

Forthcoming in: *Episteme*

Why Disagreement-Based Skepticism Cannot Escape the Challenge of Self-Defeat

Thomas Grundmann, University of Cologne

[penultimate version]

Abstract. Global meta-philosophical skepticism (i.e. completely unrestricted skepticism about philosophy) based upon disagreement faces the problem of self-defeat since it undercuts its motivating conciliatory principle. However, the skeptic may easily escape this threat by adopting a more modest kind of skepticism, that will be called “extensive meta-philosophical skepticism”, i.e., the view that most of our philosophical beliefs are unjustified, except our beliefs in epistemically fundamental principles. As I will argue in this paper, this kind of skepticism is well-motivated, does not undercut the conciliatory principle, but still poses a radical challenge to philosophy as a cognitive discipline. Moreover, I will argue that non-global skepticism that is still extensive undermines itself as well. The deeper reason for this is that this more modest kind of skepticism can only be motivated by the assumption that disagreement with philosophical peers is abundant and that we can identify peers only by relying on track-record arguments. But then one can argue for extensive meta-philosophical skepticism only if one presupposes that those philosophical beliefs that form the basis of track-record evaluations are justified. Here, the threat of self-defeat looms again. I will proceed by first defending the premises of this new anti-skeptical argument against standard objections from the literature. Second, I will show in more detail where the epistemic inconsistency arises in the argument for extensive meta-philosophical skepticism. I conclude with an assessment of the scope and the limits of my argument.

1. Introduction

There are domains of judgment in which disagreement occurs regarding almost every substantial question, not only among isolated thinkers but also among larger groups. In these domains the disagreement does not seem to be resolved over time, even if the opponents are exposed to almost the same body of relevant evidence and arguments. This kind of *ubiquitous* and *intractable* disagreement is a common phenomenon in matters of morality, religion, politics, and aesthetics and is also very familiar in disciplines like philosophy. In contrast, scientists tend to converge much more in their views, at least over time (Kornblith 2010). One kind of reaction to this observation about ubiquitous disagreement is the skeptical worry that disciplines such as philosophy provide us with little, if any, knowledge or justified belief. Although this general worry has been around for quite some time, detailed arguments for skeptical consequences that arise from disagreement have been developed only recently.

Arguments of the sort I have in mind usually start with our natural responses to ordinary yet somewhat idealized cases. Suppose you go out for dinner with some friends at a restaurant. After dinner, you decide to split the bill evenly. Since no calculator is available, you and your friend Paula do some mental math. In most cases of mental calculation, the two of you agree about the result. But in cases where you don't agree, you and Paula arrive at correct results equally often. Suppose further that in the *Restaurant Case* the two of you arrive at slightly different results about each party's share. Only one of you can be correct. Neither of you is drunk, distracted, or particularly tired, so there is no obvious reason why one rather than the other should have made the mistake. In such a situation, the natural verdict is to say that both of you should suspend judgment about the shares as soon as you realize your disagreement (Christensen 2007). Why is this reaction rationally required? A simple explanation goes as follows: Both of you have tried to find a solution to the same math problem. When you notice the disagreement you know that *at most* one of you has arrived at the correct result. But you also know that *at least* one of you has calculated incorrectly. Since you also know, independently of the present dispute, that the two of you are equally good at mental math, you have acquired new evidence that suggests that the probability of your having committed a performance error is at least 50%. This undermines the justification of your initial belief about the shares, and it does so even if you have in fact calculated correctly. Here is a more abstract characterization of your epistemic situation: If you realize that you disagree with someone who possesses the same relevant evidence as you do and who is—as assessed from your perspective, although independently of the current controversy—equally competent in her reasoning based on this evidence, then rationally, you should suspend judgment. Thus, known disagreement with a strong epistemic peer (i.e., someone who shares one's evidence and has the same domain-relevant reasoning competence) demands a strongly conciliatory reaction, i.e., suspension of belief.

However, we need not rely on recognition of strong epistemic peerness to motivate the rationality of strong conciliatory reactions to disagreement, as the following *Election Case* illustrates. Suppose there is a very close presidential race between two competing candidates, X and Y. After the people have voted, two equally renowned election forecasters using different, but equally representative samples come up with conflicting predictions regarding who won the election. Forecaster A predicts that X has won the election by a very narrow margin. Forecaster B predicts the same of candidate Y. If these facts are disclosed to A and B, the rational reaction would once again seem to be suspension of judgment for both. But in contrast to the *Restaurant Case*, A and B do not fully share the same relevant evidence. They

clearly rely on different inductive evidence, since they use different samples of voters for their predictions. It would be more appropriate to claim that they have *equally good, though not identical* inductive evidence. Why is suspension of judgment rationally required in this case? Here is a simple answer: Suppose that A's initial judgment that candidate X had won the election was justified. By recognizing that someone who has equally good evidence and who is equally reliable in her predictions made a different judgment about the results of the election, A acquires defeating evidence that is as strong as the evidence that initially justified her judgment. Since this new evidence speaks against the truth of what A believed, a fair weighing of the equally strong pro and con evidence would rationally lead A to suspend her belief. Whereas it is crucial to the Restaurant Case that you note the sufficient likelihood of your own performance error, it is crucial to the Election Case that after having been exposed to a disagreeing weak peer (someone who has equally good rather than identical evidence and who is an equally good reasoner), A has equally strong evidence for and against her initial belief.

Reflection on the two cases suggests that whenever we face disagreement with either strong or weak epistemic peers, we acquire new evidence that *defeats* our initial justification, if there was any. Whereas in cases like the Restaurant Case it is an undercutting defeater, in cases like the Election Case we acquire a rebutting defeater. But how does this strongly conciliatory position lead to extensive skepticism about philosophy? Typically, philosophers themselves characterize their discipline as a battlefield of persistent disagreement and controversy (e.g., Feldman 2006: 217, Goldberg 2013: 181–83). One might, therefore, claim that with respect to almost every philosophical belief that you take to be justified you will find some epistemic peer who disagrees with you.¹ Admittedly, it might be hard to find strong epistemic peers apart from highly idealized situations, such as the Restaurant Case. Even if philosophical opponents have been exposed to the same relevant data and arguments, they still need not share *all* of their relevant private and background evidence. However, even if strong peers can hardly be found in philosophy, weak peers may be more abundant—and, as the Election Case suggests, recognizing disagreement with weak peers seems to be sufficient to motivate suspension of belief.

