

Critical Social Epistemology of Social Media and Epistemic Virtues

Lukas Schwengerer - Lukas.Schwengerer@uni-due.de

This is a penultimate draft. Please cite the published (and properly proofed) version at <https://www.tandfonline.com/doi/full/10.1080/02691728.2024.2414446>

Abstract

This paper suggests that virtue epistemology can help decide how to respond to conflicts between different epistemic goals for social media. It is a contribution to critical epistemology of social media insofar as it supplements system-level consideration with insights from individualist epistemology. In particular, whereas the proposal of critical social epistemology of social media by Joshua Habgood-Coote suggests that conflicts between epistemic goals of social media have to be solved by ethical consideration, I suggest that virtue epistemology can also solve at least some of these conflicts fully within the epistemic realm.

Keywords: Critical Social Epistemology of Social Media; Online Intellectual Virtues; Epistemology of Social Media; Internet Epistemology; Virtue Epistemology

1. Introduction

A recent proposal in analyzing and potentially designing social media made by Joshua Habgood-Coote (Forthcoming) puts forward multiple epistemic goals for a social media platform and socio-technological systems more generally. These multiple goals can come into conflict, but any way out of such a conflict is supposed to require tools outside of the epistemic

realm. Even though I broadly agree with the identification of the epistemic goals and their potential conflicts, I suggest that more can be done with the purely epistemic toolset available than Habgood-Coote claims. In particular, I argue that looking at individual agents in a virtue epistemology framework can at least in some cases help us decide what to do when different epistemic goals for a social media system come into conflict. As such, I discuss how system-level approaches to epistemic problems of the Internet can be combined with individual approaches to epistemic conduct in the online environment. This is both a contribution to the new critical social epistemology of social media that Habgood-Coote proposes, and to the debate around how much an individualist, virtue approach can achieve in promoting an epistemically better online environment (cf. Heersmink (2018), Smart & Clowes (2021) Schwengerer (2021b)).

My plan is the following: I start with a short introductory overview of critical social epistemology of social media as proposed by Habgood-Coote and how my aims relate to his proposal. I then explain the virtue epistemic approach to online environments and the notion of online intellectual virtues. Finally, I show how that virtue approach can help us decide how to respond to the conflicts presented by Habgood-Coote. I argue that virtue epistemology can give us clear and practical answers for some of the conflicts, but not all of them.

2. Critical Social Epistemology of Social Media

Critical social epistemology of social media is an approach proposed by Joshua Habgood-Coote (Forthcoming) inspired by critical projects in technology studies, including critical information studies (Noble, 2018), critical code studies (Marino, 2020), and critical disinformation studies (Marwick, Kuo, Cameron, & Weigel, 2021). Moreover, it combines these critical projects with a system approach in social epistemology championed by Goldman (2010). Goldman evaluates

epistemic systems on their epistemic outcomes. For instance, one can evaluate how many true beliefs a system generates as a whole, rather than looking at individual instances of belief formation. Whereas Goldman only accepts one epistemic goal, Habgood-Coote (Forthcoming) chooses three dimensions of evaluations of epistemic systems: their consequences for the epistemic states of individuals, their institutional properties, and their contributions to epistemic justice.

The first dimension of evaluating epistemic systems is to look at the results for the individual agents that are part of these systems. A good epistemic system on this dimension brings about desirable epistemic states in individuals and prevents them from forming undesirable ones.

The second dimension looks at system-level epistemic properties that can come apart from epistemic properties of individual agents. Proponents of this so-called institutionalism work with the idea that the methodological prescriptions for individuals and communities diverge and what makes an individual decision epistemically good, might not do the same on the level of the system or community. This is labelled *the independence thesis* (cf. Habgood-Coote (Forthcoming); Mayo-Wilson, Zollman, and Danks (2011)) and argued for by formal means (e.g. Zollman (2007; 2013); Boyer-Kassem, Mayo-Wilson, and Weisberg (2017)), considerations in philosophy of science (e.g. Longino (1990)), and work in feminist philosophy (e.g. Longino (2022)).

The third dimension concerns problems of epistemic (in)justice in real-life systems and aims to promote design decisions that foster epistemic justice and prevent epistemic injustices. As Habgood-Coote (Forthcoming) notes, work on epistemic justice as an evaluative dimension is extremely diverse. He explicitly draws on ideas from Fricker (2007), Dotson (2011; 2014), Fallis (2003) and Coady (2010) and tries to tie their approaches together by following Anderson

(2012) in thinking of epistemic justice roughly “[...] as a feature of social systems, which obtains when those systems respect the epistemic agency of relevant individuals, distribute epistemic goods in an equitable way” (Habgood-Coote, Forthcoming).

Critical social epistemology is an approach used to understand technological systems (in this case social media) as generating problems when these evaluative dimensions come into conflict. Because these problems exist on a system level, the solutions to those problems are also meant to be on a system level, rather than particular actions or responses done by individuals. The main problem is, as Habgood-Coote argues, that by looking at design on a system level we cannot satisfy all three epistemic dimensions properly because they end up being in conflict. Hence, he suggests that ethical considerations should be used to decide which of the epistemic dimensions should be prioritized. Although we cannot satisfy all three perfectly, we can figure out by considerations beyond the epistemic which of those dimensions is more important and should guide the design of social media, even if that means worse outcomes on other dimensions. Towards the end of his paper, he writes:

How we ought to resolve conflicts between the different epistemic desiderata of social systems is a difficult question. In the absence of a compelling meta-epistemological theory about the fundamental values are in the domain of epistemology, we have a couple of options. We might think that conflict cases are irresolvable dilemmas, or we might appeal to values from outwith epistemology to resolve the conflict cases. I think that our discussion of the social epistemology of social media, and an appreciation of the political valences of technology points towards the second option. The conflict cases we have considered centre around political and ethical conflicts in which the needs and interests of individuals, different groups, and the ethical value of

knowledge-producing institutions need to be weighed up against one another.

(Habgood-Coote, Forthcoming)

I will grant Habgood-Coote that there is no single solution for all these conflicts on offer. Moreover, I grant that he is right that there is no meta-epistemological theory that helps us out. But I am less sure that we have to look for answers outside of epistemology if we want to design the epistemically best social media platforms. I want to suggest that we can do more within epistemology first. It might turn out that ultimately epistemological considerations will not be enough. Perhaps moral reasons or political rationales are required. But there is something more that epistemology can offer – and it comes from a source that plays little role in Habgood-Coote’s approach so far: individual epistemic agents. Critical social epistemology of social media is an approach that looks at the epistemic systems. However, I suggest that it should be enhanced by looking at individual epistemic agents as well. To do so I will present an approach to online epistemic virtues and then look at some of the conflicts Habgood-Coote identifies. I argue that for each of those conflicts the virtue epistemic approach can aid us to solve the conflicts.

