Title: A normative framework for sharing information online

Authors: Emily Sullivan & Mark Alfano

Abstract: People have always shared information through chains and networks of testimony. It's arguably part of what makes us human and enables us to live in cooperative communities with populations greater than the Dunbar number. The invention of the Internet and the rise of social media have turbo-charged our ability to share information. In this chapter, we develop a normative framework for sharing information online. This framework takes into account both ethical and epistemic considerations that are intertwined in typical cases of online testimony. We argue that, while the current state of affairs is not entirely novel, recent technological developments call for a rethinking of the norms of testimony, as well as the articulation of a set of virtuous dispositions that people would do well to cultivate in the capacity as *conduits* (not just sources or receivers) of information.

Keywords: social epistemology, testimony, virtue epistemology

Word count: 7550 (maximum 7500)

1. Introduction

The vast majority of what any individual knows depends on chains and networks of testimony stretching outwards socially and backwards temporally. Our knowledge has accumulated over the course of years, decades, centuries, and even millenia. Each talk exchange is an opportunity to extend and enrich this vast testimonial network, which is — with the exception of a few isolated tribes — almost fully connected around the globe. Moreover, since the invention of the internet, both the speed and the fidelity of transmission have increased dramatically. Despite these exciting developments, the epistemology of testimony has made modest progress in addressing its inherently social components. The paradigm case of testimony in the philosophical literature consists of the transmission of knowledge from *exactly one* person to *exactly one* other person, neglecting proximal and (even more so) distal social sources of knowledge. In this chapter, we draw on resources from multiple disciplines to explore the differences among various modes of communication and the social epistemic consequences of these differences. We highlight several key parameters of different testimonial regimes, including *fidelity, velocity, traceability, democratization,* and *anonymization*. Online communication is characterized by extremely high fidelity, velocity, traceability, and democratization. For this reason, the epistemic norms that ought to govern online communication may differ in interesting and important ways from the epistemic norms that have traditionally governed other testimonial regimes such as oral tradition, handwriting, and print media.

2. A brief history of testimonial networks

Few philosophers, even in social epistemology, have moved beyond the hearer-speaker dyad. For example, some have considered what it means for *groups* to testify (Tollefsen 2007; Lackey 2015, 2018). And in the philosophy of science, others have started to explore which network structures are conducive to sharing knowledge within idealized communities of scientific researchers through agent based modeling (Holman and Bruner 2015; Rosenstock et al. 2016; Weatherall et al. 2019; Zollman 2007, 2010). Little attention has been paid to the evaluation of testimonial networks consisting of both experts and non-experts in non-ideal settings. In philosophy, Coady (1992) briefly considers the influence of network structure on the transmission of historical knowledge. However, he does not identify formal network features or structures that are likely to produce epistemic goods and avoid epistemic ills. More recently, Sullivan et al. (2019, 2020) showed that some online social networks are ill-suited to deliver the wisdom of crowds.

Researchers in other fields offer a broader perspective. In sociology, Senturk (2005) addresses the structure of the *hadith* transmission network from Muhammad's contemporaries to subsequent generations. This is arguably the longest extant intergenerational testimonial network of its sort, spanning 610 CE to 1505 CE. Senturk shows that participants in this network, known as *huffaz* (roughly "ones who memorize and protect" *ahadith*), aimed to learn *ahadith* from other *huffaz* with the shortest paths back to the prophet and his companions. In the terminology we use below, this means that the *huffaz* saw themselves as playing the role of epistemic conduits, and that they attempted to minimize the length of epistemic geodesics in the *hadith* transmission network.

In anthropology, Blong (1982) shows that oral legends passed down through generations of Papua New Guineans accurately record a catastrophic volcanic eruption that occurred

centuries prior. And in literary studies, researchers have pointed out that bards (like *huffaz*) were professional memorizers and reciters of important pieces of language. Lord (1960; see also Parry 1971) argues that bards and other singers of epic poetry from Homer to the present did not memorize whole poems verbatim. Instead, they memorized themes and formulas (essentially, important, memorable phrases that employ a particular rhythmic pattern, e.g., "Ajax with his shield like a wall"), which they creatively recombined in each successive recitation of the epic. Lord contends that the transmission methods associated with oral versus written transmission make an important difference to how language is preserved over time. Essentially, the shift from the spoken to the written word increases fidelity. At the same time, this tectonic shift poses a threat to the creativity of professional memorizers and reciters. And by democratizing the transition process, writing has the potential to destroy the livelihoods of these professionals.

These processes continued to accelerate with the shift from the written to the printed word (Gutenberg). Typographical errors introduced by medieval monks could be largely eliminated by the use of the printing press, though it did not produce transmittable copies with perfect fidelity. Even more recently the invention of the internet has made it possible in many cases to quickly and automatically track back a piece of testimony to its original source. At the same time, the unregulated way in which online knowledge is recorded and transmitted has led to a seemingly intractable problem of citogenesis, in which track-backs among dynamic webpages such as Wikipedia entries lead one in a loop rather than back to any genuine original source (Sterzer et al. 2008, Saez-Trumper 2019).

