epistēme

ARTICLE

Echo Chambers and Moral Progress

Tyler Wark ⬥

University of Calgary, Calgary, AB T2N 1N4, Canada
Email: tyler.wark@ucalgary.ca

Abstract

In this paper, I argue that echo chambers pose a problem for moral progress because of their threat to moral reasoning. I argue for two theses about the epistemology of moral progress: (1) the *practical utility thesis*: moral reasoning plays an important role in improving moral judgments, and (2) the *conflictive social reasoning thesis*: the kind of moral reasoning that is important for moral progress involves social reasoning with disputants. Without some conflict, human beings will naturally reason in a biased and otherwise poor manner. Thus, good reasoning must be social so that reasoners who *disagree* can keep each other in check. These two theses explain why echo chambers are a problem for moral progress. I argue that echo chambers isolate individuals from reasoning with those they disagree with. This is because echo chambers act as a mechanism for discrediting those outside the chamber. If this is true, then the members of an echo chamber will only reason with those who agree with them. The result is that echo chamber members won't reason according to the conflictive social reasoning thesis. Reasoning will only reinforce their existing echoed beliefs rather than improve them.

1. Introduction

> Moral progress does exist. Yet it is chancier, slower, less thorough, and less systematic than it might be. One way of conceiving moral philosophy identifies the function of the subject as that of supplying tools for facilitating moral progress. (Kitcher 2021: 14–15)

I take this first claim – that there is moral progress, but it isn't as effective as it could be – to be granted by most of those writing about moral progress today. The second claim – that the function of moral philosophy is to facilitate moral progress – is not as widely accepted. Despite this, I think most philosophers would accept the weaker claim that moral philosophy *can* facilitate moral progress, even if that isn't moral philosophy's 'function'. For this paper, I will accept the first claim and the weaker version of the second. In other words, I intend to help facilitate moral progress be more effective than it currently is.

To do this, I'll try to work out a specific problem that moral progress faces: the problem posed by 'echo chambers'. I won't give a worked-out response to this problem, but as Dewey says, 'a problem well put is half solved' (Dewey 1938: 112). If I can put the problem well, we will be on our way to finding a full solution.

To illustrate the problem, I argue for two theses about moral reasoning: (1) *the practical utility thesis*[1] and (2) *the conflictive social reasoning thesis*. The first thesis says that moral reasoning is an important tool for bringing about at least some kinds of moral progress. The second thesis says that for moral reasoning to be effective, it must be done socially – I call this the *social reasoning thesis*. I then go farther and argue that people reasoning together must have some genuine disagreement – I call this expanded thesis the *conflictive social reasoning thesis*. Reasoning with other people we disagree with ensures that reasoning might change our minds rather than only confirming what we already think. By reasoning with those we disagree with, they hold us accountable and combat our biased reasoning.

The problem with echo chambers, I argue, is that they make conflictive social reasoning difficult – if not impossible – for members of an echo chamber. Members of an echo chamber discredit outsiders, especially those with conflicting views. Those within an echo chamber are inoculated from changing their moral judgments by reasoning. This means that echo chambers echoing regressive views will make members resistant to changing those views. Moral progress requires that we change regressive views. Thus, echo chambers are a problem for moral progress. I conclude by looking at the social media platform Gab and 6 January as a case study of echo chambers that enforce regressive views.

2. Reasoning and moral progress

Ever since the modern notion of progress developed during the Enlightenment, philosophers have thought there was a close connection between progress and reasoning. In the heady days of the Enlightenment, the connection was incredibly strong. Moral progress was just a matter of applying one's reason. When you reason in the right way, moral progress follows. Consider this quote from Kant:

> Reason in a creature is a faculty of widening the rules and purposes of the use of all its powers far beyond natural instinct; it acknowledges no limits to its projects. Reason itself does not work instinctively, but requires trial, practice, and instruction in order gradually to progress from one level of insight to another. Therefore a single man would have to live excessively long in order to learn to make full use of all his natural capacities. Since Nature has set only a short period for his life, she needs a perhaps unreckonable series of generations, each of which passes its own enlightenment to its successor in order finally to bring the seeds of enlightenment to that degree of development in our race which is completely suitable to Nature's purpose. (1784: 13)

There are several ideas in this passage. The first is a conception of progress that involves the development of our capacities through reason. That's what Kant means when he talks about going 'beyond natural instinct'. To put this idea in a broader way, reasoning is how we bring about moral progress – at least one way to do so. In this section I'll

---

[1]Thanks to an anonymous reviewer for suggesting the name for this thesis.

defend a version of this idea. Moral reasoning is an important tool for bringing about at least some kinds of moral progress – what I called the *practical utility thesis*. The second idea is that we develop our capacities through education between generations. We cannot simply reason on our own and completely develop our capacities. This second idea is important, and I think a version of it is important for any account of reasoning in moral progress. This is what I earlier called the *social reasoning thesis*: any effective reasoning undergirding moral progress will involve reasoning as a group. Note that this is a more general thesis than Kant's. The group Kant seems to have in mind is the entire human species, but it might be more specific than that. Kant is also concerned with intergenerational reasoning, but on the more general thesis, it could also include contemporaries and not generations. I will argue for the *practical utility thesis* and the *social reasoning thesis* in turn.

2.1. The practical utility thesis

First, I will argue for the *practical utility thesis*: the thesis that reasoning is important for moral progress. This is because reasoning is often the best way to purposefully and efficiently bring about certain kinds of moral progress.