3. Virtues Guiding Conflict Choices

Let me begin with the general strategy, before filling in the details for the necessary steps. I start with the conflicts between epistemic goals presented by Habgood-Coote and agree that they all pose questions of what to prioritize. Epistemic goods in individual agents? Desirable institutional epistemic properties? Epistemic justice? Whenever two of these dimensions are in conflict we have to give up on one to some degree to benefit the other dimension. However, Habgood-Coote does not consider whether we can compensate for any loss in such a trade-off. In general, if one has to choose between A and B, but has available a different,

independent way to get B even when choosing A, then the cost of choosing A over B becomes much more palatable. Applied to the epistemic goals that come into conflict, there might be instances in which one has to pick between maximizing informational connections of individuals and desirable institutional epistemic properties. But if there is a way to gain the desirable institutional properties through other means, then choosing to maximize connections comes at little cost. Perhaps the design of a social media platform can only foster one epistemic goal properly, but are there other ways outside of the design approach that can also lead us to the other goal that the social media platform was not designed for? The question then becomes the following: for each of the conflicts between two distinct epistemic goals for our system, can we compensate for giving up on one of them at the design level more easily?

This is, I suggest, where the qualities of individual epistemic agents come in. For many conflicts of the kind that Habgood-Coote envisions, some potential losses can be compensated by the individual agents. For the design of epistemically desirable social media platforms that entails that we design primarily with an aim for those epistemic features in individuals and systems/communities that cannot – or only with great difficulty – be achieved in other ways. Within the design stage, we give up on those epistemic goals that can be achieved outside of the design of social media. To see how this suggestion is meant to work, let me introduce the notion of an *online intellectual virtue*¹.

3.1. Online Intellectual Virtues

Online intellectual virtues are a particular implementation of the virtue epistemology framework to the case of Internet users proposed by Heersmink (2018) and Schwengerer

¹ I use epistemic virtue synonymous to intellectual virtue.

(2021a). Such virtues are understood as *“instances of general intellectual virtues that are applied in an online environment based on relevant background knowledge thereof”* (Schwengerer, 2021a, p. 313). Hence, to understand online intellectual virtues, I need some background of intellectual virtues more generally.

All theories of virtue epistemology emphasize the agent’s role in acquiring knowledge. They are primarily interested in what makes agents good or bad believers and inquirers (Battaly, 2008). Within virtue epistemic frameworks one can distinguish virtue reliabilism from virtue responsibilism. The former identifies epistemic virtues with faculties that are involved in forming beliefs (e.g. Sosa (2007; 2015), Greco (2009), Pritchard (2012)). The latter holds that the relevant epistemic virtues (and corresponding vices) are character traits (e.g. Zagzebski (1996), Baehr (2015)).

The proposal of particular online intellectual virtues falls within the responsibilist camp. Online intellectual virtues are acquired character traits that are applied to an online environment. Both Heersmink (2018) and Schwengerer (2021a) spell out the proposal of a virtue epistemology of the internet by looking at a desirable engagement with Google Search. When one uses Google Search, the results are generated according to a particular algorithm. Search algorithms usually use personal information to inform results. Simpson (2012) argues that such personalization undermines the objectivity of search results and thereby is detrimental to the user. Moreover, the algorithm can be gamed by websites being designed for the search algorithm, regardless of the actual use of the site for the user. And even the algorithm behind something so minor as the autocomplete feature has epistemically problematic consequences. Miller and Record (2013) suggest that autocompletion might nudge people to direct their search in highly problematic ways. Particularly striking examples of these problematic consequences are provided by Noble (2018), who demonstrates that

autocomplete within Google Search can promote racist and misogynist ideas. She also shows that search results themselves often reflect and promote these ideas too. Google Search seems to foster oppression.

Overall, Google Search seems to be an epistemic tool one should use with caution. This is what the virtue epistemologist takes to be an important lesson. They do not deny that the search engine itself needs to be better, nor that its function can be part of more systematic problems (e.g. white ignorance (Mills, 2007)). However, they suggest that the user of Google Search needs to know of these problems and adapt their epistemic behaviour accordingly. A good epistemic agent uses search engines with caution and in full awareness of the biases and problems that the search engine likely comes with. This is the core idea of online intellectual virtues. They are acquired character traits that help an agent to navigate the online environment properly. And they require a substantial degree of knowledge about that online environment. Hence, Heersmink's emphasis on "Internet literacy skills" (Heersmink, 2018, p. 7) and "lifelong learning" (Heersmink, 2018, p. 10) to counteract the quickly changing environment.

Even though Google Search is the example of choice in both Heersmink (2018) and Schwengerer (2021a), it is easy to see how the same concept applies to social media. Such platforms tend to foster some epistemically problematic tendencies just like Google Search. For instance, they tend to generate epistemic bubbles (Nguyen, 2020) by selecting only particular content for a user's default view. That such bubbles exist is well evidenced for Twitter/X (e.g. by looking at the connections of anti-vax and pro-vax users (Sullivan, et al., 2020)).² Users who do not know about the possibility of such bubbles, or that an algorithm

² However, Coady (2024) argues that talk of problematic echo chambers and filter bubbles is mistaken and we should avoid such talk altogether.

chooses to show the user posts they likely agree with, might draw false conclusions from the testimony seen on the social media platform. After all, if one is presented with an overwhelming amount of people testifying that p and barely any objections to that testimony, it seems rational to form the belief that p. However, the virtuous user of Twitter/X will be more cautious. They know that the posts they see are selected by an algorithm that does not even aim at providing good information. The algorithm is implemented to improve engagement. Hence, it is easy to distinguish better from worse epistemic agents in their use of social media as an epistemic tool. In other words, some users possess online intellectual virtues to a higher degree than others.

