29

KEEP QUIET INSIDE THE ECHO CHAMBER

The Ethics of Posting on Social Media

Yuval Avnur

People share news items and commentary on social media, usually with people inside their "friend" group, or those that "follow" them. In one way of framing it, there are three kinds of agents in this arrangement: (1) producers, who originally find the information and write up the news item in the first place (journalists); (2) consumers, who consume the news via their social media feeds; and (3) posters, who post, repost, or comment on news items for their "friends" or "followers" to consume. One person might play more than one of these roles. There has been much discussion lately of the effects this arrangement has on consumers, and some discussion of the duties of producers of news. But there has been relatively little discussion of the duties of those in the third group: posters.[1] This chapter makes a first step towards characterizing those duties, focusing on the act of posting news and commentary within a social network. I will argue that, with some exceptions, there is a standing moral reason to refrain from posting links to, or commenting on, news items in your social media accounts. The exceptions are significant, and my conclusion is in some ways a modest claim. There is a reason against posting, and this might often be outweighed by a reason in favour of posting. But we must begin the ethics of posting somewhere, and it is not obvious whether the standing reason against posting news is generally recognized. The problematic sort of posting is still a typical, if not predominant, form of contribution to social networks, so the standing reason defended here is, at least, relevant to our justification for typical posting behaviour.

Here is the argument:

1 It is, all else being equal, morally wrong to contribute to people's cognitive bias on matters of social importance.
2 With some exceptions, posting news and commentary on matters of social importance on social media contributes to people's cognitive bias.

Therefore,

3 It is, all else being equal, morally wrong to post news and commentary on matters of social importance on social media, with some exceptions.

DOI: 10.4324/9781003134749-35

The "all else being equal" is meant to clarify that it is not always, all things considered, wrong to post the relevant material. There are circumstances in which the reasons for posting outweigh this general, standing reason against. Here my point is to establish a general, default reason against, and whether there are stronger reasons in favour will depend on specific circumstances, some of which I discuss below. It is also worth noting that a similar argument might apply to discussing politics with friends in person, though here I am concentrating on online interactions.

The following two sections defend each premise, respectively. The third section discusses the exceptions and addresses objections.

29.1 The First Premise

Why think that it is morally wrong to contribute to people's cognitive bias on matters of social importance?

To begin, I understand cognitive bias to be an illicit influence or factor in belief-formation, which results in the belief-forming process falling short of some standards of rationality. To the extent that one's belief-forming process is cognitively biased, the belief is irrational. Everyone should accept that *some* processes result in irrational beliefs, and that these processes are individuated at least *in part* by elements of one's belief-forming process.

To the extent that one's beliefs on matters of social importance are cognitively biased, those beliefs are irrational. And, worse, they are irrational in predictable ways, often in line with one's prejudices, entrenching one's views and making them less responsive to new evidence. There are two relevant dimensions to the act of (knowingly) contributing to someone's irrationality in this way: the well-being of the other, and your attitude towards the other. On both dimensions, it is wrong to contribute to another's irrationality. This is especially so given that the people in question, those who are in your "friend" network on social media, are your friends, acquaintances, and family to whom you owe more than to randomly selected people, according to commonly held principles of partiality.

29.1.1 The Consequences for Others' Well-being

When you knowingly contribute to someone's irrationality, especially on matters of importance, you are corrupting them, predictably making them worse off than they otherwise would be. They become worse off, first, in the sense that their ability to pursue their interests rationally is degraded (whatever their interests are). This is because it is instrumentally valuable to be rational, because accurate beliefs are required for effective practical reasoning; when one becomes less rational, one becomes less effective in one's practical reasoning. For example, one's practical interests in a pandemic are, presumably, in large part to keep oneself and one's family healthy. But irrational beliefs about how to stave off infection make pursuit of that interest less effective, and sometimes even counterproductive.[2]

