

Conciliationism and Moral Spinelessness

James Fritz

Abstract This paper renews a challenge to conciliationist views of the epistemically correct response to disagreement. This challenge, which has sometimes been called the ‘problem of spinelessness’, asks conciliationists to show that their view does not require intuitively implausible reductions in confidence. Here, I raise problems for a recent attempt to solve the problem of spinelessness about moral belief. This attempt, due to Katia Vavova, crucially involves individuating disputes quite broadly. In doing so, it reduces the dispute-independent ground from which we can assess those who dispute our moral beliefs. This sort of conciliationism is initially attractive, but it cannot explain why we are rationally permitted to retain some of our most important moral beliefs. On any plausible version of such a view, the extent of moral agreement on commonsense matters provides us with a dispute-independent reason to consider our interlocutors trustworthy about even highly controversial moral questions. Given enough such interlocutors, conciliationism counsels reductions in confidence. It implies, in fact, that we will have to become much less confident in some controversial moral beliefs to which, intuitively, we are entitled.

This paper renews a challenge for conciliationist views of disagreement (collectively, *conciliationism*). This is the challenge that Adam Elga calls the “problem of spinelessness” (2007, 484): roughly, conciliationists must show that their views do not counsel intuitively implausible belief revision.¹ When steadfast belief seems clearly appropriate, they must explain why conciliationism does not recommend changes in confidence.

In a longer version of this paper, I argue that conciliationists cannot satisfactorily resolve the problem of spinelessness about moral beliefs.² Here, I will restrict my attention to the problems

¹ On my usage, to *revise* a belief is either to abandon it or to lower one’s confidence in it.

² Noncognitivist reject the assumption that moral judgments are beliefs, and some may also reject the assumption that anything like epistemic normativity applies to the formation and revision of moral judgments. Though this paper will assume that moral judgments are beliefs, it’s worth noting that this sort of noncognitivist can avoid its central tension. She can simultaneously accept conciliationism about beliefs and reject the conclusion that there is anything like an epistemic requirement to be spineless about moral judgments.

for a particular brand of conciliationism. This is the sort of conciliationism that individuates disputes widely. Katia Vavova (2014) has recently defended conciliationism from the charge of spinelessness about moral beliefs by adopting this very approach. She argues that, generally, we do not have good dispute-independent reason to trust people who reject our most cherished moral beliefs. This paper offers some reason to think that we *do* have good dispute-independent reason to trust such people, and that a conciliationism like Vavova's must therefore counsel moral spinelessness.

1. Conciliationism and Spinelessness

Conciliationist views are accounts of the epistemically rational response to disagreement. The defining trait of a conciliationist view is its acceptance of a claim that I'll be calling the *principle of independence*, or Independence for short.

Independence When considering how much epistemic weight to give your beliefs, I must assess your epistemic credentials in a way that relies only on dispute-independent reasoning.³

Perhaps the trickiest question about Independence is this: just which beliefs and reasoning are *dispute-independent*? In other words, just how broad is the scope of any given dispute? One family of conciliationist views, most notably defended by Katia Vavova, construes disputes quite broadly. Vavova claims that, when I disagree with a crystal-ball-using mystic about what will happen tomorrow, my background beliefs about what provides evidence about the future are not dispute-independent (2014, 310). According to Vavova's Independence, then, I must set these background beliefs aside when assessing the mystic's credentials.⁴

³ Cf. Vavova 2014, 308-9; Christensen 2011, 1-2; and Elga 2007, 490-1.

