



# Higher Order Evidence and Deep Disagreement

Klemens Kappel<sup>1</sup>

Published online: 3 September 2018  
© Springer Nature B.V. 2018

## Abstract

In deep disagreements local disagreements are intertwined with more general basic disagreements about the relevant evidence, standards of argument or proper methods of inquiry in that domain. The paper provides a more specific conception of deep disagreement along these lines and argues that while we should generally conciliate in cases of disagreement, this is not so in deep disagreements. The paper offers a general view of disagreement, holding roughly that one should moderate one's credence towards uncertainty in so far as disagreement with others provides undefeated higher order evidence that one might have made a mistake in one's appreciation of the first order evidence. When applying this view to deep disagreement we get that in cases of deep disagreement higher order evidence from disagreement is rebutted or undercut by the nature of the disagreement. So, in cases of deep disagreement one should not moderate one's credence. I finally argue that this gives a better general view of deep disagreement than views appealing to epistemic peers, personal information or independence.

**Keywords** Disagreement · Deep disagreement · Higher order evidence · Peer disagreement · Personal information · Independence

## 1 Introduction

The characteristic feature of deep disagreements is that local disagreements about facts in a particular domain are intertwined with basic disagreements about the relevant evidence, standards of argument or proper methods of inquiry in that domain. For convenience, we can say that deep disagreement features clashes of *perspectives*, where a perspective includes views about how we should acquire evidence, assess evidence, argue, form beliefs, or gain knowledge in the domain in question. Deep disagreement involves not only disagreement about facts in a particular domain, but clashes of perspectives, where there is no more basic level we can resort to in order to resolve the disagreement.

Arguably we see deep disagreements, or at least patterns resembling deep disagreements, in controversies over science and religion, and in many moral disagreements. Though this may be less obvious, there are, I suggest, also elements of this pattern in disagreements over GMO, climate

change, evolutionary theory, and complementary/alternative medicine (CAM). Here people disagree not only about the facts, but also about how to gather reliable evidence or gain knowledge about the disputed facts, with no commonly agreed level to resort to. Academic disciplines like economics or indeed philosophy may also be subject to deep disagreement, at least in some areas. Disagreements involving conspiracy theories will also often involve patterns of deep disagreements. The same may hold for disagreements purportedly due to ideology insofar as certain wicked ideologies systematically prevent us from seeing vital moral or non-moral truths.

Like in the frequently discussed cases of peer disagreement, there seem to be two intuitions pulling in opposite directions in cases of deep disagreement. Suppose that A and B find themselves in a case of deep disagreement and fully understand this. Should A and B remain steadfast, or should they conciliate? On the one hand, A understands how her view about the proposition in dispute depends on her wider perspective, and A can also see how someone with B's perspective will have to reject what she regards as true. But A also understands why someone with B's perspective would be inclined to acquire false beliefs in the domain in question. So, for A there seems to be no reason to be worried that B disagrees. This is the *non-conciliationist intuition*. On

---

✉ Klemens Kappel  
kappel@hum.ku.dk

<sup>1</sup> Section for Philosophy, Institute for Media, Cognition and Communication, University of Copenhagen, Copenhagen, Denmark

the other hand, it surely seems that A and B are on a par in a significant way. Their disagreement depends on different perspectives, but neither of them have independent basis for saying that their perspective is better than their opponent's. So, A cannot insist that her perspective is right and B's is wrong, but neither can B insist on this. Indeed, on what basis would A or B do that? For this reason, it might be thought, both A and B should reduce their confidence in their view about P. This is the *conciliatationist intuition*.

My aim in this paper is to discuss which, if any, of these two reactions to deep disagreement is the epistemically proper one. I will argue, roughly, that though we should in general conciliate in disagreements, this is not so when two parties understand that they are in a deep disagreement. The plan of the paper is as follows. Section 2 outlines a more detailed conception of deep disagreement. In Sect. 3, I sketch a general view of disagreement holding, roughly, that disagreement is epistemically significant to the extent that it provides us with higher order evidence that our cognition may be faulty in certain ways. In Sect. 4 I apply this view to the case of deep disagreement. In Sect. 5, I turn to some of the main ingredients in the epistemology of disagreement, peer disagreement, personal information and independence and discuss their relevance for deep disagreement. Finally, in Sect. 6, I consider a couple of worries about the proposed view.

## 2 On Deep Disagreement

The term 'deep disagreement' is clearly a philosophical term of art, but it is helpful to propose a more detailed stipulation informed by reflection on the sort of cases mentioned above.

Say that a *substantive epistemic principle* is a directive telling us that we ought to use certain *doxastic practices* for acquiring justified beliefs and knowledge in certain domains, where a doxastic practice is a way of acquiring or sustaining beliefs. I intend doxastic practices to be construed broadly so as to include trusting certain authorities or sources of evidence, using certain methods of inquiry, and analyzing and processing evidence in particular ways. So, for example, insisting on the need for randomized controlled trials for assessing the efficacy of medical interventions whether of the CAM variety or not, is adhering to a substantive epistemic principle. Trusting scientific evidence from well-functioning scientific institutions is a doxastic practice. A *truth-conducive* substantive epistemic principle picks out *reliable* doxastic practices, that is, doxastic practices that tend to yield true and avoid false beliefs when used for appropriate cognitive aims under proper circumstances.

We can think of substantive epistemic principles as involving an evaluative (or conceptual) part and a factual part. The evaluative part specifies that relying on certain

epistemic practices (say visual perception for specific purposes under proper circumstances) yield knowledge, epistemic justification, epistemic entitlement, or some other epistemic status. The factual part says or implies that the recommended doxastic practices are reliable when used for certain cognitive aims under proper circumstances.

Doxastic practices are actual psychological or social belief forming processes, so whether doxastic practices are reliable or not is a contingent empirical fact about them. Generally, there are no doxastic practices M, domains D and set of circumstances C such that it is necessary or a priori true that M is reliable with respect to domain D in circumstances C.<sup>1</sup> So, the truth-conduciveness of substantive epistemic principles and the reliability of our doxastic practices depend on contingent features of the world, features that generally are knowable only a posteriori. Accordingly, to provide evidence for the truth-conduciveness of a particular epistemic principle we need to refer to empirical facts pertaining to features of the world and our cognition.