While Heersmink (2018) and Schwengerer (2021a) provide good reasons for virtuous users in a vacuum, actually following their suggestions and developing such virtuous is easier said than done. After all, companies who provide tools such as search engines are not necessarily interested in virtuous users, but might be more interested in the company's ad revenue. It might be in the best financial interest³ of a company to make it difficult for users to be virtuous. Hence, a company might design their website or app in a way that tricks users into behaving against their own best interest (so-called dark patterns in user interface design).⁴ Users of a search engine are situated in an online environment that can be hostile to the user's manifestation of epistemic virtues and users need to be trained on how to overcome these problematic aspects of the online environment. In addition, this is a place where regulation might come into play as a limiting factor on the hostility of the epistemic environment online. What companies are allowed to do to encourage epistemically vicious behavior in users can be regulated. One example of such a regulation is the European Union's General Data

³ There might be other non-epistemic interests as well, such as political goals.

⁴ See also Schwengerer (2022) for a discussion of web design to promote intellectual vices.

Protection Regulation that limits the hidden use of tracking cookies on websites. In this sense, being able to manifest virtues is not all on the individual agents, but embedded in a larger social environment. This social environment is not only limited to the particular technology being used as one factor for individual user's ability to manifest epistemic virtues. For this reason Frost-Arnold (2023) argues that virtue theoretic considerations about Internet users need to also consider the situatedness of users with regard to their race, class, gender, sexual orientation, disabilities.⁵ Whether a user can manifest virtues will depend on how the environment treats that particular user in their social situation. To give one example, it might be more difficult for a user to manifest epistemic virtues in an environment that comes with regular harassment and such harassment disproportionately affects women, people of color, and queer users (Citron, 2014). Hence, while online epistemic virtues can function as a guiding aim, it would be overly optimistic to think that these virtues can be achieved in every social environment.

So far, I have given a short explanation of online intellectual virtues and have shown that they also apply to users of social media. I also suggested that the ease of manifesting such virtues depends on the broader social environment and particular technology used. The next step is to look at Habgood-Coote's (Forthcoming) conflicts and find out whether the choices for these conflicts can be informed by virtue epistemic considerations.

3.2. Virtues in Individualism versus Institutionalism

The first conflict I want to look at is between maximizing the positive consequences of epistemic systems for the individuals involved and maximizing the positive epistemic features of the system as a whole. I show that at least some of these conflicts can be decided by virtue

⁵ The list is not exhaustive.

theoretic considerations. Let me start by looking at Habgood-Coote's description of such a conflict. He argues, that what leads to the best epistemic outcome for individuals at times does not lead to the best epistemic outcomes for the system or community. For instance, in some cases, worse individual performance increases group performance (cf. Smart (2018), Levy & Alfano (2020)). Moreover, the Zollman Effect (Zollman, 2007) shows that in other cases what is best for individuals (as much information connections as possible), turns out to be epistemically bad for the group overall. It can be detrimental on a system level, even though the individuals involved seem perfectly rational and epistemically virtuous. Zollman shows with formal models that a system with extremely well-connected people tends to end up with false beliefs settled in the system more often than a system with less connected agents.

This is not merely a theoretical case based on an abstract model. Social media, like Facebook, aims for connectedness between users. For every individual user, it seems epistemically prudent to improve their connectedness in their quest to gain knowledge and better coverage of important information (cf. Goldberg (2010)). But, as the Zollman Effect shows, connectedness can lead to a system settling too quickly on false beliefs.

Let me look at one such case suggested by Habgood-Coote in more detail. The case can be modelled as a conflict between two goals:

1. Maximizing informational connections for the individual users of a social media platform; and
2. Maximizing the system's overall epistemic results.

The first goal is almost self-explanatory. It is epistemically beneficial for individual users to have as many informational connections available as possible. Having more sources of

information in reach is generally a good thing for an epistemic agent.⁶ Moreover, social media platforms like Facebook explicitly aim to provide connections. Habgood-Coote argues – though with a tentative spirit – that this individually desirable goal of maximizing connections for individuals might conflict with the system’s overall epistemic features. A well-connected system might settle on consensus too easily, even when the consensus is a false belief. He writes:

Facebook presents itself as promoting connectedness amongst its users, encouraging users to create connections with many others. This kind of network structure does offer some epistemic benefits. In a highly connected network, individuals have access to a greater amount of information, can expect to hear of interesting news in a timely manner, and will be able to make use of coverage-supported reasoning (see Goldberg 2010 C6). However, the fact that highly connected networks provide epistemic goods for individuals does not mean that they are without problems at the institutional level. Highly connected networks can reach consensus too quickly (Zollman 2007, 2010), and can lose reliability due to informational cascades (Zollman forthcoming). Although this would take more evidence to back up, we might speculate that the reason why high-connection sites like Facebook and Twitter remain popular despite their well-publicised problems is that they trade off individual epistemic goods against our collective epistemic interest. (Habgood-Coote, Forthcoming)

Just looking at generic agents in the system and their connectedness, it seems that Habgood-Coote is right. And his tentative application to the real-world case seems plausible as well. However, the conflict does not seem to hold for any agent. I suggest that virtuous agents can

⁶ Of course, there are cognitive limits in processing information. I bracket this issue.

have their cake and eat it too. A social media platform with virtuous users can maximize individual connectedness without settling too quickly on false beliefs.

One way to argue for that is to look more closely at Zollman's results. Zollman (2010) already acknowledges himself that limiting communication – lower connectedness – is only one way to prevent a system from settling too quickly. He shows that extreme differences in priors work as well, because they ensure enough diversity between the agents in the system for long enough, such that the system does not come to a premature consensus on a wrong result. What turns out to be the key factor is not the amount of connections itself, but rather the speed at which people come to a consensus and stop looking for another potential answer. The difference in priors already shows, that we can achieve this result by having different agents involved, not merely by reducing connections.

This points to the virtue epistemic solution: the users of social media need to have character traits that prevent them from coming to a consensus in the online environment too quickly. Muldoon (2013) suggests that stubborn scientists can help ensure that scientific consensus tracks the truth. Something similar holds for users of social media. They need to be sufficiently stubborn – or in more positive terms – cautious in their belief acquisition. A similar effect can be achieved with users who keep exploring theories and views that are not the current consensus (cf. Rosenstock et al. (2017)). Overall, it seems that at least in this case one can settle the conflict between individualism and institutionalism suggested by Habgood-Coote (Forthcoming). Our virtue epistemic approach shows that one can opt for more connectedness while achieving the desirable features on the system (or institutional) level through a different source: virtuous agents that do not settle on a belief too quickly in a social media environment. To get those we need to follow Heersmink's (2018) plea to teach online epistemic virtues and promote lifelong learning for the required background knowledge.