Second, insofar as there are objective standards of rationality, and insofar as failing to meet them is a failure, it is intrinsically better to be rational, or to meet those standards. A cognitive bias is a failure, to some extent, to meet such standards. Shared or objective standards of rationality aside, given that the aim of one's cognition is to believe truth and avoid falsehood – an aim that is plausibly constitutive of belief – cognitive bias, which introduces aims external to these, obstructs the route cognition takes towards its aims. Even

if one doesn't care whether one is rational in one's believing, one *ought* to care, and it is in one's own interest not to fall short of the standards for belief. It is controversial whether this basic interest in rationality is of moral concern, or whether it is strictly prudential. But even if it is merely prudential, it seems plausible that, in general, one is better off if one can satisfy such a prudential concern, and so, all else equal, thwarting one's effort to satisfy it (or, rather, the effort one should be making) is detrimental to one's wellbeing, and is, thus, of moral relevance. Compare: your friend has a prudential interest in paying off their house, or at least they *should*. All else equal, you have some moral reason not to thwart their efforts to pay off their house.

There is a third relevant point that goes beyond an individual's wellbeing. The foreseeable consequences of contributing to others' cognitive bias, such as decreased well-being, can affect the entire group, or society. The effects on society of widespread cognitive bias, especially as they manifest in acceptance of conspiracy theories and susceptibility to propaganda, are widely considered to be bad. This is the misinformation crisis that many commentators fear is a threat to our society. Here is one way to argue for this: Cognitive bias leads to irrational, or less rational, responses to available evidence. This applies to evidence about: political positions, facts about happenings within society (and in government), and candidates for elected office. So, to the extent that cognitive bias is more prevalent in a democratic society such as ours, political actions and electoral results are more likely to be based on irrational beliefs. But, perhaps worse, cognitive bias can be exploited, because it results in predictable tendencies to interpret information, evaluate sources, and react to news (as I explain in more detail later). A society plagued with cognitive bias is one that is ripe for manipulation.

Here, then, is the upshot for the morality of contributing to others' cognitive bias on matters of social importance, from a consequence-focused perspective: Since a person's well-being, and a democratic society's function, is degraded by cognitive bias, it is, all else being equal, wrong to contribute to that cognitive bias.

29.1.2 Respect for Others

The wrongness of contributing to the cognitive bias of others can also be appreciated when considering your attitude towards those others. Each person has a duty to respect others, and in particular to respect the autonomy and the exercise of the rational capacities of others. I take this to be, roughly, one main current in Kantian ethical theory. When one knowingly contributes to the cognitive bias of others, one is shirking this duty, and manifesting a disrespect for the others' autonomy and the exercise of their rational capacity. Cognitive bias is an influence on one's reasoning that is external to the pursuit of truth (and avoidance of error). This influence makes it more difficult, in other words, for agents to successfully think for themselves. Thus, to amplify one's cognitive bias is to intrude, negatively, on their cognitive autonomy. More directly, intervening in one's rational process *in a way that foreseeably degrades it* is patently disrespectful of that process.

To illustrate both this consideration of respect and the previous one concerning consequences, consider an analogy. Being an addict is not good for a person, in part because an addict has a misaligned will structure. An addict willingly does something that is bad for the addict, or else does so unwillingly.

Consider, first, someone of the second variety: their addiction consists partly in their wanting to, but being unable to, stop. Such a person does not endorse their addictive

behaviour. Despite their will, they persist in the addictive behaviour.[3] Now, suppose you can choose to dispose of some substance (be it a drug, or cash) in such a way that the addict has access to it. And suppose that this, predictably, will contribute to the addiction. You arguably have a duty not to dispose of the substance in this way. To dispose of it in this way is to disregard the well-being of another person (and will predictably harm that person), but in a particular way: you are contributing to the degradation of their will and autonomy, insofar as the addictive behaviour is a failure of the will and autonomy. It is wrong to contribute to this other person's addiction (all else being equal).

Consider now, instead, the addict of the first variety: They do not want to stop, so they do not display a weakness of the will. Still, given that the addiction is bad for them, and thwarts lots of their other central interests, they *should* want to stop. So, they still suffer from a problem of the will: they want the wrong things. Say, at least, that this is your considered opinion of the situation. You should then refrain from contributing to their addiction, since to do so would be to contribute to something that you believe contributes to their misguided will. It is easy to imagine a case in which one's will for an addictive substance is incompatible with, for example, one's meeting one's basic commitments to oneself and others, so that refraining from assisting one's addiction is a matter of respecting their autonomy and capability to exercise it.[4] Note that I am not suggesting here that you should refrain from helping them pursue their addiction if they sincerely ask, or that you should break some commitment to them in order to "save them from themselves." Rather, I'm making the minimal point that, all else being equal, you have some reason not to contribute.

