



Moral Disagreement and Higher-Order Evidence

Kappel, Klemens; Andersen, Frederik J.

Published in:
Ethical Theory and Moral Practice

Publication date:
2019

Document version
Peer reviewed version

Citation for published version (APA):
Kappel, K., & Andersen, F. J. (2019). Moral Disagreement and Higher-Order Evidence. *Ethical Theory and Moral Practice*, 22(5), 1103-1120.

Moral Disagreement and Higher-Order Evidence[#]

By Klemens Kappel and Frederik J. Andersen

Forthcoming in *Ethical Theory and Moral Practice*, October 2019

Abstract

This paper sketches a general account of how to respond in an epistemically rational way to moral disagreement. Roughly, the account states that when two parties, A and B, disagree as to whether p , A says p while B says not- p , this is higher-order evidence that A has made a cognitive error on the first-order level of reasoning in coming to believe that p (and likewise for B with respect to not- p). If such higher-order evidence is not defeated, then one rationally ought to reduce one's confidence with respect to the proposition in question. We term this *the higher-order evidence account* (the HOE account), and present it as superior to what we might call *standard conciliationism*, which holds that when agents A and B disagree about p , and are (known) epistemic peers, they should both suspend judgement about p or adjust their confidence towards the mean of A and B's prior credences in p . Many have suspected that standard conciliationism is implausible and may have skeptical implications. After presenting the HOE account, we put it to work by applying it to

[#] We have presented this material at workshops in London, Lisbon, and Copenhagen. Thanks to audiences there for helpful comments, in particular Julien Dutant, Alexander Heape and participants in LVU18, Lisbon.

a range of cases of moral disagreement, including those that have featured in recent debates assuming standard conciliationism. We show that the HOE account support reasonable, non-skeptical verdicts in a range of cases. Note that this is a paper on moral disagreement, not on the HOE account, thus the account is merely stated here, while defended more fully elsewhere.¹

1. Introduction

Much discussion of moral (peer) disagreement explores various versions of what we might call *standard conciliationism*, holding that when agents A and B disagree about p , and are (known) epistemic peers, they should both suspend judgement about p .² The equal weight view holds that A and B should rather adjust their confidence towards the mean of A and B's prior credences in p . For discussions of moral (peer) disagreement and implications, see e.g. (Besong 2014; Christensen 2007; Christensen 2011; Matheson 2016; McGrath 2008; Rowland 2017b; Vavova 2014). Two individuals are epistemic peers with respect to p , when they have access to the same body of evidence, are equally adept in interpreting the evidence, and have scrutinized it equally carefully. As has been observed, perhaps no individuals are epistemic peers, strictly speaking, and hence we might want to talk about

¹ (Kappel 2018).

² Note that the term 'conciliationism' is also used in a somewhat broader sense, denoting views stating that in peer disagreements one should reduce one's confidence in the disputed belief.

approximate peers instead, see e.g. (Lackey 2010). We will set these complications aside, and henceforth simply talk about conciliationism and epistemic peers, unless otherwise indicated. Much current discussion in moral epistemology explore the epistemological implications of conciliationism, which may be a thorough moral skepticism.

This paper outlines what we hope to be a more plausible general view of disagreement, *the higher-order evidence account* (HOE account). We argue that the HOE account is both better motivated than standard conciliationism, and that it supports intuitively plausible verdicts in a range of different cases of moral disagreement.

The HOE account argues, first, that disagreement generally provides defeasible higher-order evidence that one has made an error; either by following a faulty belief forming procedure for collecting or interpreting evidence (what we term *a principle error*), by making a mistake in following a good one (*a performance error*), or by failing to consider relevant evidence (*an evidence error*). Second, when such higher-order evidence is strong enough, and not defeated, one should rationally reduce confidence in one's first-order view.

After briefly outlining the HOE account (section 2), we put the HOE account to work on the problem of moral disagreement. We do so by distinguishing various cases of moral disagreement, all featuring different structures (section 3). In applying the HOE account to these cases, we will ask whether these kinds of disagreements give an agent strong undefeated higher-order evidence to the effect that she has made a mistake on her first-order level of reasoning. We hope that this

will show that the HOE account provides intuitively plausible results in various types of disagreements, and that the HOE account is overall more plausible than the standard conciliationism.³

2. Higher-order evidence and disagreement

In this section we briefly outline the HOE account, though a full defense of it is the project of other work.⁴ The general idea that disagreement counts as higher-order evidence is familiar, yet we think that our specific version of the view is new, and as will be clear shortly, various of these details matter for how the HOE account responds to moral disagreement. For similar or related views, see e.g. (Bergmann 2009; Horowitz 2014; Kelly 2010; Kelly 2013a).

To appreciate the significance of higher-order evidence it is helpful first to focus on a case that does not involve disagreement:

³ Throughout the paper, we will assume that moral intuitions and moral beliefs can be right or wrong/true or false, and that we are epistemically committed (as well as morally committed) to have true moral views. So, in effect we treat moral beliefs and moral intuitions just like ordinary factual beliefs and factual intellectual seemings, or dispositions to judge in certain ways. We also freely assume that rational credences in propositions come in quantitative degrees, i.e. a credence in p can be represented by a number between 0 and 1.

⁴ This section is based on material first published in (Kappel 2018).

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.⁵

The Pill invites the following intuition: after being informed about my ingestion of the pill, I am no longer rationally highly justified in my belief that p . I should adjust my confidence so that it reflects my evidence concerning the high likelihood that my reasoning is impaired because of the pill. It would, for example, be 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.⁶ In other words, upon being told about the pill, I should reduce my first-order credence.

Assuming that we accept this basic intuition about the case, note a couple of additional things. First, even if I did in fact get my first-order evidence right, the advent of the evidence that I have ingested the reasoning-distorting pill should lead

⁵ This case is adapted from (Christensen 2011, pp. 5-6).

⁶ Cf. Horowitz's discussion of epistemic akrasia (2014).

me to significantly reduce my confidence in my object-level belief. Suppose that the reasoning-distorting pill gives me a 90% chance of being affected and a 10% chance of being unaffected. In fact, I am lucky and my reasoning capacities are not impaired, so I correctly see that E provides very strong evidence for p . Even so, we submit, I should significantly reduce my confidence in p . Second, I should reduce confidence even if my higher-order evidence is misleading. Suppose that, quite uncharacteristically, and in ways I cannot detect and have no reason to suspect, my credible source is wrong. In fact, I have not ingested the reasoning-distorting pill. Yet, I should reduce my confidence that I am right.

