



The Surprising Truth About Disagreement

Neil Levy^{1,2} 

Received: 16 January 2020 / Accepted: 26 May 2020 / Published online: 8 June 2020
© The Author(s) 2020

Abstract

Conciliationism—the thesis that when epistemic peers discover that they disagree about a proposition, both should reduce their confidence—faces a major objection: it seems to require us to significantly reduce our confidence in our central moral and political commitments. In this paper, I develop a typology of disagreement cases and a diagnosis of the source and force of the pressure to conciliate. Building on Vavova’s work, I argue that ordinary and extreme disagreements are surprising, and for this reason, they carry information about the likelihood of error. But deep disagreement is not surprising at all, and token deep disagreements do not put pressure on us to conciliate. However, a pattern of deep disagreements points to a different concern: not the problem of disagreement but the problem of irrelevant influences. Deep disagreement constitutes some pressure to examine the foundations from which we reason, rather than to conciliate on our central moral and political claims.

Conciliationism is the view that discovering that an epistemic peer disagrees about a proposition should make agents less confident in their opinions. The conciliationist response to peer disagreement is highly intuitive with regard to a range of everyday cases. But it appears to have steep costs. Many of the issues we care about most are highly controversial: we cannot help but be aware that many others passionately disagree with us about (say) the morality of abortion, the need for gun control or the rights of asylum seekers (call these cases, on which left and right seem intractably opposed, cases of partisan dissent).¹ Unless we can find grounds for dismissing many intelligent and well-informed, disputants from peerhood, we seem to be required to substantially decrease our confidence, perhaps even becoming agnostic, on many questions central to our identities. Conciliationism thus gives rise to what is sometimes called the problem of ‘spinelessness’ (e.g. Elga 2007; Sosa 2010).

The standard conciliationist strategy for avoiding widespread scepticism has consisted in attempting to show that our partisan opponents are not our epistemic

✉ Neil Levy
neil.levy@philosophy.ox.ac.uk

¹ Department of Philosophy, Macquarie University, Sydney, NSW 2109, Australia

² Uehiro Centre for Practical Ethics, University of Oxford, Oxford OX1 1PT, UK

peers. They fail to count as peers not because they are indoctrinated, stupid or biased, but because *setting aside the dispute and the reasons it implicates*, we should not assess them as equally likely to be right about the target proposition. In this paper, I argue that this strategy fails. We cannot dismiss our partisan opponents from peerhood on the grounds its proponents adduce, at least in an uncomfortably large range of cases. Should we therefore conciliate?

I will argue that different kinds of disagreements call for different kinds of responses. I will distinguish three basic kinds of case. Disagreement provides information of a special sort to the extent that it is unexpected and the kinds can be distinguished by reference to just how unexpected the disagreement is. In *ordinary* disagreements, disagreement is unexpected but not shocking, and we should conciliate by significantly reducing our confidence in the proposition at issue (the ‘target proposition’). In *extreme* disagreements, disagreement is shocking and we should conciliate by significantly reducing our confidence that our partisan opponent is an epistemic peer. In *deep* disagreements, disagreement is not unexpected at all, and gives us no reason to conciliate with regard to the target proposition or with regard to the peerhood of our disputants. Deep disagreements constitute pressure if not to conciliate then to engage in examination of the principles and assumptions from which we reason. Any such pressure arises from a different source to the pressure to conciliate stemming from ordinary and extreme disagreements: from a suspicion of irrelevant influences.

I will begin with a brief presentation of the conciliationist response to epistemic disagreement. While non-conciliatory views remain influential, I ignore them, simply assuming the truth of conciliationism. Section 1 introduces the debate and describes its contours. Section 2 turns to the question of peerhood, on which the best-known strategy for holding fast to our moral and political beliefs turns. In that section, I will argue that our partisan opponents cannot be dismissed (sufficiently) from peerhood. Section 3 introduces and elaborates on what I will call the *surprise* account of peer disagreement, which builds on work by Katia Vavova. Section 4 applies the account to more difficult cases, in ways which conflict with Vavova’s own views. In the light of the taxonomy developed in this section, I will argue that different kinds of disagreement have different sorts of epistemic upshots.

1 Peer Disagreement and Conciliationism

Disagreement with apparent peers is a fact of epistemic life. Most philosophers agree that in some such cases, we ought to reduce our confidence in the belief we held prior to discovering the disagreement. Consider this familiar case (based on Christensen 2007):

Restaurant Check: Anika and Bindi are old friends who eat out together once a fortnight. They always split the bill. As they always do, each calculates her share on her own, dividing the check by 2 and adding 15% to the total for a gratuity. They are both pretty good at mental arithmetic, and they almost always agree on the total. When in the past they have disagreed, checking has shown that Anika is right about half the time. Tonight is one of those rare occasions when they disagree: Anika announces that each owes \$43, while Bindi comes up with the figure of \$45.

This is designed to be a case of peer disagreement. Anika and Bindi have track records that show that they are equally competent at mental arithmetic, and there is nothing about tonight that impugns either's competence. Neither (let's suppose) has drunk alcohol with their meal nor have they recently experienced any unusual or undue stress or sleep disruption. Neither has a reason to discount the judgement of the other. In cases like this, most philosophers think, each ought to reduce her confidence in her initial judgement. Due to the symmetry in their epistemic positions, neither has a good reason to prefer her own judgement to that of the other person, and they ought to conciliate. They should think that they should assign (more or less) identical probabilities to \$43 and \$45 being the right figure.

Examples like this one support the most widely discussed view on the epistemic significance of disagreement: the *equal weight view* (Christensen 2007; Elga 2007). The equal weight view can be presented as a combination of three claims (Frances and Matheson 2018):

- *Defeat*: Learning that a peer disagrees with you about p gives you a reason to believe you are mistaken about p .
- *Equal weight*: The reason to think you are mistaken about p stemming from your peer's opinion about p is just as strong as the reason to think you are correct about p stemming from your opinion about p .
- *Independence*: Reasons to discount your peer's opinion about p must be independent of the disagreement itself.

These three conditions are intimately related. Peer disagreement is a defeater because, prior to and independent of the dispute, the probability each assigns to the likelihood of the other person being right in the event of disagreement is (roughly) 50%. Anika and Bindi know that ordinary people, like themselves, sometimes make mistakes on moderately complex mental arithmetic, and each believes that she is no less (or more) likely to make such a mistake than her peer. Because the probability of being wrong in the event of dispute is roughly 50%, each should put equal weight on her own verdict and the dissenting verdict.

The justification for the Independence principle is easy to see. Flouting Independence seems to allow us egregiously to beg the question against those who disagree with us. It would clearly be improper for Anika to use her own belief as a reason to discount Bindi's: she cannot say 'you must have made a mistake, because the total is \$43'. That would be citing a claim that Bindi denies as a reason for Bindi to give up that very claim. It is important to recognize that Independence requires us to bracket not only the very proposition that is under dispute but also propositions that both sides agree entail that proposition. So Anika cannot cite the penultimate step of her arithmetic ('and $41 + 2 = 43$ ') as a reason to discount Bindi's conclusion, if Bindi agrees that $41 + 2 = 43$ but denies that Anika should have reached those figures at the penultimate stage of the argument.