3.3. Virtues in Institutional Properties versus Epistemic Justice

Habgood-Coote discusses the conflict between maximizing institutional epistemic properties and epistemic justice by looking at the problem of the second-best. In a nutshell, the best institutional arrangement in ideal conditions might not be the best one in non-ideal conditions, and might even promote epistemic injustice. The example he provides is based on Nancy Fraser's (1990) work. Fraser shows that women and minority groups face problems in real-life communicative spaces, even when they are not officially excluded from those spaces. These spaces might still make it exceedingly difficult for women and minority groups to contribute equally, give their voices less weight, restrictively conceptualize what counts as an appropriate topic or contribution, invoke male-coded norms, and so on. Moreover, she argues that in many real-life societies, these issues might even be more prevalent when no one is excluded from participation and everyone speaks in a unified public sphere⁷. In contrast, excluding some voices might help the otherwise unheard. Counterpublic spaces that are limited only to otherwise marginalized or oppressed groups allow those groups to develop their own norms of communication and contribution, without having to follow the norms set currently by the dominant group.

The argument applies straightforwardly to social media spaces. If the aim is to promote epistemic justice and allow every group to contribute properly, it might be prudent to restrict the users who can participate on a social media platform. Habgood-Coote provides a long list of work on online counterpublics supporting this idea:

Work in media studies has been heavily influenced by the public sphere tradition, and there is a rich body of work on online counterpublics, including curated blocklists on

⁷ For the notion of public sphere see Habermas (1989).

twitter (Geiger 2016), Black twitter (Graham and Smith 2016, Brock 2020 C3), hashtag networks (Jackson, Bailey, Foucault Welles 2020), blogs (Steele 2016a, 2016b, Gabriel 2015), and BlackPlanet (Byrne 2007). Some of this literature has a normative bent, arguing that Online counterpublic spaces can play a role in resisting epistemic oppression by providing safe havens and support networks (Geiger 2016, Sobieraj 2018), allowing marginalised groups to contest media narratives (Steele 2016a, 2016b, Gabriel 2015), and allowing marginalised groups to organise (Jackson, Bailey, Foucault Welles 2020). We can think of this literature as providing us with a first pass at an account of what effective online counterpublics look like. (Habgood-Coote, Forthcoming)

It seems that maximizing epistemic justice needs such counterpublics, but works against a unified public sphere. Even though from an institutional level it might be ideal for a social media platform to constitute a unified public sphere (cf. Sunstein (2017)), such unification is not only unrealistic but might be harmful in our non-ideal world. We end up with two goals that seem to be in conflict:

1. Unification as a desirable institutional, system-level property.
2. Counterpublics promoting epistemic justice.

One might ask here whether unification is even a desirable property. I cannot provide a full argument here, given that I work with the conflicts presented by Habgood-Coote. However, Habgood-Coote defends unification as a goal by pointing to Sunstein's (2017) work and giving it a Habermasian reading. Sunstein holds that an ideal deliberative democracy is a unified public sphere. In such a sphere all citizens are on equal footing, share some common experiences and encounter a diverse range of different views from other citizens. Democratic

problem-solving works best when such a diverse range of views come together on even grounds, such that a joint deliberation in the public sphere can take place. The idea in the background is that a unified public sphere allows for the maximal amount of evidence to be collected and used in deliberation. Less unification leads to groups that have their own deliberative spaces but lack some evidence that would be available in the more unified public sphere. Sunstein (2017) worries that this is exactly what happens when groups separate on the Internet to different platforms.

Suppose Habgood-Coote is right and there is a conflict between unification as an epistemic aim and epistemic justice. Is there a way to have both by looking at virtue epistemology? In an ideal world, the answer would be yes, at least to some degree. Epistemic justice within a community (or towards another person or group) can be a product of the virtue of epistemic justice in individual agents. The idea that justice in general can be a virtue of individuals has been discussed in LeBar (2020) and the specific proposal for epistemic justice as a virtue has been provided by Fricker (2007) and Kotsonis (2023). Fricker talks of the virtues of testimonial justice and hermeneutical justice, whereas Kotsonis aims to provide an account of a unified virtue of epistemic justice. The virtue is defined as follows:

“The character trait of epistemic justice is characterized by a disposition to act justly towards other knowers in matters pertaining to epistemic goods, involving a regard for agents as knowers. The disposition to act justly is an intuitive component of our concept of ‘being an epistemically just person’. (Kotsonis, 2023, p. 600)

If all users of a social media platform had this virtue of epistemic justice, then even a single unified platform without any exclusions would have little, if any effect on furthering oppression. No counterpublics would be needed. To make this idea a little more concrete I

will draw on Frost-Arnold's (2023) discussion of virtuous lurking. Lurking is here understood as the consuming of content from an online community without actively participating. Virtuous lurking is a complex activity. The virtuous lurker recognizes when lurking is appropriate, when they should rather not be present, and when they should rather participate. The virtuous lurker reads to unlearn prejudices and to experience epistemic friction – a concept introduced by Medina (2013) to capture a form of cognitive conflict that occurs when different ways of seeing the world clash. The virtuous lurker seeks out different perspectives without derailing conversations. As such, the virtuous lurker can be in a unified public sphere without pressuring other citizens to confirm to the lurker's perspective in any way. And the epistemically just person would be such a virtuous lurker. They give others in the public sphere the room when appropriate, which is why isolated counterpublics would not be necessary anymore (or at least rarely necessary). Hence, in an ideal world, we could choose institutionalism and would not suffer in the dimension of epistemic justice because the epistemically just agents would step back and fully give the public sphere to others when needed.

But of course, our world is not ideal. The virtue of epistemic justice is rare among social media users and few users are virtuous lurkers. In practice, users often invade spaces, put explanatory burdens on marginalized people (cf. Berenstain (2016)), or simply derail conversations. Working towards fostering this virtue among the public will lead to needing counterpublics less, but it is doubtful whether we can become so virtuous that those are not needed at all anymore.