In both versions of the addiction case, then, you have some reason not to contribute. Structurally, this is the same as the case of contributing to the irrationality of others because irrationality is an obstacle to autonomous and effective thinking. Some might think that I have strayed too close to a sort of paternalism, wherein you are expected to shape another's trajectory according to your own vision of what is "good" or what they "should" want. But I am not arguing here that you have a standing moral reason to go after people online who you think are reasoning in a biased way. Rather, I am making the more minimal point, that you have some reason not to contribute to others' bias, and that this reason can be discerned within the framework of respect for others' rationality.

The above constitutes an initial case for the first premise. There are objections, including the idea that contributing to cognitive bias can be justified if one does so in a way that predictably leads the other to possess true and useful beliefs, or other positive outcomes. I will take this up below, but it will be helpful to first consider the second premise.

29.2 The Second Premise

Why think that, with some exceptions, posting news and commentary on matters of social importance on social media contributes to people's cognitive bias?

In Avnur (2020) I argued in detail that using social media platforms to get news, in the usual way, typically amplifies and aggravates the news consumers' cognitive biases. In particular, one's motivated reasoning, or the influence of one's interests, fears, and desires on one's belief-formation predictably becomes greater. In summary, here is how this works. One typically shares common interests, fears, and desires – at least on important social issues – with most of one's social network (and, on some platforms, one is more likely to see posts from elsewhere that are aligned with one's interests). This shared set of common interests determines a "target" conclusion on important social issues. For example, a liberal

Democrat in 2019 will tend to prefer the conclusion that Trump committed an impeachable offence, because this would confirm the correctness of their commitments and the desirability of their pre-existing preferences. Because on social media one is more likely to see news items, or commentary on news items, from one's network, one is more likely to see material posted by those who share target conclusions. This material has predictable features: evidence for the targets (Trump committed an impeachable offence) is usually taken at face value, emphasized, or given a charitable reading, and evidence for unwelcome targets (Trump did not commit an impeachable offence) is subjected to withering scrutiny and its sources undermined. This is, in part, because journalists and other producers of content are themselves usually engaging in motivated reasoning, and in part because only such articles will be selected by members of the network of friends for re-posting.[5] As a result, routes from the available public evidence to one's preferred targets are encouraged, and routes to unwelcome targets are discouraged. The basic process of motivated reasoning is a well-documented phenomenon;[6] what I argued in the previous work is that this effect is predictably magnified by the way we gather news on social media.

The process summarized above need not result in a social media user's insulation from inconvenient evidence, as it often claimed.[7] Insulation from unwelcome evidence is usefully distinguished from illicit influence on how evidence is processed or weighted.[8] I do not claim that social media tends to insulate the consumer from unwanted information – in fact, several recent studies suggest it does not.[9] Rather, I claim that the way in which one is encouraged to process the evidence – what evidence to downplay, what evidence to make much of, and what to take the evidence itself to indicate – is predictably biased by news feeds whose main inputs are those with predictably similar dispositions or political orientations. Social media echo chambers amplify a bias of evidence processing, not necessarily evidence exposure.

So far, I have invoked and briefly described the case for the view that getting one's news on social media typically contributes to one's cognitive bias on important social issues. It follows that when one posts material of the relevant sort, one is predictably contributing to the cognitive bias of the news consumers in one's network.

Given the variety of material one can post on social networks and the many reasons one might have for posting it, the idea that one shouldn't post needs some qualification. Some kinds of posting either do not raise the ethical problem described above, or else they may in some cases outweigh the ethical problem such that, all things considered, it is permissible to post. This is the subject of the next section.

29.3 The Exceptions, Some Non-exceptions, and Objections

Let's begin with the exceptions mentioned in the second premise. Clearly, posting jokes, pictures of friends and family, and generally catching up, though these are "socially important" in some sense, are exceptions here.

There are also cases in which your audience, or those who are expected to view your post (your friends, or, on some platforms, the members of the public you expect to stumble upon your post), differ significantly with respect to your interests in a way that is relevant to the post. For example, suppose you are on a social network where most of your friends are professional colleagues, and you differ politically from most of them. Posting news, editorials, or opinions about these would not in this case raise the problem above, because it would not foreseeably contribute to the audience's biases.