Despite its attractiveness, Independence is controversial. There are cases in which it is apparently appropriate to appeal to the disputed proposition itself to downgrade the epistemic status of one of the disputants. Consider this pair of cases:

Elementary Math. Harry and Jennifer are trying to determine how many people from their department will be attending the APA. Jennifer notes that Mark and Mary

are going on Wednesday, and Sam and Stacey are going on Thursday. ‘Since $2 + 2 = 4$, there will be four members of our department at the conference’, she adds. In response, Harry asserts, ‘But $2 + 2$ does not equal 4’. (slightly modified from Lackey 2010a: 283.)

Holocaust. Deborah is very confident that the Holocaust occurred. Her evidence is extensive, ranging from the testimony of a consensus of historians through to direct testimony from survivors. David, like Deborah, has done extensive research on the Holocaust. He is equally confident that the Holocaust did not occur. Each is fully aware of the evidence the other person cites.²

Opponents of the equal weight view take these cases to challenge the view by undermining Independence. In *Elementary Math*, Jennifer’s antecedent confidence that $2 + 2 = 4$ is sufficiently strong that she can appeal to that very belief to reject Harry’s testimony. She should not conciliate at all. Similarly, Deborah can reject David’s testimony on the grounds that he has reached the wrong conclusion, without worrying that she is begging the question against him.

Since (as we shall see) Independence does a lot of work for conciliationists in staving off the threat of moral scepticism in response to partisan dispute, these apparent counterexamples seem very threatening. In *Elementary Math* and *Holocaust*, we can be confident that the other person is not an epistemic peer, because the answer they give to the question under dispute is so obviously crazy. It seems appropriate to downgrade the epistemic status of an agent who defends a beyond the pale proposition (at least their epistemic status with regard to that proposition), by appealing to our own response and the reasons implicated in it. It seems illegitimate to appeal to our own reasoning to dismiss others from peerhood in some cases, but not others; we need a principled account of when such appeal is in order.

2 Partisan Peerhood

Perhaps, however, we can resist the pressure to conciliate by adducing grounds to relegate dissenters from the status of peerhood. Some philosophers advance extremely demanding conditions for peerhood, making appeal to Independence unnecessary. Insofar as we are concerned with real-life disagreements and how to manage ourselves epistemically in their light, there is good reason to use a more relaxed account. Set the bar too high and we render the debate almost entirely irrelevant to our epistemic lives, and undermine the motivation for investigating peer disagreement (Lackey 2010a). We should not require that disputants are precisely cognitive equals. Were Bindi and Anika to discover that one of them had a slight edge over the other on mental arithmetic, being right in 51% of cases in which they disagree, both should continue to feel significant pressure to conciliate. Similarly, we do not need to establish exact evidential equality to feel pressure to conciliate. Consider another hackneyed case; *Horserace* (Elga 2007). When two friends, with (roughly!) equally good eyesight, find themselves disagreeing over which horse finished a very close race first, they should both feel under pressure to conciliate, even though it is unlikely that their viewing angles are genuinely equally good: so long as the differences are small enough, we do not regard such inequalities as sufficient to undermine epistemic peerhood.

Many readers might feel, nevertheless, that the differences between partisan disputants are often big enough for *us* to be able to relegate *them* from peerhood (see Sherman 2014 for discussion). A reader who shares the author's liberal sympathies might point out that Republican voters often get their news from sources, like Fox (and worse), which are significantly less reliable than other mainstream sources (Fox News is the only outlet which leaves viewers less well informed than those who consume no news at all, though MSNBC is not very much better (PublicMind 2012)). She might point to the fact that consumption of fake news and highly partisan information is not symmetrical. While both Democrats and Republicans accessed sites with highly slanted coverage during the 2016 election, those with pro-Trump leanings were significantly less likely to have consumed mainstream media, inhabiting an alternative media ecosystem, in which information was highly unreliable (Benkler et al. 2017). These seem to be good reasons for downgrading them from peerhood.

I do not think these kinds of facts do much to reduce the pressure to conciliate. Though it is probably true that (right now) on average those on the right tend to be somewhat less well informed than those on the left, there are plenty of well-informed people on both sides of debates. In fact, polarization is *greater* among the well-educated and well-informed than the less knowledgeable (Drummond and Fischhoff 2017). Each of us on the political left has many opinions that are rejected by large numbers of conservatives, including many who are at least as well-informed and as intellectually able as we are.³ We can dismiss beyond the pale partisan dissent, but we cannot reasonably dismiss partisan dissent across the board in this way. Perhaps dissenting partisan peers are less numerous than those who agree with us; nevertheless, on many topics, they constitute a significant pressure to conciliate, if perhaps not to split the difference between our view and theirs.⁴

On traditional accounts of what peerhood consists in, citing intellectual virtues, access to evidence, attentiveness, and so on, it is hard to dismiss partisan dissenters from being (near enough) peers. It is at this point that Independence may be invoked to allow us to resist the pressure to conciliate. Elga (2007) proposes the following test for peerhood: a dissenter is your peer if, prior to discovering the disagreement and setting aside any of the reasons implicated in the dispute, you would have regarded them as about as *likely* to reach the right answer on the issue as yourself. This test clearly yields the right verdict in many cases of peer disagreement. It rules, for instance, that Anika and Bindi rightly regard one another as peers, because prior to discovering their disagreement, they would think one another equally likely to get the right answer on the restaurant check. It may also provide us with grounds for dismissing partisan dissent.

The likelihood account of peer agreement allows us to dismiss partisan dissent because when we set aside the disagreement and the propositions that it directly implicates, we often do *not* regard our partisan opponents as equally likely to be right about the issue. There is a kind of holism of moral and political beliefs: knowing what you believe about climate change, I can usually predict your beliefs about gun control. Due to this correlation of moral and political beliefs, we do not see our partisan opponents as equally likely to be right: rather, we see one another as going wrong across a range of issues, some indirectly related to the current question, some apparently unrelated. On the likelihood account, partisan disagreements are not disagreements between epistemic peers, for reasons that have nothing to do with how well-

informed each side is. Because peer disagreements are disagreements between people who see one another as roughly equally likely to be right, conditional on disagreement, our peers are those with whom we agree on a great deal.

There is something very attractive about the idea that two people are epistemic peers only if they are equally likely to reach the right answer to a question they have not yet considered. But the proposal faces an objection: there appears to be no principled reason to individuate the dispute narrowly, in the way that Elga wants, as a dispute about, say, abortion, rather than more broadly, as a dispute about a whole set of more or less closely related claims (the existence of God, the moment at which an individual begins to exist, the role of religious belief in public life....). The fact that *I* do not regard *you* as equally likely to reach the right answer with regard to the question 'is abortion permissible' does not entail that *I* should not regard *you* as equally likely to be correct about that whole set of background issues.

Elga has a response to this objection. He agrees that you do not have a good basis to think that those with whom you are in broadly individuated dispute are less likely to get the right result about that whole set of questions. But that, he claims, is because you have no basis for *any* judgement about their likelihood of getting the right result, and therefore no basis for seeing them as a peer. Once we set aside all the issues in dispute (not to mention all the further propositions each issue directly implicates), there is no determinate answer at all to the question how likely my opponent is to reach the right conclusion on *any* of these questions. Whether we take our dispute to be narrow or broad, we ought not to see one another as peers.

Several critics have pointed out in response to Elga that the disputants he has in mind do not live in different moral worlds (Rowland 2017). Even those on opposing sides of bitter disputes about abortion, say, agree with one another in a great deal of their moral judgements, some shared more or less universally (lying is wrong), some accepted only at some places and at some times (slavery is wrong; animal suffering counts morally). This shared background might be sufficient to generate a determinate answer to questions about how likely each should think the other to reach the right answers (McGrath 2008; Fritz 2018). I am more sympathetic to Elga's claim than these critics.⁵ But the facts they point to illustrate what is fishy about this test for peerhood.