Third, and very important for our discussions below, my higher-order evidence can be rebutted or undercut (see e.g. Pollock 1970, 1986). Suppose I learn from a credible source that I have ingested a reasoning-distorting pill without noticing. I am also told, however, 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 have received evidence undercutting my higher-order evidence. Or suppose that a trustworthy source tells me that my reasoning in these kinds of cases are actually correct, despite apparently having ingested the pill, and repeatedly so. This would tend to rebut the evidence that I have actually ingested a reasoning-distorting pill, or it would undercut the testimonial evidence that this is so. Generally, 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 obliged to reduce credence in my first-order belief.

Note, finally, that the epistemic impact of the higher-order evidence in The Pill is quite significant. Intuitively, if I believe that the reasoning-distorting pill makes me just as likely to be right as wrong in my answer to the mathematical puzzle, then it seems that I should adjust my confidence in p from about full certainty to 0.5. If, however, I believe that p is just one of many possible answers to the problem, all but one is false, and I believe that the pill makes me pick a random answer, then my credence that I have the right answer should be less than 0.5. Indeed, it seems that I should be almost entirely certain that I got the wrong answer. This remarkable effect of higher-order evidence arises despite the fact that the higher-order evidence in question has no bearing on the first-order proposition. By itself, the fact that I have ingested a reason distorting pill does not speak to the truth of the mathematical proposition in question – it does not bear on the truth or falsity of p , nor the evidential relation between E and p , and it may still be true that my evidence E entails p , and further I may even have correctly apprehended that E entails p . See also the discussion in David Christensen’s paper “*Formulating Independence*” (2019).

We suggest that disagreements generally provide higher-order evidence of this kind. When I believe p , and you believe that p is false, this is generally evidence that I may have made a mistake when forming my belief that p . The fact that you reject p is *prima facie* first-order evidence that p is false, which should be weighed against my evidence suggesting that p is true. But our disagreement is also higher-

order evidence that I may have made a mistake: I may have employed a wrong method or principle when collecting or assessing the evidence (a principle error), I may have employed a right method or principle but made an error in my application of it (a performance error), or there may be relevant evidence that I should have considered (an evidence error).⁷ This is all higher-order evidence that there might be something wrong with the way I have formed my belief. As we have seen, when undefeated, higher-order evidence can significantly impact one's first-order credences. This is the HOE account in its most basic form.

2.1 The HOE account, non-conciliationism, and standard conciliationism

Though the HOE account is conciliationist in spirit, it formally counts as non-conciliationist, in so far as it admits there to be cases of peer disagreement where one should not conciliate (see discussion below). We will now briefly compare the HOE account to other non-conciliationist views. A prominent class of these are what one might call *first-person views*: I am entitled to accord my own basic moral intuitions greater epistemic weight than those of others, and thus I can sometimes

⁷ Of course, epistemic peers cannot have different evidence, and maybe peers cannot differ with respect to whether they commit a principle error. This leaves performance errors as the only option for explaining peer disagreement, which by itself seems plausible. But ordinary disagreements, of course, often involve different sets of evidence, and a general theory about disagreement should cover such cases.

remain steadfast even in the face of moral peer disagreement (Enoch 2010; Wedgewood 2010). The main problem with these views, we contend, is the difficulty of motivating that one's own basic moral intuitions are epistemically privileged vis-à-vis those of a similarly situated peer. We also find it hard to see how the higher-order evidence that one might be mistaken in one's moral intuition can be ignored. The HOE account is more plausible than the various first-person views because it avoids these problems. Other non-conciliationist views focus on the degree of epistemic justification of one's view by first-order evidence. Roughly, if A's belief that p is epistemically justified by a given pool of first-order evidence, and B's belief that not- p is not, then A should not conciliate. We count Lackey, Kelly and Weatherson (Lackey 2010; Kelly 2010; Weatherson 2019) among those championing this sort of view. Again, the main challenge, we think, to these views is that they do not acknowledge the significance of higher-order evidence. In *The Pill* I am presumably maximally warranted by my first-order evidence to believe p , and yet, intuitively because of the higher-order evidence I should lower my credence in p considerably. In recent debates many have denied, in various ways, that higher-order evidence can rationally affect what one should rationally believe on the basis of first-order evidence – the levels split up as far as epistemic rationality is concerned (Lasonen-Aarnio 2014; Titelbaum 2015; Weatherson 2008, 2019; Williamson 2011). A full defense of the HOE account needs to address this batch of objections. Unfortunately, we cannot do so here, but see discussions in (Horowitz 2014; Kappel 2019; Sliwa and Horowitz 2015a, 2015b).

Let us now briefly compare the HOE account to standard conciliationism. First, note that standard conciliationism is often motivated by considerations about the significance of higher-order evidence constituted by disagreement. However, as standard conciliationism is stated, the strength of higher-order evidence drops out of the picture; all that matters is whether one is facing an epistemic peer or not. So, standard conciliationism displays no sensitivity to whether the higher-order evidence arising from a peer disagreement is defeated or not. Instead, how one should respond to disagreement is a simple function of whether the agent one disagrees with is a (known) epistemic peer. For reasons that should become clear below, this makes the HOE account more plausible than standard conciliationism.

Second, standard conciliationism is restricted to cases of epistemic peerhood, while the HOE account is general and not tied to such a restriction; it is not even tied to cases of disagreement. Yet, disagreements with non-peers seem clearly epistemically significant. Suppose that I am a medical doctor with a known track record of getting a certain type of diagnosis right 80% of the time. My junior colleague has a known track record of 60%. It turns out that we disagree about a particular patient; I am initially quite confident that the patient has the diagnosis in question, but my colleague is sure that this is not the case. Intuitively, this should make me reduce my confidence that I am right – after all, my colleague is right 60% of the time, and this should give me some reason to think that I might have made a mistake. Yet, my junior colleague is not my peer on this question. The HOE account easily explains why I should pause: my disagreement with my colleague is higher-

order evidence that I might have made a mistake, which should lead me to reduce my confidence.

Third, standard conciliationism concerns two-person cases, leaving it open how one should respond to cases involving more than two individuals. Suppose I believe p , and realize that many reasonably competent individuals independently of one another believe not- p . For reasons originally discovered by Condorcet, this would be very impressive evidence that not- p is true, and thus very strong higher-order evidence that I have made a mistake. Likewise, if I believe p , and you believe not- p , but it turns out that many other competent individuals independently of one another side with me, then again this is strong evidence for p , even if none of the other individuals are as competent as I am with respect to the issue we disagree about.⁸ The HOE account can accommodate what seems to be epistemically significant facts about independence: suppose that I hold p , and it turns out that I disagree with many other individuals, but they do not form beliefs independently of one another – in fact, they all blindly follow just one individual. This is a very different situation from the previous one: the evidence provided from a disagreement featuring this structure is rather similar to me just disagreeing with one individual.