In some cases, you might intend to report new information, or inform your network of basic, first-order facts, rather than interpretations of this or suggestions of what conclusion to draw. This kind of posting became more common during the COVID-19 crisis, with people posting the latest information about case numbers, for example. There was usually some neutral information contained in such material, but it was often connected with some loaded conclusion: the latest outbreak was especially bad, or it showed that we have nothing to worry about and should go back to normal living, for example. Sometimes, such discussions are helpful. But other times they can contribute to a bias oriented towards the shared target conclusion of the discussants, as summarized above. Usually, it can be expected that others can soon learn the neutral information in some form anyway, since one rarely has access to information on issues of social importance that others cannot, and will not, access. Very little is likely lost if one refrains from playing "the reporter" and sharing new information with friends. So, the thrill of being the first to report some information to one's friends can hardly outweigh the potential harm of contributing to their bias. But, in cases in which you truly have unique access – for example if you are the only one in your network who has the expertise to understand scientific literature relevant to the topic – reporting new evidence might reasonably be considered to be non-bias-enhancing. Note, however, that reporting information in an entirely unbiased way, though possible, is usually nearly impossible, and very difficult to self-monitor.[10] There is, thus, usually at least *some* reason to refrain, given the aim of not contributing to others' bias.

Other situations present, not exceptions, but ways in which the standing reason identified by the argument might be outweighed. That is, sometimes there is a reason to post that outweighs the all-things-being-equal wrong of posting, since it makes things not equal. For example, consider the expressive potential of posting. Sometimes people post things to express frustrations, fears, and hopes, perhaps hoping to find support or commiseration. I grant that, insofar as such things are valuable, they constitute some reason to post. This must then weigh against the standing reason not to post. In most realistic cases, it seems implausible that the reason against posting will be outweighed in this way, because the good of expressing oneself in this way intuitively pales in comparison to the bad of contributing to the cognitive bias of others. Recall, for example, the potential harm to our political system of large scale, enhanced cognitive bias. Can the good of expressing one's attitudes outweigh this? Perhaps, if social media is one's *only* social outlet, but I take it that this is a relatively rare case (though perhaps less rare during, for example, the lockdowns associated with the COVID-19 crisis). Perhaps there are some isolated individuals, perhaps those surrounded by intolerant political adversaries in their physical lives, for whom the opportunity to vent and commiserate is too valuable to forego.

A more common example is posting on wedge issues. Sometimes in-groups, such as your friend network on a social media platform, are not entirely synchronized in their positions on politically important matters. Almost everyone's friends or followers on social media will contain *some* diversity of positions. But here I have in mind particular topics that serve as "wedge issues", as I will call them, that divide one's audience into two or more substantial or sizable groups. In my own recent experience, vaccine mandates and school closures have been a wedge issue, with some in my network being strongly in favour, and some strongly opposed. In this case, if you have an agenda on the wedge issue, you may well have a legitimate interest in swaying the opinion of others on the matter. And though your posts might contribute to some cognitive bias on the side of those who agree with you, the value of presenting your point of view, through selective posting of news, to the others in your network may well outweigh the harm.

Finally, one potentially important use for social media is recruitment and coordination for political activity, including organizing for legally sanctioned political activity such as elections, and revolutionary action. These are complicated cases, in which the details will matter a lot. For one example: you might post some articles in the hopes that you energize those who already agree – this is different from hoping to convert or win the hearts of those in the opposition on a wedge issue. This energizing might be aimed at turning out votes during an election, for example, or gearing up for a protest. In such a case, the good of attempting to organize *this group of people*, who already largely agree, should be weighed against the harm of contributing to their further, biased entrenchment into this position. In many cases, your audience is already likely to do what you are, apparently, trying to get them to do, so this is not the typical case, and usually you don't have more to gain by posting than you do to lose, morally, by contributing to their bias. It is a good question, and one that I cannot take up here, to what extent this applies to rallies and organized meetings of like-minded people. We should, at least, take seriously the question of whether such "riling up" of those who agree is all things considered good, given that it will predictably entrench everyone's positions.