In the cases of deep disagreement that we are concerned with, people rely upon or trust different modes of gathering and evaluating evidence and forming beliefs. These differences are best understood, I suggest, as differences and conflicts in substantive epistemic principles. Since the truth-conduciveness of substantive epistemic principles depends on contingent features of the world, our rational endorsement of such principles depends on our beliefs about the relevant parts of the world, even if the rationality of these beliefs themselves depends on other epistemic principles. If I believe that the features needed for some epistemic principle to be truth-conducive simply do not obtain, then there is a sense in which I am irrational in endorsing it and in relying on the doxastic practices it recommends. Thus, if we have very different beliefs about the relevant features of the world, we might well rationally have to have very different beliefs about which doxastic practices are truth-conducive. This partly explains why there can be deep disagreements despite the fact that we populate the same world and are roughly equally intelligent and receptive to evidence. Our immediate evidential input and formal constraints on what we can rationally believe underdetermine many of our beliefs about the world, and once we have different beliefs about various parts of the world, these beliefs may in turn support very different substantive epistemic principles, which may in turn sustain deep disagreements.<sup>2</sup>

<sup>1</sup> Unless, of course, proper circumstances are defined as whatever circumstances make M reliable with respect to D.

<sup>2</sup> Substantive epistemic principles might not be the only kind of epistemic principles there are. Some epistemic principles might be merely formal in nature in that they tell us only very general things such as that we should respect the evidence and avoid inconsistent beliefs. Formal epistemic principles are neutral with respect to which more

We need to be more specific about what it is for substantive epistemic principles to conflict. As a first approximation, let us say that two substantive epistemic principles are *incompatible* if and only if they rank doxastic practices differently, given a particular cognitive aim, and a fixed set of circumstances, say when one epistemic principle selects one particular doxastic practice as the preferred one, and the other principle selects another doxastic practice. If two epistemic principles rank doxastic practices differently they are incompatible: there will be possible situations in which one cannot comply with both principles. Yet, two sets of differently ranked doxastic practices may be pairwise equally reliable, or may be on average equally reliable.

But substantive epistemic principles can also conflict in a more dramatic way. Let us say that two substantive epistemic principles  $EP_1$  and  $EP_2$  are *opposed* to one another when (roughly)  $EP_1$  implies that the doxastic practices praised by  $EP_2$  are epistemically unfit for the relevant purposes, and vice versa. This is vague in certain ways, as doxastic practices clearly can be more or less unfit for a certain purpose, as they can be more or less reliable, but I will ignore this complication. As I have defined things, substantive epistemic principles can be incompatible without being opposed to one another.

In this paper I will be mostly concerned with deep disagreement that involves opposed epistemic principles, though disagreements involving incompatible but non-opposed epistemic principles may raise some of the same issues. So, deep disagreement, as I will mainly think of it, involves opposed substantive epistemic principles, where one can resort to no other substantive epistemic principle to resolve the conflict. In this sense, deep disagreement involves *basic* epistemic principles. A substantive epistemic principle being basic just means that there are no distinct epistemic principles that one can turn to for showing that the principle is truth-conducive. This can happen when the rationality of adopting a particular epistemic principle depends on holding certain beliefs, and yet the rationality of these very beliefs also depends on holding that particular epistemic principle. We can state this in terms of epistemic reasons, if we think of *epistemic reasons* in terms of arguments, or at least reasoning that can be represented as arguments: S has an epistemic reason for some proposition P iff S has access to a set of premises that constitutes a cogent argument for P, and these premises have an appropriate epistemic status for S (that is, S is justified or entitled in believing them). When a substantive epistemic principle EP is basic for a subject S, then S does not have

access to an argument for the truth-conduciveness of EP, where this argument does not somehow rely on EP itself.

It should be noted that basicness is *agent-relative* in a certain sense. What is epistemically basic for me need not be so for you, as you may have other cognitive routes available in your perspective. This complicates our description of deep disagreement, as A and B may have opposed substantive epistemic principles, where A's principles are basic for A, but the negation of A's principles need not depend on principles that are basic for B.

We can now give the following general, if still highly schematic, model of deep disagreement:

*Deep Disagreement.* A and B are in deep disagreement regarding a proposition P if and only if: (i) A and B adopt different doxastic attitudes to P, (ii) A's doxastic attitude to P depends on substantive epistemic principle  $EP_A$ , (iii) B's doxastic attitude to P depends on substantive epistemic principle  $EP_B$ , (iv)  $EP_A$  is basic for A, (v)  $EP_B$  is basic for B, (vi)  $EP_A$  and  $EP_B$  are opposed, (vii) A's rejection of  $EP_B$  depends on her reliance on  $EP_A$  and vice versa.<sup>3</sup>

A couple of features of deep disagreement should be noted. For convenience we have said that when A and B have a deep disagreement, they have a conflict of *perspectives*. We can now specify that A's perspective consists of A's system of substantive epistemic principles  $EP_A$ , her associated doxastic practices, and her beliefs and reasons sustaining the suppositions of truth-conduciveness of those principles (and the same for B).

In this paper I will assume that deep disagreements are transparent in that parties to a deep disagreement understand how their own view and that of their opponent is entangled in a wider perspective. Of course, this is not a realistic assumption. Many cases aspiring to be cases of deep disagreement, including those I mentioned in the beginning of the paper, are unlikely to be transparent, but I will set this aside.

In a deep disagreement between A and B, A's best argument for the truth-conduciveness of  $EP_A$  is *epistemically circular* (cf. Alston 1986). This implies that if A were to question the truth-conduciveness of  $EP_A$ , if only for the purpose of argument, then A should not consider herself as having a good reason to think that  $EP_A$  is truth-conducive. The same holds for B and her epistemic principle  $EP_B$ . Similarly, A

Footnote 2 (continued)

specific doxastic practices one should use, and they don't rely on assumptions about which specific doxastic practices are reliable.

<sup>3</sup> This is a revised version of my account in Kappel (2012) which is intentionally similar in spirit to Lynch's way of defining deep disagreement, though there are differences of detail, cf. Lynch (2010). Compare also to Fogelin's somewhat different conception, according to which deep disagreement occurs in argumentation that involves disagreements over the norms governing the argumentation (Fogelin 1985, p. 8).

does not have a good argument to offer to B to the effect that  $EP_A$  is truth-conducive. A's best argument for the truth-conduciveness of  $EP_A$  appeals to premises whose justification depends on relying on  $EP_A$ , and B should therefore rationally reject it since she does not accept  $EP_A$ . The same holds for B's arguments for  $EP_B$ , of course.

