You might therefore think that the whole debate about the

epistemic status of the peerness assumption is irrelevant. But that would be a mistake, too.

Justification does not play any role for the truth of premise (1''') of the skeptical argument (since

doxastic defeaters need not be justified), but it plays a crucial role for the truth of premise (2''').

To see this, assume that in philosophy everybody permanently believes that all his first-order

beliefs are disputed by peers. Further, assume that this ubiquitous second-order belief is false and

unjustified. Then it is certainly not true that philosophers cannot escape from having defeaters to

almost everything, as (2''') claims. For, if philosophers were to rationally adjust their beliefs

about peerness, they immediately would lose most of their defeaters.[34]

But then skepticism would be unreasonable. Consider, for example, the analogy of someone who

has the *unjustified* and false belief that an evil demon systematically deceives her. Although this

belief constitutes a doxastic undercutting defeater, it does not actually call for skepticism. If the

victim of the alleged deception were to rationally adjust her beliefs to the evidence available to

her, then she would no longer believe in the global deception and thereby lose her defeater. We

can conclude that defeaters are inescapable only if they are based on good reasons. In order to

---

[33] Carrey 2011, 374, is not quite that clear when he says: "whenever I disagree with someone I have good
(undefeated) reason to think is an epistemic peer with equally good evidence, I should suspend judgment on the
disputed proposition." Strictly speaking, Carrey does not claim that suspension is rational *only if* the peerness
assumption is justified. One could understand him also as making only an uncontroversial sufficiency statement. Of
course, if doxastic defeaters are possible, then the appropriate justified beliefs will also suffice to provide us with a
defeater.

[34] One might wonder whether it is always easy to find out whether one's defeaters are based on good reason. But this
is not the crucial point here. If one can argue, as I will try in the following, that there are no good reasons supporting
the peerness assumption, then this insight alone would suffice to dispense the challenge from disagreement. (Thanks
to Joachim Horvath for raising this worry.)

assess the second premise of the skeptical argument we therefore have to check whether the peerness assumption is justified in philosophy.

The skeptic about philosophy has to claim that we have *good reason* to believe that there is genuine and roughly symmetric disagreement among philosophical peers that concerns almost everything and persists through time. Certain aspects of this claim can be easily assessed on an empirical basis. We can just look at the history of a debate to find out how widespread, stable and symmetric the disagreement was. Admittedly, the question of whether a given disagreement is genuine or merely verbal is much more difficult to answer. We are still looking for a generally workable criterion of genuine disagreement. But let us put this issue aside for now and assume that most philosophical disputes are genuine. The really hard case then is assessing epistemic peerness. At the beginning of this paper, I defined epistemic peerness such that two epistemic agents, A and B, are epistemic peers with respect to a certain domain D iff (i) A and B have access to the same evidence relevant to D and (ii) A's and B's abilities to assess the relevant evidence are equally competent. Are we able to establish that disagreeing philosophers have access to the same relevant evidence? There seem to be two obstacles to the satisfaction of this condition. Even if philosophical opponents know all the relevant arguments and considerations from the literature and in addition have exchanged their personal views continuously over the years, it is highly likely that they do not share all the relevant evidence. First, there seems to be essentially private evidence that cannot be shared through communication. Assume that it is controversial between my doctor and me whether I feel a stabbing pain in my stomach. I strongly believe that I have this kind of pain because I feel it. My evidence for having the belief about my pain is my having this pain. My doctor thinks that I am either simulating to get a medical prescription or that I am a hypochondriac since he cannot find any results that would explain this

pain. Now, when I communicate my pain to my doctor he thereby does not obtain the same evidence that I have when I feel my pain. He rather obtains a mere *description* of my evidence, but not that evidence itself. I will call a piece of evidence that cannot be shared by describing it to other people 'private evidence.' In philosophy, it is widely held that intuitions and insights play a large evidential role. But this is also a kind of private evidence. We can mention it to our opponents, but they do not thereby come to share it with us. So there seem to be clear limits of sharing evidence in philosophical disputes because at least some of the relevant evidence is private. Second, there is also a practical limit to the relevant evidence that openly lies on the table for all opponents. Our judgments are always determined by explicit evidence as well as by implicit background evidence that is sometimes inaccessible even to ourselves. But then, of course, opponents eo ipso do not share their relevant background evidence because they cannot communicate to others what is even hidden to themselves.