On the basis of the above considerations one can put forward the following *argument for global meta-philosophical skepticism*:

¹ To simplify matters I ignore here the size of the relevant groups that agree with you or with your peer. Strictly speaking, these groups should be of roughly the same size to motivate a conciliatory response. See my [reference omitted].

- (P1) Whenever I have dispute-independent reasons for believing that the subject I am disagreeing with is my epistemic peer, I am rationally required to suspend judgment.
(strong conciliationism)
- (P2) For any philosophical proposition that I believe, I have dispute-independent reasons for believing that at least some of the subjects I am disagreeing with are my epistemic peers.
- (C) Therefore, I am rationally required to suspend all philosophical beliefs.

Global skepticism about philosophy is a view that is currently endorsed by some philosophers (e.g., Feldman 2006, Goldberg 2009, 2013, 2015, Brennan 2010, Kornblith 2010, Ribeiro 2011, Machuca 2013). But there are also many philosophers who try to resist the skeptical challenge by objecting to one of the argument's premises: Proponents of the Steadfast View, as e.g. Kelly 2005, 2010, Lackey 2010, Sosa 2010, dispute (P1); others dispute (P2) by either arguing that one often has symmetry-breaking information introspectively (Lackey 2010, Sosa 2010) or by arguing that as philosophers we are often agnostic with respect to whether our opponents are our peers (Elga 2007, King 2012, Grundmann 2013).

In this paper I will try out a different anti-skeptical strategy. I will argue that the argument for (meta-philosophical) skepticism based on disagreement is epistemically self-undermining. More specifically, I will show that if the conclusion of skepticism is true, one of the essential premises of the skeptical argument cannot be justified. Hence, the skeptical argument proves to be epistemically inconsistent—i.e., if its conclusion is true, one of its premises is unjustified, while if, on the other hand, all of its premises are justified, then its conclusion cannot be true. In section 2, I will consider a simple and well-known argument that supports that global skepticism is self-undermining. I will argue that this argument is not sufficient to undermine all kinds of extensive skepticism based on disagreement. In section 3, I present a new anti-skeptical argument that also addresses more moderate versions of disagreement-based skepticism. In sections 4 and 5, I defend two *prima facie* controversial premises of this argument. In section 6, I show in more detail where the epistemic inconsistency arises in the argument for extensive meta-philosophical skepticism. Finally, I conclude with some remarks about the scope and the limits of the new argument.

2. A simple argument to undermine global skepticism and why more than that is needed

Global justification skepticism is the view that there are *no* justified beliefs (in a given domain). In contrast, extensive skepticism claims that *only very few* beliefs are justified. Even a quick look at the above argument for global meta-philosophical skepticism reveals that this argument is self-undermining: If its conclusion correctly claims that I am rationally required to suspend all philosophical beliefs, then I am rationally required to suspend my belief in strong conciliationism, i.e. premise (P1) of the above argument, which clearly makes an epistemological and thus philosophical claim. Hence, the standard argument for global meta-philosophical skepticism is self-undermining. However, this problem can be fixed by revising the conciliatory principle such that it does not apply to cases of disagreement about conciliationism itself. We then may get the following revised *argument for extensive meta-philosophical skepticism*:²

(P1*) Whenever the controversy is not about epistemically fundamental principles and I have dispute-independent reasons for believing that the subject I am disagreeing with is my epistemic peer, I am rationally required to suspend judgment. (partial conciliationism)

(P2) For any philosophical proposition that I believe, I have dispute-independent reasons for believing that at least some of the subjects I am disagreeing with are my epistemic peers.

(C*) Therefore, I am rationally required to suspend all philosophical beliefs unless they refer to epistemically fundamental principles.

The revised argument is no longer self-undermining since its conclusion denies any implications for the rational belief in epistemically fundamental principles as conciliationism.³ Is the exemption of the conciliatory principle from conciliationism well-motivated and non-arbitrary? Elga (2010) argues that it is. He believes that one should entertain a dogmatic attitude towards *all* epistemically fundamental principles, rules and policies in order to avoid the possibility of incoherent advice and self-undermining. Since conciliationism is only one among many of those principles this strategy is not purely ad hoc, i.e., just designed to save conciliationism.⁴ Let us assume here that Elga's defense of partial conciliationism, i.e. a version of conciliationism that rules out its application to controversies about conciliationism

² All revisions are underlined.

³ I just mention here the worry that it may well be controversial whether conciliationism articulates a *fundamental* epistemic principle.

⁴ For criticism of Elga's strategy, see Christensen 2013 (88–89) and Decker 2014. For further defense, see Pittard 2015.

itself, is viable. We can then use the above argument to motivate extensive skepticism about philosophical beliefs—an argument that does not run the risk of blatant self-defeat. Although extensive skepticism is not absolutely global (since it does not apply to the philosophical belief in conciliationism), it is overall almost global (since all philosophical beliefs except beliefs in fundamental epistemic principles are threatened) and absolutely global for all philosophical sub-disciplines that are different from epistemology. If this is correct, then we cannot use the simple strategy of self-defeat to avoid extensive meta-philosophical skepticism. In the balance of this paper, I will argue that a subtler version of self-defeat threatens even arguments for a extensive meta-philosophical skepticism that is not absolutely global.

3. The new anti-skeptical argument

Let me start with a sketch of my anti-skeptical argument:

- (P3) Disagreement-based arguments justify extensive skepticism about philosophy only if one justifiedly believes that one's opponents are weak epistemic peers.
- (P4) In philosophy, one justifiedly believes that one's opponents are weak epistemic peers only if one justifiedly believes the premises of relevant track-record arguments.
- (P5) One justifiedly believes the premises of relevant track-record arguments only if one has many justified philosophical beliefs.
- (P6) One has many justified philosophical beliefs only if extensive skepticism about philosophy is false.
- (C1) Disagreement-based arguments justify extensive skepticism about philosophy only if extensive skepticism about philosophy is false. (epistemic inconsistency)

Before defending premises (P3) and (P4) in more detail, let me begin with some general comments. The core idea of this anti-skeptical argument is that we must rely on justified attributions of (weak) epistemic peerness when we intend to derive extensive skeptical conclusions from the observation of philosophical disagreement. But in order to attribute (weak) epistemic peerness we have to use track-record arguments that rely on many justified philosophical beliefs about the relevant domain. Hence, we must already presuppose the

falsity of extensive skepticism about philosophy in order to justify the premises of the skeptical argument.