⁸ For discussions of Condorcet's Jury Theorem, and the significance of independent belief formation and consensus, see e.g. (Lackey 2013; List 2001; Goldman 2001).

Fourth, as we think of it, the HOE account is not wedded to the *Independence Principle*. As Christensen originally stated this principle, it says:

Independence Principle. 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.⁹

The Independence Principle seems intuitively very plausible and is motivated by intuitions about question-begging. When A and B disagree about p it seems that A cannot use his belief that p to reason as follows: p , so B must be wrong in her belief that not- p , hence there is no reason to modify my own belief. Clearly, this pattern of reasoning seems question-begging. Yet, the Independence Principle remains controversial.¹⁰ We claim that the HOE view is not wedded to the Independence Principle in the form stated, though clearly some restrictions regarding question-begging patterns of reasoning will need to be adopted by any plausible view. For example, in *The Pill* I cannot use my belief that E supports p to argue that I have

⁹ (Christensen 2011, p.1).

¹⁰ For discussion, see e.g. (Christensen 2009, 2011; Lord 2014; Moon 2018; Sosa 2010). For Christensen's most recent defense of a modified version of the Independence Principle, see (Christensen 2019).

not been drugged after all. This is both because my mathematical evidence E does not bear on whether I am drugged or not, but also because using first-order evidence that p to dismiss evidence that one has misjudged the evidence that p seems question-begging.¹¹

Fifth, the HOE account does not seem to depend on whether we accept permissivism or uniqueness about first order-evidence (Kelly 2013b; Matheson 2011; Rosa 2012; White 2013). Think again about the Pill. Surely, we can assume that the mathematical first-order evidence considered on its own uniquely determines that full belief in p is the uniquely rational attitude to have. Yet, even granting this, the higher-order evidence should make one reduce one's first-order confidence. Assume then that permissivism is true, and suppose that my initial level of credence in some first-order proposition p is one of several permitted attitudes. I then receive higher-order evidence that I may have made a mistake when forming this attitude. Shouldn't that make me pause, make me less certain in my belief, even if there are other permitted attitudes?

We cannot present a full defense of the HOE account here, but we want to note again that the core idea – that higher-order evidence can impact first-order epistemic rationality even when the higher-order evidence in question does not bear on the truth of the first-order proposition – is not without prominent critics. See in

¹¹ So, we submit that the HOE account will support restrictions somewhat similar to but not identical to the Independence Principle. We cannot develop the details of this here, but see (Kappel and Andersen unpublished manuscript).

particular (Lasonen-Aarnio 2014; Horowitz 2014; Kappel 2019; Sliwa and Horowitz 2015a, 2015b; Titelbaum 2015; Weatherson 2019; Worsnip 2018).

3. Moral disagreement and higher-order evidence

We now apply the HOE account to moral disagreement. We don't assume that moral disagreement is fundamentally different from disagreements in other domains, granted of course, that moral convictions are just ordinary beliefs. However, recent discussion of moral disagreement has featured discussion of a range of cases, most notably what we below term *evil disagreement*, *comprehensive disagreement*, *deep disagreement* and *fundamental disagreement*. Evil disagreements and fundamental disagreements are, arguably, distinctive for the moral domain. Also, the current debate has featured the suspicion that standard conciliationism is implausible because it leads to skepticism, and this is in part because epistemic peerhood may be more difficult to dismiss in the moral domain than elsewhere. These are general reasons for applying the HOE account to the moral domain in the hopes of showing that it supports more plausible answers than standard conciliationism.

One important implication of the HOE account is that not all cases of disagreement will be treated equally; what basically matters is the extent to which disagreement provides undefeated higher-order evidence of a certain kind which in turn depend on the structure of the disagreements. Thus, how one should react to disagreements depend on the details of the structure of the disagreement, and not

merely on whether the disagreement is between peers. We now present a variety of types of moral disagreement, including the cases discussed in the literature.

3.1 Local benign disagreements

Consider the following:

Local benign disagreement. A clearly intuitis that criminal offenders deserve punishment, while B has an equally clear intuition that criminals do not *per se* deserve punishment. Other disagreements between A and B are either irrelevant to the issue at hand, or they are direct consequences of these conflicting moral intuitions.

It seems that some actual cases of moral disagreement have this structure. People just appear to have different moral intuitions about specific matters, though the exact extent to which this is true is of course a difficult empirical question. This case is a local disagreement in that it concerns one specific moral intuition, i.e. a moral judgement about a particular case, or as in the case above, a more or less general moral principle, where this judgement is based on an intuition. We take an intuition to be a kind of intellectual seeming; an urge to see a certain proposition as true, though not based on inference from other propositions. The case is also a benign form of disagreement because the intuitions that set A and B apart are not evil or spurious in any way, i.e. the intuitions are not by themselves evidence of moral depravation or insensitivity to moral truth (see discussion below). From A's

perspective the disagreement with B gives her higher-order evidence of some force that she has made a mistake in her first-order reasoning, most likely that she has relied on a way of forming moral intuitions which is not truth-conducive, or perhaps it is truth conducive, but then in this particular case she has made a performance error. This is A's higher-order evidence. As the case is stated, A has no evidence that rebuts or undercuts this higher-order evidence. Hence, we contend that A should reduce her confidence in the disputed proposition.

Note that for A, the fact that she intuits in a certain way, is not good evidence that rebuts or undermines her higher-order evidence stemming from the disagreement with B. A has a strong moral intuition supporting p , and thus she believes p . But A's belief that p , and her access to (introspective) evidence that she herself has the seeming that p , makes no headway against the suspicion that there might be something wrong with her belief that p , or the process upon which it is based, i.e. the moral intuition that p . So, A's own intuition is no help against the hypothesis that her intuitions are misleading her. But A could have rebutting or undercutting evidence against the higher-order evidence. This would be evidence that suggests that A's way of forming moral intuitions in the relevant domain is actually highly reliable. Suppose, for example, that an independent moral expert confirms for A that she has a remarkable accuracy in her intuitive judgements about this kind of cases. This would be evidence rebutting the higher-order evidence. Or suppose, instead, that some credible source correctly informs A that though B is generally reliable, someone has slipped a drug into her coffee such that on this

occasion B is entirely unreliable in her judgements regarding p . This would be undercutting evidence.

3.2 Local evil and bizarre disagreements

Consider then:

Local evil disagreements. A has a clear intuition that all humans deserve equal respect, whereas B thinks that only the strong and successful deserve respect, while the weak deserve the death penalty. Other disagreements between A and B are either irrelevant to the issue at hand, or they are direct consequences of these conflicting moral intuitions.