You might think that these *limits of shared evidence* clearly establish that we never have good reasons to assume that we are disagreeing with a peer.[35] But this conclusion would be too quick, for we can redefine the relevant notion of epistemic peerness in such a way that the problem dissolves. The crucial feature of epistemic peers that turns their dissenting views into rebutting defeaters is the equal reliability of their judgment about the disputed issue.[36] If an opponent uses the same evidence and her assessment of the evidence is equally conditionally reliable, then her judgment will be equally reliable simpliciter. But there are other ways of arriving at a result that is of equal epistemic quality. For example, the opponent might have evidence of the same quality,

---

[35] Compare Sosa 2010.

[36] One might be worried here that this is much too reliabilist to pass the adequacy condition of metaepistemological neutrality. Remember, however, that I have already argued for the claim that the discovery of peer disagreement can provide us with a defeater only if peerness is analyzed in terms of reliability. Moreover, when we talk about defeaters, the opponents need not be *in fact* equally reliable, but rather have to be *taken* as peers with respect to their reliablity. (Thanks to Joachim Horvath and Jens Kipper for pressing me on this point.)

although it is not strictly the same evidence.[37] Or the opponent might have evidence of a lower quality, but her ability to assess the evidence is so much better than our own that this compensates for the poor quality of her evidence. If any of these scenarios obtains, then we may have a reason to believe that we are disagreeing with a peer, even though our opponent does not strictly possess the same evidence as we do.[38]

The whole issue therefore comes down to the question whether we have good reasons to regard our opponent as someone who is of equal epistemic standing on the disputed matter. How can I justify that my opponent is epistemically on a par in this respect? There are in principle four different ways to justify the peerness assumption: (i) justification by default, (ii) the track record method, (iii) arguing from widespread background agreement, (iv) arguing from general philosophical competence. In what follows, I will argue that none of these ways is likely to succeed in philosophy.

(Ad i) Do we have a default justification for our peerness assumption? That would mean that we are justified in regarding our opponent as our peer as long as nothing speaks against her peerness. In practice, this would reverse the burden of proof. Assume that on an international conference some historian presents a lucid interpretation of some historical facts. At the end of her talk someone in the audience comments that he doubts this interpretation on the basis of having thought hard about these issues himself. He does not give any reasons for his view. Moreover, nobody knows who this person is. If default justification were in place here, the historian would

---

[37] Compare Christensen 2007, 211; Thune 2010; King 2011.

[38] However, when we deviate from the original definition of peerness in this way, it seems impossible to show that the conditions of peerness are satisfied prior to (or independent of) assessing the reliability simpliciter of the opponents. For example, how can we know that our opponent has equally good, though not identical evidence? In order to acquire this knowledge we have to presuppose that the conditional reliability of the ability to assess the evidence is roughly the same and in addition we have to know that the reliability simpliciter of the opponents' judgment is equal to ours. Only on the basis of this prior knowledge can we acquire knowledge of the equal quality of the opponents' evidence.

have to suspend her judgment on the issue. By default, she would have to assume that the commentator is her epistemic peer – at least unless she acquires evidence to the contrary. But this implication strikes me as highly counterintuitive.[39] A default justification of the peerness assumption therefore seems implausible.

(Ad ii) In many domains we assess the reliability of agents or methods on the basis of their track record. We simply check how often a person or a method was correct in the past and then inductively project this past truth ratio into the future. Here are some typical cases. When Paul and Paula are doing mental arithmetic, they can later check their reliability against the results provided by a calculator. In the sciences, theories imply predictions which can retrospectively be tested on the basis of experience. Halley, for example, predicted in 1705 on the basis of Newton's laws that the famous comet that finally got his name would return in 1758 – at that time a bold prediction that was later empirically confirmed. In this way, we also can assess the reliability of the scientific methods that gave rise to the theories in question. We do not always use independent methods to establish a track record, however. Sometimes it suffices when the opponents later agree who was right earlier. Can we apply the track record method to philosophy? This seems highly questionable for the following two reasons. First, we do not possess independent methods to assess the truth of philosophical beliefs. Second, and more importantly, philosophical disagreement is typically persistent over time. So there typically is no later point in time at which former opponents agree who was wrong earlier.

(Ad iii) There seems to be another good indicator of equal reliability. If in the relevant domain the disagreement occurs against the background of widespread agreement, then this makes an equal reliability of the controversial judgments likely. The initial appeal of this line of argument

---

[39] See also Christensen 2011, 15f.

relies on two assumptions. First, the background agreement would have to be quite extensive such that only rather local disagreements occur. Second, the agreement would have to obtain in the relevant domain of the dispute. For example, it would not help if philosophers who disagree about how objects persist through time agree on philosophical matters outside of metaphysics. Adam Elga, on the one hand, claims that disputes in philosophy are always pervasive in the relevant domain since the disputed propositions are never fully isolated.[40] Hilary Kornblith, on the other hand, insists that disputes in philosophy are much more local than Elga believes.[41] In the end, this appears to be an empirical question which clearly would deserve more attention. But what I will argue now is that even if there were a widespread background of agreement in the relevant domain, this would still not make it likely that the conflicting judgments are equally reliable.