⁴ Adam Elga, another notable conciliationist, would likely object to Vavova's broad individuation of disputes. Elga individuates disputes quite narrowly. In the context of disagreements about abortion's permissibility, for example, he contends that beliefs about when fetuses begin to feel pain can be dispute-independent (496). Elga's reason for avoiding a broad individuation of disputes is interestingly similar to Vavova's reason for adopting it: on a broad individuation of disputes, we will sometimes be required to set aside so much that we will have no independent grounds from which to assess our opponents' epistemic credentials. There is a

Independence gives rise to a worry that Adam Elga calls the “problem of spinelessness” (2007, 484). By forbidding me from appealing to some reasoning, Independence restricts the epistemically legitimate ways in which I can conclude that my opponent is untrustworthy. *Prima facie*, this seems to require me to take certain opponents quite seriously. It may even require me to lose confidence in cases where, intuitively, I am fully entitled to retain my disputed beliefs. Conciliationists usually seek to reject the charge of spinelessness; they argue that Independence does not require intuitively implausible belief revision.

Conciliationists begin to address the problem of spinelessness by distinguishing between what Vavova calls NIRP (the No Independent Reason Principle) and GIRP (the Good Independent Reason Principle)

NIRP Insofar as the dispute-independent evaluation fails to give me good reason to be confident that I am more trustworthy than my interlocutor, I must revise my belief in the direction of hers.

GIRP Insofar as the dispute-independent evaluation gives me good reason to be confident that my interlocutor is at least as trustworthy as I am, I must revise my belief in the direction of hers.⁵

Conciliationists generally reject NIRP precisely because it would lead to spinelessness.⁶ Consider a disagreement with a global skeptic. She disputes every one of my beliefs. For any proposition *p*, I have no undisputed reason to think that I am more trustworthy about *p* than the skeptic is. If NIRP is right, this means that I should revise my belief in the direction of the skeptic’s. GIRP does not

particularly common conciliationist reply to this worry: even after setting aside all of our beliefs in a particularly important “cluster” (like the cluster of beliefs around abortion), we would still be reasonable in forming different expectations about the comparative trustworthiness of different judges (Kornblith 2010, 49-50; Christensen 2011, 16; Vavova 2014, 314-5; and McGrath 2007, 105-6).

⁵ This formulation is adapted from Christensen (2011, 15). Here and throughout the paper, I use “trustworthy” as an abbreviation for the property that conciliationism asks us to look for in others. Different conciliationists offer different candidates for this property. According to Vavova (2014, 317) and Christensen (2011, 15) alike, the right sort of dispute-independent evaluation aims at determining whether an interlocutor is *well-informed* and *likely to have reasoned correctly from her evidence*. Contrast (Elga 2007, 490).

⁶ There are a few apparent exceptions to this rule; some thinkers appear to endorse a conciliationism that uses (1) instead of (2). See Sidgwick 1981, 342; King 2010, 267-9; and Feldman 2006.

have such counterintuitive results. It only requires me to revise my belief in cases where I have positive reason to take my interlocutor to be quite trustworthy.

By accepting GIRP instead of NIRP, a conciliationist begins to address the charge of spinelessness. But a crucial question remains: what would it take to have good dispute-independent reason to consider an opponent just as trustworthy as I am? According to GIRP, whenever this condition obtains, disagreement requires me to revise my belief.

This requirement seems particularly troubling in the moral domain. Certain moral beliefs are both very intuitively secure and deeply controversial. For this reason, conciliationists often attempt to explain why their theories do not recommend spinelessness about moral belief.⁷ The next section reviews Vavova's recent attempt to meet this challenge.

2. An Attempt to Avoid Spinelessness

Vavova argues that, when Independence is rightly understood, conciliationism does not require moral spinelessness. In fact, she thinks that conciliationism has attractive results in the moral domain: it only requires reduced confidence, she argues, when reduced confidence is independently well-motivated (2014, 325).

Vavova defends this claim with several complementary observations about moral thought and practice. First, she notes that there is negligibly little actual disagreement about certain moral beliefs. Even if disagreement threatens some moral beliefs, it is surely no threat to the belief that pain is bad or the belief that we should not kick puppies (2014, 326). There is simply not enough disagreement about these subjects to present a problem.