In what follows, we identify several key parameters that describe testimonial networks (e.g. fidelity and democratization), and how the online information environment influences these parameters, leading to changes in our epistemic responsibilities. We then discuss how these parameters could inform individual epistemic norms and institutional norms surrounding information sharing in the online context.

3. Key parameters of testimonial networks

There are several parameters that can describe information flow in a testimonial network. These parameters range from the speed of information flows to whether the structure of the network creates bottlenecks to information flow. In this chapter, we focus on the following five parameters and their impact on epistemic norms of testimony in the online context:

Fidelity
The similarity that message, *m`*, has to the original or transmitted message, *m*. Fidelity can be understood in terms of both semantic, linguistic, or orthographic similarity.

Velocity
The speed of message transmission.

Traceability
The ease with which it is possible for a receiver of *m`*to trace *m`* back to *m*.

Democratization
The proportion of those in the network who are able to transmit messages. The mechanism that determines who in the network is able to transmit messages.

Anonymization
The degree to which the identity of individuals transmitting messages in the network is publicly known to each other member of the network.[1]

3.1 Fidelity

As information moves through a network, it can get more or less distorted or noisy along the way (Shannon 1948). Fidelity measures the extent to which noise or decay occurs. The way information is shared, or the medium that information is shared through, can impact the type of information stored and shared. For example, greater processing power allows for sending complex messages that otherwise would have needed to be condensed.[2] As such, technology along with the type of testimonial network can impact fidelity. As mentioned above, the printing press increased fidelity by eliminating many typographical errors when transmitting messages. Social media platforms increase fidelity even further through their share functions that create digital copies of the original message.

The content of a message can be more or less similar to the original message both in terms of its semantic content and its linguistic features. For example, a message may have the same words as another message but have different semantic content, or a message may have different words but expresses the same semantic content. For example, a message translated into a different language will not have linguistic fidelity but may hold onto its semantic fidelity (and such translation services are increasingly being automated online). Shortened messages can also have semantic fidelity. In the context of online information sharing, messages are broadcasted to large numbers of people very quickly. While linguistic fidelity remains high, increased spread allows for context shifting where the meaning and function of the message evolves as it moves through the network, resulting in semantic fidelity decay while the linguistic fidelity remains intact (Sullivan 2019). The epistemic climate on social media platforms, Twitter in particular, becomes a series of over-hearers engaging in context shifting due in part to short messages, but also due to the sheer number of people that messages can travel through. Linguistic fidelity can be easy to measure with the right tools, but the degree of semantic fidelity is considerably more difficult because the pragmatic context of utterance matters (Lackey 2006). The pragmatic context can sometimes be difficult enough for the average person to identify, and is an especially complicated problem for automated algorithms. As a result, algorithms that aim to filter messages are often trained on linguistic similarity which can miss important trends in semantic fidelity.

It can seem as though increased fidelity in a network makes the network epistemically better than one with decreased fidelity. After all, higher fidelity means less room for miscommunication and misunderstandings. However, sometimes message decay can have epistemic benefits. For example, information decay can be useful in collective systems by disrupting gridlock and promoting cooperation (Shirado and Christakis 2017). There is even evidence in schools of fish that decreased fidelity can prevent the spread of false alarms (Rosenthal et al. 2015). If we take the spread of misinformation and disinformation as a type of

---

[1] While other network parameters may also be important for norms of testimony, for the purposes of this paper, we restrict our scope to focusing on the five identified here.

[2] See Szathmary and Smith (1995) for an example from biology.

false alarm, then perhaps it is the increased fidelity of messages in the online environment, that is in part to blame for our current epistemic landscape where rumors and misinformation spread faster than truth (Vosoughi et al. 2018).

3.2 Velocity

Message fidelity is also impacted by *velocity*, that is how quickly messages are transmitted throughout the network. Testimonial networks found in early human communication involved slow and noisy travel. As a result of a low velocity, information degraded or changed as it moved through the network. For example, early battlefield communication suffered from low velocity, and armies with more efficient means of sending messages had an advantage over their enemy while operating under the fog of war (Clausewitz 1832 / 1989). The internet affords almost instantaneous transfer of messages. This speed can help preserve message fidelity since there is no time in between the sharing and receiving of a message for the message to change.

Increased velocity not only benefits message fidelity; there are other benefits to the faster spread of messages. In an emergency situation such as a wildfire or flood, the speed with which messages can be propagated can save lives. High velocity networks also level the epistemic playing field among agents in the network. Information can be powerful, and in a slow-moving network, those who get information first can have advantage over those who receive it later. Thus, a network with increased velocity creates a more egalitarian epistemic environment in some respects.