Moreover, it is doubtful whether fostering virtues can always be done without counterpublics in the first place. The development of epistemic virtues is not done outside of a social environment and not every such environment will be equally suitable to foster virtues. It may

be the case that epistemic justice is difficult to achieve without having counterpublics that facilitate a productive kind of epistemic friction. Counterpublics can facilitate epistemic friction insofar as they allow members of a counterpublic to discuss and establish their own perspective on the world outside of the dominant perspective held by the larger public. As such, a perspective developed in a counterpublic can clash with the perspective of the larger public. Following Medina these conflicts can be particularly productive in developing a virtue of epistemic justice, because epistemic friction allows agents to see that their epistemic perspective is one of many. Friction is no guarantee for epistemic justice, but it can enable it. And if counterpublics play a role in enabling friction, it seems that the suggestion that counterpublics would not be (or would be less) needed with epistemically just users neglects how important counterpublics are for developing the virtue of epistemic justice in the first place. As such, one has to think of users “[...] as moral agents with varying degrees of virtue that can *change and grow over time*” (Frost-Arnold, 2023, p. 194; my italics). When one uses a virtue approach in looking at the conflict between institutional properties and epistemic justice, one has to consider the development of virtues – and that development will be harmed without counterpublic spaces.

Overall, from this perspective, virtue epistemology can provide a theoretical solution to the conflict, but that solution has practical limits. It shows where to go, but there is no quick and easy way to get there. Only through slowly fostering the virtue of epistemic justice the conflict between institutional properties and epistemic justice can be resolved. In the current, non-ideal world counterpublics are required while working on becoming virtuous agents and that needs to be respected in designing social media. Moreover, some counterpublics might always be needed to facilitate the development of epistemic justice in agents themselves. Hence, there is a limit to unification.

Perhaps there is another option. So far, I have explored whether maximal unification in social media is possible while compensating for potential consequences for epistemic justice with virtuous agents. Is there a way to maximize epistemic justice while still gaining unification through virtuous agents? I am unsure what such a way might even look like. Giving up on unification in social media in favour of epistemic justice comes with the creation of counterpublics online. It is difficult to see how unification could be achieved through other means once we have counterpublics. With choosing unification over epistemic justice, it at least seems possible in the ideal world to achieve said epistemic justice anyway through other means (fostering virtues) that are not in conflict with unification. But choosing against unification cannot be compensated in the same way. It seems that the only option is to choose unification while fostering the virtue of epistemic justice. And that option is certainly limited in our non-ideal world. Hence, the conflict remains, although the virtue epistemic position tells us at least in which direction to go.

3.4. Virtues in Individual Epistemic Properties versus Epistemic Justice

The final conflict highlighted by Habgood-Coote (Forthcoming) arises when we compare the benefits of a social media platform for individual epistemic goods and the effects for epistemic justice. He writes:

[...] consequentialist structure of individualist approaches to social epistemology means that individualism is unable to account for the importance of equality of epistemic outcomes and certain kinds of epistemic rights. Conflicts between majority epistemic interests and epistemic justice can in principle break either way, and there may be some cases in which the epistemic interests of the majority justify some undermining of epistemic agency. However, in the majority of realistic cases the

badness of epistemic oppression and unequal distributions of epistemic outcomes trump maximising aggregate epistemic outcomes. (Habgood-Coote, Forthcoming)

The worry is that looking for the best outcome for individual epistemic agents on average, often comes at the cost of bad consequences for a minority of agents within the system. Habgood-Coote illustrates this with a couple of examples, some of which I take to be better than others. For instance, the use of images without automatic descriptions in text form seems one that seems rather easily solved in a way that promotes epistemic justice without sacrificing the average individual epistemic gain.⁸ Images (can) have positive epistemic effects, but by relying on images without any alternative description those with difficulties seeing the images are pushed out. In this case, epistemic goods are distributed unequally, even though the average individual agent might benefit overall. However, adding those automatic text descriptions does not seem to come with significant costs for individual epistemic success.

A better example is the anonymity of users in social media. One commonly discussed way to hold users in social media accountable is to de-anonymize users. Being able to identify and significantly sanction or punish people for unwelcome actions online can be a deterrent to sharing misinformation. To a lesser degree we can also hold pseudonymous users accountable, by calling them out, and providing tools in the social media platform to flag unreliable sources (cf. Rini (2017)), but de-anonymized users are the paradigmatic example. A social media platform without anonymity might do better in terms of misinformation shared on the platform, but it will do a lot worse on the epistemic justice axis. For instance, it becomes increasingly risky for marginalized and oppressed groups to contribute on such a platform, for they easily could become a target – both on social media and in their normal life. Anonymity

⁸ See also Costanza-Chock (2020) for wider considerations of design justice.

counteracts (among other things) what Dotson (2011) calls testimonial smothering: a form of self-censorship to avoid being perniciously misunderstood and targeted. Without anonymity some users will therefore be excluded from having a proper voice in discussions. This might even question whether de-anonymization is the best choice if one is solely focused on maximizing individual epistemic goals. If such an accountability practice pushes out marginalized voices everyone on the social media platform will lose out on some knowledge. Nevertheless, suppose that an accountability practice using de-anonymized names does prevent the sharing of misinformation to a degree that overcompensates the loss of knowledge from pushing out marginalized voices. Even if that is the case, it is clearly not an epistemically just practice, because it excludes particular social groups from an epistemic practice. Can virtue theoretic considerations tell us what to do in this conflict?

I suggest they can. There are two different approaches available. Suppose one is in the process of designing a social media platform and wonders whether it should allow anonymous users. Let me assume the choices are then between the following:

1. Maximizing individual epistemic gains by choosing no anonymity.
2. Maximizing epistemic justice by choosing anonymity.

Again, the virtue theorist should ask, which of these options can be more easily achieved by other means, even if it is not chosen. I suggest that both paths seem possible, but one appears slightly better than the other.

First, if one takes option (1), then one has to create epistemic justice in the community (or system) by other means. As shown in the last section, from the virtue epistemologist's point of view this can be done by looking at the corresponding virtue of epistemic justice that individual users can have. In an ideal world, every user has acquired a character trait of

epistemic justice generally, and online epistemic justice in particular. Such online epistemic justice is the general virtue applied to the environment provided by the social media platform. Epistemically just users know of potential problems for users from marginalized communities and help create an environment in which no user experiences threats or pressure not to contribute. Epistemically just users also know when to step away from a space (cf. Frost-Arnold (2023) on virtuous lurkers). As such, just users work against the conditions that lead to testimonial smothering. In a world full of epistemically just users, anonymity would not be required and tools that sanction the sharing of misinformation could be implemented with fewer (or perhaps no) bad side-effects.⁹ However, as already mentioned earlier, a world full of epistemically just users is unrealistic. At best it can serve as an ideal goal to aim for, but it seems difficult to reach.