Second, Vavova notes that some "apparent moral disagreements aren't really disagreements about moral matters" (313). We might agree about whether graffiti artists should be imprisoned,

⁷ Indeed, Adam Elga's seminal discussion of the "problem of spinelessness" primarily focuses on an imagined disagreement about the permissibility of abortion (2014, 492-3).

for instance, while disagreeing about whether Jones should therefore be imprisoned. Our disagreement would come down to a disagreement about a non-moral question: whether Jones is a graffiti artist. Conciliating here would not amount to spinelessness about our moral commitments.⁸

If conciliationism requires spinelessness about our moral commitments, then, it will only be in a restricted class of disagreements: widespread controversies that “aren’t, at root, non-moral” (324). Even in these crucial cases, Vavova argues, disagreement cannot have much defeating power. When my interlocutor has a dramatically different moral worldview, our disagreements will not be limited to only a few topics; they will arise in “clusters” of related beliefs. On Vavova’s understanding of dispute-independence, this means that fewer of my beliefs will be dispute-independent. Now, recall that a disagreement cannot have defeating power unless it leaves me with dispute-independent reason to trust my interlocutor.⁹ So the more I have to set aside, the less epistemically powerful a disagreement is. This, Vavova argues, is appropriate: moral disagreement from the average political opponent should be more epistemically worrisome than moral disagreement from a sociopath (314-5).

By individuating disputes widely, Vavova appears to offer a tidy solution to the problem of spinelessness. Disagreements with people who share our “basic worldview” (324) take place against a backdrop of shared moral beliefs. So we have plenty of dispute-independent reason to trust these opponents. Conciliationism will sometimes call on us to lose confidence in these cases, but this is not intuitively implausible (322). We do *not*, on the other hand, have good dispute-independent reason to trust people whose moral judgments differ from ours systematically. These disagreements require us to set aside too much; they “go too deep” (323) to have defeating power.

⁸ There is an interesting question as to whether Vavova’s conciliationism would require objectionable spinelessness about *non-moral* beliefs. But this question lies beyond the scope of this paper. For some promising first moves, see Decker and Groll 2014.

⁹ Or, more accurately, it cannot have defeating power for reasons specified by a conciliationist view that, like Vavova’s, opts for (2) over (1). Vavova acknowledges that some disagreements about which her conciliationism is silent may have epistemic defeating power for distinct reasons (2014, 320).

Despite this initial promise, conciliationists cannot solve the problem of spinelessness by individuating disputes widely. As we'll see in the next section, we plausibly do have dispute-independent reason to trust the moral judgment of some people who reject some of our most treasured moral beliefs. When there are enough such people, conciliationism counsels spinelessness.

3. The Problem

I've just claimed that, even if we individuate disputes widely, we have dispute-independent reason to trust the moral judgment of a great many people—including some who reject some of our most cherished moral beliefs.

This may seem like a preposterous claim. Recall GIRP:

Insofar as the dispute-independent evaluation gives me good reason to be confident that my interlocutor is at least as trustworthy as I am, I must revise my belief in the direction of hers.

Isn't this an extremely high standard? Though I know plenty about my own trustworthiness, I know far less about the trustworthiness of the people who reject my most treasured moral beliefs. Most of these people are strangers. Wouldn't I need to know an enormous amount of personal information about someone before I could have reason to be confident that he is at least as trustworthy as I am?

This line of thinking is misleading. We could, of course, read "reason to be confident that my interlocutor is at least as trustworthy as I am" in a way that requires specific evidence about a particular person's track-record. But this is not the reading that the conciliationist should accept, precisely because even the disagreement of total strangers can clearly be an epistemic defeater.

Say that you are at a racetrack, watching the end of a horse race. You watch horses A and B cross the finish line, and you form the belief that horse A narrowly won. Now, say that a stranger beside you claims that horse B won. Intuitively, this disagreement has some defeating power.

Perhaps not very much; depending on the details, you may only be required to reduce your confidence a bit. But the notion that this person's disagreement just cannot rationalize belief revision, merely because you know your own track-record better than his, is clearly mistaken.