It may seem implausible that all disagreement-based arguments for extensive skepticism about philosophy are committed to the (justified) assumption of weak peerhood (see P3). It is true that in my informal motivation for this kind of skepticism, I relied on assumptions of peerhood. But aren't peerhood-independent roads to disagreement-based skepticism available in the literature? And if not, why do I focus on *weak* peerhood only? Let me say two things in response: First, you will indeed find skeptical arguments from disagreement that are claimed to not rely on any peerhood assumption. I hope to persuade you in the next section that these arguments are either unsuccessful or tacitly rely on the assumption of peerhood after all. Second, as I indicated above, strong epistemic peerhood (identical evidence and equal reasoning competence) implies weak peerhood (equally good evidence and equal reasoning competence), but not the other way around. So, in order to show that any skeptical argument based on peerhood assumptions fails, we only need to show that the assumption of weak peerhood spells trouble when used in an argument for extensive skepticism.

But why do we have to rely on track-record arguments in order to attribute weak peerhood to our opponent? Recall the intuitive argument from weak peer disagreement as it was extracted from the Election Case. My evidence for believing that *p* is outweighed by my evidence that a peer disagrees with me only if I assume that this peer is equally likely to be right about propositions in the disputed domain. If epistemic peerhood were disconnected from equal truth-likeness and, e.g., involved only equal subject-relative rationality, then recognizing peer disagreement would provide me with no reason to revise my initial belief. I could then think that whereas *for me*, the belief that *p* is justified, *for her*, the belief that not-*p* is justified. My opponent's disbelief provides counterevidence for *p* only if I assume *from my perspective* that my peer is equally likely to be right. Now, one basis for comparative judgments about truth-likeness is a track-record argument. Track-record arguments justify a conclusion about the general truth ratio for a method (or cognitive capacity) by counting how often this method (or cognitive capacity) delivers correct results in particular cases. So, if we want to find out how likely someone is to be right when employing the method of mental math, we may first check how often past results of her mental calculations were correct, generalize from this and apply it to the situation. One can easily see how one can justify judgments about comparative truth-likeness in a similar way.

One crucial feature of track-record arguments is that in order to support an assessment of a method, they have to rely on many justified beliefs about the method's domain. Hence, to assess the truth-likelihood of someone's ability to do mental calculation, one has to rely on one's own mathematical beliefs (which in turn may be based on either one's own calculation or testimony or one's calculator). This explains why (P5) is trivially true. It follows directly from the definition of track-record arguments: In order to assess the truth-likelihood of philosophers by track-record arguments we have to rely on many justified philosophical beliefs. (P6) is also analytically true. Extensive skepticism about philosophy conceptually entails that not many philosophical beliefs are justified. Hence, if there are many justified philosophical beliefs, extensive skepticism about philosophy is false (by contraposition).

4. Why all disagreement-based arguments must rely on epistemic peerness

In this section, I will defend premise (P3) of my argument in more detail against two alleged counterexamples from the literature. To begin with, you might think that there is a quick and straightforward argument from *systematic disagreement* in philosophy to a general skeptical conclusion that does not rely on the assumption of epistemic peerness. In what follows I will use the phrase "systematic disagreement" if and only if the following three conditions are satisfied: (i) the disagreement is *global*, i.e., in the relevant domain there is no uncontroversial belief, (ii) the disagreement is *symmetric* in the sense that the sizes of the disagreeing parties are approximately equal, and (iii) the disagreement is *persistent* and stable over time, even if the relevant evidence has been fully disclosed to all opponents.⁵ Now, when there is systematic disagreement within a discipline such as philosophy, this might indicate that the discipline's underlying method is generally unreliable. This, in turn, would constitute an undercutting defeater for the whole practice of philosophy. The previous informal skeptical considerations can be put in the form of the following *argument from unreliability* (for similar arguments see Brennan 2010, Goldberg 2009, 2013, Ribeiro 2011):

(P7) If one has sufficient reason to believe that there is systematic disagreement in philosophy, then justification of one's philosophical beliefs is generally undermined.

(P8) One has sufficient reason to believe that there is systematic disagreement in philosophy.

⁵ For a similar characterization, compare Goldberg (2013: 169–170).

(C2) Justification of one's philosophical beliefs is generally undermined.

Let me begin by saying a few things about (P8). It is a claim about the extent, social distribution, and historical dynamics of philosophical disagreement. As such it is a fully empirical claim. I will not defend (P8) here. For the sake of argument, I assume that it is basically correct and can be defended empirically.

Typically, (P7) is motivated by something like the following consideration: If there is systematic disagreement in a discipline such as philosophy, then probably no more than 50% of the beliefs in that discipline are correct.⁶ Knowledge of this fact generates an undermining defeater for the whole practice of philosophical belief formation.

If this is the true motivation behind (P7), it is not a good one. It would only work if every philosophical belief relied on the very same method. If the discipline of philosophy employs many different methods, we cannot simply infer the unreliability of philosophical methods from widespread error in philosophy. It is, indeed, extremely plausible to assume that philosophers use many different methods even within one domain. Not only is it the case that different philosophical schools use different methods; on a more fine-grained scale, there are also methodological differences depending on the weight that is given to various theoretical virtues such as simplicity, fruitfulness, or parsimony. Thus, there should be no doubt that in philosophy many different methods are often applied to a particular type of question. Without the assumption of a homogeneous method we cannot project the truth ratio for the discipline onto the reliability of the applied methods. Compare the following analogous case: You know regarding a certain basketball team that its scoring rate is 50%. If this is the only thing you know about the team, you certainly cannot claim that every player of the team has a scoring rate of 50%. The poor performance of some players might be compensated by the extremely strong performance of other players on the same team.⁷

One might think that there is a related motivation for (P7) that works even under the assumption that there are many, though not too many, methods used in philosophy. If the truth ratio of philosophical beliefs is in general not higher than 50% and if not too many methods are used in philosophy, then at least a significant number of philosophical methods will be unreliable. But then, one might argue, it is highly likely that the method currently used by

⁶ Most likely, fewer than 50% of the beliefs are correct in this case, since often we have many conflicting parties with contrary rather than contradictory beliefs such that at most one of many beliefs is correct (see Goldberg 2013).

⁷ [Reference omitted]

oneself is unreliable, and recognition of this fact undermines one's prima facie justification for one's belief (for a similar argument see Goldberg 2009). However, this is a non-sequitur! From the fact that many (or even most) philosophical methods are unreliable, it does not follow that the philosophical method that I am currently using is most likely unreliable. It would follow only if a further premise were added, namely, that my opponents are my epistemic peers. Without this further assumption, we simply cannot tell how likely it is that our own method is among the unreliable ones. Interestingly, Sanford Goldberg uses this additional premise explicitly in his more recent papers on the topic (Goldberg 2013:178, 2015). The revised argument from (recognized) unreliability therefore depends essentially on a peerness assumption:

(P7*) If one has sufficient reason to believe that there is systematic disagreement *among epistemic peers* in philosophy, then the justification of one's philosophical beliefs is generally undermined.