A very plausible explanation for *why* we agree on so much cites cognitive capacities we share, combined with our shared evidential state, and these facts strongly suggest that in an important sense, we are peers. These kinds of considerations might reasonably motivate us to reject likelihood accounts of peerhood, in favour of accounts which understand peers as those with access to the relevant evidence and who are equally intelligent, conscientious, attentive and so on.

But if we cannot downgrade partisan disputants from peerhood, the problem of spinelessness looms large once more. In what follows, I will argue that we can see off the problem of spinelessness by recognizing that deep disagreements, of the kind that give rise to the apparent problem, are qualitatively different from the kinds of disagreement that place rational pressure on us to conciliate. I will argue that each kind of disagreement carries information of a distinctive sort, and homing in on that information provides a guide to how we ought to respond: not (directly) by updating our beliefs or lowering our confidence, but by engaging in an appraisal of the influences to which they point.

3 What We Learn From Disagreement

As Vavova (2014b) forcefully argues, the pressure to conciliate is strongest when disagreement is most *surprising*. In fact, this is a specific case of a more general truth: we should become more or less confident in our beliefs in proportion to how surprising new evidence is. Vavova illustrates the point with an example involving a person drawing a marble from an urn, knowing that the urn contains both black and white marbles. Should she change her opinion about the composition of the urn on drawing (say) a black marble? That depends on what her (appropriate) expectations were. If she believed that the urn contained mostly black marbles, or a roughly equal mix of black and white marbles, she should be unsurprised to draw a black marble and she has no reason to change her opinion. If she believed that the urn contained predominantly white marbles, then drawing a black marble is surprising, and constitutes pressure to decrease her confidence in the belief. The higher the proportion of white marbles she believed the urn to contain, the more unexpected it is to draw a black marble and the stronger her reason to rethink. Similarly, Vavova claims, when we expect someone to agree with us on a proposition, their disagreement is surprising, and puts pressure on us to revise our belief.

Though Vavova does not put it like this, one way to understand the claim that the more unexpected the disagreement is the more it constitutes pressure to conciliate is to think of such disagreement as carrying information. Unexpected disagreement is just a special case of unexpected input, and unexpected input carries information for the agent. Unexpected input is evidence that the source of the input differs from our prior model of it, and therefore provides us a reason for updating our model.⁶ If I believe that my dog is lying on the couch in front of the TV, I do not expect to see him chasing a squirrel in the yard, and if I do see him in the yard, that unexpected input is information in the light of which I must update my beliefs. Similarly, if Anika thinks that the total she and Bindi owe is \$43, and that Bindi is a reliable mathematician, then she expects Bindi to agree with her. If Bindi does not agree with her, she is under pressure to update her belief in the light of this unexpected input.

Thinking of disagreement in this way provides a satisfying explanation of the pressure to conciliate and gets the cases right (as Vavova shows). Consider, first, cases like *Restaurant Check*. Anika expects Bindi to agree with her, and vice versa. The disagreement is unexpected. But it is not all *that* unexpected. When we do arithmetic in our heads, we sometimes make mistakes. When someone we take to be a peer disagrees with us, that is evidence that at least one of us has made a mistake. Since we have no (or at any rate very little) reason to think that it is more likely that the other person has made the mistake than ourselves, the disagreement carries information that we may have erred, and we should lower our confidence accordingly.

At first glance, it might seem that understanding the pressure to conciliate as arising from surprise gets cases like *Elementary Math* wrong: Jenny is *extremely* surprised by Harry's assertion that 2 plus 2 does not equal 4, but here the unexpected disagreement does not provide her with a reason to conciliate. As Lackey herself indicates, however, cases like this *do* provide us with a strong reason to update our beliefs. The best explanation for the disagreement is not that one or other has made a mistake in arithmetic (the sum is too easy for that to be plausible) but rather that Harry is either insincere or delusional. In the face of unexpected input, we must update our model of

the world, but there are many ways to do so that accommodate the input. The model we should adopt is the one that best preserves our prior web of beliefs, with some beliefs weighted more strongly, so that they serve as (more or less) fixed points which are abandoned only in extremis. Since Jennifer is so strongly committed to the belief that $2 + 2 = 4$, she preserves her web of beliefs in the face of this highly unexpected input by changing her beliefs about Harry.⁷

This response (with its appeal to the best explanation of the disagreement, given our priors) echoes Christensen's (2007), and therefore may seem vulnerable to the objection Lackey gives to that account. Lackey argues that if such an account is not to collapse back into her justificationist view (according to which we should hold fast when we are highly justified in believing the proposition), it 'opens the door to rampant bias and dogmatism' (Lackey 2010b: 323). For agents may have priors that are highly irrational: an opponent of abortion might regard his opponents as insane, for example. I am unmoved by this objection. It is true that if we begin with bad priors, we may not succeed in improving our beliefs by way of reasoning procedures, but that is simply the epistemic condition. We cannot do any better in our attempts to track truths than use the tools we have to hand.

Before concluding this section, it is worth pausing to see how the surprise account (as I now dub it) handles cases of *agreement*. Several philosophers (e.g. Matheson 2015) have pointed out that agreement can provide a reason for belief update too. How can this be, on the account offered here? After all, if I believe p , then a peer's agreeing with me is exactly what I expect, and therefore does not offer me a reason to update. Again, though, the account gets the right result: agreement provides me with a reason to update just to the extent to which it is unexpected. Consider *Restaurant Check*. Anika believes that the right result is \$43, and because she believes that Bindi is reliable, she believes that Bindi will agree. But she is not *very* confident that Bindi will agree (mental arithmetic is not a highly reliable procedure after all). Because she is only somewhat confident that Bindi will agree, her actual agreement provides Anika with information, in the light of which she should increase her confidence that the answer is \$43. The more confident she is both of the answer and of the peerhood of the other agent, the less unexpected is agreement and, correlatively, the less effect it should have on her confidence. Harry's (counterfactual) agreement with Jennifer that $2 + 2 = 4$ should not increase Jennifer's confidence in the proposition to any discernible extent. Agreement provides evidence for the target belief in precisely those situations in which disagreement would provide evidence against the target proposition (and not the competence or sincerity of the other person): when we are not so confident in the belief that disagreement would not shift us. When we are so confident, agreement is entirely expected. In those cases, *disagreement* is shocking and provides us with a reason for doubting the sincerity or competence of the dissenter, rather than conciliating on the target proposition.

4 Disagreement: Ordinary, Extreme and Deep

In the last section, we began to sketch how the surprise account handles cases of disagreement and agreement. The account itself fleshes out Vavova (2014b); though I hope that I have succeeded in demonstrating its attractions in a somewhat different way

to her, the credit is largely hers and not mine. In this section, I put the account to work in a way that goes beyond (and directly conflicts with some of) her claims. I will argue that the surprise account can provide a satisfying solution to the problem of spinelessness in cases of partisan dissent. Disagreement typically provides us with a reason to reduce our confidence in a proposition: either in the disputed belief or in our belief that our disputant is an epistemic peer (whichever of the two we had less confidence in prior to the disagreement). In these cases, disagreement introduces incoherence into our belief set, and we restore it by conciliating. But partisan disagreement does not introduce incoherence and (therefore) does not place pressure on any particular belief(s). Rather, it is at most a spur to further investigation.