How can the conciliationist explain this?¹⁰ What gives you a reason to consider the stranger trustworthy? Well, plausibly, you are warranted in making certain assumptions about strangers based on what you know people are generally like. In other words, justified beliefs about the base rates in a population can underwrite justified beliefs about the epistemic credentials of individuals in that population. We all have a history of interactions that justifies us, for instance, in thinking that most people are good at perceiving medium-sized objects. So even when your interlocutor at the racetrack is a stranger, you still have *some* dispute-independent reason to expect that he perceives horse races well. The conciliationist should say that this provides dispute-independent evidence that the stranger is at least as perceptually trustworthy as you are.

Of course, this conclusion only threatens moral spinelessness if we have dispute-independent reason to consider strangers just as morally trustworthy as we are. But, plausibly, we do have such a reason: the enormous backdrop of moral beliefs that almost all people share. Vavova herself argues that, when it comes to most cases, people are in broad moral agreement. And this seems right. Although the practical importance of coordinating behavior tends to draw our attention toward areas of dispute, most people probably share the vast majority of their moral beliefs. Nearly everyone agrees that pain is bad and that we shouldn't kick puppies. Nearly everyone agrees that we should generally keep our promises, that it is generally better to compliment someone than to stab her, and so on. The consensus extends to a huge subset of the moral propositions that we are capable of entertaining. When we note that other people are reliable

¹⁰ First, the conciliationist faces an interpretive hurdle: she must use the phrase "reason to be confident that my interlocutor is at least as trustworthy as I am" in a way that suggests I could possibly have some dispute-independent reason to be confident that the stranger at the racetrack is at least as trustworthy as I am. For the purposes of this paper, I'll simply grant that such an interpretation can be offered. For some first moves in defending this sort of talk, see Christensen 2011, 16 and 2007, 212.

about all these moral questions, we gain inductive support for the expectation that their judgment will be accurate in other cases too.

Even when moral disputes require us to set quite a bit aside, then, we have good dispute-independent reason to trust our opponents. Now, this may not mean that any given person's disagreement will call for much belief revision. But, as in the horse race case above, the numbers matter.¹¹ When enough opponents independently arrive at beliefs that contradict mine, and those opponents are to some extent independently credible, my confidence should drop dramatically.

To see why this result is troubling, consider a case of actual disagreement. It seems likely that an enormous amount of people now and throughout history have considered homosexual behavior *pro tanto* morally impermissible. Conciliationism tells me to assess these people's epistemic credentials in a dispute-independent way. Even after setting aside all relevantly disputed beliefs, however, I have some reason to take these people to be as trustworthy as I am. Specifically, I have evidence that they believe the truth about most everyday moral questions. This means that, if my opponents are sufficiently numerous, Vavova's conciliationism will call on me to reduce my confidence in the permissibility of homosexuality.¹² But, intuitively, this is not the right result. If I know any moral truths at all, I know that homosexual behavior is morally permissible. To the extent that I am ever warranted in steadfastly retaining my current level of confidence about a moral belief, I am entitled to steadfastly retain my level of confidence about this one. This is not, in other words, a case in which a reduction in confidence is "independently warranted" (Vavova 2014; 302).

Given the argument of the last two sections, the conciliationist has an argumentative burden. She must reject the claim that agreement about commonsense moral matters gives us dispute-independent reason to trust our opponents in more controversial cases. The next section explains why this task will be so difficult.

¹¹ Cf. Vavova 2014, 313; Elga 494.

¹² Note that I may not have to reduce my confidence much. Depending on the number of my interlocutors, I may be left with an outright belief, albeit a more tentative one, in homosexuality's permissibility.