It may be difficult to point to actual cases of evil disagreements, but moral disagreement over racist, misogynist and supremacist views may come to mind. However, any candidate case of evil disagreement is likely to be contestable, not only because some might not find the views evil or abhorrent, but also because one might wonder whether candidate views are moral views at all, as distinct from some preferences of a non-moral kind.

Again, we can say that this disagreement provides A with higher-order evidence that she may have made a mistake in forming her intuition.¹² However,

¹² Disagreements of a similar sort is discussed by Richard Rowland (2017a) and Kieran Setiya (2012). Setiya writes about moral monsters, i.e. individuals who have horrible moral beliefs. Clearly,

B's intuition is evil, in the sense that it not only seems to support a false moral judgement, but one that is malignant as well. How should this affect A's rational response to the disagreement? For A it is not just that B's intuition differs from A's, thus it must be wrong. Rather, B's intuition has other features that qualifies it as evil. Assume that, independently of A's particular intuition about equal respect, A has reason to think that B's intuition is evil. Suppose further that, by default, evil intuitions are had by evil individuals and that evil individuals are not truth-sensitive (reliable) in the relevant moral domains. Then, it is not just that A and B disagree, it is also the case that the very content of B's view constitutes a reason to suspect that B is entirely wrong in her intuition and moral judgment. In effect, this is undercutting evidence. The higher-order evidence provided by the disagreement is undercut by the content of B's belief. Thus maybe, even if A's reason for thinking that B's intuition is evil in part relies on A's assumption that all humans deserve equal respect, one might suggest that A should remain fairly steadfast. This is at least comparable to a revised version of Christensen's restaurant case.¹³ In order to see what amount each of us needs to pay we add 15% and then divide the bill by 5. I say that the correct number is 43. You say that it is 583, which is more than the

on behalf of conciliationism it might be objected that moral monsters are not peers. One problem with this response is whether the judgement that moral monsters are not peers is permitted given that one accepts the Independence Principle. See also our brief mention of the weak and strong evidence requirement below.

¹³ Discussed by (Christensen 2007, p. 199).

total bill. I realize that not only do we disagree, but there must be something entirely wrong with your reasoning – you have made a mistake out of the ordinary. Christensen uses this case to show that one can, in some cases, remain steadfast in a disagreement, and disregard the opinion of one's peer, without violating the Independence Principle. This of course presupposes that when A is evaluating B's moral response, A is permitted to rely on her general moral beliefs and other considerations, relative to which B's particular intuition comes across as evil. But on the HOE account, there is nothing to suggest that A could not do so. The HOE account does not commit to or rely upon the Independence Principle.

Reflection on this case suggests that B's intuition need not be evil for A to be in a position to discount it, i.e. for the higher-order evidence provided by the disagreement to be undercut or rebutted. Consider a variant of the case:

Local bizarre disagreement. B strongly intuits and believes that drinking coffee on Mondays is intrinsically morally better than drinking coffee on Tuesdays. A intuits no such thing and believes that there is no difference in intrinsic moral goodness to drinking coffee on different weekdays. Suppose that, apart from this odd disagreement, A and B have very similar moral intuitions, very similar general moral views, and that their views are equally coherent (as similar as they can be).

Should A reduce her confidence? It seems not, though B's intuition is by no means an evil intuition, and even if in other cases A has reason to believe that B's capacity

of forming moral intuitions is highly reliable, A would have some reason to believe that B's intuition is off track. We submit that this is because B's view is bizarre in the sense that it is very difficult to imagine an account of right and wrong making-features that would make B's intuition about the ethics of coffee drinking come out true.¹⁴ In the absence of this, A should discount B's intuition as misleading, and A should consider the evidence provided by the disagreement as equally misleading. So, A's considerations about the possibility of giving a coherent account of how B might be right in his moral intuition undercuts the higher-order evidence provided by the disagreement.

Again, one might say that this reasoning violates the Independence Principle, if part of A's reason for thinking that B is wrong rests on A's intuition about the moral neutrality of coffee on different weekdays. This is using evidence that is not entirely independent of the disagreement itself, but the HOE account does not accept the Independence Principle.¹⁵ Again, to see the intuitive plausibility of this

¹⁴ When is a moral view bizarre? Often professional philosophers claim to find their colleagues' views bizarre, but professional philosophers' judgement that some otherwise widely shared view is hard to make sense of is not quite what we are after. One might also think that if a moral intuition is widely held among otherwise reasonable people this counts against it being bizarre. On the other hand, given that we generally adopt moral beliefs that are entrenched in our culture, we might expect that a moral belief may become widespread even if it is bizarre in some sense, including our sense. Thanks to a reviewer for bringing up this issue.

¹⁵ Note that it is not obvious that A's reasoning needs to cite A's specific intuition – it may be enough that A reasons from the apparent difficulty of making sense of B's bizarre intuition.

response to bizarre disagreement, compare to a bizarre perceptual disagreement. I see someone from about 50 meters away, and I am confident that it is my colleague John. You, who are otherwise quite accurate in these matters, to my astonishment, say that it is not John, but a huge elephant. In this case, your belief is so far off track that I should consider the higher-order evidence undercut – since my visual perception is generally accurate, there must be some other explanation of your belief (one that does not presuppose that the accuracy of this perceptual belief of yours is within the normal range). Again, even if I suspend my belief that it is John rather than someone else, say Peter, that I see, I can set aside the possibility that it is a huge elephant. In these cases, it is difficult to see why I could not argue as follows: I saw John walking at a distance, but my friend thought he saw an elephant, hence I immediately suspected that something was wrong with him. So, my belief that I see John at a distance may be part of my reason to believe that you are off track. This, we suggest, is what the HOE account supports.¹⁶

¹⁶ As the cases are described, one may wonder, as a reviewer did, whether local evil and bizarre disagreements provide even *prima facie* higher-order evidence in the first place. In response, one can modify the cases and add that the individuals in question are generally intelligent and thoughtful, or that we know that they form competent moral beliefs in other cases. It would then be natural to think that the disagreement yields some higher-order evidence, which is then in turn undercut. But of course, there is a general issue about distinguishing between cases in which higher-order evidence is undercut, and cases in which we might prefer to say that there is no such evidence in the first place. Note that the implications for the HOE account would be the same for local evil and local

3.3 Comprehensive disagreements

Consider now cases in which the disagreement at hand is not about a singular moral intuition, but a wide range of moral beliefs:

Comprehensive disagreement. A believes that p , where this belief is part of A's much wider web of moral beliefs. B disagrees about p , but also disagrees with many parts of A's web of moral beliefs.