The other alternative – choosing option (2) – seems more attractive. One can maximize epistemic justice by allowing anonymous users to contribute while promoting epistemic virtues that reduce the epistemic harm that anonymous users can do by sharing misinformation. This can be done in two ways: on the contributor side and the consumer side.

On the contributor side, the virtue epistemologist urges to teach such epistemic virtues that prevent users from sharing misinformation. The conceptual tools for such virtues are already available in the discussion of other-regarding epistemic virtues, such as honesty, (Kawall, 2002) and other-regarding epistemic vices, like dishonesty (Harris, 2022) and epistemic apathy (Schwengerer & Kotsonis, Forthcoming). For the virtue epistemologist, the goal is to teach users of social media how to be good contributors to a social media platform by fostering other-regarding epistemic virtues like honesty. Moreover, the virtue epistemologist can

⁹ Note that if we think of users as individuals there could still be epistemically unjust agents in other forms, e.g. corporate agents or government agents.

encourage interventions into the social media platform itself that make it easier for users to develop and manifest these other-regarding epistemic virtues. For instance, social media platforms could introduce pop-ups that help recall people their responsibilities in sharing content.¹⁰ Something similar already happens with the use of tracking-cookies on websites and their acknowledgement following the European Union's General Data Protection Regulation. Schwengerer (2022) argues that "by being presented with a pop-up pointing to the tracking-cookies the mechanisms behind the website become more salient to users with enough background information" (Schwengerer, 2022, p. 306), which allows users to manifest virtues more easily. While Schwengerer (2022) is interested in solutions to make consumer side virtues manifest more easily, something similar might work for the sharing of information on social media.

On the consumer side, online epistemic virtues help users to better disentangle the misinformation from good information available on social media. A virtuous consumer of social media content knows the limitations and problems with online communications and considers those in their judgments. They will be appropriately careful in the social media environment. They might not be immune to misinformation, but more resistant. As such, social media with virtuously careful users can maximize epistemic justice by allowing anonymity and still achieve desirable epistemic properties for the users on an individual level. It does so because the epistemic virtuous can compensate negative aspects of anonymity. And again here, the virtue epistemologist is not only limited to facilitating the acquisition of epistemic virtues, but can also look at the manifestation of such virtues in the online

¹⁰ Thank you to a reviewer for suggesting this example.

environment. A social media platform can be designed in a way to make the responsibilities of a consumer more obvious to the user.

Neither contributor-side virtues nor consumer-side virtues are easy to acquire nor manifest. However, especially the consumer side virtues seem magnitudes easier to achieve than a user base full of epistemically just users. It seems comparatively easier to teach users how to avoid the pitfalls of misinformation than to be epistemically just. After all, being a virtuous consumer of social media can be motivated out of pure epistemic self-interest, whereas acting epistemically just likely appears to be costlier for the acting individual. Being a virtuous consumer of social media fits with the instrumental epistemic goals users already have.¹¹ As such, maximizing for (2) and compensating the loss in (1) by fostering epistemic virtues seems the more practical way to go. Promoting the relevant epistemic virtues is not easy, but achievable.

4. An Objection: Are these Orthogonal Issues?

At this point I have shown how the virtue theoretic approach can help us – at least sometimes – decide between available options when we are confronted with a conflict between different broadly epistemic goals of designing a social media platform. The general strategy was to look at which available choice could be compensated more easily if it was not chosen. One might wonder here whether that is something that Habgood-Coote (Forthcoming) should even be thinking about. Habgood-Coote is concerned with how to design a social media site, whereas the virtue theoretic approach concerns who we ought to be online. Perhaps these are two completely different questions, one might worry.

¹¹ There are some limitations here. For instance, some self-serving biases might work against being a virtuous consumer on social media by promoting denial of presented facts. Cf. Bardon (2019, Chapter 1)

As often with topics of internet epistemology, these questions cannot be kept neatly apart. When one thinks about how to design social media sites, one cannot do so without taking the users into consideration. Users themselves are a moving part that change the epistemic properties of a social media site. Hence, the question I am interested in is how to design a social media site to give us the best epistemic results, given that the users using the site can also change and be changed. In a sense, I add one more moving part to Habgood-Coote's approach and then ask, given these moving parts, how should a social media site look like. Accepting the users as a moving part is important, because otherwise one might end up with suboptimal social media systems in the long run. How a good social media site looks like will change depending on the users. Once one recognizes the users as one moving part (among many) that should be taken into considerations for designing social media, one has to decide whether users are the right part to tweak and change to get better results. One has to decide which broadly epistemic problems of social media should be best approached by attending to which moving part.¹² In case of Habgood-Coote's conflicts tweaking the design part of a social media site cannot solve all problems. Hence, one needs to ask which of the conflicting aims should be achieved by changing a social media platform's design and which aims should be achieved by other means. This, I suggest, is best done by also looking at a virtue theoretic approach in addition to more systemic features such as connectedness, content moderation, features allowing for easier generation of counterpublics, and so on. By looking at users from a virtue theoretic point of view I can decide which aims should be achieved by design choices and which can be achieved in another way. This is still a question of how to design social media

¹² This is similar to Frost-Arnold's (2023) approach, who draws on many different frameworks (Feminist accounts of Objectivity, Veritism, Ignorance, Virtues, and Injustice) to capture different moving parts. I am suggesting here that critical social epistemology of social media should pay more attention to the virtuous or vicious users as a moving part.

sites. However, it includes a prior question on what should or should not be aimed for with design alone to get the overall best outcome. Hence, critical social epistemology of social media should also look at virtue epistemology.

5. Conclusion

After looking at these conflicts, it seems clear that virtue epistemic considerations can sometimes help us decide how to choose within these conflicts. In the example of a conflict between individualism and institutionalism virtue epistemology provides an argumentative ground to choose connectedness in social media, even if that would lead to the epistemic system potentially settling too quickly. There are ways to compensate for the problems of the system settling too quickly by promoting epistemic virtues in the users of the social media environment.

In the example case of a conflict between the system-level (or institutional) property of unification and epistemic justice, the results were less promising. Even though in an ideal world the virtue of epistemic justice in individual agents could be paired with a choice of maximizing unification, in the non-ideal real world this is hardly achievable. Moreover, maximizing unification seems to have problematic consequences for the development of intellectual virtues, including epistemic justice itself. Counterpublics perform a valuable role in providing epistemic friction that fosters the virtue of epistemic justice. With a fully unified public sphere such friction might be lost and it becomes more difficult for agents to acquire the virtue of epistemic justice. Hence, virtue epistemology provides an ideal goal at best, but no proper answer to the conflict. Non-epistemic considerations are required.