4. Responses

This section considers several ways in which the conciliationist might attempt to block the inference from awareness of commonsense moral agreement to the conclusion that other people are trustworthy about controversial moral questions. Broadly, these approaches fall into two categories: those that lean on differences between moral beliefs, and those that lean on connections between moral beliefs.¹³

4.1 Differences

It's surely an oversimplification to treat all moral thinking in the same way. Consider an analogy with mathematical thinking. The fact that people are generally reliable about two-digit addition does not give us good enough reason to think that any given stranger will get a calculus problem right.¹⁴ This analogy suggests two ways in which the conciliationist might try to protect controversial moral beliefs.

¹³ It's worth mentioning a third approach: couldn't the conciliationist argue that the mere fact *that there is controversy about p* renders irrelevant our inductive grounds for expecting people to get *p* right? After all, even if humans are generally reliable about a certain class of moral judgments, that reliability has its limits. And in the context of any given widespread moral disagreement, we have sufficient reason to think that we are outside the limits: the human population at large is *not* reliable about this question. This line of reasoning is initially attractive, but it proves too much. If the conciliationist adopts it, she will allow objectionably dogmatic belief in even non-moral cases. In the horse race example, for instance, your knowledge of others' general perceptual reliability is a dispute-independent reason for you to take your opponent's disagreement seriously. And you cannot dismiss this reason simply by noting that, between the two of you, there is controversy. Even though you have very strong reason to think that one of the two of you is mistaken in the case at hand, your belief in others' general perceptual reliability continues to give you dispute-independent reason to trust your interlocutor.

The conciliationist might be tempted to argue that, though the mere fact that there is controversy in a group of two does not screen off evidence of general reliability, controversy in a larger population can do so. But this fix also has unacceptable results. Intuitively, widespread cases of disagreement have (*ceteris paribus*) even more defeating power than two-person ones. It would be preposterous to say that, although a single stranger's disagreement about a horse race can call on me to reduce my confidence, the disagreement of half of the stadium cannot. Even when I know that there is massive controversy about *p* among all the people who have any belief at all as to whether *p*, then, I can have sufficient dispute-independent reason to take disagreement about *p* seriously.

¹⁴ Cf. Rotondo 2015, 262-3.

First, she might note that some moral questions are more difficult than others. The mere fact that most people generally get the *easy* moral questions right may give us no reason to expect that they will get the *hard* moral questions right.¹⁵

This approach is not a promising general solution to the problem of spinelessness. The *difficult* moral questions are not, intuitively, the ones which we are most entitled to retain in the face of disagreement.¹⁶ If we become convinced that a moral disagreement is explained by the difficulty of thinking clearly about the topic at hand, we thereby gain a reason to worry about *anyone's* trustworthiness—including our own.¹⁷ Vavova acknowledges this point. She writes that, when it comes to more “difficult moral questions, we should expect the rational confidence level to be lower” (325). So conciliationists should not attempt to avoid spinelessness by appealing to difficulty.

On to a second approach. Rather than emphasizing the *difficulty* of certain moral questions, the conciliationist might claim that we reach our controversial moral beliefs through a different *kind* of thinking than we use to form our generally shared moral beliefs. Perhaps, for instance, there is an important difference between the process that delivers moral beliefs about particular cases (e.g. John was wrong to lie to his mother) and the process that delivers beliefs in general moral principles (e.g. the principle of utility).¹⁸ If so, widespread agreement about particular cases may provide no support for the conclusion that people are trustworthy inquirers into the more general truths of ethics.

There is a problem for this approach as well: some of our most controversial moral beliefs are just as fine-grained as, and reached in the same way as, the beliefs about which there is most

¹⁵ Cf. Parfit 2011, 554.

¹⁶ Of course, any conciliationist will admit that, if we have some dispute-independent reason to think that we are more trustworthy about difficult questions than the average person, this symmetry-breaker can justify steadfast moral belief. But we do not usually have this sort of symmetry-breaker.

¹⁷ Cf. Kornblith 2010, 45.