In this section, I will distinguish three kinds of disagreement: Ordinary, Extreme and Deep.⁸ I will argue that the surprise account gives a plausible and satisfactory account of all three. It also provides a satisfying account of the epistemic significance of agreement. Since deep disagreements are those that give rise to the problem of spinelessness, they will receive by far the most attention. These kinds of disagreement are ideal types, of course: in the real world, matters can be messy. Far from casting doubt on the usefulness of the typology, however, it allows us to predict and make sense of the ways in which it can be expected to be messy.

(1) Ordinary disagreement.

In *ordinary* disagreements (like *Restaurant Check*), A and B disagree with one another whether p and the disagreement is unexpected but not shocking. When A and B disagree with one another in this kind of case, the disagreement is evidence of a *reasoning error*: that is, the most likely explanation for the occurrence of the disagreement is that one or both agents has made a mistake.

Unexpected but not shocking disagreement arises when generating the right answer requires reasoning that is easy enough for disagreement to be unexpected (if the reasoning is very difficult, generating discordant answers is expectable and should not lower our confidence very much; note that in such cases, our confidence should be quite low to begin with) but not so easy that disagreement is shocking. Since the best explanation for the disagreement is that one or both agents have made a mistake, each should conciliate, lowering their confidence that p (in such cases, *agreement* is expected, but since error is not unlikely, the agents should not be very confident in this expectation and its occurrence should cause them to increase their confidence that p).

Note that ‘ordinary disagreement’ is a term of art. While many of the disagreements that routinely occur in the course of everyday life are ordinary disagreements, not all of them are (as we shall soon see).

(2) Extreme disagreements

In *extreme* disagreements (like *Elementary Math*), A and B disagree with one another over p and the disagreement is very unexpected. When A and B disagree whether p and the disagreement is very unexpected, the best explanation of the disagreement is that one or other is not sincere or is not competent (in this domain). Extreme disagreement is also evidence of error: it is evidence that the apparent peer is deeply in error in a way that impugns their competence, or that it is an error to take them seriously.

Extreme disagreement comes in two forms: expert and non-expert. The standard examples of extreme disagreements are non-expert: they involve one party asserting a belief that is obviously unacceptable to anyone who considers the issue: asserting that $2 + 2$ does not equal 4, or cases like Christensen's (2007) *Extreme Restaurant Check* (imagine that Anika asserts that the share each owes is \$43, while Bindi asserts it is actually \$476, a figure that is higher than the bill total). Expert cases arise in specialized domains, when one party asserts a claim that is unacceptable to experts, though it may sound plausible or at least arguable to outsiders. Real-life examples include pseudo-scientific claims, like those made by proponents of 'intelligent design', vaccine scepticism and climate change denial. Expert cases themselves divide in two, the first closely analogous to non-expert. These extreme disagreements arise when the relevant problem is easy enough (for experts in the domain) for them to be highly confident of agreement. Disagreement is highly unexpected, and the best explanation cannot cite a simple reasoning error (because the problem is too easy for that). Instead, the disagreement forces the experts to change their credences with regard to the dissenter: the symmetry between them is broken and the dissenter is relegated from the status of peer (in such cases, agreement is very confidently expected; so confidently that its actual occurrence does not carry significant information and should not lead us to adjust our credences).

The second type of expert extreme disagreement is probably more common than the first. Interestingly, this kind of expert disagreement builds on some of the features of ordinary disagreement: it arises in cases in which the relevant experts expect others *already* to have conciliated in the light of something very like ordinary disagreement. In these cases, the relevant problem is *not* easy. Think of climate change: the processes involved are highly complex and multifactorial, and the evidence comes from multiple sources. Detecting signal in all that noise takes hard work and a great deal of analysis. The problem is difficult enough that any researcher or research group tackling it on their own might easily make a mistake. So far, so parallel to ordinary disagreements. These disagreements are extreme, however, because there is a consensus on the issue, and researchers are highly confident of the claim as a consequence. Because the problem is difficult, errors in generating the solution are highly expectable. But because errors are expectable, experts have another expectation as well: they expect other researchers who reach solutions contrary to the consensus to lower their confidence in the solution—that is, to conciliate—and take the disagreement as a strong reason to check for errors (understood widely, to include checking that data is representative, looking for alternative explanations, and so on). They expect this conciliation and checking to occur *prior* to publication of the results. In light of the consensus, they are confident that rechecking will likely expose errors, or that conciliation will have occurred such that responsible researchers will not take their findings to conflict with the consensus (they will present them as anomalies to be explained, at most). While an initial disagreement is expectable, a published disagreement is highly unexpected, given that conciliation is assumed already to have taken place.⁹ Because it is highly unexpected, it is highly informative, and the best explanation for it cannot cite a simple reasoning error (given the assumption that such errors would have been detected at the rechecking stage). Instead, the disagreement forces experts to change their credences with regard to the dissenter, relegating them from the status of peers. Once again, *agreement* is highly expected and uninformative.

Before turning to our third kind of case, let me make good on the claim that the typology provides resources for making sense of more messy cases. Obviously, ordinary and extreme disagreements are ends of a continuum, because the difficulty of problems is continuous. It follows, therefore, that the degree to which a dispute is surprising is itself continuous. Cases midway between ordinary and extreme disagreements are therefore to be expected. In ordinary disagreements, because the disagreement is only somewhat surprising, we best restore coherence in our beliefs by lowering our confidence in the proposition over which we conflict. In extreme disagreements, our confidence in that proposition is far too high for it to be easily shaken, but we are less confident that our disputant is a peer, and we best restore coherence by reducing our confidence in peerhood. In messier cases, we may be faced with a choice: we might restore coherence either by reducing our confidence in the target proposition, or in peerhood, or by making adjustments in both.¹⁰

(3) Deep disagreement

Partisan disagreements—those that paradigmatically give rise to the problem of spinelessness—are *deep* disagreements. Unlike ordinary and extreme disagreements, deep disagreements are *not*, or at any rate need not, be unexpected.¹¹ Because these disagreements are not unexpected, by themselves (the need for this qualification will become apparent in due course) they do not carry information. *A fortiori*, they do not carry information about the likelihood that one of us has made an error or is not competent or sincere. Because they do not carry this kind of information, they do not constitute pressure to revise our beliefs about the target proposition or about peerhood.

Locke (2017) defines deep disagreements as disagreements that arise out of a disagreement over some principle. Empirically, these kinds of disagreements may usually arise from the different emphasis that the opposing sides place on what Jonathan Haidt and colleagues (Graham et al. 2013; Graham et al. 2016) call the *moral foundations*: the basic values or principles by reference to which individuals make token moral judgements. While their theory was, as the name suggests, developed to explain differences in moral judgements, it appears to explain much of the partisan divide. Differences in scores on the moral foundations predict allegiance to political parties and (even controlling for such allegiances) attitudes to polarized issues. For example, scores on the purity/degradation foundation predict attitudes toward euthanasia, abortion and pornography (Koleva et al. 2012; again, see Graham et al. 2013 for detailed review).¹² In any case, whatever the explanation, partisan disagreement is often expected, and individual cases therefore carry no additional information. Exposure to partisan disagreement is something akin to drawing a black marble from an urn that I believe contains a roughly equal mix of black and white marbles. It does not give me a reason to change my views about the target proposition, any more than I should change my view about the composition of the urn.¹³

Because I expect deep disagreement whether *p* to occur, and expected it *before* it occurred, (token) disagreement does not put me under pressure to conciliate with regard to the target proposition. Ought I to change my view about the sincerity or competence of the disputant? There are of course disputes on which there is a partisan split and with regard to which such a response is warranted (think of ‘pizzagate’, or claims that the Sandy Hook massacre was a ‘false flag’ operation; both theories found almost

exclusively on the right). The best explanation of this kind of partisan dissent may indeed impugn either sincerity or competence. Many of those who advocate for these views are trolls, or what Pritchard (2018) calls ‘dialectical poseurs’, and do not genuinely disagree with us. Those who are sincere are often in the grip of some poisonous ideology that distorts their capacity to process evidence. These cases are better seen as cases of extreme, rather than deep disagreement, given that the falsity of the claims should be obvious to everyone, whatever their ideology; correlatively, the best explanation for such disagreements impugns the sincerity or the competence of the disputants.