Yet, as many have pointed out, including Alston, even if A does not have a non-circular argument for the truth-conduciveness of  $EP_A$ , relying on  $EP_A$  may nonetheless in some externalist sense be epistemically permissible or epistemically rational for A. So, while deep disagreement is symmetrical in some epistemic respects, it can be asymmetrical in other epistemically relevant ways, notably in the sense that A is right about P, while B is not, and  $EP_A$  can be truth-conducive while  $EP_B$  is not, and that A is in some externalist sense right in relying on  $EP_A$ , whereas B is not similarly properly relying on  $EP_B$ . On familiar externalist accounts of epistemic justification, this might imply that A, but not B, knows or is epistemically justified in the relevant beliefs. For these reasons, we can accept that we ultimately rely on basic epistemic principles without giving in to the skeptic.

One may wonder whether deep disagreement between two conflicting but equally coherent perspectives is really a coherent possibility. Suppose, for example, that B rejects scientific standards of inquiry and evidence, but only for the purpose of defending her belief in creationism, not in other domains. Shouldn't we charge B of incoherency? I think that once we appreciate that deep disagreement involves disagreements over substantive epistemic principles whose truth-conduciveness in turn depends on facts of the world, we can see that B's perspective need not be incoherent. As noted, which substantive epistemic principles B should rationally accept depends on B's beliefs about the world. B might think that the nature of the world licenses believing that while many questions about the world can be answered by scientific methods, questions about the origins of life cannot—they are too profound or too fundamental, or concern features of the world that are not amenable to scientific or naturalistic treatment.

### 3 The Higher Order Evidence View of Disagreement

Turn now to the question of what we should rationally do in cases of deep disagreement. This covers at least three distinct more detailed questions: (1) Are deep disagreements rationally resolvable? Say that a disagreement between A and B about P is rationally resolvable iff there is a set of mutually accessible epistemic reasons available to A and B (that is, a set of premises that both A and B should rationally accept, and an inference from these premises) determining a common doxastic attitude. Without elaborating, I assume

that the above account of deep disagreement suggests that they are not rationally resolvable (cf. Lynch 2010). (2) Even if deep disagreements are not rationally resolvable, there might well be some doxastic attitude we ought to have in deep disagreements. So, the question arises what doxastic attitude (or range of attitudes) we are rationally obliged or permitted to have in deep disagreements? (3) When we are in deep disagreement about factual or moral matters that are important for public policies or decisions affecting us all, how should we deliberate with one another, and how should we identify a legitimate policy or decision? (see Lynch 2010; Kappel 2012, 2017a).

In this paper, I will focus solely on the second question. Any answer to this question will depend on a general view about the epistemology of disagreement. I cannot fully defend a general view here, but I want to propose a particular way of thinking of the epistemology of disagreement in terms of higher order evidence. This idea is not new, of course, though I want to suggest that the plausibility of the view depends on certain details that are novel (for similar though importantly distinct views, see Kelly (2010, 2013), Bergmann (2009)).

The epistemic significance of higher order evidence can be exemplified by a case adapted from Christensen (2011, p. 6):

The Pill. Suppose I consider a mathematical problem on the basis of some evidence E. Suppose that E entails that P is the correct answer to the mathematical problem. After careful scrutiny I come to believe that the correct answer is P. I am then told by a credible source that without noticing I have ingested a reason-distorting pill that makes me completely unreliable with respect to those kinds of mathematical problems, though this is not in any way perceptible to me.

The Pill invites the following intuition: after being informed about the pill, I am no longer rationally highly justified in my belief that P. It would, for example, seem highly irresponsible of me to bet my fortune on the truth of P, or to regard P as true without qualification in my theoretical reasoning. In other words, upon being told about the pill, I should reduce my first order credence.

Note a couple of things about this case. First, even if I did in fact get the first order evidence right, the advent of the evidence that I have ingested the reason-distorting pill should lead me to significantly reduce my confidence in my object level belief. Second, it seems that what matters is my level of propositional justification in the higher order belief in question. It does not matter whether I actually have this higher order belief—it seems enough that I am propositionally justified in this belief. To illustrate, suppose that the story is as above, but for some reason I am psychologically blocked from believing that I might be under the influence

of a reason-distorting drug. It seems that I should still reduce my confidence in the first order belief, though I wouldn't actually do so.

What we have seen so far can be accounted for in the following way. Say that a *negative higher order belief* is a belief to the effect that a first order belief is epistemically flawed in one of two distinct ways. First, a belief can be epistemically flawed by being based on a non-truth conducive epistemic principle, or a doxastic practice that is not sufficiently reliable. Call this a *principle error*. Second, a belief can be epistemically flawed when it is the result of a *performance error*, that is, when it is based on a truth-conducive epistemic principle and a reliable doxastic practice, and yet something has gone awry in the execution. In the Pill, my suspicion that I have made a performance error should be raised, but we can imagine other cases that invite the thought that I have used a non-truth conducive principle, and therefore am guilty of a principle error.

Beliefs about these two forms of errors are obviously related to defeaters. There are slightly different ways of defining defeaters. One standard way makes a distinction between undercutting and rebutting defeaters. Relative to some body of evidence E indicating the truth of a proposition P, an *undercutting defeater* removes the evidential relation between E and P, where a *rebutting defeater* indicates the falsity of P. When I receive evidence that I have ingested a reason distorting pill, this is *not* evidence that the mathematical proposition P is false, nor is it evidence that the evidential relation between the mathematical evidence E and the correct answer is absent after all. So, my evidence that I ingested a reason distorting pill is strictly speaking neither an undercutting nor a rebutting defeater in the sense just defined. Rather, it is evidence that I may have made a performance error, and this affects my ability to correctly appraise the evidence in the case—I can no longer take myself to have correctly understood the evidential import of E (and in this slightly different sense, the higher order evidence may be said to be an undercutting defeater). Similarly, if I were to believe that I am relying on a flawed epistemic principle or a non-reliable doxastic practice, this is evidence that I am doing something wrong, and need not concern the actual evidential relation between some evidence and a target proposition, though this depends on how evidence is construed.