¹⁸ Vavova provides some support for this proposal. She claims that, when it comes to morality, we agree on “a lot of the cases” but disagree more as matters become “more theoretical” (326).

plausibly near-universal agreement. We do not simply agree about the rightness or wrongness of act-tokens; we also agree a great deal about the defeasible rightness or wrongness of act-types.¹⁹ And some of the controversial beliefs about which we should most clearly be steadfast—for instance, the belief that homosexual behavior is *pro tanto* permissible—are themselves claims about the defeasible rightness or wrongness of act-types. The conciliationist, then, cannot protect these beliefs by appealing to their level of generality.²⁰

Nor can she do so by appealing to the way in which these beliefs are formed or maintained. Just like the beliefs about which there is near-universal agreement, these beliefs are generally introduced to children through testimony from adults; as those children grow, they question, revise, and maintain them in the light of their intuitions and experiences. Of course, there are certain important disanalogies between sorts of moral inquiry. But it seems quite unlikely that there is a consistent difference between the process that delivers our shared moral beliefs and the process that delivers our controversial ones.

4.2 Connections

The conciliationist might attempt to avoid moral spinelessness by noting connections between moral beliefs. Here's the idea behind this strategy: sometimes, when I learn that a person does not share one of my moral beliefs, this gives me powerful evidence that our disagreement is, in some sense, *systematic*. Once I know this, I either *lose* my dispute-independent reason to consider my interlocutor trustworthy or *gain* a dispute-independent reason to consider her *untrustworthy*.

¹⁹ In fact, it seems plausible that there is even *more* agreement about the *pro tanto* rightness or wrongness of act-types than there is about the rightness or wrongness of any given act-token. Real life is more complicated than thought experiments are.

²⁰ Vavova suggests another compelling reason not to defend our controversial beliefs by appealing to their level of generality: intuitively, our epistemic position with respect to general moral principles is worse than our epistemic position with respect to particular cases (2014, 326; cf. Wedgwood 2014, 36-8). This means that the conciliationist would get intuitively backward results by blocking the inference from others' reliability about particular cases to an expectation of reliability about general principles. Such a conciliationism would call for *more* belief revision about particular cases than about general principles.

Conciliationists who use this strategy must explain how disagreement about a particular moral question provides evidence of systematic disagreement. Let's consider two such explanations.

First, the conciliationist might claim that connections between beliefs are part of the very structure of moral thought. She might, for instance, argue that people who form judgments about some case-type are thereby committed, in some sense, to certain moral beliefs about many other case-types. Vavova appears to pursue this line of argument. She claims, for example, that "if the reason we disagree about abortion turns out to be *that we disagree about what morality requires*, then we have less independent ground than we thought" (2014, 323, emphasis original). The idea, I take it, is that a person's moral beliefs about particular cases or case-types are often indicative of that person's commitment to quite general principles. And when we disagree about general principles, we must set aside beliefs that follow from them. In such cases, we have scant ground from which to evaluate our opponents' credentials.

This strategy relies on an implausibly systematic picture of moral thinking. To borrow Vavova's phrase, we can disagree about *what morality requires* in a particular case (or case-type) without disagreeing about *what morality requires* generally. Two people who disagree about whether morality requires strong constraints against killing fetuses, for instance, could in principle agree about almost every other moral subject. Now, granted, there are certain moral theories that offer only a single fundamental moral principle (e.g. the principle of utility). When we learn that a person is committed to a principle like this, we can justifiably infer a great deal about her moral thinking. But this is not how typical moral disagreements work. To say otherwise would be to attribute an implausibly systematic approach to morality to the average person.

Perhaps the conciliationist could do better by appealing to a less controversial point: as a matter of sociological fact, certain moral beliefs tend to accompany others. Given what I know about

distribution of moral beliefs, the fact that a person disagrees with me about abortion gives me evidence about her views on, e.g., gun control or environmental policy.