Increased velocity can also come with epistemic downsides. There is evidence linking high-speed message flow with overwhelming cognitive processes or leading to less accurate decisions (Tversky and Kahneman 1974; Chittka et al. 2009). When there are too many messages coming in quickly, it is not possible for each one to be fact-checked or even paid attention to. There needs to be some mechanism for filtering important messages from unimportant messages. Unfortunately, left to our own devices, we humans often rely on “system 1” processing that relies on heuristics and biases that can lead us astray (Kahneman 2011). High velocity itself also puts pressure on the network not to slow down. This can lead to a lack of due-diligence in fact-checking or tracing back the message to authenticate it. Agents may feel rushed in making their decisions on whether to spread a piece of information to keep up with underlying flow trends. Utilizing automated algorithms as a means of filtering information can reduce cognitive overload, but algorithms come with their own biases, which often recapitulate or even exacerbate human biases, resulting in systematic bias across the whole network. In the case where individuals in the network filter for themselves, there may be an increased likelihood that individuals introduce message diversity due to different interests and perspectives; however, algorithms run the risk of shutting out important messages entirely.

3.3 Traceability

Traceability concerns the ease with which it is possible for a receiver of $m`$ to trace $m`$ back to m. Traceability is important for verifying the accuracy and authentication of messages, contributing to the epistemic health of a network. The more difficult a message is to trace back to its original source, the harder it will be for the message to be assessed on its epistemic credentials. In networks where there are few actors that are able to transmit messages — especially if the network has low velocity — traceability is likely to be easier to achieve. However, in a high

velocity network where there are numerous active information-sharers, the thread of a message can get lost in the sea of interconnections (especially when screen-capping, photoshop, blocking, and other technologies enter the scene). Amplification farms that repost content from other pages, either with or without the original source name, have significantly decreased traceability on social media platforms. An increasing number of these amplification farms repost the same content under a different source name, creating the appearance that the content has originated from many places instead of just one (Burke 2018; Nyhan 2019). In regular offline environments, if the same message comes from a variety of sources, that signals its importance and reliability; however this epistemic heuristic only makes sense if the messages are genuinely sprouting up from diverse sources. Increased traceability allows for knowing whether our epistemic heuristics about sources continue to hold true in the current network, or need to be revised in some way.

While social media technology has in many ways hindered traceability, there are possible technological fixes. For example, introducing technology architectures that authenticate and keep a record of message sharing can increase traceability. For example, blockchain architectures provide an immutable and decentralized ledger with an authentication process. Blockchain is used primarily in currencies and banking transactions, like Bitcoin (Nakamoto 2008). However, the underlying concept of an immutable ledger for message sharing may be attractive for epistemic reasons.

3.4 Democratization

Who in the network is able to transmit messages? How much training/certification is required for someone to be given the ability to transmit messages? Democratization tracks the proportion of those in the network who are able to transmit messages, and the ease with which they are able to gain this status. Historically, technology has played a role in increasing the democratization of testimonial networks. The printing press not only increased the fidelity of messages flowing through contemporary testimonial networks, but it also increased democratization by expanding who could own *and write* books. However, even though more people could own and write books, the publishing system was not fully democratized since gatekeepers prevented most from getting their books published and only would give certain people or organizations the rights to distribution, limiting access to many.

Social media in turn has created a highly democratized system of message sharing by eliminating gatekeepers. Accounts can be created in minutes with very little oversight of the content of shared messages. We all now serve as editors, choosing which messages to pass on and influencing the content that those in our immediate network see. The public acting as editor has even begun to bleed over into more traditional media outlets. For example, the *New York Times* in 2017 decided to eliminate their public editor. A public editor acts as a liaison to the public and seeks to uphold ethical journalistic practices at the paper, balancing the interests of the public with the interests of journalists. The NYTimes decided instead of a public editor that the public itself could directly be the judges on quality and informativeness. Comment boards were opened up and journalists were advised to read these comments to better their job. So while there is still editorial oversight for the content at the paper, it is the public at large that votes on what is worth reading and what is considered newsworthy with their shares and other forms of engagement. Critics argue that this move opens up the NYTimes to promoting more click-bait and lower quality articles (Lind 2017).

As with the other network parameters discussed above, there are epistemic advantages and disadvantages to the increased democratization of our testimonial networks. A more egalitarian system for information sharing allows those from historically disenfranchised groups to contribute to the shared knowledge of the network. It can also increase the epistemic diversity of the topics that are shared. If more people are engaged in content creation and content sharing, then a variety of content should make it through at least portions of the network, since each person has his or her own editorial sense. However, there are also costs to increased democratization. When everyone is an editor, there is no centralized check or filter to content. This means there can be a flood of disinformation and conspiracy theorizing, since there is no central body fact-checking or eliminating such information. The repercussions of a lack of epistemic oversight is showcased across social media with the increased prevalence of flat earthers and anti-vaccine advocates. Importantly, disinformation in a highly democratized network can come both from bad actors (e.g., spooks, grifters) and those who are well-intentioned but lack the required expertise (e.g., aunts, uncles, tech bros on Medium).

Moreover, because of the aforementioned problems stemming from high velocity and the sheer number of messages propagating throughout contemporary networks, along with increased democratization, the need for message filtering becomes even more pressing. Each individual cannot possibly attend to all the messages in a careful way. Instead of using more traditional gatekeeping methods for quality checks, such as needing to be a member of a professional organization or a newspaper editor, the network relies on other ways of filtering content. For example, celebrities have enjoyed a greater degree of epistemic power in social networks online (Archer et al. 2020). People rely on those whom they trust or are in the public eye to determine what is important information for them. However, celebrities are not necessarily epistemically qualified in determining what is important or truthful (Dennis, Alfano, & Archer under review). Again, we see the negative side-effects celebrity power has with regards to the anti-vaccine movement (Freed et al. 2011).