Ordinary disagreements provide information: because generating the right response in these cases is moderately difficult, and the other person is a peer, disagreement provides evidence that at least one of us has failed to reason properly. Extreme disagreements also provide information: because generating the right answer in these cases is easy, or because there is a consensus, disagreement provides evidence that the other person is not a peer by casting doubt on their sincerity or their competence. In both cases, disagreement points to a significant probability of *error*: error in my reasoning or error in taking my disputant to be a competent and sincere reasoner. But deep disagreements do not provide information about the possibility of error. Deep disagreements do not arise because one or the other side has made an error in reasoning (typical partisan disputes are after all tokens of a type: *my* dispute with *you* over gun control or prison reform is more or less representative of *my* side’s dispute with *yours* on these polarized issues; but whereas it is plausible that I or you may have made a mistake in reasoning on this occasion, it is scarcely credible that all of us and all of you have made such a mistake on most occasions). Nor do they arise because one side or the other is incompetent or insincere (again, the fact that our dispute is a token of a type serves to rule that out: your side has its experts, after all).

Deep disagreements therefore do not provide us with any new information. Given what I know about you, or what I can infer, I expected you to disagree with me about climate change or gun control. Neither you nor I learn anything new from these expected disputes: they provide no information about our reasoning or that of the other side. They do not challenge our causal model of the world or of each other; rather they confirm it. Because *token* deep disagreements do not inform us either about the possibility of error or about the competence or sincerity of the disputants, they do not place pressure on the disputants to conciliate on the target proposition.

Might not this merely push the question back, though? While the token disagreement does not indicate that either you or I have made an error in reasoning, given our different starting points, is not the very fact that we expected the disagreement indicative that at least one of us has made a much more fundamental error? We expected the disagreement because it is one of a series of such disputes—*our* side disagrees with *theirs* time and time again—and this pattern of disagreements surely suggests that at least one of us had made an error, if not in the reasoning which leads us to our conclusions, then in the premises from which we reason? Surely, then, we are now under pressure to conciliate with regard to these foundations and principles, rather than the target propositions or the capacities of agents?

Surprise carries information: it provides us with a reason to update our causal model of the world. In its light, we ought to reduce our confidence in the disputed belief, or in our belief that the disputant is a peer. The fact that we expect deep disagreements

indicates that this disagreement is predictable given our causal model. Given further assumptions (such as uniqueness—the view that the evidence licenses a single doxastic attitude), this in turn entails that we *do not* see our disputants as likelihood peers: that is, agents as likely to reach the conclusions on these matters as we are.¹⁴ However, while these facts show that deep disagreements do not have the epistemic significance of ordinary and extreme disagreements, they seem troubling nevertheless. The other side has (in many cases, at any rate) the markers of peerhood: intelligence, education, thoughtfulness and so on, and access to the same evidence as we do. Does not the fact that we reach opposing conclusions place pressure on us to conciliate with regard to our starting points?

I will argue that the deep disagreements have a different kind of epistemic significance than ordinary and extreme disagreements. They give rise to concerns about irrelevant influences, not error.¹⁵

5 The Real Epistemic Significance of Deep Disagreements

In ordinary and extreme disagreement, dissent is unexpected and provides information. At least one of us has made a mistake or is—unexpectedly—not a peer. Deep disagreements do not provide either kind of information. They do not provide evidence that at least one of us has made a mistake in reasoning. They are precisely what we should expect, given competent reasoning *from our divergent principles*. The token disagreement provides us with no new information at all, while the pattern helps constitute and provides us with confirmation that we reason from such different principles. Any pressure to conciliate therefore attaches to these principles, not reasoning or peerhood.

Empirical evidence (from moral foundations theory) supports the commonsensical hypothesis that these differences in basic premises stems from differences in enculturation (as well, perhaps, as differences in genetics). In asking whether deep disagreement provides us with reason to conciliate, we ask about our confidence in these starting points, and whether they constitute irrelevant influences on our beliefs.¹⁶ An influence is irrelevant if it does not bear on the truth of the belief (Vavova 2018). If my belief that *p* is counterfactually dependent on factors that do not bear on its truth, then I should lower my confidence in *p*. But that is not because I have made an error: it is because of the irrelevant influence.

It is at this point that my account of the epistemic significance of deep disagreements diverges most sharply from Vavova's. She, too, thinks that partisan disagreement is to be expected. For her, partisan disagreement is expected *because the issues are so difficult*: given how complex these questions are, we should have low confidence in our answers, and therefore we should not be surprised when disputes arise. For her, these disputes do not place us under pressure to conciliate, not because the disputants are not peers, but because our confidence in our answer should *already* be low. The reasons for low confidence and for expecting disagreement are one and the same—the question is difficult.¹⁷ The view put forward here is not committed to the claim that these disputes are difficult.

Some of the questions on which partisans divide *are* difficult. Vavova gives as an example the question whether death penalty is an effective deterrent. Answering that

question involves a wide variety of factors, data from multiple jurisdictions and difficulties in controlling for confounds. But these kinds of difficulties do not prevent us from having justified high confidence in other areas: climate change is equally subject to partisan dispute, equally (if not more) difficult, and yet one side (alone) has the right to high confidence in its beliefs. Other deep disagreements may not be especially difficult at all. Conservatives and liberals continue to be deeply divided on the acceptability of homosexuality (Brenan 2018), but the question really does not seem all that difficult. In the recent past, of course, there were deep partisan disagreements on a range of questions that we now see as obvious: the equality of women, for example. We can have justified high confidence in some of our partisan beliefs, either because they are easy for those of us lucky to have the right foundations, or because relevant experts can assess them despite their difficulty, and transmit their findings to us via testimony. We should lower our confidence only when further investigation reveals that the belief forming strategies we competently deploy are not reliable.

The second way in which my account diverges from Vavova's is more significant. While she distinguishes irrelevant influence from disagreement, noting that not all cases of the latter involve the former, she maintains that evidence of irrelevant influence is a genus of the same species as evidence of disagreement, and that species is error. I have suggested that deep disagreement turns on irrelevant influences, and that it does not constitute evidence of error. At first pass, the difference may appear merely terminological: Vavova uses 'error' to encompass all factors that lead to mistaken conclusions, including those that are typically at issue in irrelevant influence cases (such as bias and inculcation (2018: 141)), whereas I use 'error' more narrowly, to pick out those mistakes that can be rapidly corrected on the basis of new information (by lowering our confidence in either the disputed belief or in the belief that an agent is a peer; of course, such correction is only a first response but counts as a genuine correction insofar as it restores coherence to our beliefs). But the dispute is more than merely terminological.