(P8*) One has sufficient reason to believe that there is systematic disagreement *among epistemic peers* in philosophy.

(C2) The justification of one's philosophical beliefs is generally undermined.

One might think that there is another direct argument from systematic disagreement in philosophy to general skepticism looming one level up. Over and above first-order disagreement about philosophical issues, there is also systematic disagreement about which methods are permissible in philosophy. Given this observation, one might argue that at least the choice of one's philosophical methods is always unreliable or, at least, arbitrary.⁸ Let me discuss both worries in due order. Does systematic disagreement about the permissibility of philosophical methods suggest that the selection of any one of them would be unreliable? Not necessarily. Sure, it would if only one single method were available to decide about the permissibility of philosophical methods. But as we have seen above, the best explanation of systematic disagreement is that the opponents already rely on different methods when they make up their minds about this question. But then one can use the same strategy here. Given that various methods such as, e.g., relying on intuitions, empirical research or reflective equilibrium are already in play, systematic disagreement about whether particular philosophical methods are permissible does not establish that one's own method is very likely unreliable. What about the arbitrariness worry? If any particular selection of philosophical

⁸ Thanks to an anonymous reviewer for raising this objection.

methods is already based on a background method, then the agent has already some motivation for her selection. This motivation might be epistemically circular (if the agent argues for the permissibility of a particular method by using that very method); and it probably will not influence the opponent's choice; but this does not make the choice arbitrary. Nevertheless, there is the pressing question of how one should rationally respond if systematic disagreement persists even about philosophical methods. There seem to be two options here. First, one might be able to ascent to a level of methodological abstraction where one will agree with one's opponent on the permissible methods. In this case, there is some hope that one might be able to dissolve all the lower level disagreements by relying on shared abstract methods. However, I am quite skeptical that it will always be possible to find methodological agreement at some level; and I am also quite skeptical that even if one discovers such an agreement, one will always be able to dissolve all the lower level disagreements. In particular, I do not see why we should be able to dissolve disagreement about more concrete methods by relying on some shared abstract method. According to the second option, both sides do not share methods at any level and hence cannot overcome their disagreement. This clearly indicates an obstacle to establishing agreement, but it does not show that neither party can acquire justified beliefs by relying on her favorite method. This certainly requires epistemically circular reasoning at some level. However, there is nothing objectionable about this kind of reasoning. For the externalist, all that matters is whether the method that is used is in fact reliable. For the mentalist, one can justify a method by relying on one's mental evidence for the method's reliability. Nothing rules out that this evidence can be delivered by the target method itself. As it seems to me, neither of the indicated options poses any additional skeptical problem for philosophy.

In recent literature, one further attempt to argue for skeptical consequences of disagreement without relying on any peerhood assumption has been promoted by King (see King 2012). His argument utilizes an epistemic principle according to which we must suspend our belief regarding any controversial issue unless we have sufficient reason to believe that we are epistemically superior to our opponents. In what follows, I will call this principle *Sidgwick's principle*, since Henry Sidgwick was the first to endorse this principle explicitly in his argument for suspending belief in the face of disagreement (see Sidgwick 1907: 342, also McGrath 2007: 91–92). King (2012) uses Sidgwick's principle in an argument that I will call *Sidgwick's argument from disagreement*:

(P9) Whenever I believe that p and recognize that my opponent holds a conflicting belief, I am rationally required to suspend my belief unless I have sufficient reason to believe that I am epistemically superior to my opponent with respect to whether p is the case. (Sidgwick's principle)

(P10) In philosophy, I typically encounter opponents with respect to whom I cannot reasonably claim epistemic superiority.

(C2) In philosophy, I typically am rationally required to suspend belief.

(P10) is plausible because in philosophy, almost everything is disputed by people who are not clearly inferiors. Sidgwick's principle (P9) is more controversial. In the literature, some philosophers have disputed this principle just because it has excessively skeptical consequences (see Christensen 2011: 15–16; Vavova 2014). I believe that there is an even better reason for rejecting it.

Sidgwick's principle licenses verdicts that are intuitively too strong. Consider the following case of an *Epistemic Troublemaker*: Suppose you are giving a public lecture at a foreign university to an unspecific audience. Students might sit next to faculty or interested laypeople from town. Apart from the colleague who invited and introduced you, all people are unknown to you. In your talk you defend the proposition that p on a topic from normative ethics. At the end of your talk, someone completely unknown to you stands up and asserts with a serious tone " p is false". Without saying anything else she leaves the room. If Sidgwick's principle were correct, the troublemaker's intervention would be sufficient to remove justification from your belief that p . This is so because you have no reason to assess the completely unknown opponent as epistemically inferior to you. But this consequence seems absurd. Justification cannot be lost so easily.

In this section, I looked more closely at two disagreement-based arguments for skepticism that seem to work without relying on any peerhood assumption. It turned out that the argument from unreliability is promising only if we add the assumption of epistemic peerhood as a further premise. In contrast, King's argument for skepticism relies on an epistemic principle, namely Sidgwick's principle, that is intuitively too strong and hence implausible. The moral seems to be that any promising disagreement-based argument for skepticism requires the assumption of epistemic peerhood. Therefore, (P3) seems to be correct.

5. Why reliance on track records is the only trustworthy method of detecting epistemic peers in philosophy

Let us now turn to premise (P4). Why do I claim that we must use track-record arguments when we attribute weak epistemic peerness to our philosophical opponents? In the introduction, I argued that since strong peerness implies weak peerness, it is sufficient to address the problems of attributing weak peerness in order to attack all kinds of peerness attributions. In general, two people are weak peers with respect to proposition *p* if their beliefs about whether *p* is true have equally high probabilities. We can understand these probabilities in two fundamentally different ways. First, they can be understood in a subject-relative way such that the subjectively permissible credence is equally high on both sides. Alternatively, we can understand the probabilities as degrees of truth-likelihood. If we understand weak peerness in the first sense, the recognition of peer disagreement does not generate any rational pressure to revise one's initial belief. It is perfectly consistent to claim, on the one hand, that one's own belief is probable relative to one's own perspective and to admit, on the other hand, that a conflicting belief is probable relative only to someone else's perspective. In contrast, if we understand peers as believing truly with an equal likelihood, it is quite easy to see how disagreement with people whom one is justified to take as one's peers generates rational pressure to revise one's initial belief. If I take my opponent to be as likely to be right about *p* as I am, then I acquire equally strong counterevidence for *p* when I realize that she is disagreeing with me about *p*.