So, negative higher order beliefs concern performance errors or principle errors. The epistemological import of negative higher order beliefs might be captured, I suggest, in the following rough way:

(D) To the degree that S is propositionally justified in a negative higher order belief, S should moderate her credence in the first order belief towards uncertainty.

Note that while my first order rational credence may be strongly affected by a propositionally justified negative

higher order belief, there is no similar influence from my first order belief to my higher order belief. Consider again the Pill, and suppose that I happen to process the first order evidence correctly. Surely, this fact alone does not make me justified in thinking that I didn't ingest the pill after all, or that my friend misled me when he said so. Rationally appreciating the evidence that one has ingested the reason-distorting pill should make one believe that this is so, and this should in turn make one less rationally confident in the correctness of one's object level belief. Yet, correctly appreciating the first order evidence should not make one rationally believe that one has not ingested a reason-distorting pill after all, or that one's reasoning isn't affected by the pill. So, there is a distinct asymmetry—higher order evidence concerning a particular first object level proposition affects object level rational credence in a way that object level evidence does not affect rational credence in higher order beliefs (for a fuller discussion see Kappel (2017a, b)).

I suggest that the general epistemic significance of disagreement can be accounted for by these features. In general, when I believe that P and encounter others who think not-P, this is typically some evidence that P is false. But disagreement is also higher order evidence that I might be subject to a performance error or a principle error. If this higher order evidence is strong enough, and not defeated, then I should moderate my credence in the disputed proposition. We get stronger higher order evidence when we disagree with peers than when we disagree with our epistemic inferiors. But clearly, disagreement with inferiors can also provide decisive higher order evidence. A senior medical doctor might get a certain kind of diagnosis right 80% of the time, while his junior colleague is only right 60% of the time. Surely, when the senior doctor notes that his junior doctor disagrees with him about a particular patient, this should make him consider the possibility that he has made a mistake, and reduce his credence. Numbers as well as patterns of dependency matter. If I hold P, but disagree with a large number of people who independently of one another believe not-P, this might be very massive evidence that I have made a mistake, even if none of these other people qualify as my epistemic peer. If, however, they are not fully independent of one another, say because they influence one another, or tend to derive their beliefs from the same source, then I should assign less weight to their number (see Goldman 2001).

The significant property, I suggest, is the level of propositional justification for negative higher order beliefs, that is, the balance of evidence for these negative beliefs, as (D) says. Clearly, higher order evidence about performance error and principle error can itself be undercut or rebutted. Suppose I learn from a credible source that I have ingested a reason-distorting pill without noticing. But I am also told, from another credible source, that I happen to be one of very few people on whom the active ingredient has no effect. Or



suppose that I learn that, uncharacteristically, the source telling me about the pill has a strong incentive to mislead me (it turns out that he will win a large prize if he succeeds in making me waver in my answer to the mathematical problem). In these cases I receive undercutting evidence, not of my first order evidence, but of my higher order evidence. Or suppose that a trustworthy source tells me that my calculations are actually correct, despite apparently having ingested the pill, and repeatedly so.<sup>4</sup> This would tend to rebut the evidence that I have actually ingested a reason distorting pill, or it would undercut the testimonial evidence that this is so. When higher order evidence that would otherwise make me propositionally justified in a negative higher order belief is undercut or rebutted, I am not propositionally justified in that higher order belief, and thus not by (D) obliged to reduce credence in my first order belief.

Mostly, higher order evidence for negative higher order beliefs should make us reduce our confidence in our first order beliefs, but not always. Undercutting and rebutting evidence for negative higher order beliefs explain why we should sometimes not conciliate in disagreements. Suppose I believe that  $2 + 2 = 4$ , but to my astonishment you believe it makes 5. Our disagreement does provide higher order evidence that I have made a mistake, but this evidence is normally rebutted by other higher order evidence that I possess for the very high reliability of my elementary arithmetic skills. By contrast, small children, who don't have such rebutting higher order evidence, should reduce confidence when facing disagreement over simple math with their young peers. Compare to Christensen's restaurant case, where we disagree about how to split the bill after doing the calculations in our heads. Even if I know that I am pretty good at doing these calculations, the fact that you disagree is higher order evidence that I might have made a mistake, and in this case I possess no rebutting or undercutting evidence.

For convenience I will refer to the view I have sketched as the Higher Order Evidence Account of the epistemology of disagreement, or just the HOE account, though admittedly this label is a bit broad. The HOE account is a version of the total evidence view—both first order and higher order evidence matters for how we should respond to disagreement. But my version of the HOE account assumes a distinctive way in which first order evidence and higher order evidence come together and affect the rational credence of our beliefs. Higher order evidence affects first order credences as asserted in (D), but first order evidence for a particular proposition does not similarly affect the propositional justification of higher order beliefs regarding beliefs in that proposition. So, the HOE account asserts a distinctive asymmetry in the direction of influence between first order and

higher order levels (for discussion and further references, see Kappel (2017b)).

#### 4 Why the HOE Account Implies that We Should Not Conciliate in Deep Disagreement

Return now to deep disagreement, and let us apply the HOE account. It will be helpful first to elaborate the details of a case of deep disagreement, albeit still a fictitious and schematic one<sup>5</sup>:

Suppose that Ann is relying on scientific methods and standards of evidence for addressing the question of the age of the earth. So, Ann thinks that the earth is billions of years old. Her friend Beth, on the other hand, relies on biblical evidence concerning this very question, and on these grounds she thinks that the earth is much younger, merely some thousand years.

Consider first Ann's perspective. Ann relies on certain epistemic principles that support the use of various scientific methods and standards of reasoning used in physics and geology to determine the age of the earth. Call these principles  $EP_A$ . Ann thinks that there is a general explanation why  $EP_A$  are truth-conducive, which also explains why the methods and procedures  $M_A$  they pick out are reliable as to answering questions such as the age of the earth. This explanation, naturally, in part, appeals to features of the natural world and the way our cognition works. These further facts we know from scientific investigations, which in turn rely on epistemic principles and doxastic practices that are similar to  $EP_A$  and  $M_A$ . These are all parts of Ann's scientific perspective, and for Ann there is no escaping this perspective when she tries to argue for the various principles and doxastic practices she adheres to.