Finally, consider a couple of objections:

Objection 1: *Perhaps the above holds for some people, but I possess the truth and understand how the evidence supports it in a completely clear and unbiased way. So, when I contribute to someone's cognitive bias, I do so in a way that tilts their reasoning towards the truth, which isn't a bad thing.*

I have two, independent replies.

First, increasing the cognitive bias of others, thereby making their beliefs less rational, for the purpose of making them believe true things is an instance of manipulation and paternalism. You are, in that case, failing to respect their autonomy as thinkers in their own right. Granted, some are comfortable with some amount of paternalism. But, to the extent that you are "doing others' thinking for them", by facilitating their arriving at the "truth" through an irrational route, this seems plausibly to land on the bad end of the paternalism spectrum. Furthermore, your opponents are likely just as confident about their view as you feel about your own, and, presumably, you should hold that their manipulation of their audience is permissible, as yours is. Or, at least, you are committed to the view that, as far as *they* can tell, they are doing something permissible. So, the foreseeable consequence of this policy is the weaponization of information, where people on either side of a debate attempt to manipulate their respective audience more effectively than the other. (Some would argue this has already happened in our own society, and that this is a sad state of affairs.) The result is widespread irrationality, and the success of whichever side has greater skill or resources to manipulate and corrupt their respective audience.

Second, as confident as you are that you are not biased, it is risky to stake actions that potentially harm others on it. Since cognitive bias works under the surface, so that the biased individuals do not know, or believe, that they is biased, it will seem perfectly reasonable for everyone on every side of an issue to believe that this exception applies to them. Thus, if it seems morally wrong to you that your opponents attempt to manipulate public opinion, and if it seems to you that the *apparent* reasonableness of their attempt does not exculpate them, this should give you pause: from the inside, their position looks just like yours does to you. If you are biased, it would feel just like it does to you now.

Note that, given that one's audience in social media typically largely agrees with one (or else one of the exceptions above applies), the objection here cannot be that the matter is of such great importance that people need to become convinced of the truth. What one does when one further entrenches one's audience's views – when one "preaches to the choir" – is to harden their views and make them potentially less responsive to new evidence and rational debate. But one does not typically convert anyone in this way, because one's audience already largely agrees. This is potentially especially bad for them because, on the off chance that you're all wrong, they won't be as responsive to new evidence as they were before you contributed. As Mill rightly noted, even when we all agree, we all have interest in considering objections and dissenting views, in order to make our position stronger. But contributing to cognitive bias does the opposite.

You're likely not changing any minds, in most cases, by posting news on social media. Rather, you are entrenching, and (further) biasing, those who agree with you. Except for the exceptions above, all else being equal, this is a bad thing to do, and thinking you are right (as one always does, after all) does not make things less equal.

Objection 2: *Refraining from posting is futile. Everyone keeps posting, and often posting more flagrantly biased material than anything I would post. So, does it really help if I stop posting?*

This is a good objection, and it is similar to other collective action problems. Consider, for example, the classic rebuttal to the consequentialist case for vegetarianism: it won't lead to any animal's suffering if I buy this one piece of already-slain flesh at the store, so what reason have I to refrain? Replies to this abound.[11] Their basic gist is to appeal to one's responsibility for taking part in a group's action. Although your individual action in one instance may not make much difference, you thereby participate, and are partially responsible for, the actions of a group – the posters. But, further, the above arguments were not solely consequentialist. We have a duty to each other regardless of the consequences, at least according to some reasonable ethical theories. The question whether such theories are correct is of course a great one, and not one that I can settle here. But I take it that if the reason against posting is on a par with the reason against purchasing meat, according to standard vegetarian arguments, I have succeeded at least in putting a starting position in the ethics of posting on the philosophical map.