The overall truth-likelihood of an epistemic agent with respect to a proposition *p* is a product of two separate factors. It depends on (i) the quality of the evidence and (ii) the (conditional) reliability of the reasoning competence that takes the evidence as input. To identify a peer one need not necessarily rely on a comparative track record. Typically, one is using indirect indicators of epistemic peerness. For example, when one encounters situations like the Restaurant Case and one's friend is recognized as a cognitively normal person, one is justified in assessing her as one's peer. One certainly need not rely on a detailed calculation of who was right how often in relevantly similar situations in the past. The same is true when someone who is obviously a normal perceiver makes a diverging observation, e.g., whether the lights were already red, when the car passed or who won a tight race. One may have reason to attribute peerness without considering who was right how often in the past about cases that are perceptually hard to judge. We need to keep in mind that the epistemic peerness relevant in this context is defined as equal truth-likelihood with respect to the target

proposition's relevant domain. The following list of criteria might indirectly indicate epistemic peers:

- a. equal general intelligence, rationality, and thoughtfulness (Kelly 2005: 168, Feldman 2006: 218, Christensen 2007: 188);
- b. equal familiarity with the relevant evidence and arguments (Goldman 2001, Kelly 2005: 175, Feldman 2006: 218, Christensen 2007);
- c. absence of bias (Kelly 2005: 175);
- d. equal training and equal reputation (Goldman 2001);
- e. equal dialectic abilities (Goldman 2001);
- f. equal creativity and potential to build theories in the field;
- g. normality of both subjects.

None of these criteria taken alone seems to be sufficient to indicate (equal) truth-likelihood. That much should be obvious. For example, even people who are generally equally intelligent will often differ radically with respect to more domain-specific competences. Similarly, people who are equally well trained might still differ significantly with respect to their respective skills since they are more or less talented. But the defender of indirect indicators need not give up so easily. She can admit that a track-record argument would be the best indicator of reliability. Nevertheless, she might argue that the satisfaction of *all or many of the above listed criteria* taken together makes equal truth-likelihood sufficiently likely. This seems to be true at least in many real-world cases.

I will now argue that this method is not applicable to standard cases of philosophical disagreement. Consider two philosophers who are equally intelligent, rational, and thoughtful *and* who are familiar with the same relevant evidence and arguments *and* who have the same training and reputation in the relevant field of the discipline *and* who are equally good at explaining and defending their views in discussion *and* who are equally creative *and* who seem to be pretty normal both. Isn't it sufficiently likely that they are equally likely to be right as well?

I do not think so. Here is why. Consider two philosophers, Yoko and Zoe, who disagree about modal realism. Yoko accepts modal realism of a Lewisian kind, whereas Zoe believes in some kind of modal actualism of an Armstrongian flavor. They are both highly intelligent and thoughtful. For years, they have both been reading all the relevant literature about the metaphysics of modality, so that they both know all of the arguments for and against

modal realism and actualism. They are both assistant professors in their early career at philosophical departments of renowned colleges. Yoko has published three papers on the metaphysics of modality in *Nous*, *Philosophical Quarterly*, and *Australasian Journal of Philosophy*. Zoe has published two papers in *Philosophical Studies* and one in *The Philosophical Review* on the same topic. Yoko and Zoe often discuss the metaphysics of modality with each other. For an external observer, it is hard to say who is better in defending her view. Both are extremely sharp and quick in discussion. But they both stick to their guns. Yoko's and Zoe's dissertations not only contain critical assessments of the current debate, but also include some constructive theory building. Both seem to be normal members of society's set of highly talented young researchers. In short: Yoko and Zoe are perfect equals according to criteria (a)–(g). That Yoko and Zoe continuously disagree about the truth of modal realism although they have been discussing this issue for quite some time suggests that neither of them has made a simple mistake. A simple mistake would have been discovered already. The best explanation for their persistent disagreement is that they weigh theory virtues such as simplicity, fruitfulness, and parsimony differently. Whereas Yoko gives priority to fruitfulness and favors modal realism because of its explanatory strengths, Zoe gives priority to ontological parsimony and therefore finds it hard to believe that possible worlds exist in the same way as our actual world. Now, three things are revealing about this case. First, the fact that Yoko and Zoe give different weights to theory virtues makes it likely that they use different methods of belief formation with different degrees of reliability. Although one cannot completely rule out the possibility that different methods that license widely incompatible verdicts may after all be equally (un)reliable, most likely such conflicting methods will not be equally reliable. From this a second point follows: Criteria (a)–(g) do not sufficiently discriminate between different degrees of truth-likeness in the case of philosophical disputes. The case at hand is the norm rather than the exception. Hence, indirect indicators of peerhood are not reliable within the domain of philosophy. Third, one cannot simply add the criterion of equal weighing of theory virtues as a further entry in the above list. This is because Yoko and Zoe disagree only in virtue of their unequal weighing of theory virtues. If one takes away the latter, their disagreement would go away as well. So, one could not use the extended list to attribute peerhood among opponents.

The upshot is that, although using any item from the list of suggested indirect indicators may work in many domains of real-world disagreement, it does not work in domains in which it is highly probable that opponents who satisfy all of these criteria are not

equally reliable.⁹ We know from cases like the above that, in philosophy, the equal truth-likelihood of disagreeing philosophers is not sufficiently correlated with these criteria. They often fulfill all of them equally well, but still use different methods that are in all likelihood not equally reliable. As soon as one recognizes that these indirect indicators are disconnected from reliability, this insight defeats one's justification for the peerhood assumption to the extent that it relies on indirect indicators. One therefore has acquired evidence that undermines one's rational use of these indirect indicators.¹⁰

Let me address three objections to the above line of argument.¹¹ First, one might hesitate to generalize the case of Yoko and Zoe. In particular, one might argue that if this case were generalizable, we should expect to see agreement among those philosophers who share their methodology. Is this what we find in reality? It seems to me that this kind of agreement among philosophers with shared methods can in fact be observed. Consider the controversy about utilitarianism. Whereas proponents typically reject the epistemic value of intuitions about cases, opponents typically rely on the method of cases. Or consider the controversy about physicalism. Whereas proponents typically object to the epistemic value of modal intuitions, opponents as, e.g., Kripke, Jackson or Chalmers, typically rely on them. One might still hesitate to generalize the case of Yoko and Zoe since there seems to be an explanation of philosophical disagreement that does not refer to any methodological differences. According to this alternative explanation, philosophical opponents share their basic methodology but apply it inconsistently in ways that are shaped by peculiarities of taste or circumstance. If this explanation applies, opponents use the very same method in unreliable ways. In response, let me point out that there are two different readings of the phrase "inconsistent application of a method." According to the first reading, each philosopher is applying the method at hand in a specific but robust way. On that reading, there is no principle difference between applying the same method in two robustly different ways and applying two different methods. Hence this does not provide a substantial alternative to my initial explanation. According to the other reading, each philosopher is applying the same method in different ways on different occasions. This reading does not involve the use of different methods since methods are