There are other means of filtering messages that rely less on heuristics and shortcuts made by individuals. Filtering algorithms are prevalent on social media platforms; however, the way they work is typically opaque to those in the network. Often, these algorithms are not optimized to promote truth or other epistemic values, but rather are developed to promote message engagement to drive advertising revenue. Facebook, for example, even goes so far as to actively avoid making epistemic decisions on content (e.g., refusing to fact-check political advertisements (Ortutay and Anderson 2020)). Facebook and Twitter do not see themselves as media or news providers but as a social network connectors akin to the community bulletin board in a café. Unfortunately, given the monopolistic place that social media platforms have in our economy, they are unlikely to change their business model on their own. As of the writing of this chapter, just three companies account for nearly 90% of global market share. Facebook alone enjoys 64.22%, plus another 7.05% from Instagram (which Facebook owns). Twitter enjoys 13.96%, and YouTube 3.79% (Alfano and Sullivan *forthcoming*). Regulation could help by classifying social media companies as news agencies and holding them to the standards that other news agencies face, in regards to fact-checking content, or by forcing them to move away from the current attention-economy business model towards one that is more consistent with epistemic norms. However, notice that any centralized mechanism that restricts the messages within the network cuts down on democratization. A tension thus exists between epistemic values and network democratization.

3.5 Anonymization

Anonymization measures the degree to which the identity of individuals transmitting messages in the network can be determined by others. It ranges from no anonymization to pseudonymization, to anonymization. Anonymization, as with the other network parameters, has both epistemic upsides and downsides. The upsides are similar to those seen with democratization. If sources of information or those who merely transmit messages can remain anonymous, then sensitive information, such as whistleblowing, leaking to the press, and behind-the-scenes gossip can be introduced in the network that might otherwise be kept silent (Alfano & Robinson 2017). People may fear retaliation for spreading certain true information and thus refrain from sharing it even if they have the ability to do so.[3] Thus, a level of anonymization can improve the epistemic standing of the network by introducing important information that would otherwise never be reported.

The downsides of anonymization also mirror the downsides of democratization. Anonymization reduces the social costs of producing low quality information or sharing misinformation. If there is no social risk to transmitting a message, then there is less need to be epistemically diligent. Thus, the removal of filters and a high-level of democratization coupled with a high level of anonymization can perpetuate the rapid distribution of false information in a network (Veliz 2019).

4. Translating Parameters into Epistemic Norms

It should be clear from the previous section that we can consider each network parameter on its own, but they also influence each other to a large extent. Velocity can impact fidelity, anonymization has an impact on traceability and democratization, and so on. Due to the interconnected nature of these network parameters, we propose a virtue-theoretic approach to epistemic norms for online information sharing. However, we should note that this approach does not presuppose that knowledge can be defined or analyzed in terms of epistemic virtues, and much of the following could, we suspect, be made compatible with an alternative approach.

4.1 Conduits and Epistemic Norms

One aspect that is different in epistemic networks that we find online compared to ordinary cases of testimony often considered by epistemologists is that there are more than two testimonial roles. In an ordinary case of testimony, there is the *source* of the message and the *receiver* of the message. We submit that there is a third possible epistemic role, someone simply passing messages across the network, i.e., a *conduit*. Conduits pass on information from sources (or from other conduits).[4] Thus, they are not content creators or sources. Moreover, conduits are not mere receivers of information. A receiver of information is someone who listens for incoming

[3] An example of retaliation for whistleblowing might be seen in the developing case of three doctors mysteriously falling out of windows in Russia after speaking out about the medical conditions there surrounding COVID-19 (Ilyushina, 2020).

[4] The line between a conduit and a source is not always clear cut. If a conduit substantially changes a message before sending it along, then the conduit could also be considered a source. In this chapter, we stick to cases of conduits that do not intentionally change the messages they receive.

messages and perhaps is looking for a specific message, but a receiver need not pass on this information. Conduits, on the other hand, receive messages *and* pass them on (Sullivan and Alfano *forthcoming*, Sullivan et al. 2020).

The question concerning epistemic norms and responsibility of online information sharing is largely a question concerning the epistemically responsible behavior of a conduit. It is a question about passing information along through the network, not necessarily about content creation or merely being on the receiving end of messages. Thus, our focus here is on the epistemic responsibility of conduits. In the next section, we look at more institutional norms that can also mitigate some of the epistemic problems of the online information landscape.

Maintaining knowledge in a social epistemic environment requires distinct considerations. First, once we move beyond the speaker-hearer dyad, network structure plays a role in the way information is spread throughout the network. There has been interesting work in agent-based modeling that looks at testimonial network structures at the level of the whole network, suggesting that certain types of structures are epistemically better than others, at least for certain purposes (Citation Zollman and others). However, individual agents in the network also need to bear in mind their surrounding network structure in order to be epistemically responsible conduits (Sullivan and Alfano *forthcoming*, Sullivan et al. 2020).