Finally, in the example case of a conflict between the benefits of individual epistemic agents and epistemic justice, both choices are available. I have considered the example of online

anonymity in social media. In this case, it turns out that choosing epistemic justice is the better choice because the epistemic goods for individual agents can be achieved comparatively easier independently by fostering online epistemic virtues. Overall, the considered cases show that at least in some conflicts the better course of action can be evaluated purely on epistemic grounds.

The lesson here is not that Habgood-Coote (Forthcoming) is wrong. I believe he is mostly right. We need to look at multiple epistemic goals when designing how social media should look like and those goals often come in conflict. Sometimes to find a good way out of the conflict we need to look for answers outside of the epistemic. But Habgood-Coote does not bring enough attention to solutions within epistemology yet. This is partially a result of only looking at the conflicts from a system-level approach and partially the result of being the first contribution to the critical social epistemology of social media. His discussion so far does not engage with approaches focusing on the individual epistemic agents involved. However, if we consider broadly individualist epistemological approaches there are tools available that are already developed to deal with problems related to the conflicts mentioned.

Habgood-Coote points in his conclusion to the need for different epistemological methodologies, and his analysis so far is somewhat programmatic and leading to future development. My proposal of online intellectual virtues is such an additional methodology – a tool that can be used to decide conflicts fully within the epistemic domain. Some choices in these conflicts come with costs that can be potentially more easily compensated by virtuous epistemic agents than other choices. And the effort required for such compensation can help decide how to deal with the conflicts. As such, the proposal is a suggestion to combine approaches looking at individual epistemic agents and approaches looking at epistemic systems. Not every conflict will be decidable by epistemic considerations alone, regardless of

whether those are at a system level or the individual level. Some will still need ethical considerations. Hence, Habgood-Coote's proposal is still a good one and the inclusion of virtue epistemology should be considered a fruitful addition.

Acknowledgements

Work on this paper has been funded by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Projektnummer 462399384. Thank you to two reviewers of *Social Epistemology* who helped improve this paper significantly, both in structure and content.

References

- Anderson, E. (2012). Epistemic Justice as a Virtue of Social Institutions. *Social Epistemology*, 26(2), pp. 163-173.
- Baehr, J. (2015). *Cultivating Good Minds: A Philosophical and Practical Guide to Educating for Intellectual Virtues*. Retrieved from <https://intellectualvirtues.org/why-should-we-educate-for-intellectual-virtues-2-2/>
- Bardon, A. (2019). *The Truth About Denial: Bias and Self-Deception in Science, Politics, and Religion*. Oxford: OUP.
- Battaly, H. (2008). Virtue Epistemology. *Philosophy Compass*, 3(4), pp. 639-663.
- Berenstain, N. (2016). Epistemic Exploitation. *Ergo: An Open Access Journal of Philosophy*, 3(22), pp. 569-590.
- Boyer-Kassem, T., Mayo-Wilson, C., & Weisberg, M. (2017). *Scientific Collaboration and Collective Knowledge*. New York: OUP.

- Brock Jr, A. (2020). *Distributed Blackness: African American Cybercultures*. New York: NYU Press.
- Byrne, D. N. (2007). Public Discourse, Community Concerns, and Civic Engagement: Exploring Black Social Networking Traditions on BlackPlanet.com. *Journal of Computer-Mediated Communication*, 13(1), pp. 319-340.
- Citron, D. K. (2014). *Hate Crimes in Cyberspace*. Cambridge, MA: Harvard University Press.
- Coady, D. (2010). Two Concepts of Epistemic Injustice. *Episteme*, 7(2), pp. 101-113.
- Coady, D. (2024). Stop Talking about Echo Chambers and Filter Bubbles. *Educational Theory*, 74(1), pp. 92–107.
- Costanza-Chock, S. (2020). *Design justice: Community-led practices to build the worlds we need*. Cambridge, MA: MIT Press.
- Dotson, K. (2011). Tracking Epistemic Violence, Tracking Practices of Silencing. *Hypatia*, 26(2), pp. 236-257.
- Dotson, K. (2014). Conceptualizing Epistemic Oppression. *Social Epistemology*, 28(2), pp. 115-138.
- Fallis, D. (2003). Social Epistemology and the Digital Divide. *ACM International Conference Proceeding Series*, 101, pp. 79-84.
- Fraser, N. (1990). Rethinking the public sphere: A contribution to the critique of actually existing democracy. *Social Text*(25/26), pp. 56-80.
- Fricker, M. (2007). *Epistemic Injustice*. Oxford: Oxford University Press.

- Frost-Arnold, K. (2014). Trustworthiness and truth: The epistemic pitfalls of internet accountability. *Episteme*, 11(1), pp. 63-81.
- Frost-Arnold, K. (2023). *Who Should We Be Online?: A Social Epistemology for the Internet*. Oxford: OUP.
- Gabriel, D. (2015). Alternative Voices, Alternative Spaces: Counterhegemonic Discourse in the Blogosphere. In E. Thorsen, D. Jackson, H. Savigny, & J. Alexander (Eds.), *Media, Margins and Civic Agency* (pp. 15-28). London: Palgrave Macmillan.
- Geiger, S. R. (2016). Bot-based collective blocklists in Twitter: the counterpublic moderation of harassment in a networked public space. *Information, Communication & Society*, 19(16), pp. 787-803.
- Goldberg, S. C. (2010). *Relying on Others: An Essay in Epistemology*. Oxford: OUP.
- Goldman, A. (2010). Systems-oriented social epistemology. *Oxford Studies in Epistemology*, 3, pp. 189-214.
- Goldman, A. I. (1999). *Knowledge in a Social World*. Oxford: Oxford University Press.
- Graham, R., & Smith, S. (2016). The Content of Our #Characters: Black Twitter as Counterpublic. *Sociology of Race and Ethnicity*, 2(4), pp. 433-449.
- Greco, J. (2009). Knowledge and Success From Ability. *Philosophical Studies*, 142(1), pp. 17–26.
- Habermas, J. (1989). *The Structural Transformation of the Public Sphere*. Cambridge: Polity Press.