This sort of sociological evidence cannot help the conciliationist to justify a steadfast response in the cases that concern us. To see this, first note that your most-likely-disagreeing-with-me about a matter related to p is only relevant to our disagreement about p if the disagreement about the related matter would be evidence of something about your thinking with respect to p . Now, it's entirely plausible that some disagreements do work this way. If I learn that you disagree with me about gun control and environmental policy, I often thereby get evidence suggesting that you were raised in a certain way, that you rely on certain sources of moral testimony, and so on. But this is the only way that related disagreements can give me reason to question your judgment about abortion. They must suggest that there is some bad-making feature of your moral thinking about abortion and other matters alike.

The problem is this: this crucial belief is usually not dispute-independent. So it is usually illegitimate for me to rely on this belief in assessing your epistemic credentials.²¹ Recall Vavova's example of disagreement with the mystic. In the context of this disagreement, my belief about what counts as good evidence for the future must be set aside. In just the same way, when I disagree with someone whose moral beliefs seem reprehensible to me, I am not entitled to rely on my belief (even if justified) that he was raised by a morally bankrupt family. After all, he would surely not grant as much.²²

²¹ There may be some features that are generally understood, in a dispute-independent way, to be bad influences on moral inquiry. This cannot solve the problem of moral spinelessness: we would not be justified in assuming, when it comes to the widespread disagreements that are most troubling, that there is always some uncontroversial bad-making feature of any opponent's moral inquiry.

²² There is, moreover, good reason for the conciliationist to require us to set aside these contested beliefs. If I can appeal to contested beliefs about bad influences on my opponent's thinking, I can justify steadfast belief in even some cases where conciliation seems appropriate. Consider an example: two groups of detectives, who otherwise take one another to have spotless track records, tend to reach contradictory conclusions in neighborhood X . Say that the evidence suggests that this is no mere coincidence; everyone involved believes that, for some reason, at least one of the two groups does detective work more poorly in X . Given their knowledge of past disagreements in X , is Group 1 entitled to approach Group 2's inquiry in X as distorted? If so, they may be justified in steadfastly ignoring any given new conclusion that Group 2 reaches in X . But this

Vavova's conciliationism, then, asks us to set aside disputed beliefs about what has gone wrong with our opponents' inquiry. But the fact that our opponents disagree with us in multiple ways entitles us to consider them untrustworthy only if it gives us dispute-independent reason to believe that something has gone wrong with their inquiry. So conciliationists cannot solve the problem of spinelessness by appealing to multiple instances of disagreement.²³

5. Conclusion

Conciliationists cannot solve the problem of spinelessness by individuating disputes widely. Even after we set a great deal aside, our history of commonsense moral agreement gives us a reason to trust others' moral judgment. This means that, given enough opponents, we must revise even our most cherished moral beliefs.

But just how bad is this result? Grant that we have *some* reason to trust strangers' moral judgment. Even so, the conciliationist might argue, our backdrop of shared moral beliefs might not be a *good* reason to trust our opponents. And conciliation only calls for belief revision when we have *good* dispute-independent reason to trust others. Conciliationists would go a long way toward solving the problem of spinelessness, then, if they could claim that widespread moral agreement is not a good reason to trust others' moral judgment.

But any claim to this effect is likely to be objectionably *ad hoc*. We've already seen reasons to think that moral agreement can be epistemically significant. For one, the extent of moral agreement is vast and impressive. Moreover, a background of impressive agreement is just the sort

seems quite contrary to the spirit of Independence. The right conciliationist verdict here is more moderate: in the absence of a dispute-independent tiebreaker, members of Group 1 should respect the possibility that they are the biased ones.

²³ This is not to say that the clustering of moral disagreements has no epistemic relevance at all. When I set aside more and more disputed beliefs, I have less and less dispute-independent reason to consider any given person trustworthy. But that is not enough to defeat the challenge raised in sections 2 and 3. Those sections suggested that our general background of commonsense moral agreement gives us good reason to trust our opponents. This case for trusting strangers is not affected by a difference in our ability to rely on *controversial* moral beliefs.

of thing that, in other domains, can give us sufficient reason to conciliate—recall the case of perceptual disagreement with a stranger.