In previous work, we identified three central classes of virtues that are required to be a good epistemic agent in a social network: monitoring, adjusting, and restructuring (Sullivan and Alfano *forthcoming*). These virtues also apply specifically to conduits within networks on social media platforms concerning when to virtuously share information. *Monitoring* virtues cover what is needed for agents to attend to the structure of their network. It is only through monitoring paths of information flow that one can see whether the sources they are relying on are independent from each other or not. For example, if a piece of information is shared by multiple people in your network, you have *prima facie* evidence that this information is important and reliable. However, if you trace back the origin of this message as coming from only one original biased source, then the fact that multiple people amplified the message should not increase your confidence in the information. On the contrary, it should make you question whether your network structure has enough independent sources (Sullivan et al. 2020). A conduit blindly contributing to the amplification of information from a single biased source reduces the epistemic quality of the network overall.

One reason for monitoring one's network is so that one can more accurately judge how to calibrate and adjust one's credence in sources and pieces of information, not only in order to adjust belief formation, but also in order to decide when to take on the conduit role and pass the information along. *Adjusting* virtues are those that govern this decision. One specific factor in the decision to become a conduit should be based on how independent your network is. For those that find themselves in a network structure that relies on only one or two sources, it may be foolhardy to amplify these messages, if there are no countervailing reasons to suspect the information is reliable. As such, monitoring virtues also involve monitoring the epistemic roles and track-records of those in the network. Do my sources have a reliable track-record, or do they often provide false or misleading information? Are my sources independent, or are they themselves conduits simply amplifying the messages of others? Adjusting virtues in turn govern your sharing behavior based on this information. In the former case, there is good reason to share information, but in the latter case refraining from passing along the message could be the more responsible action.

Being a responsible conduit also means taking action to *restructure* your network, if the current one is too severely flawed. This could involve seeking out new sources, no longer listening to sources one had previously trusted, or effecting more distal changes in the structure of the network. Doing this well depends on sufficiently successful monitoring and the motivation and capacity to identify efficient and effective changes that one has the power to enact. It might be that the nature of the network is so far gone that the epistemically responsible thing to do is either remove oneself from the network (Levy 2017) or to no longer engage in the conduit role. Moreover, while in an ideal world, conduits would do their epistemic due-diligence in developing the dispositions necessary to share information in a epistemically responsible way, in reality it can involve a lot of time and effort, such that only a small number of actors in the network are able to become responsible conduits. This creates a tradeoff between democratization and epistemically responsible agents. If only a small number of actors function as conduits in an epistemically responsible way, then the network cannot be fully democratic and epistemically successful at the same time.

Another feature of online testimonial networks is that epistemic virtues should not just be understood as localized within the individual. As with ethical virtues, epistemic virtues can be *other-regarding* (Kawall 2002; Fricker 2007). After all, the role of the conduit is to pass on information in the network to the benefit or hindrance of others. Other-regarding considerations should both be directed toward specific individuals (e.g. the conduit's inner circle) and the health of the network as a whole. Indeed, there are some things that individuals can do to help mitigate the epistemic problems that come with fidelity, velocity, traceability, democratization, and anonymization. One of the problems with maintaining fidelity in online networks is that the same message can take on different semantic meanings because of context shifting. One response to this problem is for conduits to refrain from passing along messages that would change the semantic fidelity of the original message when sharing in a particular context, and also refraining from sharing messages with a low level of semantic fidelity to their originals. For example, satirical posts are likely candidates of context shifting. Some people will interpret the message as being informative instead of non-informative.[5] Conduits should take care that they are not sharing a satirical message in a way that could be interpreted by others as informative. Conduits then might even have the responsibility to clearly state that the message is not informative. This responsibility can increase for a specific conduit depending on the position the specific conduit has in the network. Do their messages tend to propagate far through the network? Or do the messages tend to spread only to a few? However, this heuristic regarding *current* network position can only take us so far. The viral nature of platforms like Twitter means that any message can propagate far through the network even if there was no reasonable indication it would do so beforehand. If a given network has a low level of traceability, then fidelity checks will be hard to perform. In the absence of sufficient traceability, conduits again should think twice about passing along a message without qualification.

As discussed above, high velocity networks come with the drawback of eliminating the space for fact-checking and can lead to cognitive overload and poor decisions about what to believe and what to share. Individual conduits can help mitigate these issues by intentionally slowing down the speed of information flow simply by waiting before sending. A conduit can take time to trace back the message, checking source reliability as a way of delaying. Delaying

[5] See, for example, url = < https://literallyunbelievable.tumblr.com/ >.

messages could be especially important for breaking news based on unverified sources. Retraction in general can be ineffective (Ecker et al. 2011), especially in online environments. In such an environment, conduits have a larger responsibility to get things right the first time, whereas in other social network environments where retraction is effective, delaying messages might not be as important.