29.4 Conclusion

Let us sum up. We have seen that consuming news and opinions on social media predictably leads to motivated reasoning and cognitive bias. So, posting news and opinions predictably contributes to others' cognitive bias. There is a standing reason (all else equal) not to contribute to the cognitive bias of others, due to both the foreseeable consequences to their wellbeing, and to one's duty to respect their autonomy and rationality. So, there is a standing moral reason not to post news and opinions on social media. It is often said that these days we get our news and opinions in an "echo chamber", meaning that we hear too much of our own perspective reflected back at us. Our friend groups and the ways in which news and opinions appear online act as mirrors, and this amplifies our biases. But we are not merely *passive* consumers. We can choose not to consume our own news in this way, we can actively seek out diverse sources. I have not argued here for that policy, though I do think it

is wise. We are also active in that we can choose to post, or not to post, onto these news feeds. When we post, we add to the cacophony of like-minded opinions and perspectives, because our audience is typically like-minded people. So we should typically choose to stay silent and not add to the cacophony. The online world doesn't need any more echoes.

Notes

1 One recent exception is Marin (2021), though my account here differs substantially from hers.
2 Some may have interests that make biased reasoning attractive. For example, one might want to have one's views confirmed, and one might even want to be irrational. In that case, contributing to their bias helps them to achieve *those* interests. But those desires are, arguably, not good for them, and not in their interest in any objective sense. I discuss cases like thisin Section 29.3.
3 See Levy (2014) for an illuminating discussion of addiction and belief.
4 See Figdor (forthcoming) for an account that takes this to be more than just an analogy.
5 See Worsnip (2019).
6 e.g. Kunda (1990).
7 e.g. Sunstein (2001, 2009) and Pariser (2011).
8 Nguyen (2020) helpfully marks the distinction by calling insulation from evidence an "epistemic bubble" and biased (or otherwise compromised) weighting of evidence an "echo chamber", though the latter is characterized solely in terms of asymmetric trust in sources, whereas I recognize other ways in which motivated reasoning, or influence on the processing of evidence, can occur.
9 Guess et al. (2018) and Dubois & Blank (2018).
10 See Worsnip ibid. Section 29.2 and Anderson (2011).
11 One classic is Singer (1998).

References

Anderson, Elizabeth (2011). Democracy, Public Policy, and Lay Assessments of Scientific Testimony. *Episteme* 8 (2):144–164.

Avnur, Yuval (2020). What's Wrong with the Online Echo Chamber: A motivated Reasoning Account. *Journal of Applied Philosophy* 37 (4):578–593.

Dubois, Elizabeth & Blank, Grant (2018). The Echo Chamber Is Overstated: The Moderating Effect of Political Interest and Diverse Media. *Information, Communication & Society* 21(5):729–745.

Figdor, Carrie (Forthcoming). Doxastic Addiction and Effective Interventions. In Jennifer Lackey & Aidan McGlynn (eds.), *The Oxford Handbook of Social Epistemology*. Oxford University Press.

Guess, A. M., Lyons, B., Nyhan, B, & Reifler, J. (2018). 'Avoiding the echo chamber about echo chambers: Why selective exposure to like-minded political news is less prevalent than you think (Knight foundation report)': Retrieved from https://kf-site-production.s3.amazonaws.com/media_elements/files/000/000/133/original/Topos_KF_White-Paper_Nyhan_V1.pdf

Kunda, Ziva (1990). The Case for Motivated Reasoning. *Psychological Bulletin* 108: 480–498.

Levy, Neil (2014). Addiction as a Disorder of Belief. *Biology and Philosophy* 29 (3):337–355.

Marin, Lavinia (2021). Sharing (mis) Information on Social Networking Sites. An Exploration of the Norms for Distributing Content Authored by Others. *Ethics and Information Technology* 23 (3):363–372.

Nguyen, C. Thi (2020). Echo Chambers and Epistemic Bubbles. *Episteme* 17 (2):141–161.

Pariser, Eli (2011). *The Filter Bubble: What the Internet Is Hiding from You*. Penguin.

Singer, Peter (1998) A Vegetarian Philosophy. In Sian Griffiths & Jennifer Wallace (eds.), *Consuming Passions*. Manchester. pp. 66–72.

Sunstein, Cass (2001) *Echo Chambers: Bush v. Gore, Impeachment, and Beyond Princeton*. Princeton University Press.

——— (2009). *Going to Extremes: How Like Minds Unite and Divide*. Oxford University Press.

Worsnip, Alex (2019). The Obligation to Diversify One's Sources: Against Epistemic Partisanship in the Consumption of News Media. In Carl Fox & Joe Saunders (eds.), *Media Ethics: Free Speech and the Requirements of Democracy*. Routledge. pp. 240–264.