Here's an argument for the *practical utility thesis*:

1) Humans make moral judgments.
2) These moral judgments are often poor.
3) Moral progress often involves improving moral judgments.
4) Often the best way to purposefully improve moral judgments is through moral reasoning.
C) Moral reasoning is often the best way to bring about moral progress.

Note that by moral judgments, I mean a judgment about whether a behaviour or practice is right or wrong. Moral judgments are improved just when the agent moves from an incorrect or inappropriate judgment to a correct or appropriate one. For example, someone judges that slavery is morally permissible and then changes their judgment to think it is morally impermissible.

I take (1) and (2) to be obvious. There are two reasons to think that (3) is true: first is that improving moral judgments seems like a kind of moral progress (see Buchanan and Powell 2018: 56–57). Second, improving moral judgments can bring about moral progress in other ways. Take the British abolition of slavery as an example of moral progress. This was a case of an immoral practice being ended, and it's plausible that a change in moral judgment was necessary for that change in practice. If the British public and power brokers, by and large, judged slavery to be permissible, then slavery probably wouldn't have ended. There must have been some point where moral judgments about slavery had to be changed. This is not to say that a change in judgment would be sufficient since it's plausible that the British widely regarded slavery as wrong well before abolition (Tam 2020: 82). Despite this, we shouldn't underestimate the importance that they did think that slavery was wrong.

We should believe (4) for several reasons: first, it looks like it would be best if moral progress is something we purposefully bring about rather than leaving it up to blind chance. Insofar as reasoning can help bring about moral progress and do so purposefully, that's good. Second, it doesn't look like there are many other ways to purposefully change our moral judgments. One other candidate is to try to appeal to people's

sentiments. Think of how showing pictures from factory farms of animals in horrible conditions might cause people to change their minds about the morality of factory farming. This sort of strategy is entirely compatible with reasoning and should probably be used alongside it since people aren't always moved by appeals to sentiment. Further, it isn't clear that we should expect our sentiments to guide us towards better moral judgments since they don't aim at improving our judgments. Our sentiments are just sentiments – they just push us towards certain judgments because we have them. By contrast, *good* moral reasoning is aimed at improving our judgments by looking for good reasons to make our judgments. Kumar and May put the point nicely:

> If people are to make better moral judgments and engage in more virtuous behavior, it seems likely that moral reasoning must play a central role. Emotions and other psychological processes, by themselves, are as likely to distort moral judgment as improve it. (Forthcoming: 1)

Appeal to sentiments has their place, but this strategy doesn't undermine the importance of moral reasoning.

Another strategy would be to change our practices.[2] To continue with the factory farming example, an individual might lead by example by adopting a vegan diet which doesn't rely on factory farming. On encountering such a person, a meat eater would have to reckon with the possibility that this alternative practice could be a better one. The consideration of such a possibility may very well lead to a change in judgment. By showing that a practice like factory farming is not inevitable, it opens the possibility that it may be wrong. However, it's unclear how someone moves from exposure to an alternative practice to a change in moral judgment. Most people today have been exposed to vegan lifestyles in some way or another, yet comparatively few have changed their judgment on factory farming.

What connects alternative practices to a change in moral judgment? I see two possibilities: sentiments or moral reasoning. A meat eater might have some guilt or other negative feelings about factory farming and meat consumption, yet ultimately still think it's permissible. When they see the – presumably guilt free – alternative practice they judge that this is the moral way. Alternatively, the meat eater may see the alternative practice and begin to question the reasons for the current practice. If they find the reasons for the current practice lacking, and reasons for the alternative practice satisfactory, they change their judgment. Encountering alternative practices is a prompt for moral reasoning. Thus, the lead by example strategy must be combined with one of the other two strategies. You need to appeal to sentiments such that a subject will see how the alternative practice is morally superior. (This faces all the same problems the sentiments strategy alone faces.) Or else you need to generate reasons for that the alternative practice is more moral to convince the now reasoning subject. You might leave it up to the subject, but – as we'll see – the subject is likely to just confirm their own existing judgment. Moral reasoning is still going to play an important role in changing moral judgments even if there are models for alternative practices. The point here is that just showing an alternative possibility isn't enough, you also need to illustrate why the alternative is better. This will be done either by sentiments or reason.

---

[2]Thanks to an anonymous reviewer for this suggestion.

Therefore, because moral reasoning is often better than the sentiments in improving moral judgments, moral reasoning is often the best tool for the job when we're trying to make moral progress. Despite this, moral reasoning faces problems in bringing about moral a change in moral judgment. I'll now explain two major ones.

The first problem: grounded in empirical evidence from moral psychology, Jonathan Haidt (2001) argues that 'moral reasoning does not cause moral judgment; rather, moral reasoning is usually a post hoc construction, generated after a judgment has been reached' (814). For Haidt, moral judgments are primarily driven by rapid, automatic, intuitive and emotional responses, which we only afterward justify with reasons. He puts forward his 'social intuitionist model' of moral reasoning to better model this. Reasoning plays no role in forming our moral judgments, but only confirming them. Call this the *problem of otiosity*.