The accusation that someone believes something due to an irrelevant influence may sometimes come as a surprise to the accused but it need not; importantly, it retains whatever force it has as an epistemic challenge even after it comes to be expected.¹⁸ The fact that these accusations do not owe their epistemic significance to being unexpected entails that they do not function in the same way as ordinary and extreme disagreements. Unexpected input is evidence that my causal model of the world is in some way faulty (you are not my peer; I have made a mistake in my reasoning; the beer is not where I thought it was; and so on). Expected input is consistent with my causal model of the world; if anything, it should raise my confidence that things are as I thought. The fact that deep disagreement is unsurprising demonstrates that it does not have the same kind of epistemic significance as other kinds, and aligns it with accusations of irrelevant influence. Of course, these challenges are indeed species that belong to the same genus, if that genus 'challenges to the reliability of our beliefs'. But there are important differences between them (and indeed other members of that genus, such as sceptical challenges) and it is better to think of them as belonging to quite different categories.

I do, however, join with Vavova (2018) in thinking that irrelevant influences sometimes and only sometimes weaken the epistemic credentials of our beliefs. One question is whether the case is a permissive one, and whether permissivism might

plausibly stretch to encompass the competing opinions. There may be symmetry breakers between influences. It is an open question whether our moral foundations are genuinely irrelevant influences and whether there might be such symmetry breakers between those who emphasise different foundations. Haidt sometimes suggests that there are such symmetry breakers: conservatives are in a better epistemic position than liberals because they are sensitive to a broader range of foundations. Of course, by itself that is question-begging. Assessing the epistemic credentials of these competing visions is obviously extremely difficult. We might resort to pragmatic criteria (which set of foundations better conduces to the welfare of the members of a society, for instance)? Perhaps we might invoke the epistemic perspective of people who have experienced ‘conversions’ from one outlook to the other; perhaps they have some epistemic privilege on the question. Perhaps we may be able to undermine one outlook or the other (or both) on the kinds of grounds invoked in the large literature on evolutionary debunking.¹⁹ Any of these strategies might provide us with a reason to lower our confidence in our beliefs. But they do so by providing *additional* evidence: evidence beyond the fact of deep disagreement itself. Since these disagreements are (to repeat) expected or expectable, they do not require updating in the absence of such additional evidence.

While deep disagreement constitutes a spur to further investigation—to the search for such additional evidence—there is no *a priori* reason to think that the investigation will reveal that our belief forming processes are not reliable. There is no reason to think that our epistemic situation is akin to the paradigms featuring in the literature on irrelevant influences, such as the person who has just taken a drug that distorts certain kinds of reasoning 20% of the time (Christensen 2016) or who has participated in an experiment in which half the subjects were subliminally primed to believe something (Vavova 2018). If we know that there is some significant probability we are subject to such distorting influences, we should significantly lower our confidence in the affected beliefs, but we have no special reason to think that is true of us just because we are disputants in a deep disagreement. But we cannot ignore these disputes either, refusing to see them as reasons to engage in further investigation. If we are not in an epistemic situation akin to those just mentioned, we are also not in the position of the person who contemplates whether she is a brain in a vat. Unlike the first kind of case, we *lack* evidence that there is some significant probability that we are subject to a distorting influence, but unlike the second, we *possess* evidence that we are subject to an influence that is possibly distorting. We seem to have reason to conciliate, pending further investigation, but to a far lesser extent than those who have evidence of error, or of irrelevant influence.

6 Conclusion

Epistemologists talk about the epistemic significance (singular) of disagreement. If the view put forward here is correct, that is something of a mistake. Different kinds of disagreement have different kinds of significance. I have developed a typology of disagreement, distinguishing *ordinary*, *extreme* (which itself comes in two kinds) and *deep* disagreement. I have argued that the first two kinds of disagreement provide information, because they are (to some degree) unexpected. Disagreement of these

kinds inform us of the likelihood of error on one or both sides. When we have no less reason to think the error is on our side than the other, we ought to conciliate. When we have reason to think that the error is on the other side, we ought to relegate the disputant from the status of peerhood.

Deep disagreements—*persisting* deep disagreements—are beasts of a crucially different kind. They are not evidence of error, either in our reasoning or in identifying peers. Rather, they point to differences in the principles or worldview from which we argue. Surprise is the hallmark of ordinary and extreme disagreements, and surprise points to a need to update our causal model of the world. But deep disagreements are not surprising; at least, they do not owe their significance to surprise. (Of course, we sometimes learn of a new disagreement that turns out to be deep, and that is surprising. But when that happens we mistakenly assimilate it to extreme or ordinary agreement; we take it to be evidence of error. It takes time to learn that the disagreement is deep and that those who disagree with us are, in some sense, our peers, if not our likelihood peers).

Rather than pointing to error, deep disagreement points to worries about the starting points from which we reason. They provide us with grounds for worrying that these starting points are arbitrary: that they are irrelevant influences on belief. They do not provide us with reasons to conciliate, since we have no reason to think those with whom we disagree in these cases are likelihood peers. Rather, they provide us with a reason to investigate these starting points. This is not a reason to conciliate because the evidence of irrelevant influences is weak. I do not gain evidence that one of us has made a mistake and that it might just as likely be me as them. Rather, I gain evidence that one of us might have made a mistake and it just might be me. Perhaps that is a reason to lower my confidence a little, but not to anything like the degree we see in ordinary and extreme disagreements. (Any such conciliation should take place early, but deep disagreements persist. Encountering yet another token of the type is entirely expectable and should not lead to any further reduction in confidence).

Whereas ordinary and extreme disagreements are the province of epistemology, traditionally conceived, deep disagreements require us to move beyond epistemology. Deep disagreements are the product of social facts, and they are the province of social epistemology. Their assessment may take us further afield, too: into sociological and psychological inquiry into the nature of the mechanisms involved. The significance of deep disagreement is quite different from that of the other kinds, and calls for a different response and different kinds of inquiry.²⁰

7 Notes

1. While my focus is on these kinds of moral and political disputes, conciliationism might entail agnosticism about religious beliefs too. See De Cruz (2018) for discussion and proposals for leveraging religious diversity to better justify religious belief.
2. This example is based on the infamous case of David Irving, who (unsuccessfully) sued Deborah Lipstadt for calling him a Holocaust denier. Unlike the Holocaust cases which feature in the disagreement literature (see,

for instance, Christensen 2018), we cannot dismiss Irving from peerhood on the basis of lack of acquaintance with the evidence.