Due to the high level of democratization on social media platforms, conduits need to take their role of public editor seriously. They should work to cultivate virtues that are conducive to responsible filtering. For example, not giving celebrity voices undue epistemic power and making sure that the expert voices amplified are actually experts in the specific topic of the message seems prudent. Even though social media platforms are egalitarian in terms of how message sharing operates (i.e., everyone shares a message in the same way), these platforms are not egalitarian in terms of epistemic power. People who are seen as celebrities or experts have more influence on others with more weight attached to the messages they share. Moreover, those in the network whose messages are more likely to spread far throughout the network because of their network position have more epistemic power. Following up on sociological work by Conway & Oreskes (2010), network simulations have shown that a determined propagandist in a testimonial network is often capable of shifting the opinions of many others (Weatherall et al. 2019). So even in the absence of a centralized system that controls epistemic power in online networks, unequal power structures can emerge. Conduits should do well to monitor these power relations and do what they can by choosing which messages they share in a way that would distribute epistemic power along epistemic lines instead of perpetuating epistemically poor patterns of information flow.

4.2 Institutional Norms

Many of the epistemic problems associated with the network parameters discussed in section 3 cannot merely be solved by individuals developing conduit virtues. Change must be made on an institutional level in the form of regulation, oversight, education, and developing new technological tools to mitigate some of the problems. Rini (2017) likewise argues that individual epistemic dispositions and virtues are inadequate to address the problem of misinformation, and that targeting the structure of epistemic institutions such as social media platforms is necessary.

It is the design of social media platforms themselves that can play a role in perpetuating false and misleading information. Platforms play an active role in perpetuating some of the network parameters discussed in this paper in a way that hinders epistemic goals. The functionality that allows immediate and quantifiable social feedback through likes and shares can amplify the velocity of sharing and works to prioritize certain types of content. For example, Bradey et al. (forthcoming) discuss the propensity of sharing morally charged content and Vosoughi et al. (2018) discuss the propensity of people sharing news that tends to inspire fear, disgust, and surprise, which is often associated with false news. This problem is exacerbated by another common feature of social media platforms: displaying the number of times an article has already been shared by others. Research suggests that people are up to seven times more likely to share content online when they see others are already doing so (Bakshy et al. 2012). Platforms are also known to delay the like and share count of particular posts to give the appearance of continued engagement throughout the day (Morgans 2017). This has the added benefit that users check the platform more regularly.

Social media platforms are also designed to encourage as many connections as possible (Arfini et al. 2018), which means that any time an individual shares something it will be spread further than if the connections were more limited. In network simulations of highly connected networks, failures of epistemic competence are demonstrated (Hahn et al. 2018) as well as bandwagoning effects (e.g. O'Connor & Weatherall 2018; Zollman 2007). So while individual epistemic dispositions matter for sharing information online, several of the epistemic problems with social media platforms come from their business model and the way they are structured. For example, Alfano and Sullivan (*forthcoming*) argue that social media platforms are natural monopolies, which explains the business models of these companies and the reluctance to adapt to growing ethical and epistemic concerns.

In the remainder of this section, we articulate institutional changes affecting the five network parameters discussed above, with the hope of mitigating some of the epistemic problems with online information sharing. Recall that one of the problems with message fidelity in online environments is that semantic fidelity is hard to measure and yet can change rapidly in a highly connected and high-velocity network. Automated algorithms, if developed to identify semantic fidelity, may help to mitigate this problem. However, since algorithms can be biased, these algorithms should be subject to oversight. Users of the platforms should have some understanding of how information is flagged and filtered. Regulatory agencies should have access to algorithms to oversee whether they respect privacy and protected classes. Social media platforms also have a responsibility to track linguistic similarity. Sometimes small changes to the words of a message can alter meanings, giving the message, at first glance, the appearance of being reliable even if it was changed to exhibit a falsehood. As of this writing, we are not aware of any pressure on social media platforms to track message fidelity.

The central epistemic problem arising from high-velocity networks is that they can overwhelm cognitive processes and leave little room for fact-checking and proper gatekeeping. Again, there are institutional changes that could mitigate these issues. For example, platforms already engage in altering the actual timeline of posts, likes, and shares. However, they engage in these delaying tactics in order to increase engagement and their profit margins. Delaying messages that are indicative of rumors or false and misleading information could help to reduce the spread of such problematic messages. The delaying tactic could be done while the platform is able to engage in more direct fact-checking or authentification.[6] It is important to note, though, that fact-checking can be difficult during emergencies and with developing stories. Care must be taken when engaging in delaying tactics that vital life-saving information is not delayed.

Traceability is also a big issue in large interconnected networks. Social media platforms have well-known problems with amplification and troll farms that distort where a message originated from. If platforms were able to make progress on fidelity, it could come with gains in traceability. Moreover, more vigilant oversight of amplification and troll farms can help to remove bad actors from the network. Given the monopolistic character of these platforms, such oversight might need to come from regulatory agencies and not simply be done internally by the

[6] Content moderation in itself is not an easy problem to solve. A recent report found that content moderators for Facebook often end up believing the very conspiracy theories they are meant to screen (Newton 2019). Moreover, the working conditions for moderators is poor with workers suffering from PTSD given the type of content they need to screen (Newton 2019). Developing labor regulations around content moderation is an important first step.

company. Another avenue for a technological solution, as mentioned above, could include the introduction of an immutable ledger that would keep track of messages and how they changed. Including an authentication mechanism to the ledger could also make strides toward identifying amplification farms.