The second problem: what moral reasoning faces is bias. For example, we engage in 'satisficing', meaning instead of looking for the best arguments, we look for many good enough arguments. Further, we engage in 'confirmation bias'. We usually only produce those arguments in favour of our own position rather than considering any arguments against it (Mercier 2011). There's also the effect that power and interests have on our moral reasoning. People are generally biased towards their own interests while reasoning, but power exacerbates this. Elizabeth Anderson identifies two pervasive biases among the powerful: arrogance and ignorance. They are arrogant in the sense that they think their power gives them moral authority. They are ignorant in the sense that their power insulates them from the accountability that would inform them of how they're infringing on the interests of others (2014: 8). This bias is particularly antithetical to moral reasoning's power to change social norms for the better since the powerful are in the best position to change norms. Call this the *problem of bias*.

To see how to combat these problems I'll now move on to the *(conflictive) social reasoning thesis*.

2.2. The social reasoning thesis

Kant's account of progress (and enlightenment theories of progress in general) assumes a pure or ideal sort of reasoning; we can reason apart from everyday concerns. Most contemporary authors think this is too hopeful. Humans living in non-ideal circumstances tend to be biased in their reasoning, especially regarding their power and interests (Anderson 2014). For this reason, many contemporary authors understand moral reasoning as being importantly social. I take these authors to be committed to what I called the *social reasoning thesis*: any effective reasoning undergirding moral progress will involve reasoning as a group. Specifically, the *social reasoning thesis* in contemporary guise is often used as an antidote to poor moral reasoning.

I will now argue that the *social reasoning thesis* is true. I argue in this way: The effectiveness of moral reasoning in changing moral judgments faces several problems – *otiosity* and *bias*. These problems, I argue, primarily arise with non-social reasoning. Thus, the solution to these problems is to reason socially. Effective moral reasoning, then, must be social – which is just to say that the *social reasoning thesis* is true.

To begin, note that even Haidt, who I've set up as a critic of moral reasoning, gives it a place in his 'social-intuitionist model of moral judgment'. Post-hoc moral reasoning can be verbalized to other people to change their moral judgments. Despite this, he says: 'reasoned persuasion works not by providing logically compelling arguments but by triggering new affectively valenced intuitions in the listener' (819). Notably, this is

just a prediction of the social-intuitionist model, not something Haidt has illustrated. My point is just that even Haidt can admit that moral reasoning can play a role in changing moral judgments – even if it isn't in the way we might hope. Importantly, moral reasoning is good for changing *other people's* moral judgments, it's a social role for moral reasoning.

Hugo Mercier (2011) offers an alternative model to Haidt's: the 'argumentative theory of reasoning'. This model gives reasoning the job of giving and evaluating arguments in an argumentative context. Importantly, such a context is always social – it is always several people arguing among themselves (135).

When people make a moral judgment, they are almost always prepared to defend it by producing arguments. Even Haidt admits this, but he thinks these arguments are both rationalizations and inevitably biased. That might be fine when we initially give arguments to support our own view. However, when evaluating arguments, it's important to re-emphasize the social component of argumentation. If someone is faced with resistance to their own view, they will produce counterarguments. These will be biased, just like when people initially rationalize their own view, thus biasing the evaluation of the counter view. What we need is someone who holds the opposite view to answer the counterarguments. Mercier thinks that when faced with someone who disagrees with our view and will call us out on our biased reasoning, we can temper our bias. When we temper this bias, we are reliable evaluators.

Elizabeth Anderson (2014, 2015, 2016) offers a pragmatist democratic model of moral reasoning as an answer to concerns of bias and power. This is a dialectical model of moral reasoning that centred around argument and counterargument. The most important aspect of Anderson's model is *who* is part of the dialectic. Reasoning needs to be democratic, meaning:

> All sides to a moral dispute … manage to participate on terms of equality in contention over the principles governing their claims, and do so in ways the others cannot ignore or dismiss but must address in their own terms. (2016: 93–94)

Democratic dialogue isn't enough; the dialogue must also be 'contentious'. Contention is 'practices in which people make claims against others, on behalf of someone's interests' (2015: 32). Moral reasoning is just one version of contention. For Anderson, the status quo can be contested by moral reasoning involving a democratic dialectic of argument and counterargument. Democratic contention, when done well, forces the powerful to acknowledge the moral authority of the oppressed (2014: 8). When this happens, the powerful can no longer merely rationalize their own views. While facing opposition, they must engage in genuine moral reasoning (2015: 39). It also means the powerful cannot simply ignore the interests of others, which counteracts bias towards their own interests. In other words, much like Mercier, Anderson thinks that disagreement is important for moral reasoning.

My point here is that the *social reasoning thesis* acts as a means of making moral reasoning effective. If the *practical utility thesis* is true, then we should also accept the *social reasoning thesis*. Now I don't want to say that any of the models of moral reasoning I've mentioned are true, but I take it that social reasoning of some sort will be required to mitigate the problems of *otiosity* and *bias*. We can see that mere social reasoning is not enough since disagreement is also required. You need someone who will hold you accountable for your poor reasoning to avoid bias. Someone who agrees with you would share your biases and so would probably let poor reasoning slide. We saw

this in both Mercier and Anderson's accounts of moral reasoning. Further, it looks like *mere* disagreement will not be enough. Two subjects might disagree, but one or neither of the disputants will treat the disagreement seriously. If the disputants do not respect their interlocutor as serious, intelligent enough, informed enough, etc., then they will not care that they disagree. An environmental scientist (rightly) will not care about a climate change denier who disagrees with them about climate change – assuming the climate change denier is (at least from the scientist's point of view) ignorant of the relevant science. Such a disagreement would not lead the scientist to change their reasoning (nor should it). Such a disagreement could make the climate change denier reason better, but only if the climate change denier recognizes the scientist as worth taking seriously. In other words, only the scientist can hold the denier accountable for their reasoning since the denier takes the scientist seriously and not the other way around. In moral reasoning, disputants need to take each other seriously so that they can hold each other accountable.