Consider now Beth's perspective. Beth thinks that there are truth-conducive epistemic principles that imply that relying on biblical evidence about the age of the earth is a reliable doxastic practice. Call these epistemic principles  $EP_B$ , and Beth's preferred practice  $M_B$ . Like Ann, Beth can offer a general explanation of why biblical evidence is reliable in such matters, that is, why  $EP_B$  is truth-conducive, and why  $M_B$  is reliable. Assume that Beth's explanation ultimately refers to certain religious assumptions: the scripture is a reliable source of information in these matters because, essentially, it provides divine testimonial evidence telling us the truth, and that truth is that the earth was created by divine fiat some thousand years ago.

Note that it is difficult for both Ann and Beth to go beyond defending their preferred epistemic principles by

<sup>4</sup> Thanks to a reviewer for bringing up this case.

<sup>5</sup> Thanks to a reviewer for prompting this.

appeal to principles that are both more fundamental and rationally acceptable for both of them. What would these principles be? How can Ann appeal to epistemic principles that (together with facts about the world) both support  $EP_A$ , and are at the same time rationally acceptable to Beth? This is hard to see. Conversely, it is difficult to see how Beth could provide a more fundamental defense of  $EP_B$ , which would appeal to premises that Ann could accept.

Consider now how Ann's and Beth's perspectives relate to one another. Ann's general considerations about reliable ways of amassing evidence about questions such as the age of the earth will not support the idea that biblical evidence has any merit whatsoever. On the contrary, the reasons supporting the truth-conduciveness of  $EP_A$  lends no support to thinking that taking the testimony from one select human religion among many is a reliable way to discover the truth of intricate questions like the age of the earth. So, according to Ann's perspective, Beth's preferred method of relying on biblical evidence will be seriously misleading and typically lead to false beliefs about the world. This in turn suggests that Ann should consider her own and Beth's preferred methods as not merely different, but also opposed. Similarly, given Beth's view about why  $EP_B$  is truth-conducive in this domain, it seems that Beth must hold that  $EP_A$  cannot be truth-conducive as well. If relying on biblical evidence is reliable in the way that Beth thinks, then attending to the evidence suggesting that the earth is billions of years old will be seriously misleading; going by this evidence takes us away from the truth. So, again Beth should view Ann's preferred methods as not only incompatible, but also opposed.

Could Ann and Beth adopt a more ecumenical perspective, such that they regard both scientific evidence and biblical evidence as bona fide forms of evidence?<sup>6</sup> On this construal, all they disagree about is whether one type of evidence outweighs the other in particular instances. Beth would fully accept the scientific evidence, but merely consider it outweighed by the biblical evidence, whereas Ann would consider the scientific evidence more weighty than the biblical evidence in the question at hand. For the reasons already given, I find this ecumenical perspective hard to make sense of, though there is no space for a detailed discussion. As indicated above, the obvious general explanation why biblical evidence might be reliable regarding questions such as the age of the earth seem to imply that scientific evidence is misleading concerning this question (though not necessarily in general).

Suppose now that Ann and Beth encounter one another and realize that they disagree about the age of the earth, and that the structure of their disagreement is as outlined above. Suppose also that they both understand how their

disagreement is based on their acceptance of very different epistemic principles and doxastic practices. How should Ann and Beth react to this disagreement? According to the HOE account, the decisive question will be whether the disagreement constitutes evidence that they have made a mistake, either by relying on a non-truth-conducive epistemic principle, or by making a performance error. Upon encountering the disagreement, would Ann or Beth become propositionally justified in believing either of these possibilities?

Consider first Ann's perspective. Should Ann think that her disagreement with Beth is evidence that she has made a performance error? Well, it is hard to see. For Ann, Beth's dissenting view is entirely due to Beth's adherence to doxastic practices that Ann considers wildly inappropriate, and highly likely to generate false beliefs about the true age of the earth. Why would Ann think that this is any indication that she, Ann, might have made an error when employing an otherwise reliable doxastic practice? For Ann, there is a perfectly good explanation of why Beth disagrees with her, and it has to do with Beth's adherence to certain epistemic principles  $EP_B$ , which, again, Ann considers wrongheaded. Similarly, why should the disagreement with Beth make Ann doubt the truth-conduciveness of her own epistemic principles? Why should the fact that someone else uses a wrongheaded epistemic principle and reaches a false belief be evidence that one's own epistemic principle is not truth-conducive? To illustrate with a crude, though hopefully valid, analogy: Suppose I believe, and take myself to have good reasons to believe, that your watch is quite inaccurate. I glance at my watch and form the belief that the time is 10 pm. You then tell me that the time is more like 8.15 pm. Why should this disagreement make me doubt that my watch is accurate, or that I have made an error in reading it?

From Beth's point of view, the story is similar. She takes herself to have reasons to believe that Ann's preferred epistemic principles are not truth-conducive when applied to questions such as the age of the earth, though scientific standards and methods may be fine in other domains. Since this is so, Beth has no particular reason to take her disagreement with Ann to indicate that her own epistemic principles might be at fault, or that she might have made a mistake when applying them. So, on the HOE account of disagreement, Beth should not reduce her confidence in her belief about the age of the earth.

In general terms the HOE accounts says the following about deep disagreement. In the general case, disagreement is prima facie higher order evidence of performance errors or principle errors, and in response to this one should modify one's credence towards less certainty. However, in deep disagreement, this higher order evidence is undercut by evidence stemming from knowledge of the nature of the disagreement. In so far as deep disagreements are symmetrical, neither party should conciliate.

<sup>6</sup> Thanks to a reviewer for raising this question.

The case above is, as I said, detailed but still schematic, and also hypothetical. Like the other examples mentioned, I intend this case to illustrate a case of deep disagreement, and make no claim that any actual case instantiates deep disagreement as defined here. Clearly, it would be preferable to work with an unequivocal detailed real-life case of deep disagreement. However, finding such a case is difficult for reasons we can clearly appreciate: deep disagreement is defined by the nature and epistemic status of the epistemic principles underlying disagreement. In general, however, we don't have easy empirical access to how these matters stand in actual cases. Hence, any claim that a particular case is an instance of deep disagreement as defined here is going to be hard to justify. I assume, however, that the phenomenon of deep disagreement is nonetheless of both practical and theoretical interest. First, it seems to me that some real-life cases are similar enough to deep disagreement to warrant discussing them under the assumption that they are indeed instances of deep disagreement, even if we don't have solid empirical evidence that they are. Second, cases of deep disagreement are of theoretical interest in the epistemology of disagreement. A full theory of the epistemology of disagreement should account for deep disagreement, as well as other kinds of disagreement, say peer disagreement, or disagreement between groups of individuals and so on.<sup>7</sup>

## 5 Peers, Personal Information, and Independence

I will now discuss certain aspects of the HOE account by relating it to three common themes in the disagreement debate: epistemic peerhood, personal information and independence.