The kind of disagreement one finds between a prototypical conservative and liberal may fit this pattern. They disagree not only about some insulated issue, but about a whole cluster of views such as abortion, gay rights, gender equality, environmental protection, nuclear energy, equality, desert, individual freedom, the role of the welfare state, the use of military power.

Suppose that the agents' moral systems of beliefs differ in comprehensive ways. There are many instances where the verdicts of A and B differ: they have different moral intuitions at important junctures, as well as different general moral beliefs. Yet, assume that their systems are roughly equally coherent and free of internal conflict. Suppose further that both systems are reasonable: it is not that

bizarre disagreements even after the suggested modifications. According to the HOE account, if there is no *prima facie* higher-order evidence to begin with, we should not conciliate.

from A's perspective, B is a depraved person, who would be grossly insensitive to moral truth, or that B's system involves bizarre moral beliefs or intuitions that cannot be made sense of. Of course, A does believe that B is wrong, but not in the way of someone who is morally blind, perverted or evil. A has no particular explanation of why B is wrong – it is just that B seems to see a number of things differently. In this case, it seems that A should reduce her confidence (and likewise for B). After all, at each point of disagreement, at least one agent is mistaken. From A's point of view, B's perspective seems to be one that morally and epistemically well-functioning people could have, and still, A may believe that she got it right and B got it wrong, while having no particular reason to rule out that she might be wrong herself. So, the disagreement provides higher-order evidence that A might have made mistakes in her first-order moral reasoning and there is nothing to rebut or undercut that evidence. This is what the HOE account implies about comprehensive moral disagreements, when they are not evil or bizarre (for more caveats, see below).

Conciliationism's response to comprehensive disagreements has been discussed assuming the Independence Principle. Adam Elga argues that comprehensive disagreements are common, but that the Independence Principle helps showing why conciliationism does not lead to wholesale skepticism in such cases (Elga 2007, pp. 492). His argument is that we have no reason (based on independent evidence) to assume that our interlocutors are epistemic peers in comprehensive disagreement, and hence we are not obliged to change our view. As

has been pointed out by others, this inference turns on which of the following we accept:¹⁷

The strong evidence requirement. When A and B disagree about p , then A should conciliate (be agnostic), if and only if A has independent reason to believe that B is an epistemic peer.

The weak evidence requirement. When A and B disagree about p , then A should conciliate (be agnostic), unless A has independent reason to believe that B is A's epistemic inferior.¹⁸

When combined with the weak evidence requirement, conciliationism is susceptible to skepticism as we rarely have independent reasons to believe that others are epistemic inferiors, but when combined with the strong evidence requirement conciliationism is less prone to skepticism, as we should only conciliate when we can positively affirm that others are epistemic peers.

We suggest that the HOE account offers a more plausible response to comprehensive disagreement, and to the problem of skepticism. The HOE account cannot accept anything like the Strong evidence requirement, so a non-skeptical

¹⁷ See e.g. (Christensen 2011; Vavova 2014; Rowland 2017a).

¹⁸ Here we arbitrarily state the evidence requirements from A's point of view.

response cannot rest on this. Yet, we contend that the HOE account does not lead to wholesale skepticism, or at least not a variety of skepticism that seems intuitively implausible. This is, in part, because the conciliatory response advocated by the HOE account is limited to comprehensive disagreements that do not involve bizarre or evil views. It seems likely that A and B cannot have a comprehensive and non-bizarre (and/or non-evil) disagreement, unless they also share at least some substantive moral views. Assume that A and B share general principles such as: harming innocent individuals is *pro tanto* wrong, promoting well-being is *pro tanto* good etc. This overlapping consensus might be in terms of general foundational principles (principles that pick out moral factors), or mid-level principles (principles that specifies practical guidelines for decision-making), or perhaps the consensus encapsulates classes of moral judgements about particular situations. A and B may justify mid-level principles, or classes of particular judgements in different ways (the justification may invoke different foundational principles), or if they agree on a general foundational principle, they may prefer different interpretations of it. So, there is a variety of different ways in which A and B's moral systems could overlap, it would seem. If there are not one or more of these kinds of overlap between A and B it becomes difficult to believe that their disagreement is still non-bizarre. But given this, though A and B should reduce their confidence in the non-overlapping parts of their views, they could remain confident where their views converge. Indeed, they might have a reason, also supported by

the HOE account, to become more confident where they converge.¹⁹ So, if this is right, comprehensive disagreements as defined should not lead to wholesale moral skepticism. While the conservative and the liberal disagree about many things, they also have many points of convergence.

3.4 Deep disagreement

Consider next a more complicated form of disagreement which might also (but need not) be comprehensive:

Deep disagreement. A and B disagree about p . However, A's belief that p is grounded in a perspective, which includes basic methodological assumptions about which traditions, institutions, sources of evidence, patterns of reasoning to rely upon when forming moral beliefs. B's perspective rejects the methodological assumptions that A asserts, and there is no common ground which A and B can resort to in order to resolve the dispute.

So, A and B disagree about a range of issues, but these disagreements largely turn on a more fundamental disagreement about how to gather evidence or form beliefs about moral matters. Actual examples of this may be moral disagreements about abortion, homosexuality, contraception etc. based on different views about the role

¹⁹ We elaborate the epistemic significance of convergence in ethics in (Kappel, Andersen et al. unpublished manuscript).

of religion in morality, or in particular, the status of sacred texts, or religious institutions' authoritative interpretations of the moral content of such texts.

To make matters more concrete, suppose that A relies on his considered moral judgements about abortion, contraception and homosexuality when forming his moral beliefs, whereas B relies on an interpretation of a sacred book provided by a certain clerical authority. Of course, B also has moral intuitions, but where they conflict with the authorized interpretation of the scripture, B relies on the latter. We can also imagine that A and B both accept some sort of broad coherency requirement in their thinking, and that they roughly agree about the nature of these (for more on reflective equilibrium, see below). So, what sets A and B apart is that they radically disagree about which non-derivative moral beliefs should enter moral reflection about abortion, contraception and homosexuality, and this is because they radically disagree about how to gather and assess evidence for moral views. So, A and B in part disagree about which epistemic principles should govern our moral thinking (here we take an epistemic principle to be a norm for how to gather evidence, choose methods, or doxastic practices). A relies on one epistemic principle, whereas B relies on an incompatible one. Imagine that A and B's systems of moral beliefs are on a par with respect to formal features such as coherency and consistency. So, from A's perspective, the epistemic principle that B relies on is likely to lead B to false moral views, whereas B sees A in the same way. Imagine

also that there is no more fundamental principle or procedure that A and B can use to settle this disagreement about epistemic principles.²⁰

Notice that deep disagreements are clearly different from comprehensive disagreements (discussed above) in that A and B, as a part of their perspective, have a theory about why their opponent tend to get the evidence wrong. As we defined comprehensive disagreements they do not involve assumptions about epistemic principles governing moral belief.