Let's change the *social reasoning thesis* to the *conflictive social reasoning thesis*: any effective reasoning that undergirds moral progress will involve reasoning as a group where there is significant conflict or disagreement of views.

The best way to understand the *conflictive social reasoning* thesis and the *practical utility thesis* is as complimentary. One criticism of the *practical utility thesis* is that moral reasoning cannot play a role in moral progress because moral reasoning is impotent. It cannot change our judgments but only confirm the ones we have. In other words, moral reasoning is really moral rationalizing. The *conflictive social reasoning* illustrates how moral reasoning can have utility despite the criticism of rationalization.

Now I want to argue that if the *practical utility thesis* and the *conflictive social reasoning thesis* are both true, then echo chambers pose a significant problem for moral progress.

3. Echo chambers

Echo chambers are a problem for moral progress when they make bad moral norms and other moral regressions stick. Echo chambers keep homogenous groups, homogenous – and so echo chambers also keep members radical.

One view of echo chambers is that they are just places where people only encounter views that agree with their views, where this involves filtering out other views to ensure things stay that way. Nguyen (2020) – developing on the work of Jamieson and Cappella (2008) – denies that insulation from other views and filtration are sufficient for a community to be an echo chamber. Rather, these practices only constitute *epistemic bubbles*, which he defines as 'a social epistemic structure which has inadequate coverage through a process of exclusion by omission' (143). The reason Nguyen doesn't consider epistemic bubbles to be echo chambers is that epistemic bubbles don't explain 'the apparent resistance to clear evidence found in some groups' (142). The point is that you can break up an epistemic bubble just by bringing alternative points of view to the attention of those in an epistemic bubble. There's nothing about an epistemic bubble that would make someone resistant to countervailing views when they encounter them (145). Those in an epistemic bubble are merely avoiding views which disagree with them.

Epistemic bubbles are not a problem for moral progress, or at least a minor enough problem that we shouldn't worry about them too much. The reason is that epistemic bubbles just generate the *problem of bias*. The issue with epistemic bubbles is that they let your reasoning simply confirm what you already think. As we saw, the solution

to the *problem of bias* was social reasoning with those who you disagree with. The same solution applies to epistemic bubbles. As Nguyen notes, you can break up an epistemic bubble just by bringing alternative points of view to the attention of those in an epistemic bubble. Social reasoning does exactly this. Thus, epistemic bubbles don't pose a special problem for moral progress in the way that, as I'll argue, echo chambers do. In contrast to epistemic bubbles, Nguyen defines an echo chamber as:

> an epistemic community which creates a significant disparity in trust between members and non-members. This disparity is created by excluding non-members through epistemic discrediting, while simultaneously amplifying members' epistemic credentials (146).

Nguyen also identifies several features that echo chambers exhibit: for an echo chamber 'general agreement with some core set of beliefs is a prerequisite for membership, where those core beliefs include beliefs that support that disparity in trust' (146). Another feature of echo chambers is that they act as a 'disagreement-reinforcement mechanism'. This is when 'members can be brought to hold a set of beliefs such that the existence and expression of contrary beliefs reinforces the original set of beliefs and the discrediting story' (147). Roughly, what happens is that the discrediting story predicts that those outside the echo chamber will contest the views of those in the echo chamber. When the prediction comes true, this gives members of the echo chamber more reason to trust the views within the echo chamber. They also continue to ignore disagreement since those views were, from the chamber's point of view, discredited already.

I'll assume this framework for thinking about echo chambers for the rest of the paper.

4. The problem of echo chambers for moral progress

Now I'll consider what problems echo chambers might pose for moral progress, given what has been discussed so far. I first consider several problems which already appear in the literature, but I'll argue that, though these are problems, there's also a novel problem which is more fundamental than the others: echo chambers corrupt the ability of members to engage in social reasoning in the way needed for moral progress. Before I get to that, I'll talk about the other less fundamental problems.

The first thing you might worry about is truth: in an echo chamber, members are inoculated from the truth; members may become resistant to giving up falsehoods echoed by the chambers and believing the truth. Echo chambers, however, don't have to echo falsehoods. Lackey (2018) argues that echo chambers are only bad when the views being echoed are false ones. It might be if you're stuck in an echo chamber where true beliefs are echoed that might be fine. Lackey says that echoing falsehoods is the *only* problem with echo chambers, but I'm not committing to this. I believe there can be other ways for echo chambers to be harmful, e.g. as I'll argue, by impeding moral progress. The point is instead that echoing falsehoods is one problem or harm of echo chambers, but it isn't an intrinsic one. You may worry that Lackey has a different conception of echo chambers from Nguyen, but the point still stands if Ngyuen's echo chambers can echo the truth – which they can.[3] So, there's no reason to think that

---

[3]See Fantl (2021: 1–2) for a similar argument that Ngyuen's echo chambers aren't necessarily bad, but bad when they echo falsehoods.

an intrinsic problem with echo chambers is the beliefs held by those in them. Consider a group of Holocaust historians who only read academic historians on the horrors of the Holocaust. This person also discredits all extra-academic sources as untrustworthy. Let's stipulate that they have the correct views too. This is an echo chamber, but it seems perfectly fine. This suggests that echo chambers are a problem when the chamber's views are false, but not necessarily when they echo truths.