### 5.1 Epistemic Peers

It is sometimes suggested that deep disagreements cannot be peer disagreements, and that peer disagreements cannot be deep (e.g. Siegel 2013).<sup>8</sup> Let us now consider whether this is true, and what significance it would have.

There are two main notions of epistemic peerhood that should concern us. On the first, two individuals A and B are epistemic peers with regard to a proposition P iff they have the same evidence pertaining to P, are equally competent in assessing the evidence, and have considered the evidence with the same care and attention (call this *evidential peers*). On a different notion, A and B are epistemic peers iff they

are equally likely to be right about the target proposition P (call this *probabilistic peers*). Clearly, when two individuals are evidential peers they would seem to be probabilistic peers as well, given a natural way of understanding what competence in assessing evidence is, but two individuals can be probabilistic peers without being evidential peers.<sup>9</sup>

Consider first probabilistic peers. Two individuals may accept *incompatible* substantive epistemic principles, and yet they are equally likely to be right, and thus probabilistic peers. This is possible when the substantive epistemic principles involved support doxastic practices that are less than perfectly reliable, in which case incompatible doxastic practices can be equally reliable, allowing for probabilistic peerhood. Two individuals can even, it seems, accept *opposed* substantive epistemic principles, and yet they are probabilistic peers, though if they understand that their disagreement involves opposed epistemic principles, they cannot regard one another as probabilistic peers.

Consider then evidential peers. Whether deep disagreement is compatible with evidential peerhood depends on details about what we take evidence to be. On an ordinary way of speaking about evidence, evidence consists of externally observed facts or reports of such facts that pertain to the truth of the target proposition (call this the *externalist notion of evidence*). Two individuals in a deep disagreement may then have the same evidence and thus be evidential peers if they are equally competent in scrutinizing the evidence and do so with equal care. If competence is couched in terms of the tendency to form correct beliefs given a body of evidence, then two individuals who accept radically different ways of analyzing evidence can still be equally competent. Thus, two individuals accepting incompatible or even opposed epistemic principles could be evidential peers, if we go by the externalist notion of evidence.

Often, however, evidence is thought of as an internalist notion. If two individuals observe the same set of external facts, but interpret or evaluate these facts in different ways, say because they have different background assumptions or consider a different set of explanations of their observations, then they have different evidence. It follows that they are not evidential peers, even if they do consider the same set of external facts with equal care, and are equally likely to be right about what those facts indicate about the relevant proposition. Thus, on the internalist notion of evidence, deep disagreement is generally not compatible with evidential peerhood, and one may question whether evidential peerhood can occur at all except in philosophical thought experiments.

<sup>7</sup> Thanks to a reviewer for raising these concerns.

<sup>8</sup> Note that Siegel's remark concerns deep disagreement in Fogelin's sense, not strictly the sense I have defined here.

<sup>9</sup> For both notions of peers, we should of course distinguish between someone being an epistemic peer and believing or knowing that someone is a peer.



Now consider the epistemic significance of this. As just noted, even if A and B are not evidential peers nothing follows about their probability of being right, or about the truth-conduciveness of their substantive epistemic principles or the reliability of the doxastic practices they rely on. Even individuals who are not evidential peers may still rely on equally reliable doxastic practices. This is significant, as it implies that we cannot assume that when A and B are not evidential peers then one of them is epistemically worse situated than the other. If A and B share the same set of substantive epistemic principles, then it seems reasonable to make an inference from lack of epistemic peerhood between A and B to one of them being epistemically worse off. However, if A and B fail to be evidential peers as a result of their acceptance of different substantive epistemic principles, then it is not clear that one must be epistemically worse off than the other. Even opposed epistemic principles can be equally good in terms of truth-conduciveness and reliability of underlying doxastic practices.

Note also that discussions of the significance of epistemic peerhood often presuppose that disputants have access to dispute-independent reasons to believe that someone is, or is not, one's peer. For example, Christensen's restaurant case specifies that independently of our disagreement, we have reason to take the other as at least our approximate peer. In part the significance of epistemic peerhood depends on this assumption. However, in deep disagreement we normally do not have dispute-independent reasons to say that those we disagree with are not our peers. In such cases, I might rationally believe you are neither my evidential nor my probabilistic peer, but this is not a dispute-independent verdict, but one that depends on my perspective being different from yours. Thus, attaching great significance to you not being my epistemic peer, but rather my epistemic inferior, is close to simply declaring that my perspective is right and yours is wrong. If we merely focus on the notion of epistemic peers, we cannot settle whether this inference is reasonable or not.

Finally, setting aside evidential and probabilistic peerhood, note that deep disagreement clearly involves a distinct form of epistemic parity. A and B's disagreement results from a clash of perspectives, but both perspectives involve basic substantive epistemic principles, and for both A and B it is true that independently of their perspectives they don't have any particular reasons to think they are better situated than their opponent. This parity seems significant, irrespective of the fact that it may not be captured in terms of evidential or probabilistic peers. Moreover, the parity of perspective invites the conciliationist intuition, that I mentioned in the introduction. Yet, the HOE account defended here suggests that the conciliationist intuition is a mistake. What matters on the HOE account is whether a case of disagreement generates strong, undefeated evidence for a negative higher order belief. This can be true of peer-disagreements,

but as we have seen, it might as well be true of non-peer disagreements. A deep disagreement displays a certain form of parity of perspective. Yet, cases of deep disagreement fail to provide a reason to conciliate, and this is because the higher order evidence provided by the disagreement is undercut or rebutted.

## 5.2 Personal Information

Consider then *personal information asymmetries*, another regular in the peer disagreement literature. It is often suggested that personal information may be a tie-breaker in peer disagreements. Imagine that you and I are roughly epistemic peers who disagree about some matter. I may know that I am not emotionally distressed, intoxicated, or cognitively impaired in various ways, but I don't know this about you, or at least I cannot be as certain. So, there is a relevant asymmetry that sets us apart, and on some views, this asymmetry may justify a non-conciliatory response on my part.