⁹ In philosophy, we know this from indirect evidence. The best explanation of persistent disagreement is that the parties use different methods with different degrees of reliability. Similarly, in chess, we have strong direct evidence that the indirect criteria are not reliable. The criteria on the list (with the possible exception of "reputation in the field") can be equally well satisfied while the Elo ratings indicate that the strength of the players varies radically.

¹⁰ Notice that for the sciences the situation seems to be different since factors such as reputation in the field properly reflect the reliability of scientists' forecasting. In philosophy, the situation is different, given that there are only very few philosophical propositions that are generally accepted.

¹¹ These objections were all raised by one of the anonymous reviewers.

individuated as types rather than tokens of belief-formation. However, it seems very unlikely that a well-trained philosopher uses the same method in an unstable and inconsistent manner. This is what good training is expected to rule out. Therefore, the phenomenon of inconsistent applications is either reducible to methodological difference or it is extremely unlikely to occur among well-trained philosophers.

Secondly, let us suppose that we are persistently disagreeing with another philosopher who is pretty similar to us with respect to all the above indirect criteria (i.e., intelligence, familiarity with relevant evidence, absence of bias etc.) and who, moreover, seems to share her method with us. In such cases, which are clearly possible, we can hardly deny that we have at least some degree of justification that we are confronted with a philosophical peer; and this is so although we do not rely on any track-record argument. But then there seems to be a way of justifying peerness that bypasses any track-record information. In response, two things can be said. First, even if it is possible that all indirect criteria of peerness are satisfied and one *seem* to share the relevant method with someone who is persistently disagreeing, it might still not be possible that there is persistent philosophical disagreement with someone who *in fact* shares one's relevant methods. The deeper reason here is that this kind of persistent disagreement cannot be explained except by referring to different methods being used. One should add here that methods that are in fact different might easily look similar if we abstract away from relevant details. Second, even if it is possible that persistently disagreeing well-trained philosophers are peers, real cases would be rare and hard to identify. Beliefs in philosophical peer disagreement would thus even in the best cases be only weakly justified. This cannot motivate a general kind of meta-philosophical skepticism.

Finally, one might argue that if philosophical opponents do not share their methods and thus will not be peers, one is compelled to the highly counterintuitive conclusion that one's philosophical opponents are always epistemically inferior to oneself. Here is a sketch of such an argument: since my philosophical opponent is not my peer, she can either be my superior or my inferior. Since agnosticism about the relative epistemic quality of my opponent is initially rational I need not weaken any of my philosophical beliefs in response to this disagreement but can remain steadfast. On the basis of my own philosophical beliefs my opponents appear to have false beliefs whenever they disagree with me. Therefore, opponents (including, e.g., David Lewis) must be generally assessed as inferiors. This result is not only implausible as such, it also contradicts my initial assumption that agnosticism about the relative quality of my opponent is rationally required.

I agree that this epistemic situation seems to generate a paradox of inconsistent doxastic attitudes. Since there is no sufficient evidence to assess the relative epistemic quality of my philosophical opponent, agnosticism seems initially required. But then we can use our own justified philosophical beliefs to argue that our opponent is wrong whenever she disagrees with us and thus is epistemically inferior. This leads to the inconsistency between rationally required agnosticism about the relative epistemic quality and the rationally permissible claim that our opponent is inferior to us. This inconsistency is structurally similar to the lottery paradox (or the preface paradox). In the case of the lottery paradox, you know that, say, 999 tickets will lose and only one will win. Hence, the probability is very high for each ticket that it will lose. So, you are justified in believing that it will lose. But the same applies to all the other tickets. If justification is closed under competent deduction, you can derive the justified conclusion that all tickets will lose. But this conflicts with your prior knowledge that exactly one ticket will win. According to some of the standard solutions to the lottery paradox, the justification of each of your beliefs about the losing individual tickets can be made consistent with your knowledge that one ticket will win. Whatever saves this consistency in the lottery paradox, for example giving up on closure for justification (Kyburg 1961), will also be suitable to defend the consistency of the steadfast view with peerhood agnosticism here.

It turns out that, in philosophy, there is no way to know whether one's opponents are epistemic peers other than relying on track-record arguments. In order to find out whether our opponents perform as reliably as we do, we must also determine who is correct in the disagreement range and how often.¹² Such an inductive track record argument would look like this:

(P11) I believed that p whereas my opponent S believed that not-p; and later <u>it turned out that p</u> .	(P12) I believed that p* whereas my opponent S believed that not-p*; and <u>later it turned out that not-p*</u> .
(P13) ...	(P14) ...
(P15) ...	(P16) ...

¹² Strictly speaking, this claim has to be relativized to methods. More precisely, in order to count as an epistemic peer with respect to some currently controversial proposition p the opponents must use *methods* that produced an equal ratio of true beliefs in the past. It is hence not sufficient to count the true and false beliefs on both sides, we also have to establish that we compare only beliefs that are based on those methods that are currently used by the opponents. To avoid too much complexity, I leave out this qualification in the main text. Thanks to an anonymous reviewer for drawing my attention to this fact.

(C3) Therefore, my opponent and I are epistemic peers.

If the ratio of my true beliefs approximately matches the ratio of my opponent's true beliefs within the disagreement range, then this is good inductive evidence for believing that we are peers. But such a track-record argument can justify (C3) only if every premise is justified. In order to do that, we need to justify many beliefs about the domain in question, as the underlined parts of (P11) and (P12) indicate.

6. The epistemic inconsistency in the argument for global meta-philosophical skepticism

On closer inspection, the premises of my argument have turned out to be either uncontroversial (P5 and P6) or clearly defensible, such as the premise that all *prima facie* plausible skeptical arguments from disagreement rely on a justified peerness assumption (P3) and that one can justify peerness in philosophy only by using a track-record argument (P4). Taken together, this justifies the conclusion that disagreement-based arguments for extensive skepticism about philosophy are epistemically inconsistent. But what is the epistemic significance of an argument that some conclusion is epistemically inconsistent, i.e., inconsistent with the justification of the conclusion? It seems clear that being epistemically inconsistent has negative consequences for the justificatory status of the position in question. Whenever an argument *prima facie* justifies a conclusion that is in tension with the justification, it thereby generates an undermining defeater that removes the *prima facie* justification. The moral is that meta-philosophical skepticism cannot be *ultima facie* justified in this way.