The problem for moral progress is when people make false moral judgments.[4] That's not a problem specific to echo chambers – people make bad moral judgments all the time. What is especially bad about echo chambers is that those bad judgments become entrenched. This entrenchment is a problem, but I think it's more of a symptom than the real issue. The real issue is the explanation of why moral judgments become entrenched. Thus, while echoing falsehoods is a problem of echo chambers, there's a deeper problem.

Relatedly, you might think the 'inoculation' effect I mentioned – that members of an echo chamber become resistant to changing the views echoed by the chamber – is itself harmful. You might think this effect makes members of an echo chamber act closed-mindedly. I would maintain that this is not necessarily a harm. I've so far only talked about inoculation from the truth, but inoculation cuts both ways. Echo chambers can also make us resistant to accepting falsehoods. This can be seen in the Holocaust historian case. In this way, echo chambers can actually do some epistemic good (though it's still possible for them to be overall bad). At the very least, this suggests that echo chambers aren't necessarily always harmful in every way. My point in this paper is not to show that echo chambers are always harmful. Instead, I want to explore what kind of harm echo chambers might cause relative to moral progress.

One worry you might have about the above analysis is that Nguyen sometimes seems to imply that echo chambers are intrinsically bad, such as when he calls them 'perversions' of good epistemic practices (2020: 148). He even calls them 'malicious' (147 note 5) since 'the most plausible explanation for the particular features of echo chambers is something … malicious' (149). Thus, you might take issue with putting him into conversation with Lackey, who takes a neutral view. However, Nguyen doesn't offer an argument that echo chambers are necessarily 'malicious', only that they can be used maliciously. He even says that echo chambers can form unintentionally (149). Further, I see nothing in Nguyen's definition that would entail that echo chambers are intrinsically bad. Thus, I am happy to accept that echo chambers are intrinsically neutral but sometimes have contingent harms (such harms are what this essay is about).

A second possible worry about echo chambers is the lack of diversity in opinion within an echo chamber. As we saw, being part of an echo chamber means sharing particular views with the rest of the chamber. You might think that engaging with diverse opinions is good, and insofar as echo chambers are antithetical to this, they're problematic. However, a lack of diversity is not itself antithetical to moral progress. Many scholars think convergence on certain views is moral progress (Huemer 2015). The ideal end of moral progress means everyone agrees on the 'correct' moral judgments. Any disagreement could lead to moral regression. For example, we don't need a diversity of views on whether antisemitism is bad. Rather, what we need is everyone to agree that it is bad.

---

[4]Many of those in the moral progress literature would not want to commit to the view that there are true or false moral judgments. However, most would say that there are good or bad moral judgments, not all moral judgments are equal. So, instead of true or false you can just plug in whatever term of approval or disapproval you want.

A more sophisticated version of this worry is that echo chambers make moral reasoning less democratic since they exclude relevant voices rather than just other views. The problem with echo chambers is something like what Philip Kitcher (2021) calls 'exclusion'. Like Anderson, Kitcher thinks that moral reasoning needs to be democratic to be effective. Moral reasoning about a moral problem is democratic when it involves a conversation among all relevant voices. Kitcher identifies relevant voices as stakeholders – stakeholders are those who have stakes in a solution to the moral problem. Exclusion occurs when certain people whose voices should be heard are not heard (Kitcher 2021: 33). For example, in Antebellum America, slave owners largely ignored the voices of slaves. Instead, it took white non-slaves to get the wheels in motion for abolition. Notice that this is consistent with what Ngyuen called epistemic bubbles – which, again, he defines as 'a social epistemic structure in which some relevant voices have been excluded through omission' (2020: 142). The exclusion described by Kitcher could be the result of an epistemic bubble, not just echo chambers. Echo chambers pose an even worse and unique problem. The problem isn't just that voices are excluded but that they are discredited. You might think the solution to exclusion is to include the excluded voices by bringing them into the conversation and letting them be heard. However, if they are discredited, then the excluded will be ignored and practically still excluded, no matter what. Echo chambers not only cause exclusion, but they also make exclusion entrenched. Members of an echo chamber aren't just excluding people by ignoring them. Instead, they're actively discrediting the excluded, making meaningful conversation impossible.

I think something like this is correct, but there's a deeper and more general problem with echo chambers than being exclusionary or anti-democratic. This deeper and more general problem arises because of the *conflictive social reasoning thesis*. For moral reasoning to change our moral judgments, reasoners need to come into conflict with those they disagree with while reasoning. Notice that for conflict to reduce bias in the way discussed before, the conflict must force those in the dialectic to reconsider their own position. In other words, it must put the reasoner on the defensive. Their biased reasoning won't work in response to their opponent since they'll point out the bias. The reasoner must come up with better reasons to answer their opponent. If they cannot, then they will, ideally, change their judgment. Again, the reasoner must take their opponent seriously since it's not enough to just face someone who disagrees with you. When you find someone who disagrees with you, you must take it seriously enough to warrant a response. This response should generate reasons which answer the challenge their disagreement poses and aren't vulnerable to counterarguments from the opponent.