Above, we discussed how highly-democratized networks take away gatekeeping mechanisms. In highly-connected and high-velocity networks some form of filtering or gate-keeping is necessary. Without a centralized mechanism, individuals will use their own heuristics, which may not be epistemically virtuous. For example, celebrities have large amounts of epistemic influence on social media platforms. Platforms also contribute to promoting content by allowing sponsored posts and advertising. For example, on Instagram, advertising can go unnoticed because so-called influencers are paid behind-the-scenes to promote a product, which gives a misleading impression to other users. So while everyone in the network has the same ability to post and share information, power imbalances do emerge. Not everyone in the network has equal access to amplification. Amplification algorithms are designed with the platform's business model in mind. Requiring social media platforms to be transparent about their amplification methods, and classifying social media companies as news agencies, forcing them to move away from the current attention-economy business model would be a start to mitigating these issues.

Lastly, consider anonymization. Some level of anonymization is desirable in any epistemic network because it creates the environment for information to be shared that might not otherwise be shared (e.g. whistleblowing). However, allowing unlimited anonymous or pseudonymous accounts can degrade the epistemic quality of the network, due to the lack of reputational risk. Veliz (2018) suggests a variety of approaches to pseudonyms that could ensure that reputation still does its essential work. One such suggestion is for people to have a certain number of pseudonyms that are traced by a third party in the event of either ethical or epistemic misconduct. How to best handle trade-offs between the freedom and epistemic benefits of anonymization, while still securing reputational accountability is a difficult problem. One thing is clear though, there need to be institutional wide solutions that go beyond shutting down one troll / bot at a time.

5. Conclusion

In this paper we identified five network parameters that can be used to describe online testimonial networks: fidelity, velocity, traceability, democratization, and anonymization. We explored the epistemic benefits and drawbacks of increasing each of these measures. We then used these insights to help inform individual and institutional epistemic norms for online information sharing. In the case of the individuals, we focused on what it would mean to be an epistemically responsible *conduit*, that is, someone who takes on the role of message passing in a social network. In the case of institutional norms, we identified that social media platforms could institute institutional wide changes in order to combat false and misleading information. However, given their business structure and profit motive, it seems that such institutional changes need to come from regulation.

References

Alfano, M. & Robinson, B. (2017). Gossip as a burdened virtue. *Ethical Theory and Moral Practice*, 20: 473-82.

Alfano, M. & Sullivan, E. (*forthcoming*) "Online Trust and Distrust" In Michael Hannon and Jeroen de Ridder (eds) *Routledge Handbook of Political Epistemology*. Routledge.

Archer, A., Cawston, A., Matheson, B., & Geuskens, M. (2020). Celebrity, democracy, and epistemic power. *Perspectives on Politics*.

Arfini, S., Bertolotti, T., & Magnani, L. (2018). The diffusion of ignorance in on-line communities. *International Journal of Technoethics*, 9(1): 37-50.

Blong, R. J. (1982). *The Time of Darkness: Local Legends and Volcanic Reality in Papua New Guinea*. University of Washington Press.

Brady, W., Crocket, M., & Van Bavel, J. (forthcoming). The MAD model of moral contagion. The role of motivation, attention and design in the spread of moralized content online. *Perspectives on Psychological Science*.

Burke, T. (2018) "How America's Largest Local TV Owner Turned Its News Anchors Into Soldiers In Trump's War On The Media" *Deadspin* https://theconcourse.deadspin.com/how-americas-largest-local-tv-owner-turned-its-news-anc-1824233490, Accessed May 2020.

Chittka L, Skorupski P, Raine NE (2009) Speed–accuracy tradeoffs in animal decision making. *Trends in Ecology & Evolution* 24(7):400–407.

Clausewitz, C. (1832 / 1989). *Vom Kriege / On War*. Princeton University Press.

Coady, C. A. J. (1992). *Testimony: A Philosophical Study*. Oxford University Press.

Conway, E. & Oreskes, N. (2010). *Merchants of Doubt*. Bloomsbury.

Dennis, M., Alfano, M., & Archer, A. (under review). The uses and disadvantages of celebrity for pandemic response.

Ecker, U. K., Lewandowsky, S., Swire, B., & Chang, D. (2011). Correcting false information in memory: Manipulating the strength of misinformation encoding and its retraction. *Psychonomic Bulletin & Review*, *18*(3), 570-578.

Freed, Gary L., Sarah J. Clark, Amy T. Butchart, Dianne C. Singer, and Matthew M. Davis. 2011. "Sources and Perceived Credibility of Vaccine-Safety Information for Parents." Pediatrics 127(Supplement 1): 107–12.

Fricker, M. (2007). Epistemic injustice: Power and the ethics of knowing. Oxford University Press.