Here is a first worry. Take a case of pretty broad agreement in the relevant philosophical domain. Assume that two philosophers, A and B, agree in 70% of their answers to philosophical questions in this domain. Further assume that only 60% of these answers are in fact correct. But whenever A and B disagree on questions in that domain, B will give the correct answer. Therefore, whereas A's judgments in this domain have an overall reliability of 42%, the reliability of B's judgment amounts to 72%. But now we have a situation in which, although A and B share a background of wide agreement in the relevant domain, their degree of reliability differs to such an extent that B but not A is basically reliable, at least if we assume that the threshold for basic reliability is .7.

---

[40] Elga distinguishes between what he calls 'the clean cases' and 'the messy real-world cases.' In the clean cases, as the Restaurant case, there is enough agreement in the background to assess the oponent to be one's peer, whereas in philosophy "one's reasoning about the disputed issue is tangled up with one's reasoning about many other matters." See Elga 2007, 492.

[41] Kornblith takes the example of epistemological disputes to show that although in philosophy disagreement is always broader than in the Restaurant case it is typically not global. "Even these disagreements (on epistemological internalism and externalism, TG), broad as they are, take place against a background of very broad agreement, agreement about important issues in epistemology, about which positions are worth taking seriously, about what counts for and against various views, and so on." (Kornblith 2010, 51)

Since there can be differences in basic reliability even if there is broad agreement in the relevant domain (as the above example illustrates), such a difference is even more likely for lower degrees of agreement.

Here is another worry. Even if widespread agreement were a good indicator of the agent's equal reliability in the relevant domain, it would not tell us anything about the reliability of the specific methods that were used in answering the question at hand. It is surely plausible that one and the same philosopher applies different methods with different degrees of reliability to the same domain of philosophical questions. As long as the individuation of philosophical methods remains unclear, we do not have the means to assess their reliability on the basis of the general truth-ratio of beliefs about the relevant domain.

(Ad iv) There is a further way to defend the peerness assumption that looks quite attractive at first glance. It goes like this: whenever you discover that a philosophical view that competes with yours is held by someone who works as hard and carefully on the subject matter as you do, and who is as intelligent, logically skilled and philosophically creative as you are, and who has access to roughly the same evidence as you have, and who is as intellectually honest as you are, then all of this seems to be an excellent indicator that she is your epistemic peer.[42] Is there anything wrong with considerations of this kind? I have two worries here. First, remember that the epistemically relevant notion of a peer is highly specific. Someone is my epistemic peer only if she applies an equally reliable method in the relevant domain. Now, working hard and carefully on philosophical matters, possessing general intelligence and being intellectually honest are much too coarse-grained to be sufficient for the required more specific competencies. Second, and even more importantly, the skills mentioned above are not all clearly veritistic competencies. One

---

[42] Compare Goldberg 2009, 106 for a similar formulation.

might be extremely honest, intelligent and creative without being particularly reliable with respect to the disputed proposition. These worries both emphasize that general philosophical competence along the lines above is not a good indicator of the relevant reliability

Two examples may illustrate this. First, consider the British empiricist George Berkeley. Nearly everybody takes him to be a highly competent philosopher – especially with respect to those skills listed above as general philosophical competencies. At the same time, almost nobody believes that Berkeley's view about reality was mostly correct. In contrast, most people find his phenomenal idealism radically mistaken. From their perspective, it is highly unlikely that Berkeley was reliable on philosophical matters. Second, consider two competing philosophers of modality, namely David Lewis and David Armstrong. Experts in the field would judge both of them as philosophically highly competent. But they entertain radically conflicting views about modal reality. If one of them gets modal reality basically right, then the other will be radically mistaken.[43] But then the following must be true: although they are peers in terms of general philosophical competence they cannot be peers in terms of veritistic competence, i.e. reliability. These two cases illustrate that we are committing the *fallacy of confusing competencies* when we infer from being generally philosophically competent to possessing a high veritistic competence.