When someone is part of an echo chamber, they won't face disagreement from within the echo chamber. At least if we're talking about one of the views that are constitutive of the echo chamber. For example, consider an echo chamber of Nazis who think that the Holocaust was a good thing. Part of that echo chamber is that to gain membership, you must think that the Holocaust was a good thing. It's an echo chamber, so the idea that the Holocaust was a bad thing has been discredited as, say, a Jewish conspiracy which all the academic historians are in on. Now consider a Nazi who is part of this echo chamber and believes that the Holocaust was good. They encounter someone who thinks that the Holocaust was bad. What do they do? They would probably dismiss this person as falling for a Jewish conspiracy. Whatever reasons they produce will just be talking points of the echo chamber. When their disputant puts forward counterarguments, our Nazi would dismiss them as a lost cause to the Jewish conspiracy. In fact, it would just count as more evidence for the Jewish conspiracy. They haven't changed their moral judgment. Instead – if anything – disagreement has made their

judgment stronger. That's a toy example, but assuming that Nguyen is right about the nature of echo chambers, that's roughly what we should expect to happen.

What echo chambers do is take advantage of the fact that much of our reasoning must be social and then corrupts it. Someone might get brought into an echo chamber because they are looking for people to reason with, even if it's people they already agree with. If things go badly, this person will buy into the claim that those outside the echo chamber are unreliable. When this happens, they can no longer engage in moral reasoning with those they disagree with, which, as I've argued, is required for unbiased thinking. Those within an echo chamber can't reliably reason, they can, at best, offer rationalizations.

If the *practical utility thesis* and the *conflictive social reasoning thesis* are true, then echo chambers can pose a real threat to moral progress. Those in echo chambers cannot engage in conflictive social reasoning and thus cannot effectively engage in moral reasoning. As per the *practical utility thesis*, moral reasoning is an important activity for achieving moral progress. Thus, echo chambers can act as an obstacle to moral progress.

One objection to my analysis here is that I have primarily talked about the harms of echo chambers and little about potential benefits. Perhaps the benefits of echo chambers outweigh the harms when it comes to moral reasoning. For example, Cass Sunstein (2007) argues that it can be beneficial for marginalized groups to engage in 'enclave deliberation' – defined as 'deliberation that occurs within more or less insulated groups, in which like-minded people speak mostly to one another' (77). Such deliberations are useful for marginalized groups since it gives members a chance to develop ideas 'that would otherwise be invisible, silenced, or squelched in general debate' (77). Deliberation within an echo chamber would be enclave deliberation, so perhaps some echo chambers aren't so bad? Note it would only be certain echo chambers that would have this benefit. The ones where the ideas discussed would otherwise be silenced. Also, note that not all enclave deliberation happens in echo chambers. Enclave deliberation wouldn't require the discrediting of outsiders, only the exclusion of them from the enclave. Just because marginalized groups engage in enclave deliberation, it doesn't rule out deliberation of other kinds. After engaging in enclave deliberation, an enclave will probably want to take their ideas to the public. If they want to convince others of their conclusions, they must engage in social reasoning to justify their conclusions to the public. As per the *conflictive social reasoning thesis*, for such reasoning to be effective, it will involve reasoning as a group, where there is significant conflict or disagreement of views. If the enclave is an echo chamber, then this social reasoning is ruled out. This means the reasoning is likely biased, and when presented to outsiders, they'll point out as much. This is what distinguishes cases of moral reasoning from the case of the Holocaust historians mentioned earlier. Those engaged in moral reasoning towards the end of moral progress at some point would need to change the minds of others, whereas Holocaust historians can do their work without ever engaging a Holocaust denier. My point is twofold: first, you can gain the benefits of enclave deliberation without forming an echo chamber, and second, enclave deliberation in an echo chamber would be ineffective in producing moral progress. Thus, even if echo chambers do act as a place for marginalized groups to float new ideas, at least when it comes to moral progress, this doesn't compensate for biased moral reasoning.

Another way of understanding this point is through the concept of epistemic friction. Epistemic friction is the resistance our beliefs and reasoning run up against, things which unsettle or constrain our thinking. Disagreement – at least when we take the

disputant seriously – is an example of epistemic friction. The *conflictive social reasoning* relies on the production of epistemic friction to constrain reasoning. Echo chambers can be understood as a way of dealing with and resolving epistemic friction. By reducing the amount of friction reasoning faces, echo chambers reduce its effectiveness. Things are more complicated, however. José Medina (2013) distinguishes beneficial from detrimental epistemic friction. Epistemic friction is beneficial when it 'forc(es) one to be self-critical, to compare and contrast one's beliefs, to meet justificatory demands, to recognize cognitive gaps, and so on' (50). This is roughly what the *conflictive social reasoning* aims to do. By contrast, epistemic friction is detrimental when 'censoring, silencing, or inhibiting the formation of beliefs, the articulations of doubts, the formulation of questions and lines of inquiry, and so on' (50). Echo chambers may be beneficial to marginalized groups who would be otherwise silenced, to continue with our example from before, because they can disarm the epistemic resistance which silences them. Still, while this would be beneficial for articulating views and exploring beliefs, it isn't good for moral reasoning. The attempt to win the minds of others will have to take place outside of the echo chamber if it's to be effective. Even if an echo chamber is beneficial in some ways, to take the ideas beyond the in-group the echo chamber would need to be shed and friction faced.