- Habgood-Coote, J. (Forthcoming). Towards a Critical Social Epistemology of Social Media. In J. Lackey, & A. McGlynn (Eds.), *Oxford Handbook of Social Epistemology*. Oxford: OUP.
- Harris, K. R. (2022). Outward-facing epistemic vice. *Synthese*, 200(516).
- Heersmink, R. (2018). A Virtue Epistemology of the Internet: Search Engines, Intellectual. *Social Epistemology*, 32(1), pp. 1-12.
- Jackson, S. J., Bailey, M., & Foucault Welles, B. (2020). *#Hashtag Activism*. Cambridge, MA: MIT Press.
- Kawall, J. (2002). Other-Regarding Epistemic Virtues. *Ratio*, 15(3), pp. 257-275.
- Kotsonis, A. (2023). On the Virtue of Epistemic Justice and the Vice of Epistemic Injustice. *Episteme*, 20(3), pp. 598-610.
- LeBar, M. (2020). *Justice as a Virtue*. (E. N. Zalta, Editor) Retrieved from The Stanford Encyclopedia of Philosophy: <https://plato.stanford.edu/entries/justice-virtue/>
- Levy, N., & Alfano, M. (2020). Knowledge From Vice: Deeply Social Epistemology. *Mind*, 129(515), pp. 887-915.
- Longino, H. E. (1990). *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry*. Princeton: Princeton University Press.
- Longino, H. E. (2022). What's Social About Social Epistemology? *Journal of Philosophy*, 119(4), pp. 169-195.
- Lynch, M. (2016). *The Internet of Us: Knowing More and Understanding Less in the Age of Big Data*. New York: W.W. Norton and Company.
- Marino, M. C. (2020). *Critical Code Studies*. Cambridge, MA: MIT Press.

- Marwick, A., Kuo, R., Cameron, S. J., & Weigel, M. (2021). *Critical Disinformation Studies: A Syllabus*. Retrieved from The Bulletin of Technology & Public Life: <https://citap.pubpub.org/pub/tifer5w4>
- Mayo-Wilson, C., Zollman, K. J., & Danks, D. (2011). Mayo-Wilson, Conor ; Zollman, Kevin J. S. & Danks,. *Philosophy of Science*, 78(4), pp. 653-677.
- Medina, J. (2013). *The Epistemology of Resistance: Gender and Racial Oppression, Epistemic Injustice, and the Social Imagination*. Oxford: OUP.
- Miller, B., & Record, I. (2013). Justified Belief in a Digital Age: On the Epistemic Implications of Secret Internet Technologies. *Episteme*, 10(2), pp. 117-134.
- Mills, C. W. (2007). White Ignorance. In S. Sullivan, & N. Tuana (Eds.), *Race and Epistemologies of Ignorance* (pp. 13-38). Albany, NY: State University of New York Press.
- Muldoon, R. (2013). Diversity and the Division of Cognitive Labor. *Philosophy Compass*, 8(2), pp. 117-125.
- Nguyen, T. C. (2020). Echo Chambers and Epistemic Bubbles. *Episteme*, 17(2), pp. 141-161.
- Noble, S. U. (2018). *Algorithms of Oppression*. New York: New York University Press.
- Pritchard, D. (2012). Anti-Luck Virtue Epistemology. *Journal of Philosophy*, 109(3), pp. 247-279.
- Rini, R. (2017). Fake News and Partisan Epistemology. *Kennedy Institute of Ethics Journal*, 27(2S), pp. 43-64.
- Rosenstock, S., Bruner, J., & O'Connor, C. (2017). In Epistemic Networks, Is Less Really More? *Philosophy of Science*, 84(2), pp. 234–252.

- Schwengerer, L. (2021a). Online Intellectual Virtues and the Extended Mind. *Social Epistemology*, 35(3), pp. 312-322.
- Schwengerer, L. (2021b). Revisiting Online Intellectual Virtues. *Social Epistemology Review and Reply Collective*, 10(3), pp. 38-45.
- Schwengerer, L. (2022). Promoting Vices: Designing the Web for Manipulation. In F. Jongepier, & M. Klenk (Eds.), *The Philosophy of Online Manipulation* (pp. 292-310). Routledge.
- Schwengerer, L., & Kotsonis, A. (Forthcoming). On the Intellectual Vice of Epistemic Apathy. *Social Epistemology*. doi:10.1080/02691728.2024.2356574
- Simpson, T. W. (2012). Evaluating Google as an Epistemic Tool. *Metaphilosophy*, 43(4), pp. 426-445.
- Smart, P. (2018). Mandevillian Intelligence. *Synthese*, 195(9), pp. 4169-4200.
- Smart, P. R., & Clowes, R. W. (2021). Intellectual Virtues and Internet-Extended Knowledge. *Social Epistemology Review and Reply Collective*, 10(1), pp. 7-21.
- Sobieraj, S. (2018). Bitch, slut, skank, cunt: Patterned resistance to women's visibility in digital publics. *Information, Communication & Society*, 21(11), pp. 1700-1714.
- Sosa, E. (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge* (Vol. 1). Oxford: Oxford University Press.
- Sosa, E. (2015). *Judgment and Agency*. Oxford: Oxford University Press.
- Steele, C. K. (2016a). The Digital Barbershop: Blogs and Online Oral Culture Within the African American Community. *Social Media + Society*, October-December, pp. 1-10.

- Steele, C. K. (2016b). Signifyin', Bitching, and Blogging: Black Women and Resistance Discourse Online. In S. U. Noble, & B. M. Tynes (Eds.), *The Intersectional Internet: Race, Sex, Class, and Culture Online* (pp. 73-94). New York: Peter Lang.
- Sullivan, E., Sondag, M., Rutter, I., Meulemans, W., Cunningham, S., Speckmann, B., & Alfano, M. (2020). Can Real Social Epistemic Networks Deliver the Wisdom of Crowds? In T. Lombrozo, J. Knobe, & S. Nichols (Eds.), *Oxford Studies in Experimental Philosophy Vol.3* (Vol. 3). Oxford: Oxford University Press.
- Sunstein, C. R. (2017). *#Republic: Divided democracy in the age of social media*. Princeton: Princeton University Press.
- Zagzebski, L. (1996). *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge: Cambridge University Press.
- Zollman, K. J. (2007). The communication structure of epistemic communities. *Philosophy of Science*, 74(5), pp. 574-587.
- Zollman, K. J. (2010). The Epistemic Benefit of Transient Diversity. *Erkenntnis*, 72, pp. 17–35.
- Zollman, K. J. (2013). Network Epistemology: Communication in epistemic communities. *Philosophy Compass*, 8(1), pp. 15-27.