How should A and B respond to this kind of disagreement? According to the HOE account, the decisive question will be whether the disagreement constitutes undefeated higher-order evidence for the occurrence of a mistake on the first-order level (a principle error, a performance error, or an evidence error).²¹

Consider first A's perspective. Should A think that her disagreement with B is evidence that she has made a performance error? Well, it is hard to see. For A, B's dissenting view is entirely due to B's adherence to belief forming practices that A considers wildly inappropriate and highly likely to generate false beliefs about the moral issue in question. Why would A think that this is any indication that she, A, might have made an error when thinking about the issue? For A, there is a perfectly good explanation of why B disagrees with her, and it has to do with B's adherence to epistemic practices that A considers wrongheaded. Similarly, why should the disagreement with B make A doubt the propriety of her own practices?

²⁰ For more on this conception of deep disagreement see (Lynch 2010, 2016; Kappel 2012, 2018).

²¹ This and the following paragraphs are revised material from (Kappel 2018).

Why should the fact that someone else uses a wrongheaded practice and reaches a false belief be evidence that one's own epistemic principle is not truth-conducive? From B's point of view, the story is similar, of course. B has no particular reason to take her disagreement with A to indicate that her own epistemic principles might be at fault, or that she might have made a mistake when applying them. Thus, in general terms the HOE account suggests that disagreement is *prima facie* higher-order evidence of performance errors, principle errors, or evidence errors, and in response to this one should reduce one's credence. However, in cases of *deep disagreement*, this higher-order evidence is undercut by evidence stemming from knowledge about the nature of the disagreement. In so far as deep disagreements are symmetrical, neither party should conciliate.

3.5 Fundamental disagreements

Richard Rowland briefly discusses a type of disagreement which is different from both comprehensive disagreements and deep disagreements as defined here:

We frequently find ourselves in deep seemingly irresolvable disagreements about moral issues. [...] Many of these disagreements seem to be deep because they seem to be disagreements that are due to disagreements in fundamental non-derivative moral judgments. For instance, many pro-lifers think that foetuses have a moral status, many pro-choicers disagree and there is no more fundamental moral judgment that these conflicting moral judgments derive from. And many of these deep moral disagreements can also seem to be

irresolvable. For it is hard to imagine that (all or most) pro-choices and pro-lifers, libertarians and socialists, and act-consequentialists and Kantians could come to agree about the moral matters that they disagree about by changing their fundamental moral judgments.²²

We take fundamental disagreements to concern conflicts in fundamental moral principles, say as in Rowland's example, the principle that a human fetus has moral status versus a principle denying such a status. A disagreement between a Kantian and a Utilitarian would also qualify as fundamental in this sense, in that the Kantian accepts (say) some version of the categorical imperative, and thinks that other less fundamental principles or specific judgments are supported by the categorical imperative, whereas the Utilitarian accords some version of the principle of utility this foundational role. A moral principle is fundamental when it, in some sense, does not derive from other principles, while it still provides support for less fundamental principles or judgements.

To make the notion of a fundamental disagreement a bit more precise, let us say that a moral principle T is *non-derivative* for subject S when there is no other moral principle T' that S could be epistemically justified in accepting and which

²² (Rowland 2017a, p. 1).

would support S in believing T. Given this, it is natural to define what we call fundamental moral disagreement as follows:

Fundamental disagreement. A believes a moral proposition p , and B believes not- p , and this is because (i) A accepts a moral principle T_A and B accepts an incompatible principle T_B , (ii) T_A epistemically supports p , while T_B epistemically supports not- p , (iii) T_A (T_B) is non-derivative for A (B): there are no other principles that A (B) could be epistemically justified in accepting and which would support A (B) in believing T_A (T_B).

So, fundamental disagreements are basically disagreements over moral principles that are non-derivative and occupy a certain foundational justificatory role in one's set of moral beliefs. Fundamental disagreements need not be comprehensive – they can involve only few, isolated moral beliefs. Further, fundamental disagreements need not be deep – A and B can have different moral outlooks that involve a non-overlapping set of non-derivative moral principles in foundational roles without fundamentally disagree about how one should reason about moral beliefs. Fundamental disagreement is not about which epistemic principles that govern moral evidence and beliefs. A and B may, for example, accept different fundamental moral principles, but both accept that our moral beliefs should answer to both considered moral judgments and to coherency requirements, as is usually assumed in the method of reflective equilibrium.

Fundamental disagreements are, Rowland says, “seemingly irresolvable”. Given our stipulation above, we can give a bit more content to this. Assume that a disagreement about p depends on principles T_A and T_B that are both fundamental for two subjects A and B, yet incompatible. It follows that there is no other set of principles that A and B could be epistemically justified in accepting, and which supports either T_A and T_B or their negations.²³ Clearly, this would mean that the disagreement is dialectically irresolvable in a certain sense: there is no more fundamental moral principle that A can appeal to in order to provide a rationally compelling argument to B showing the falsity of the fundamental principle that B accepts, and the same is true for B with respect to A.²⁴

Given the HOE account, how should one respond to fundamental disagreements? To make progress with this question, we need to make some assumptions about how we may become justified in accepting fundamental ethical principles. Assume some version of reflective equilibrium, according to which A

²³ Note that a principle T may be foundational in this sense for A, while the negation of T is not foundational for another subject B. We will ignore this complication here.

²⁴ Note that this dialectic impediment is compatible with the disagreement being resolvable in a different sense. Suppose that both A and B accept some version of reflective equilibrium (RE) and suppose A can make clear to B that A’s view is better in terms of RE than B’s, i.e. A’s view instantiates the good-making epistemic features of RE to a higher total degree than B’s view. If we assume that RE is (also) a theory about what theories to choose in ethical theory, then this could be a rationally convincing reason for B to change his view.

and B may be epistemically justified in accepting some non-derivative ethical principles in so far as these principles are supported by considered moral intuitions and are part of a system of beliefs that answers to the theoretical desiderata such as simplicity, generality and explanatory power, as well as the consistency of the belief system as a whole. Assume that A and B both reasonably insist that their view is not inferior to the other in terms of the theoretical desiderata involved in reflective equilibrium.