So, whatever the right story is about the epistemic weight of background agreement, commonsense moral agreement seems likely to give us plenty of reason to take certain moral disagreements seriously. The conciliationist who wants to avoid this conclusion must offer a positive proposal about the epistemic weight of agreement. Given the considerations above, we should worry that any such story will have to take a gerrymandered, unprincipled approach to agreement across different subject matters.²⁴

Could the conciliationist shrug off the charge of spinelessness in a different way? Perhaps conciliationists can grant the surprising result that we must revise a few deeply-held moral commitments, like the commitment to homosexuality's permissibility. How bad would this be?

Well, nothing like moral skepticism is lurking here: indeed, this paper has insisted that an enormous amount of our everyday moral thinking is not affected by disagreement. But to the extent that conciliationists were right to worry about spinelessness in the first place, they should also worry about this local sort of spinelessness. First, conciliationism seriously threatens the justification of unpopular moral beliefs generally, from abolitionism in the 1800s to beliefs about the moral status of animals today. Second, a local spinelessness might be even more strikingly counterintuitive than wholesale moral skepticism. I would be very surprised to learn that my moral beliefs are globally unjustified. But I would be far more surprised to learn that, though most of my moral beliefs are adequately justified, my belief about homosexuality's permissibility is not. If I know anything at all about morality, I know that homosexuality is *pro tanto* permissible; I cannot imagine losing confidence in this belief without also losing confidence in *all* my moral thinking.

²⁴ Thanks to Declan Smithies and Tristram McPherson for encouraging me to make these points explicit.

In addressing the problem of spinelessness, conciliationists attempt to show that, wherever their theory calls for belief revision, that belief revision is intuitively plausible. But the approach considered in this paper cannot deliver this result.

References

- Christensen, David. 2007. "Epistemology of Disagreement: The Good News." *Philosophical Review* 116, no. 2: 187-217.
- . 2011. "Disagreement, Question-Begging, and Epistemic Self-Criticism" *Philosopher's Imprint* 11, no. 6: 1-22.
- Decker, Jason and Daniel Groll. 2013. "The (In)significance of moral disagreement for moral knowledge." In *Oxford Studies in Metaethics*, vol. 8, edited by Russ Shafer-Landau, 140-168. Oxford: Oxford University Press.
- Elga, Adam. 2007. "Reflection and Disagreement." *Nous* 41, no. 3: 478-502.
- Feldman, Richard. 2006. "Epistemological Puzzles about Disagreement." In *Epistemology Futures*, edited by Stephen Hetherington, 216-236. Oxford: Oxford University Press, Oxford.
- Kornblith, Hilary. 2010. "Belief in the Face of Controversy." In *Disagreement*, edited by R. Feldman and T. Warfield, 29-52. Oxford: Oxford University Press.
- McGrath, Sarah. 2007. "Moral Disagreement and Moral Expertise." In *Oxford Studies in Metaethics*, vol. 3, edited by Russ Shafer-Landau, 87-108. Oxford: Oxford University Press.
- Parfit, Derek. 2011. *On What Matters*. Oxford: Oxford University Press.
- Rotondo, Andrew. 2015. "Disagreement and Intellectual Scepticism." *Australasian Journal of Philosophy* 93, no. 2: 251-271.
- Setiya, Kieran. 2012. *Knowing Right from Wrong*. Oxford: Oxford University Press.
- Vavova, Katia. 2014. "Moral Disagreement and Moral Skepticism." *Philosophical Perspectives* 28, no. 1: 302-333.
- Wedgwood, Ralph. 2014. "Moral Disagreement among Philosophers." In *Challenges to Moral and Religious Belief: Disagreement and Evolution*, edited by M. Bergmann and P. Kain, 23-39. Oxford: Oxford University Press.