There is a general worry about the strategy of this paper that needs to be addressed here. One might admit that in the case of disagreement-based skepticism the skeptic's attempt to justify directly that philosophical beliefs are generally unjustified leads to self-defeat. However, a more powerful strategy seems to be available to the skeptic.¹³ She might understand her argument as a *reductio ad absurdum* of the following type: you start with the *assumption* that a certain method M is reliable, then you apply M to the evidence that is uncontroversially available to you, which in turn leads you to the result that M is unreliable. If additionally you assume that reliable methods justify beliefs you can derive that you are justified in believing that M is unreliable. Strictly speaking, there is no contradiction between

¹³ The following objection was raised by one of the anonymous reviewers. Compare also Foley 2003.

the assumption that M is reliable and the conclusion that you are justified in believing that M is unreliable. However, if the assumption of M's reliability leads you to the conclusion that you are justified in believing that M is unreliable, M will be self-undermining *even under the assumption of its reliability* and thus cannot justify any beliefs. Pessimistic meta-induction is a paradigm case of this skeptical strategy. One starts with the assumption that induction is a reliable method. Then, one takes into account that all inductively justified scientific theories of the past have been falsified later. By applying induction to these results one can now conclude that induction is unreliable. If we assume that reliable methods justify beliefs, we can derive the conclusion that we are justified in believing that induction is unreliable from the assumption that induction is reliable.

It is tempting to apply this reductio strategy to the case of disagreement-based skepticism: suppose I am reliable at assessing philosophical claims. Then, my track-record assessments in philosophy are justified. But we are assuming that if my track-record assessments are justified, then my peerhood judgments are justified, in which case I am justified in believing that I am not justified in my philosophical beliefs. As in the case of pessimistic meta-induction the assumption is self-undermining rather than leading to a contradiction. Under the assumption of being a reliable philosopher I can argue that I am justified in believing that I am not. Can we thus derive skeptical consequences from disagreement in philosophy without relying on any *justified* philosophical beliefs? That would be an attractive route for the skeptic.

I do not dispute that the reductio strategy *can* be a good tool in the hands of the skeptic, but it is not always applicable. It is important to keep in mind what exactly is established by the reductio strategy. The reductio does not unconditionally justify the conclusion of the argument. It rather shows that *under the optimistic (i.e. non-skeptical) assumption* of reliable (or trustworthy) target beliefs these beliefs cannot be justified. Now, if there is no optimistic alternative to the optimistic assumption, then the skeptical conclusion can be generally established. In the case of basic methods no optimistic alternative to the assumption of their reliability (or trustworthiness) is available. But the case of disagreement-based skepticism is different. In that case, we must start with the assumption of one's *general* philosophical reliability in order to motivate that one's track-record assessments are justified which in turn is needed to motivate the skeptical conclusion. However, in this case there are clearly optimistic alternatives available that do not motivate the skeptical conclusion. Assume, e.g., that only a proper subset of your philosophical beliefs are reliably formed. This clearly

can be true since you do not always use the same method when you form philosophical beliefs. But then the reliable subset of your philosophical beliefs might suggest that in most cases of disagreement it is you rather than your opponent who is correct. This would be an optimistic (non-skeptical) assumption that does not lead to skeptical consequences. Since in the case of disagreement-based arguments not every non-skeptical assumption leads to a skeptical conclusion, the reductio strategy does not successfully apply to this case.

We can gain a deeper understanding of what is going on here by identifying the exact point at which the inconsistency occurs in the revised argument for meta-philosophical skepticism. Here is the argument again:

- (P1*) Whenever the controversy is not about epistemically fundamental principles and I have dispute-independent reasons for believing that I am disagreeing with an epistemic peer, I am rationally required to suspend judgment. (partial conciliationism)
- (P2) For any philosophical proposition that I believe, I have dispute-independent reasons for believing that I am disagreeing with some epistemic peers.
- (C*) Therefore, I am rationally required to suspend all philosophical beliefs unless they refer to epistemically fundamental principles.

By adding the insight that (P2) can be supported only by a track-record argument, we get:

- (P1*) Whenever the controversy is not about epistemically fundamental principles and I have dispute-independent reasons for believing that I am disagreeing with an epistemic peer, I am rationally required to suspend judgment. (strong conciliationism)
- (P2) For any philosophical proposition that I believe, I have dispute-independent reasons for believing that I am disagreeing with some epistemic peers.

Inductive justification (of P2):

- (1) I believed that p and my opponent believed that not-p, but later it turned out that p.
- (2) I believed that p* and my opponent believed that not-p*, but later it turned out that not-p*.

.....

(C**) My opponent is my epistemic peer (in the relevant domain of philosophical propositions).

(C) Therefore, I am rationally required to suspend all philosophical beliefs unless they refer to epistemically fundamental principles.

But now we can easily see the conflict:

Inductive justification (of P2):

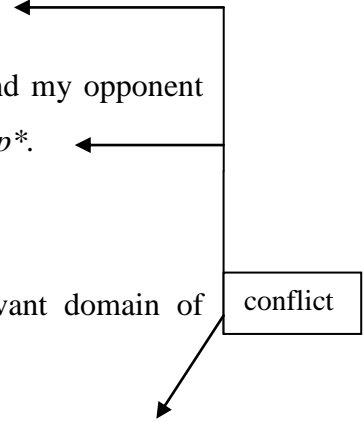
(1) I believed the philosophical proposition that p and my opponent believed that $\text{not-}p$, but later *it turned out that p* .

(2) I believed the philosophical proposition that p^* and my opponent believed that $\text{not-}p^*$, but later *it turned out that $\text{not-}p^*$* .

.....

(C**) My opponent is my epistemic peer (in the relevant domain of philosophical propositions).

(C*) Therefore, *I am rationally required to suspend all philosophical beliefs* unless they refer to epistemically fundamental principles.



In the sub-argument that supports (P2), premises (1) and (2) make first-order philosophical claims. But at the same time, the conclusion prohibits that kind of philosophical claim.