How, then, should A and B respond, according to the HOE account? Clearly, both A and B should take the disagreement to provide higher-order evidence that they have made mistakes in the process of reaching a view in reflective equilibrium. The process of adjusting a view to reach a reflective equilibrium endpoint involves making trade-offs between desiderata such as simplicity, generality and explanatory power against the intuitive acceptability of various implied judgements. It is not at all clear, in general, how this is best done. Thus, A and B cannot be very certain that they, rather than their opponent, have stricken the right balance between the various theoretical virtues. And if the fundamental disagreement is neither evil nor bizarre, it is hard to see how A or B could have rebutting or undercutting evidence defeating the higher-order evidence they get from the disagreement.

Note again the difference to deep disagreement. In deep disagreement, A has a defeater of the higher-order evidence because A believes that B is forming beliefs in what A takes to be an unreliable way. In fundamental disagreement, by contrast, A does not have a defeater of the higher-order evidence. A and B assume that they are both using the same general methods for generating moral beliefs, including

beliefs in fundamental principles, and yet they arrive at different conclusions. This makes fundamental disagreement similar to local benign disagreement, except that the disagreement concerns moral principles that play a certain kind of foundational justificatory role in one's system of moral beliefs. So, we contend that the HOE account recommends conciliation in fundamental disagreements.

Note again that this is not a wholesale skeptical conclusion, for the reasons discussed earlier. Even if A and B have fundamental disagreements, they are likely to have many areas of overlap as well, or likely to converge in mid-level principles, even where they accept incompatible fundamental principles. Thus, even when a fundamental disagreement forces A and B to reduce confidence in their incompatible non-derivative principles, they could still have high confidence in their overlapping consensus, including the mid-level principles or particular judgments supported by both T_A and T_B if such judgments exist.

3.6 Fully theorized philosophical disagreement

What we have called fundamental disagreements may just be a special case of a what we will refer to as:

Fully theorized philosophical disagreement. A has the moral belief that p and has a clear philosophical grasp of which basic moral intuitions, normative assumptions and arguments her position depends upon. Further, A is fully aware of the total dialectic of possible counter-examples or counter-arguments to these normative assumptions, and of the counter-arguments to

these counter-arguments etc. So, for any putative first-order moral argument against her view that A is aware of, she also knows how to deflect or rebut that argument. B disagrees with A regarding p but is otherwise in the same position as A, i.e. B is fully able to defend her view against every move that A might make.

As we have defined it, fully theorized philosophical disagreements are an idealization, of course, but we contend that mature philosophical disagreements tend to approach it (e.g. philosophical debates over equality, freedom, abortion, animal rights, criminal justice etc.). As competing philosophical positions get elaborated, defended and better understood they tend to become closer to being fully theorized philosophical disagreements. This is the ideal, of course, but not always what happens – sometimes positions are abandoned as a result of careful scrutiny, and sometimes ill-defended theories persist.

Fully theorized philosophical disagreements are different from mere comprehensive disagreements in that A and B have counterarguments against one another's position, and not just different views. Fully theorized philosophical disagreements need not be deep, i.e. they need not involve any disagreement about epistemic principles governing moral belief. Although the interesting fully theorized philosophical disagreements may often involve fundamental disagreements, they need not. And conversely, fundamental disagreements need not be fully theorized.

According to the HOE account, how should one respond to fully theorized philosophical disagreements? Assume that the positions held by A and B are in part defended using some version of reflective equilibrium, i.e. by arguing that a particular view is better than others in terms of desiderata such as simplicity, generality, and explanatory power, and by being supported by our considered moral intuitions. So, the relevant views are upheld not only because they are defensible in the sense defining fully theorized disagreements, but also because their proponents think that their view results from a proper trade-off between desiderata in reflective equilibrium such as simplicity, generality, explanatory power and intuitive acceptability. Clearly, our disagreements with professional peers can be considered higher-order evidence that we might have opted for the wrong trade-offs in our efforts to identify a fully theorized moral view. Clearly, none of us understands the method of reflective equilibrium well enough to rule this possibility out, thus we do not generally have evidence that defeats this higher-order evidence. For similar views, see e.g. (Frances 2013; Goldberg 2013).²⁵

²⁵ In our discussion of fundamental disagreement and fully theorized philosophical disagreements we assumed that agents are relying on reflective equilibrium as their method of reasoning. A reviewer questioned this assumption. We make this assumption to make the discussion manageable and more concrete, and it seems a far from problematic assumption as reflective equilibrium is probably the most widely used method of moral reasoning. Moreover, it seems to us that as long as two agents are using (are known to be using) the same methods, then the HOE account will recommend conciliation in both fundamental disagreements and in fully theorized disagreements.

4. Summary and concluding remarks

We have outlined a general account of the epistemology of disagreement, the HOE account. The HOE account is a generic account the epistemology of disagreement, which applies to moral disagreement as well as to disagreements in other domains. The core idea of the HOE account is that disagreement provides higher-order evidence that one might have made a mistake in one's first-order reasoning, and when this evidence is strong enough, and undefeated, one should respond by reducing one's confidence in one's affected first-order belief. We have argued that various types of moral disagreement of interest in recent debates differ with respect to whether they do generate undefeated higher-order evidence that one has made a mistake. As a consequence, the HOE account recommends conciliation in local benign disagreements, comprehensive disagreement, and fundamental disagreements, but not deep disagreements, bizarre disagreements and evil disagreement. The HOE account is generally conciliationist in spirit, but it doesn't recommend conciliation in all cases, and, importantly, it avoids the skeptical implications that standard conciliationism is often thought to give rise to. Unlike standard conciliationism, the HOE account does not hinge on the Independence Principle, on the notion of an epistemic peer, or on the thorny issue of permissivism versus uniqueness. Further, the HOE account is motivated by strong general intuitions about the significance of higher-order evidence, and this sets the HOE account apart from other non-conciliationist views.

While we have not stated or defended the HOE account in all details, we have suggested that the HOE account is well-motivated and intuitively plausible. When

applying the account to cases of moral disagreement, we find that the HOE account is consistent with relevant and intuitively plausible distinctions between a variety of different forms of moral disagreement, and that it supports intuitively plausible verdicts about these cases.

Obviously, we do not want to suggest that the types of cases we have discussed constitute an exhaustive typology, and neither have we considered cases that combine the features that we have identified. Yet, it is too much to ask for any account of moral disagreement that it issues a clear verdict on any case of disagreement, no matter how complex it is. What we might reasonably hope for is an account that identifies and motivates some general principles for rational responses to disagreement, not an algorithm that determines a precise answer to any disagreement, however complicated.