**

I conclude with wrapping up the main results of this paper. It has turned out that the challenge from peer disagreement should be best understood in terms of acquiring a rebutting defeater. Although defeaters in general need not be justified in order to remove justification, the discovery of disagreement provides us with an invulnerable defeater only if the peerness assumption is reasonable. If, on the other hand, defeaters are escapable we are not forced into the modal claim

---

[43] Here I can stay neutral on the question of whether modal realism or actualism is the correct modal view.

of skepticism, namely that we *cannot* have justified beliefs. The skeptical challenge to philosophy therefore depends on our ability to justifiably attribute epistemic peerness to philosophical opponents. Arguably, the only plausible basis for such an attribution is the track record method. Whereas in the sciences such track records are often available, and usually no challenge to their epistemic standing, they are typically out of reach in philosophy. Interestingly, it is the very same feature of philosophical disagreement, namely its stability, which makes the situation look dramatic and, at the same time, shields us from its skeptical consequences. Widespread and stable disagreement in philosophy, therefore, does not directly lead to skepticism about philosophy. In the face of controversy, competing philosophers both may rationally retain their initial judgments.[44]

ReferencesAnderegg, W.; Prall, J.; Harold, J. & Schneider, S. 2010, "Expert Credibility in Climate Change," *Proc. Natl. Acad. Sci. U.S.A* 107: 12107-9.

Brenan, J. 2010, "Scepticism about Philosophy," *Ratio* 23: 1-16.

Carey, B. 2011, "Possible Disagreements and Defeat," *Philosophical Studies* 155: 371-381.

Chalmers, D., Manley, D. & Wassermann, R. 2009, *Metametaphysics. New Essays on the Foundations of Ontology*, Oxford: Oxford University Press.

Christensen, D. 2007, "Epistemology of Disagreement: The Good News," *Philosophical Review* 116: 187-217.

Christensen, D. 2009, "Disagreement as Evidence. The Epistemology of Disagreement," *Philosophy Compass* 4/5: 756-67.

Christensen, D. 2011, "Disagreement, Question-Begging and Epistemic Self-Criticism," *Philosopher's Imprint* 11:1-22.

Elga, A. 2007, "Reflection and Disagreement," *Nous* 41: 478-502.

Enoch, D. 2010, "Not Just a Thermometer: Taking Oneself Seriously (but not Too Seriously) in Cases of Peer Disagreement," *Mind* 119: 953-97.

Feldman, R. 2006, "Epistemological Puzzles about Disagreement," in *Epistemology Futures*, ed. S. Hetherington, Oxford: Oxford University Press, 216-36.

Garwood, C. 2007, *Flat Earth: The History of an Infamous Idea,* New York: Macmillan.

Goldberg, S. 2009, "Reliabilism in Philosophy," *Philosophical Studies* 142: 105-17.

Goldberg, S. forthcoming, "Disagreement, Defeat, and Assertion," in *The Epistemology of Disagreement*, eds. J. Lackey, D. Christensen, Oxford: Oxford University Press.

Goldman, A. 1979, "What Is Justified Belief?" in *Justification and Knowledge*, ed. G. Pappas, Dordrecht: Kluwer, 1-23.

Grundmann, T. 2010, "Defeasibility Theories," in *The Routledge Companion to Epistemology*, eds. S. Bernecker, D. Pritchard, London: Routledge, 156-66.

Hume, D. 1975, *Enquiries Concerning Human Understanding and Concerning the Principles of Morals*, ed. by L. Selby-Bigge, Oxford: Clarendon Press.

Kelly, T. 2005, "The Epistemic Significance of Disagreement," *Oxford Studies in Epistemology* 1: 167-96.

King, N. 2011, "Disagreement: What's the Problem? Or a Good Peer is Hard to Find," *Philosophy and Phenomenological Research:* 1-23.

Kornblith, H. 2010, "Belief in the Face of Controversy," in *Disagreement*, eds. R. Feldman & T. Warfield, Oxford: Oxford University Press, 29-52.

Lackey, J. 2008, *Learning from Words. Testimony as a Source of Knowledge*, Oxford: Oxford University Press.

Matheson, J. 2009, "Conciliatory Views of Disagreement and Higher-Order Evidence," *Episteme* 6: 269-79.

Sosa, E. 2010, "The Epistemology of Disagreement," in *Social Epistemology*, eds. A. Haddock, A. Millar, D. Pritchard, Oxford: Oxfford University Press, 278-97.

Thune, M. 2010, "'Partial Defeaters' and the Epistemology of Disagreement," *The Philosophical Quarterly* 60: 355-72.

van Inwagen, P. 1996, "Is It Wrong Everywhere, Always, and for Anyone to Believe Anything on Insufficient Evidence?" in *Faith, Freedom, and Rationality: The Philosophy of Religion Today*, eds. J. Jordan, D. Howard-Snyder, Lanham/London: Rowman & Littlefield, 137-53.