3. It is worth noting that switching the focus from disagreement between isolated epistemic peers to disagreement between large numbers of people, or between individuals who recognize themselves to be representative of such groups, helps to avoid the worry is that it rarely or never is the case that two people have access to the same set of evidence. When we focus on real-world disagreements, there is no real question whether one side has access to evidence that the other lacks. When a dispute is years or even decades old, we can be confident that all evidence has been shared. Similarly, we can often rule out the influence of certain kinds of ‘personal information’ (Lackey 2010a), like the fact that I know myself to be sincere and unintoxicated, which are often cited as symmetry breakers. I cannot reasonably speculate most partisan disputes arise from these kinds of factors (as a reviewer points out, however, there are exceptions. The most obvious are conspiracy theories. Partisans are often obsessive consumers of information about their favoured conspiracy, so cannot be accused of lack of evidence, but there is a strong case for thinking they can be dismissed from peerhood by reference to personal information, such as our confidence we are sincere and not suffering from certain psychological dysfunctions).
4. Note that in saying that partisan dissenting peers may constitute pressure to conciliate that falls short of splitting the difference, I do not beg the question against the equal weight view. The numbers count; pressure to conciliate will be proportional to the numbers on each side (perhaps weighted for expertise). Taking the numbers into account is a way of giving each (equally expert) dissent equal weight.
5. In response to Vavova’s (2014a) argument, expanding on Elga’s, that widely individuated disputes do not provide pressure to conciliate because we should not see such disputes as pitting peers against one another, Fritz objects that the view depends on ‘an implausibly systematic picture of moral thinking’ (110). As Fritz cashes out ‘systematicity’, he is quite right: people don’t in fact rely on a ‘very systematic first-order moral theory’, from which they infer moral conclusions. But there is every reason to think that moral (and nonmoral) beliefs hang together in a mutually reinforcing, if not entailing, web. It is the web of belief, and not inference from first principles, that those who urge the no determinate answer view presuppose.
6. This way of describing unexpected input and how it causes belief update is inspired by predictive processing models of cognition (see Hohwy 2013; Clark 2016). Other elements of the predictive processing framework, such as the significance of the precision of expectations and of unexpected input, could also be incorporated smoothly into a model of disagreement, though I will not attempt to fill in these kinds of details.
7. Cases like *Holocaust* can be handled in analogous ways. The best explanation for David’s denial that the Holocaust occurred, given the overwhelming evidence for it, is that he is blinded by ideology. This highly unexpected input calls for revision, but not revision of the belief *that the Holocaust occurred*.
8. This typology builds on, but goes considerably further than, Locke (2017).

9. It is important to note that ‘publication’ (literal, or the public defence of a view) is expected to occur some significant time after rechecking: it is only persistent disagreement that is shocking. To see this, consider a counterexample to conciliationism put forward by Decker and Groll (2013). They cite Marilyn vos Savant’s solution to the Monty Hall problem. The solution she put forward was correct, but widely rejected by ordinary people and experts alike (including, Paul Erdős). Intuitively, Savant was correct to hold fast in the face of such widespread and expert dissent. But *first* she should lower her confidence and recheck her answer. At t_1 , she should be surprised but not shocked that there is such a dispute and its existence informs her that she or her opponents have made an error (though she put forward her solution after careful consideration, so did her opponents, and the problem is difficult enough that mistakes may occur). At t_2 , however, a time that occurs after she has carefully checked her answer and those of her opponents (hours or days after t_1), she should return to her earlier confidence. But at t_3 (perhaps months or even years down the track), she should conciliate if she has not succeeded in convincing her opponents. She should think she has made an error or that she has lost her competence in this domain. Of course, that would entail losing confidence in the correct view, but we are fallible epistemic beings, and misleading evidence has the power to lead us astray. It is, nevertheless, evidence.
10. In her discussion of what she calls ‘astonishing reports’, Jones (2002) points out that what counts as astonishing for an agent is coloured by their backgrounds and by sometimes their biases. As a consequence, they might relegate someone from the status of peerhood on illegitimate grounds. She recommends we treat the credibility of the testimony and of the informant independently. However, her worry seems to be the product not of (in my view correctly) relegating an informant from peerhood on the basis of her testimony, but a kind of double counting: when we take the incredibility of the report as a basis to discount the informant and *then* refer to the unreliability of the informant to further discount the testimony. That is clearly illegitimate. In standard extreme disagreements (standard in the literature, that is, if not in real life) if anything the pressure is in the other direction: to give the testimony more credibility than it deserves on the basis of the prior perception of the informant’s credibility. Jones is surely right that we sometimes illegitimately relegate someone from peerhood on the basis that we find her report incredible; that is a sad fact of epistemic life and of our fallibility as agents.
11. Of course, deep disagreements may be unexpected: encountering a partisan of a particular view for the very first time is not all that unusual (a reviewer for this journal gives the example of encountering in graduate school people who believe that every logically possible world exists). Such cases illustrate how a disagreement may belong to more than one category at once. When we first encounter such a partisan, our surprise indicates that our model of the world requires revision. When we have worked through them, we come to see that the disagreement calls neither for a reduction in our confidence in our reasoning (assuming that the disagreement survives debate) nor in our perception of the disputant’s competence or reliability: they call for revision of my belief that you share my (metaphysical) principles.

12. Summarising a large body of evidence, and abstracting from a number of complications and subtleties, those on the political left tend to emphasise the care and fairness foundations to a greater extent than those on the right, while those on the right place more emphasis on the loyalty, authority and purity foundations. The differences are very significantly driven by the latter three foundations: the extent to which the left puts a greater emphasis on care and fairness is relatively small, but the right place very much more emphasis on the latter three than the left (indeed, those on the left may see deference to authority as a vice, rather than a virtue). It is easy to see how MFT helps explain why Republicans and Democrats diverge on some of the issues that characterise the partisan divide. For instance, differences in attitudes to illegal immigration are surely at least partially explained differences in concerns for group loyalty, as well as respect for authority.
13. Hallsson (2019) puts forward an alternative view that entails that partisan disagreements do not provide unexpected information about the target proposition or the disputants. He cites empirical evidence that more able partisans have more extreme beliefs, and suggests that this shows that the more able we are, the more we are susceptible to bias. This entails, he argues, that able partisans on both sides are unreliable. Disagreement should give such partisans no reason to conciliate, but they should not be confident in any case. I think Hallsson mischaracterizes the empirical evidence and its import. The polarization he points to does not arise from a boost in our capacity to confabulate provided by sophistication: rather, it arises from the facility it provides with regard to the capacity to recognize *genuine* problems with claims to which we are ill-disposed, while failing to subject our own views to much scrutiny (see Levy 2019 for discussion). Hallsson also overplays the evidence: the phenomenon is genuine, but not so powerful that real expertise cannot swamp it. Where are the genuinely sophisticated climate science denialists he seems to predict? The consensus among climate scientists demonstrates that when problems are empirically tractable, disputes can be resolved despite our biases.
14. Uniqueness is widely held, but it is of course controversial (see Kopec and Titelbaum 2016). But weaker and much less controversial theses yield the same result in many deep disagreements. We need accept only *constraint*: a particular body of evidence licenses at most only a certain range of credences; moreover *constraint* might be restricted to certain propositions.
15. Several philosophers have argued that the problem of irrelevant influences just is a way of making disagreement salient to us (see, for instance, White 2010). I think this is a mistake. There need be no disagreement at all for the challenge from irrelevant influences to arise: we might *all* be subject to an irrelevant influence (that's one way of seeing the challenge on which evolutionary debunking focuses). Moreover, it overlooks something distinctive about the problem of irrelevant influences: when we raise this problem, we point to a troubling correlation between a seemingly irrelevant factor and a belief; we don't just highlight disagreement.
16. As a reviewer for this journal points out, it is unlikely that all our deep disagreements arise out of differences in the weight we give the moral foundations. For example, disputes in metaphysics or indeed epistemology may go deep (think of disputes between naturalists and non-naturalists). While I think that most of the

- deep disagreements that give rise to concerns about spineless are moral or directly or indirectly related to and constrained by moral principles, deep disagreements surely arise out of conflicts over other principles or differences in world view.
17. At least, this is how Vavova (2014b) understands these disputes. Her other 2014 paper advances an account of why we should not conciliate that is much closer to Elga's.
 18. DiPaolo and Simpson (2016) argue that irrelevant influence is worrying insofar as it is evidence for indoctrination, understood as the bypassing or inhibition of the capacity for rational thought. While that may often be true, indoctrination is clearly not necessary for deep disagreement. The thoughtful conservative or liberal may be aware of all the sociological facts concerning how partisan beliefs arise. There is no need for them to be cut off from evidence or argument for the etiological challenge to confront us.
 19. Kumar and May (2019) suggest that such targeted debunking strategies might be more successful than the sweeping strategies more common in the literature. McKenna (2019) argues that empirical evidence about the belief formation processes involved in the formation of partisan beliefs should lead both sides to lower their confidence in their beliefs; elsewhere (Levy 2019) I have argued that the same set of evidence he discusses has different implications for the different sides.
 20. I am extremely grateful to two reviewers for this journal for extremely helpful comments. Work leading up to this paper was funded by a generous grant from the Australian Research Council (DP180102384).