References

- Bergmann M (2009) Rational disagreement after full disclosure. *Episteme* 6 (3): 336-353.
- Besong B (2014) Moral intuitionism and disagreement. *Synthese* 191: 2767-2789.
- Christensen D (2007) Epistemology of Disagreement: The Good News. *Philosophical Review* 116: 187-217.
- Christensen D (2009) Disagreement as Evidence: The Epistemology of Controversy. *Philosophy Compass* 4/5: 756-767.
- Christensen D (2011) Disagreement, Question-Begging and Epistemic Self-Criticism. *Philosophers' Imprint* 11.
- Christensen D (2019) Formulating Independence. In Rasmussen M S and Steglich-Petersen A (eds.) *Higher-Order evidence: New Essays*. Oxford University Press.
- Elga A (2007) Reflection and Disagreement. *Noûs* 41 (3): 478-502
- Enoch D (2010) Not Just a Truthometer: Taking Oneself Seriously (but not Too Seriously) in Cases of Peer Disagreement. *Mind* 119 (476): 953-997.
- Frances B (2013) Philosophical Renegades. In Christensen D and Lackey J (eds.) *The epistemology of disagreement: New essays*. Oxford University Press, pp 121-166.
- Goldberg S (2013) Disagreement, Defeat, and Assertion. In: Christensen D and Lackey J (eds.) *The epistemology of disagreement: New essays*. Oxford University Press, pp 167-189.

- Goldman A (2001) Experts: Which Ones Should You Trust?. *Philosophy and Phenomenological Research* 63: 85-110.
- Horowitz S (2014) Epistemic Akrasia. *Noûs* 48 (4): 718-744.
- Kappel K (2012) The problem of deep disagreement. *Discipline Filosofiche* 22 (2): 7-25.
- Kappel K (2018) Higher Order Evidence and Deep Disagreement. *Topoi*: 1-12.
- Kappel K (2019) Escaping the Akratic Trilemma. In Rasmussen M S and Steglich-Petersen A (eds.) *Higher-Order evidence: New Essays*. Oxford University Press.
- Kelly T (2010) Peer Disagreement and Higher-Order Evidence. In Feldman R and Warfield T (eds.) *Disagreement*. Oxford University Press, 2010, pp 111-174.
- Kelly T (2013a) Disagreement and the Burdens of Judgment. In Christensen D and Lackey J (eds.) *The epistemology of disagreement: New essays*. Oxford University Press, pp 31-53.
- Kelly T (2013b) Evidence can be permissive. In Steup M, Turri J and Sosa E (eds.) *Contemporary Debates in Epistemology*. Wiley-Blackwell, pp 298-312.
- Lackey J (2010) A Justificationist View of Disagreement's Epistemic Significance. In Haddock A, Millar A and Pritchard D (eds.) *Social Epistemology*. Oxford University Press, pp 298-325.
- Lackey J (2013) Disagreement and Belief Dependence: Why Numbers Matter. In Christensen D and Lackey J (eds.) *The Epistemology of Disagreement*. Oxford University Press, pp 243-268.
- Lasonen-Aarnio M (2014) Higher-order evidence and the limits of defeat. *Philosophy and Phenomenological Research* 88(2): 314-345.

- List C (2001) Epistemic Democracy: Generalizing the Condorcet Jury Theorem. *Journal of Political Philosophy* 9 (3): 277–306.
- Lord E (2014) From Independence to Conciliationism: An Obituary. *Australasian Journal of Philosophy*, 92 (2): 365-377.
- Lynch M (2010) Epistemic circularity and epistemic incommensurability. In: Haddock A, Millar A and Pritchard D (eds.) *Social epistemology*. Oxford University Press, pp 262-277.
- Lynch M (2016) After the spade turns: Disagreement, first principles and epistemic contractarianism. *International Journal for the Study of Skepticism* 6: 248–259.
- Matheson J (2011) The case for rational uniqueness. *LOGOS & EPISTEME* II (3) (2011): 359-373.
- Matheson J (2016) Moral caution and the epistemology of disagreement. *Journal of Social Philosophy* 47 (2): 120-141.
- McGrath S (2008) Moral Disagreement and Moral Expertise. In: Shafer-Landau R (ed.) *Oxford Studies in Metaethics: Volume 4*. Oxford University Press, pp 87-108.
- Moon A (2018) Independence and New Ways to Remain Steadfast in the Face of Disagreement. *Episteme* 15 (1): 65-79.
- Pollock J (1970) The Structure of Epistemic Justification. *American Philosophical Quarterly*, monograph series 4: 62-78.
- Pollock J (1986) *Contemporary Theories of Knowledge*, Totowa, NJ: Rowman and Littlefield.
- Rosa L (2012) Justification and the uniqueness thesis. *LOGOS & EPISTEME* III (4): 571-577.

- Rowland R (2017a) The Epistemology of Moral Disagreement. *Philosophy Compass* 12 (2).
- Rowland R (2017b) The Significance of Significant Fundamental Moral Disagreement. *Noûs* 51(4): 802-831.
- Setiya K (2012) *Knowing Right from Wrong*. Oxford University Press.
- Sliwa P and Horowitz S (2015a) Respecting All the Evidence. *Philosophical Studies* 172 (11): 2835-2858.
- Sliwa P and Horowitz S (2015b) Respecting All the Evidence. *Philosophical Studies* 172 (11): 2835-2858.
- Sosa E (2010) The Epistemology of Disagreement. In: Haddock A, Millar A and Pritchard D (eds.) *Social Epistemology*. Oxford University Press, pp 278-297.
- Titelbaum M G (2015) Rationality's Fixed Point (or: In Defense of Right Reason). In: Gendler T S and Hawthorne J (eds.) *Oxford Studies in Epistemology* 5 (253), pp 253-294.
- Vavova K (2014) Moral Disagreement and Moral Skepticism. *Philosophical Perspectives* 28 (1): 302-333.
- Weatherson B (2008) Deontology and Descartes' Demon. *Journal of Philosophy* 105(9): 540–569.
- Weatherson B (2019) *Normative Externalism*. Oxford University Press.
- Wedgwood R (2010) The Moral Evil Demons. In Feldman R and Warfield T (eds.) *Disagreement*. Oxford University Press, pp 216-246.
- White R (2013) Evidence cannot be permissive. In Steup M, Turri J and Sosa E (eds.) *Contemporary Debates in Epistemology*. Wiley-Blackwell, pp 312-323.

Williamson T (2011) Improbable Knowing. In Doherty T (ed.) *Evidentialism and Its Discontents*. Oxford: Oxford University Press, pp. 147–164.

Worsnip A (2018) The Conflict of Evidence and Coherence. *Philosophy and Phenomenological Research* 96 (1): 3-44.