Now, we can assume that personal information asymmetries are not relevant in cases of deep disagreement, since such differences are not their defining characteristics. Yet, it is worth noting that the HOE account supports a specific explanation of the relevance of personal information for disagreement. It is not that personal information as such tips the balance making one agent's epistemic position stronger than the other. Rather, personal information works by defeating higher order evidence, that is, by constituting rebutting evidence or undercutting evidence, as we have seen above. To illustrate, I might know that my elementary arithmetic skills are normal and that I am dead earnest in my belief that  $2 + 2$  makes 4. When you claim that it sums up to 5, I am not certain whether you are playing a prank on me to test my patience or philosophical acumen. So, the higher order evidence otherwise provided by the disagreement tends to be rebutted by my knowledge that my skills are normal, and tends to be undercut by my suspicion that the disagreement is not genuine.

## 5.3 Independence

Consider finally *Independence*. As Christensen originally stated this principle, it says (Christensen 2011, p. 1):

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

When applied to cases of deep disagreement, Independence seems to imply that A should bracket her perspective where this conflicts with B's. Moreover, A should bracket her perspective not because of the specific features

distinctive of deep disagreement (e.g. the parity or the lack of non-circular reasons in favor of one's perspective), but simply because A has come across someone who disagrees. Many have worried that Independence is implausibly strong in cases of comprehensive disagreements (when we disagree about many things) or repeated disagreements (when we repeatedly disagree with the same individuals) as it seems to require that we bracket very many beliefs, even to the extent that evaluation of others' expressed beliefs make no sense.

Christensen offers a couple of reasons why Independence is not counter-intuitive, despite appearing so at first sight (Christensen 2011). Suppose that you and I disagree about some matter in a particular domain, and that our disagreement is *comprehensive*: we disagree not only about P, but also about propositions Q, R, S, T and so on in the same domain, where our beliefs in these propositions are relatively independent of one another. Now, how should I respond, once I realize that our disagreement is comprehensive? If I bracket all these beliefs, how can I even evaluate your performance in the domain? Christensen offers an answer to this problem. Suppose that I know that there are a few cranks around, that is, people who confidently assert things in the domain, but who are completely incompetent, but I also know that I am not one of them. It is overwhelmingly unlikely that a comprehensive disagreement like this one could occur between two fairly competent individuals. It can only arise when one encounters a crank. Since I am not a crank, you must be, and I should not adjust my beliefs. This line of reasoning respects Independence.

While plausible, it seems to me that this defense of Independence does not work in deep disagreement. Here we might of course also have disagreements covering many different beliefs in a domain, but the beliefs that we disagree about are not independent. When you and I deeply disagree about a range of beliefs, the correctness of your swath of beliefs is conditional on the soundness of your perspective, and so is mine. Of course, it is still true that if I can assume that I am fairly competent in the domain in question, then I can infer that you are not. But this inference seems to come down to the assertion that my perspective is correct, and yours is not. It is not clear that this complies with Independence.

Christensen also notes that Independence only says that one's evaluation should proceed independently of the disputed matter, but not anything about how one should revise one's beliefs. There are many options here. One is that:

(A) Insofar as the dispute-independent evaluation fails to give me good reason for confidence that I'm better informed, or more likely to have reasoned from the evidence correctly, I must revise my belief in the direction of the other person's. (Christensen 2011, p. 15)

Clearly, this will imply that one should often revise one's credence extensively in comprehensive disagreements and in deep disagreement. But instead of (A) we might accept

(B) Insofar as the dispute-independent evaluation gives me good reason to be confident that the other person is equally well-informed, and equally likely to have reasoned from the evidence correctly, I must revise my belief in the direction of the other person's. (Christensen 2011, p. 15)

The point is that (B) entitles one to a non-conciliatory response as long as one does not have dispute independent reasons to believe that one's opponent is as epistemically well situated as oneself. So, (B) is less prone to widespread skepticism when applied to deep disagreement.

This moves us closer to what the HOE account says, but there are still significant differences. First, the applicability of (B) is triggered by the mere fact that we disagree. The set of beliefs singled out for bracketing is determined by what others happen to disagree about. The ease with which one's beliefs can be up for bracketing is, in part at least, what seems implausible about Independence. On the HOE account, by contrast, the decisive feature is undefeated higher order evidence of sufficient strength, not the mere fact of disagreement. Sometimes we should bracket beliefs in our reasoning, say if we have reason to suspect that a belief is the result of a performance error, an unreliable process, or if the content of our belief is the subject of our inquiry. Merely disagreeing with someone is not enough. The HOE account suggests that unlike many other cases of disagreement, deep disagreement does not provide undefeated higher order evidence that my first order beliefs are epistemically flawed, so there is no pressure to revise credence in first order beliefs.

Second, while (B) gives the intuitively more plausible results, there is something unsatisfactory about it. It says that when I have no dispute-independent evidence showing that you are worse situated than me, I can remain unaffected. But why not say that, in such cases, for all I know, you might as well be just as well placed as I am, and shouldn't this make me pause? The HOE account offers a slightly different story: it is not disagreement or epistemic parity as such that matters, but the force of undefeated higher order evidence about principle errors or performance errors that does the work, and there is a general explanation of why such higher order evidence forces reductions of first order credence.

## 6 Further Reflections

I will end with a brief reflection on certain residual worries that the HOE account of deep disagreement may leave us with. Shouldn't one be moved when confronted with others who have an entirely different perspective leading them to

form very different beliefs about the world? I will briefly consider two ways of characterizing this residual epistemic angst.

One suggestion would be to think of deep disagreement in terms of debunking arguments.<sup>10</sup> Suppose that certain contingent features *F* (such as upbringing, embedment in a particular cultural or intellectual environment, or simply evolution) have made us adopt particular substantive epistemic principles. Suppose that it is also true that *F* is not correlated with the truth-conduciveness of the chosen principles. Or suppose that the causal story involving *F* even tells against the truth-conduciveness of the epistemic principles. Thus, knowing about the causal history of our acceptance of the epistemic principles in question leads to their being debunked. We can view this debunking as a special case of higher order evidence to the effect that certain substantive epistemic principles are not truth-conducive, or that we lack reasons to think that they are truth-conducive. Clearly, as I have stipulated the notion, deep disagreements do not support debunking arguments. When *A* and *B* are in a deep disagreement it is not the case that *A* receives higher order evidence similar to what happens in a debunking argument. Nothing in a deep disagreement provides *A* with a debunking explanation of her basic substantive epistemic principles.