Holman, B., & Bruner, J. P. (2015). The problem of intransigently biased agents. *Philosophy of Science*, *82*(5), 956-968.

Kahneman, D. (2011). *Thinking, fast and slow*. Macmillan.

Kawall, J. (2002). Other–regarding epistemic virtues. Ratio, 15(3), 257-275.

Lackey, Jennifer. 2006. "The Nature of Testimony." *Pacific Philosophical Quarterly* 87 (2): 177-197.

Lackey, J. A. (2015). A deflationary account of group testimony. In *Essays in collective epistemology*. Oxford University Press.

Lackey, J. (2018). Group Assertion. *Erkenntnis*, *83*(1), 21-42.

Lind, D. (2017). "The New York Times is getting rid of its public editor for exactly the wrong reasons." *Vox*

https://www.vox.com/2017/5/31/15719278/public-editor-liz-spayd-new-york-times. Accessed May 2020.

Lord, A. (1960). *The Singer of Tales*. Harvard University Press.

Morgans, J. (2017). "Your Addiction to Social Media Is No Accident" *Vice News,* https://www.vice.com/en_us/article/vv5jkb/the-secret-ways-social-media-is-built-for-addiction. Accessed May 2020.

Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system.

Nyhan, B. (2019). "Americans Trust Local News. That Belief Is Being Exploited." *New York Times,* https://www.nytimes.com/2019/10/31/upshot/fake-local-news.html. Accessed May 2020.

Newton, C. (2019) "The Trauma Floor" *The Verge* https://www.theverge.com/2019/2/25/18229714/cognizant-facebook-content-moderator-interviews-trauma-working-conditions-arizona. Accessed May 2020.

Ortutay, B., and Anderson, M. (2020). "Facebook again refuses to ban political ads, even false ones" *AP news*, https://apnews.com/90e5e81f501346f8779cb2f8b8880d9c. Accessed May 2020.

Parry, M. (1971). *The Making of Homeric Verse: The Collected Papers of Milman Parry*. Ed. A. Parry. Oxford University Press.

Rini, R. (2017). Fake news and partisan epistemology. Kennedy Institute of Ethics Journal, 27(S2): 43-64.

Rosenthal, S., Twomey, C.R., Hartnett A.T., Wu H.S., Couzin I.D. (2015). "Revealing the hidden 732 networks of interaction in mobile animal groups allows prediction of complex behavioral contagion." *Proceedings of the National Academy of Sciences* 112 (15).

Rosenstock, S., Bruner, J., & O'Connor, C. (2017). In Epistemic Networks, Is Less Really More? *Philosophy of Science*, *84*(2), 234-252.

Saez-Trumper, D. (2019). Online disinformation and the role of Wikipedia. Url = < https://arxiv.org/abs/1910.12596 >. Accessed 5 May 2020.

Senturk, R. (2005). *Narrative Social Structure: Anatomy of the Hadith Transmission Network, 610-1505*. Stanford University Press.

Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, 27(3): 379-423.

Shirado H. and Christakis N.A. (2017). "Locally noisy autonomous agents improve global human coordination in network experiments." *Nature* 545 (7654) : pp. 370–374.

Sterzer, M., McDuff, P., & Flasz, J. (2008). Note to file - the challenge of centralized control faced by the intelligence function in Afghanistan. Canadian Army Journal, 11(2): 96-100.

Sullivan, E., Sondag, M., Rutter, I., Meulemans, W., Cunningham, S., Speckmann, B., & Alfano, M. (2019). Can real social epistemic networks deliver the wisdom of crowds? In T. Lombrozo, J. Knobe, & S. Nichols (eds.), *Oxford Studies in Experimental Philosophy*. Oxford University Press.

Sullivan, E., Sondag, M., Rutter, I., Meulemans, W., Cunningham, S., Speckmann, B. & Alfano, M. (2020). Vulnerability in social epistemic networks. *International Journal of Philosophical Studies*.

Sullivan, E., and Alfano, M. (forthcoming) "Vectors of Epistemic Insecurity" In Ian James Kidd, Heather Battaly & Quassim Cassam (eds.), *Vice Epistemology: Theory and Practice*, Routledge.

Tollefsen, D. (2007). Group testimony. *Social Epistemology*, 21(3): 299-311.
Tversky A, Kahneman D (1974) Judgment under Uncertainty: Heuristics and Biases. *Science (New York, N.Y.)* 185(4157):1124–31.
Veliz, C. (2019). Online masquerade: Redesigning the internet for free speech through the use of pseudonyms. *Journal of Applied Philosophy*, 36(4): 643-58.
Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, 359:1146-51.
Weatherall, J., O'Connor, C., & Bruner, J. (2019). How to beat science and influence people: Policymakers and propaganda in epistemic networks. *British Journal for the Philosophy of Science*.
Zollman, K. J. (2007). The communication structure of epistemic communities. *Philosophy of science*, *74*(5), 574-587.
Zollman, K. J. (2010). The epistemic benefit of transient diversity. *Erkenntnis*, *72*(1), 17.