Another worry is that you might take the above discussion to assume moral cognitivism: specifically, I interpret moral judgments as moral beliefs. Moral non-cognitivism takes moral judgments to be some non-cognitive mental state and moral statements to express a non-cognitive attitude.[5] To some extent, this is true. This paper has been written from a cognitivist point of view and my framing reflects that. Despite this I take it as an open question whether the argument – or at least the core of it – is compatible with a version of non-cognitivism. A non-cognitivist would just have to accept these claims: (1) moral progress can occur via the changing of moral judgments; (2) moral reasoning works in the way I have described; and (3) there can be echo chambers involving non-cognitive moral judgments.

Claim (1) seems difficult to deny and I believe most non-cognitivists would be willing to grant it. Claim (2) is perhaps less clear. I will say I don't believe that the account of moral reasoning outlined above relies on cognitivism. It is only a description of how people actually do reason. A non-cognitivist will either be able to account for this description in non-cognitivist terms or they can't. If they can, they can accept (2). If they can't, then my view entails their non-cognitivism is false. Since my view of moral reasoning doesn't rely on cognitivism as an assumption, the burden to refute my description of moral reasoning is on any non-cognitivist who wishes to object on grounds of non-cognitivism. The worry for claim (3) is that Ngyuen defines echo chambers in terms of beliefs: 'echo chambers are such that general agreement with some core set of beliefs is a prerequisite for membership, where those core beliefs include beliefs that support that disparity in trust' (2020: 146). I have made it sound like the echo chambers I have in mind are ones constituted by moral beliefs. Yet a non-cognitivist could identify other beliefs relevant to shared non-cognitive attitudes which constitute echo chambers. Consider a non-cognitivist moral judgment such as the attitude 'I approve of slavery'. This approval will be associated with several other beliefs potentially including that slavery is the most efficient economic system, slaves are better suited to

---

[5]Note that some non-cognitivists only accept one of these conjuncts. If a non-cognitivist only accepts the second claim about moral semantics, then their view would be more or less compatible with everything I have said. The problem is that I take the changing of moral judgments to be the changing of moral beliefs.

doing manual labour, etc. Importantly, these would include some beliefs which 'support disparity in trust'. Thus, an echo chamber where members share an approval attitude will also share several beliefs. Thus, I take it to be possible for a non-cognitivist to accept the core of my argument. For brevity's sake I haven't included much detail and I leave it up to the non-cognitivist reader to decide whether my arguments are compatible with their view.

One might further object to my view by arguing that this problem is too abstract to be of concern. Even if echo chambers can be a problem for moral progress, this isn't something that really happens. To soothe this worry, I'll end my paper with a case study.

5. Gab and 6 January

The following argument will maintain that echo chambers have played a causal role in the 6 January storming of the U.S. Capital building. I take this to be a case of moral regression and political violence. It was far-right groups on Gab – the social media platform – which did much of the planning for 6 January. I then argue that these communities should be thought of as echo chambers. The upshot of this is that echo chambers have played a role in a case of moral regression. You might hope that we could combat these echo chambers by use of moral reasoning. My analysis shows that this is unlikely to do anything. Worse my analysis predicts that these groups – having already engaged in political violence in the name of their moral views – will continue to hold their views. My analysis, therefore, does not just tell us something in the abstract but can illuminate a concrete case.

On 6 January 2021, a group of former president Donald Trump supporters – and a mish-mash of various far-right groups – gathered in Washington D.C. to protest the results of the 2020 presidential election, which Trump had lost to Joe Biden. The protesters claimed the election was unfairly won; their stated goal was to repair what they took to be an injustice. The protest escalated into a riot when some protesters breached the security barricades and stormed the U.S. Capitol building, interrupting a joint session of Congress that was certifying the electoral college results. During the storming, some of the protesters engaged in acts of violence, including assaulting police officers, damaging property and looting. Law enforcement officers responded with force, including tear gas and pepper spray, and eventually cleared the building. Several people were injured and five died.

Many would consider such an attack a case of – or at least a sign of – moral regression. A group of people were trying to overturn a just democratic election to unjustly install a president they thought better reflected their values.

Much of the planning for 6 January took place on Gab – a far-right social media platform. Thus, Gab played a significant role in causing moral regression. I'll now argue that Gab is also an echo chamber in the sense that I've used throughout this paper. The point being that if Gab is an echo chamber and has played significant causal role in a case of moral regression, then an echo chamber has played a role in a case of moral regression. Thus, my analysis of echo chambers is illuminating for a real case of moral regression.

If Gab is an echo chamber, we should expect three things: (1) moral homogeneity of members, (2) discrediting of outside sources and (3) conflicting views are used to support the views within the echo chamber.

While there might not be a single moral vision on Gab, there are certainly clusters of users who do share homogenous values. This can be seen empirically: Atari *et al.* (2022)

found clusters of users with homogenous moral views (see also Lima *et al.* 2018). It's these clusters of moral views that I'll suggest are echo chambers.

Discrediting the 'mainstream media' is a mainstay of the far-right playbook (Fawzi 2019; Freelon *et al.* 2020; Haller and Holt 2019) – this strategy continues to be used even on Gab (Peucker and Fisher 2022). Peucker and Fisher not only found users discrediting mainstream media but that far-right users would also use articles from the mainstream media to forward their own views: 'any news – favourable, neutral or critical – seems to be good news for some of these far-right groups' (2022: 355). Users on Gab will use critical news articles as a means of affirming their own views, for example:

> In October 2020 … a Gab user posted a partisan feature article from The Guardian about Trump's 'extremist rhetoric' and his 'refusal to condemn white supremacy' and more specifically the actions of US far-right groups such as the Proud Boys. The post rejected the article's critical stance on Trump and expresses an ideologically oppositional views of white victimhood, calling the media 'anti-white corporate parasites'. (Peucker and Fisher 2022: 365–66)

This seems like a case of a disagreement-reinforcement mechanism in action, which is what we would expect in an echo chamber.