6. Conclusion

If the new argument against skepticism from philosophical disagreement works, then the argument for *extensive* meta-philosophical skepticism, like the argument for *global* meta-philosophical skepticism, is self-defeating. This is surely an interesting result, and it puts rational pressure on the skeptical position. But this cannot be the end of the story, since we also want to understand *whether* skepticism is false. It can only be false if at least one of premises of the argument for meta-philosophical skepticism is false. The anti-skeptical argument that I presented in this paper is completely silent on this issue. However, as indicated in the introduction to this paper, both premises of the argument for meta-philosophical skepticism might be mistaken.

(P1*) might be false, or at least too strong, for a number of independent reasons. First, one might object to the independence condition that is characteristic for the so-called ‘Equal Weight View’ (Kelly 2010). Second, one might argue that the evidence of peer disagreement constitutes a *partial* rather than a full defeater such that full suspension of the belief would be an implausibly strong reaction (Thune 2010).

There are many reasonable objections to premise (P2), as well. First, it might simply be a massive exaggeration that all philosophical propositions are controversial. There may be more common ground in philosophy than is usually granted by the skeptic—in particular, if we stop counting something as controversial when there are only a few opponents. Second, in practice it is often much harder to identify a peer than the clean and idealized cases we started with suggest (King 2012). There are also cases of deep and radical disagreement. These occur when different basic methods are used or when we are confronted with someone who disagrees with almost everything we believe. In these cases, there seems to be no independent basis that permits us to justify the belief that our opponent is as reliable as we are (Elga 2007, Vavova 2014). Third, sometimes there is some dispute-independent information that allows us to demote the alleged peer in the concrete situation (Lackey 2010, Sosa 2010, Christensen 2011).

As these remarks suggest, there is some hope that we will be able to provide an explanation of what goes wrong in the premises of the argument for meta-philosophical skepticism. In any case, such a diagnosis over and above the argument offered here seems required. One should also keep in mind that I have been arguing only against the viability of extensive skepticism. This leaves the possibility (and reasonableness) of local skepticism on the basis of peer disagreement completely untouched.

References

- Brennan, Jason 2010, “Skepticism about Philosophy”, in: *Ratio* 23: 1–16.
- Christensen, David 2007, “Epistemology of Disagreement: The Good News”, in: *Philosophical Review* 116: 187–217.
- Christensen, David 2011, “Disagreement, Question-Begging and Epistemic Self-Criticism”, in: *Philosopher’s Imprint* 11: 1–22.

- Christensen, David 2013, “Epistemic Modesty Defended”, in: D. Christensen & J. Lackey (eds.), *The Epistemology of Disagreement. New Essays*, Oxford: Oxford University Press, 77–97.
- Decker, Jason 2014, “Conciliationism and Self-Incrimination”, in: *Erkenntnis* 79: 1099–1134.
- Elga, Adam 2007, “Reflection and Disagreement”, in: *Nous* 41: 478–502.
- Elga, Adam 2010, “How to Disagree About How to Disagree”, in: R. Feldman & T.A. Warfield (eds.), *Disagreement*, Oxford: Oxford University Press, 175–187.
- Feldman, Richard 2006, “Epistemological Puzzles about Disagreement”, in: S. Hetherington (ed.), *Epistemology Futures*, Oxford: Oxford University Press: 216–236.
- Foley, Richard 2003, “Three Attempts to Refute Skepticism and Why They Fail”, in: S. Luper (ed.), *The Sceptics: Contemporary Essays*, Abingdon/New York: Ashgate Publishing: 61–74.
- Goldberg, Sanford 2009, “Reliabilism in Philosophy”, in: *Philosophical Studies* 142: 105–117.
- Goldberg, Sanford 2013, “Disagreement, Defeat, and Assertion”, in: D. Christensen & J. Lackey (eds.), *The Epistemology of Disagreement*, Oxford: Oxford University Press: 167–189.
- Goldberg, Sanford 2015, *Assertion: On the Philosophical Significance of Assertoric Speech*, Oxford: Oxford University Press.
- Goldman, Alvin 2001, “Experts: Which Ones Should You Trust?”, in: *Philosophy and Phenomenological Research* 63: 8–110.
- Grundmann, Thomas 2013, “Doubts about Philosophy? The Alleged Challenge from Disagreement”, in: Tim Henning & David Schweikard (eds.), *Knowledge, Virtue, and Action: Putting Epistemic Virtues to Work*, New York: Routledge: 72–98.
- Kelly, Thomas 2005, “The Epistemic Significance of Disagreement”, in: *Oxford Studies in Epistemology* 1: 167–196.
- Kelly, Thomas 2010, “Peer Disagreement and Higher Order Evidence”, in: Richard Feldman & Ted Warfield (eds.), *Disagreement*, Oxford: Oxford University Press: 111–174.
- King, Nathan 2012, “Disagreement: What’s the Problem? Or a Good Peer is Hard to Find”, in: *Philosophy and Phenomenological Research* 85: 1–23.
- Kornblith, Hilary 2010, “Believe in the Face of Controversy”, in: Richard Feldman & Ted Warfield (eds.), *Disagreement*, Oxford: Oxford University Press: 29–52.
- Kyburg, Henry 1961, *Probability and the Logic of Rational Belief*, Middletown: Wesleyan University Press.

- Lackey, Jennifer 2010, “A Justificationist View of Disagreement’s Epistemic Significance”, in: A. Haddock, A. Millar & D. Pritchard (eds.), *Social Epistemology*, Oxford: Oxford University Press: 298–325.
- Machuca, Diego 2013, “Conciliationism and the Menace of Scepticism”, in: *Dialogue* 54: 469–488.
- McGrath, Sarah 2007, “Moral Disagreement and Moral Expertise”, in: *Oxford Studies in Metaethics* 4: 87–108.
- Ribeiro, Brian 2011, “Philosophy and Disagreement”, in: *Critica* 43: 3-25.
- Pittard, John 2015, “Resolute Conciliationism”, in: *The Philosophical Quarterly* 65: 442-463.
- Sidgwick, Henry 1907, *The Methods of Ethics*, 7th ed., Indianapolis/Cambridge: Hackett Publishing Co.
- Sosa, Ernest 2010, “The Epistemology of Disagreement”, in: A. Haddock, A. Millar & D. Pritchard (eds.), *Social Epistemology*, Oxford: Oxford University Press: 278–297.
- Thune, M. 2010, “‘Partial Defeaters’ and the Epistemology of Disagreement”, in: *The Philosophical Quarterly* 60: 355–372.
- Vavova, Katia 2014, “Moral Disagreement and Moral Skepticism”, in: *Philosophical Perspectives* 28: 302–333.