Another suggestion is to think of deep disagreements in analogy to skeptical scenarios. Skeptical scenarios are hypothetical possibilities carefully crafted to so that no possible evidence can rule out their actuality. Similarly, deep disagreements indicate the possibility that one's perspective is entirely wrong, yet one does not have evidence to rule out this possibility. Unlike skeptical scenarios, cases of deep disagreement are assumed to be real; they are assumed to be cases in which we actually confront someone who has a radically different perspective. Being in a case of deep disagreement is like encountering an actual skeptic, not merely contemplating a skeptical scenario.

Shouldn't this affect us? For *A*, the fact that *B* deeply disagrees is a coherent possibility, but if what was argued above is correct, there is nothing to this sort of case that provides undefeated higher order evidence that *A* has committed a performance error or a principle error. So, deep disagreement is not a reason why *A* should reduce her first order credence. In response it might be said that the analogy to a skeptical argument is different, and goes as follows. *A*'s encountering *B* in a deep disagreement indicates that the falsity of *A*'s perspective is a coherent possibility, and yet there is no cogent evidence available to *A* that can rule out the possibility that *A*'s perspective is wrong. Cogent evidence here means evidence that is sufficiently independent of *A*'s perspective. But since *A* cannot rule out the possibility that

her perspective is false, she should reduce confidence that her perspective is correct, and reduce confidence that her beliefs depending on this perspective are correct.

We can resist this line of reasoning if we assume that having high credence in beliefs that depend on a particular perspective does not require independent evidence ruling out the possibility that the perspective is false. Exactly this is what we must assume anyway to avoid skepticism, or so it is commonly thought at any rate. So, we can agree that deep disagreement exemplifies the features of skeptical scenarios, which are cleverly thought out to leave us with no evidence to rule them out. We should nonetheless resist drawing skeptical conclusions from deep disagreement.

## 7 Summary

A brief summary of the main line of argument in the paper may be useful. I have stipulated that *A* and *B* are in a deep disagreement when they disagree about some proposition *P*, where this disagreement depends on *A* and *B* relying on opposed substantive epistemic principles that are also basic. Substantive epistemic principles are opposed when one principle implies that the doxastic practices recommended by the other are seriously misleading and unsuitable for the cognitive task in question. Substantive epistemic principles are also basic where there are no other principles that one can resort to argue that the principles is truth-conducive. We might think of patterns of deep disagreements as being involved in disagreements over science and religion, climate change and other polarized disputes, though it is hard to substantiate such empirical claims.

How should we rationally adjust our doxastic attitudes in a deep disagreement? I have suggested that in general disagreement provides higher order evidence that we might be relying on non-truth-conducive epistemic principles (principle error), or might have made an error in relying on a truth-conducive epistemic principle (performance error). When we have sufficiently strong such evidence, we should move towards uncertainty in our first order belief. This is why we should often conciliate in disagreements, whether they are peer disagreements or not. In deep disagreements, however, the higher order evidence provided by the disagreement is undercut: provided we understand the structure of the disagreement, the higher order evidence yielded by the disagreement does not provide strong, undefeated evidence that we have made a principle error or a performance error, even if we have. Hence, we should not conciliate in deep disagreements. Finally, I have argued that this account of the epistemology of deep disagreement is more plausible than what one would get by relying on the common notions featuring in the disagreement debate: epistemic peerhood, personal information and independence.

<sup>10</sup> Thanks to Martin Kusch for suggesting this.

**Acknowledgements** Work on this paper was supported by the Independent Research Foundation Grant ‘Complex Disagreement’. Earlier versions was presented at workshops in Copenhagen and St. Andrews. Thanks to Alex France, Bjørn Hallsson, Fernando Broncano-Berocal, Frederik J. Andersen, Giacomo Melis, Jeroen de Ridder, Thirza Lagewaard, Jesper Kallestrup, Jessica Brown, Jon Matheson, Josefine Pallavicini, Katherine Hawley, Martin Kusch, Michael Lynch, Scott Aikin, Thomas Raleigh for useful comments. Warm thanks also to David Gooden and Pat Bondy and anonymous reviewers for this issue of *Topoi* for the useful suggestions.

**Funding** This study was funded by Danish Research Councils (4001-00059BFKK).

### Compliance with Ethical Standards

**Conflict of interest** Author Klemens Kappel declares that he has no conflicts of interest.

**Research Involving Human and Animal Participants** This article does not contain any studies with human participants or animals performed by any of the authors.

### References

Alston WP (1986) Epistemic circularity. *Philos Phenomenol Res* 47(1):1–30

- Bergmann M (2009) Rational disagreement after full disclosure. *Episteme* 6(3):336–353
- Christensen D (2011) Disagreement, question-begging and epistemic self-criticism. *Philos Impr* 11(6):1–22
- Goldman AI (2001) Experts: which ones should you trust? *Philos Phenomenol Res* 63(1):85–110
- Fogelin R (1985) The logic of deep disagreement. *Informal Logic* 7(1):3–11
- Kappel K (2012) The problem of deep disagreement. *Discip Filo* 22(2):7–25
- Kappel K (2017a) Fact-dependent policy disagreements and political legitimacy. *Ethical Theory Moral Pract* 20(2):313–331
- Kappel K (2017b) Bottom up justification, asymmetric epistemic push, and the fragility of higher order justification. *Episteme*. <https://doi.org/10.1017/epi.2017.19>
- Kelly T (2010) Peer disagreement and higher order evidence. In: Feldman R, Warfield TA (eds) *Disagreement*. Oxford University Press, New York
- Kelly T (2013) Disagreement and the burdens of judgment. In: Lackey J, Christensen D (eds) *The epistemology of disagreement: new essays*. Oxford University Press, New York
- Lynch M (2010) Epistemic circularity and epistemic incommensurability. In: Haddock A, Millar A, Pritchard D (eds) *Social epistemology*. Oxford University Press, New York
- Siegel H (2013) Argumentation and the epistemology of deep disagreement. *Cogency* 5:135–170