Compliance with Ethical Standards

Conflict of Interest The author declares that he has no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Benkler, Y., Faris, R., Roberts, H. & Zuckerman, E. 2017. Study: Breitbart-led right-wing media ecosystem altered broader media agenda. *Columbia Journalism Review*, March 3. <https://www.cjr.org/analysis/breitbart-media-trump-harvard-study.php>. Accessed 26 Jan 2019.
- Brenan, M. (2018). Above all issues, abortion divides liberals, conservatives. *Gallup*, June 13. <https://news.gallup.com/poll/235640/above-issues-abortion-divides-liberals-conservatives.aspx>. Accessed 19 Dec 2019.
- Christensen, D. (2007). Epistemology of disagreement: the good news. *Philosophical Review*, 116, 187–218.
- Christensen, D. (2016). Disagreement, drugs, etc.: from accuracy to akrasia. *Episteme*, 13, 397–422.
- Christensen, D. (2018). On acting as judge in one's own (epistemic) case. *Proceedings and Addresses of the American Philosophical Association*, 92, 207–235.

- Clark, A. (2016). *Surfing uncertainty: prediction, action, and the embodied mind*. Oxford: Oxford University Press.
- Decker, J. & Groll, D. 2013. The (In)significance of moral disagreement for moral knowledge. In Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics, Volume 8*. Oxford University Press.
- De Cruz, H. (2018). *Religious disagreement*. Cambridge: Cambridge University Press.
- DiPaolo, J., & Simpson, R. M. (2016). Indoctrination anxiety and the etiology of belief. *Synthese*, 193, 3079–3098.
- Drummond, C., & Fischhoff, B. (2017). Individuals with greater science literacy and education have more polarized beliefs on controversial science topics. *Proceedings of the National Academy of Sciences*, 114(36), 9587–9592.
- Elga, A. (2007). Reflection and disagreement. *Noûs*, 41, 478–502.
- Frances, B. & Matheson, J. (2018). Disagreement. *The Stanford Encyclopedia of Philosophy* (Spring 2018 Edition), ed. Edward N. Zalta. <https://plato.stanford.edu/archives/spr2018/entries/disagreement/>. Accessed 6 Nov 2019.
- Fritz, J. (2018). Conciliationism and moral spinelessness. *Episteme*, 15, 101–118.
- Graham, J., Haidt, J., Koleva, S., Motyl, M., Iyer, R., Wojcik, S. P., & Ditto, P. H. (2013). Moral foundations theory: the pragmatic validity of moral pluralism. *Advances in Experimental Social Psychology*, 47, 55–130.
- Graham, J., Haidt, J., Motyl, M., Meindl, P., Iskiwitsch, C., & Mooijman, M. (2016). Moral foundations theory: on the advantages of moral pluralism over moral monism. In K. Gray & J. Graham (Eds.), *The Atlas of Moral Psychology: Mapping Good and Evil in the Mind*. New York: Guilford Press.
- Hallsson, B. (2019). The epistemic significance of political disagreement. *Philosophical Studies*, 176, 2187–2202.
- Hohwy, J. (2013). *The predictive mind*. Oxford: Oxford University Press.
- Jones, K. 2002. The politics of credibility. In L. Antony and C Witt (Eds.), *A mind of one's own: feminist essays on reason and objectivity*. Boulder, Co.: Westview Press, pp. 154–176.
- Koleva, S., Graham, J., Haidt, J., Iyer, R., & Ditto, P. H. (2012). Tracing the threads: how five moral concerns (especially purity) help explain culture war attitudes. *Journal of Research in Personality*, 46, 184–194.
- Kopec, M., & Titelbaum, M. G. (2016). The uniqueness thesis. *Philosophy Compass*, 11, 189–200.
- Kumar, V., & May, J. (2019). How to debunk moral beliefs. In J. Suikkanen & A. Kauppinen (Eds.), *Methodology and Moral Philosophy* (pp. 25–48). New York: Routledge.
- Lackey, J. (2010a). A justificationist view of disagreement's epistemic significance. In A. Haddock, A. Millar, & D. Pritchard (Eds.), *Social Epistemology* (pp. 298–325). Oxford: Oxford University Press.
- Lackey, J. (2010b). What should we do when we disagree? In T. S. Gendler & J. Hawthorne (Eds.), *Oxford Studies in Epistemology* (3rd ed., pp. 274–293). Oxford: Oxford University Press.
- Levy, N. (2019). Due deference to denialism: explaining ordinary people's rejection of established scientific findings. *Synthese*, 196, 313–327.
- Locke, D. (2017). The epistemic significance of moral disagreement. In T. McPherson & D. Plunkett (Eds.), *The Routledge Handbook of Metaethics* (pp. 499–518). New York: Routledge.
- Matheson, J. (2015). *The epistemic significance of disagreement*. London: Palgrave Macmillan.
- McGrath, S. (2008). Moral disagreement and moral expertise. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics, volume 3* (pp. 87–108). New York: Oxford University Press.
- McKenna, R. (2019). Irrelevant cultural influences on belief. *Journal of Applied Philosophy*, 36, 755–768.
- Pritchard, D. (2018). Disagreement of beliefs and otherwise. In C. R. Johnson (Ed.), *Voicing dissent: the ethics and epistemology of making disagreement public* (pp. 22–39). New York: Routledge.
- PublicMind. (2012). What you know depends on what you watch: current events knowledge across popular news sources. <http://publicmind.fdu.edu/2012/confirmed/>. Accessed 11 Dec 2019.
- Rowland, R. (2017). The epistemology of moral disagreement. *Philosophy Compass*, 12, e12398.
- Sherman, B. 2014. Moral disagreement and epistemic advantages: a challenge to McGrath. *Journal of Ethics and Social Philosophy* 8(3). <https://doi.org/10.26556/jesp.v8i3.82>.
- Sosa, E. (2010). The epistemology of disagreement. In A. Haddock, A. Millar, & D. Pritchard (Eds.), *Social Epistemology* (pp. 278–297). New York: Oxford University Press.
- Vavova, K. (2014a). Moral disagreement and moral skepticism. *Philosophical Perspectives*, 28, 302–333.
- Vavova, K. (2014b). Confidence, evidence, and disagreement. *Erkenntnis*, 79, 173–183.
- Vavova, K. (2018). Irrelevant influences. *Philosophy and Phenomenological Research*, 96, 134–152.
- White, R. (2010). You just believe that because. *Philosophical Perspectives*, 24, 573–615.