On Gab we find clusters of users with homogenous views, we find discrediting of outside sources, and conflicting sources are used to reinforce views. Nguyen's account of echo chambers predicts all three of these things would occur in echo chambers, thus it's plausible that there are echo chambers on Gab. These echo chambers on Gab have played a pivotal role in a case of moral regression. Hence, we have a real case where an echo chamber has led to moral regression. Worse, if there are echo chambers on Gab, then we should expect these radicalized groups to continue to hold their regressive views which – as seen on 6 January – they have tried to violently enforce them on others. How to overcome these sorts of echo chambers is unclear, but hopefully, in describing the problem, we can begin to search for a solution.

References

Anderson E. (2014). *Social Movements, Experiments in Living, and Moral Progress: Case Studies from Britain's Abolition of Slavery*. Lawrence: University of Kansas, Department of Philosophy.

Anderson E. (2015). 'Moral Bias and Corrective Practices: A Pragmatist Perspective.' *Proceedings and Addresses of the American Philosophical Association* **89**, 21–47.

Anderson E. (2016). 'The Social Epistemology of Morality: Learning from the Forgotten History of the Abolition of Slavery.' In M.S. Bradley and M. Fricker (eds), *The Epistemic Life of Groups: Essays in the Epistemology of Collectives*, pp. 75–94. New York: Oxford University Press.

Atari M., Davani A.M., Kogon D., Kennedy B., Ani Saxena N., Anderson I. and Dehghani M. (2022). 'Morally Homogeneous Networks and Radicalism.' *Social Psychological and Personality Science* **13**(6), 999–1009.

Buchanan A. and Powell R. (2018). *The Evolution of Moral Progress: A Biocultural Theory*. New York: Oxford University Press.

Dewey J. (1938). *Logic: The Theory of Inquiry*. New York: Henry Holt and Co. Reprinted in *The Collected Works of John Dewey: The Later Works, 1925–1953, volume 12*.

Fantl J. (2021). 'Fake News vs. Echo Chambers.' *Social Epistemology* **35**(6), 645–59.

Fawzi N. (2019). 'Untrustworthy News and the Media as 'Enemy of the People?' How a Populist Worldview Shapes Recipients' Attitudes Toward the Media.' *The International Journal of Press/Politics* **24**(2), 146–64.

Freelon D., Marwick A. and Kreiss D. (2020). 'False Equivalencies: Online Activism from Left to Right.' *Science* **369**(6508), 1197–201.

Haidt J. (2001). 'The Emotional Dog and Its Rational Tail: A Social Intuitionist Approach to Moral Judgment.' *Psychological Review* **108**(4), 814–34.

Haller A. and Holt K. (2019). 'Paradoxical Populism: How Pegida Relates to Mainstream and Alternative Media.' *Information, Communication & Society* **22**(12), 1665–80.

Huemer M. (2015). 'A Liberal Realist Answer to Debunking Skeptics: The Empirical Case for Realism.' *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* **173**(7), 1983–2010.

Jamieson K.H. and Cappella J.N. (2008). *Echo Chamber: Rush Limbaugh and the Conservative Media Establishment.* New York: Oxford University Press.

Kant I. (1784/1963). *Idea for a Universal History from a Cosmopolitan Point of View.* Lewis White Beck (trans.), in *On History.* Indianapolis: The Bobbs-Merrill Co.

Kitcher P. (2021). *Moral Progress.* New York: Oxford University Press.

Kumar V. and May J. (Forthcoming). 'Moral Reasoning and Moral Progress.' In D. Copp and C. Rosati (eds), *The Oxford Handbook of Metaethics*, pp. 1–17. New York: Oxford University Press.

Lackey J. (2018). 'True Story: Echo Chambers are Not the Problem.' *Morning Consult.* https://morningconsult.com/opinions/true-story-echo-chambers-not-problem/

Lima L., Reis J.C.S., Melo P., Murai F., Araujo L., Vikatos P. and Benevenuto F. (2018). 'Inside the Right-Leaning Echo Chambers: Characterizing Gab, an Unmoderated Social System.' *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM).*

Medina J. (2013). *The Epistemology of Resistance: Gender and Racial Oppression, Epistemic Injustice, and Resistant Imaginations.* New York: Oxford University Press.

Mercier H. (2011). 'What Good is Moral Reasoning?' *Mind & Society* **10**, 131–48.

Nguyen C.T. (2020). 'Echo Chambers and Epistemic Bubbles.' *Episteme* **17**(2), 141–61.

Peucker M. and Fisher T. (2022). 'Mainstream Media Use for Far-Right Mobilisation on the Alt-Tech Online Platform Gab.' *Media, Culture & Society* **45**(2), 354–72.

Sunstein C. (2007). *Republic.com 2.0.* Princeton: Princeton University Press.

Tam A. (2020). 'Why Moral Reasoning Is Insufficient for Moral Progress.' *Journal of Political Philosophy* **28**(1), 73–96.

Tyler Wark is a PhD student at the University of Calgary. His interests include social and individual epistemology. Email: tyler.wark@ucalgary.ca