Vulnerability in Social Epistemic Networks

Emily Sullivan

Max Sondag

Ignaz Rutter

Wouter Meulemans

Scott Cunningham

Bettina Speckmann

Mark Alfano

1 Introduction

The growing use of online media has brought the problems of filter bubbles (Pariser 2011), echo chambers (Nguyen forthcoming), and group polarization (Sunstein 2017; Pallavicini et al. forthcoming) into focus. Recent journalism has labeled the current era a time of “epistemic crisis” (Roberts 2017a) in which “tribal epistemology” dominates (Roberts 2017b). Levy (2017) goes so far as to suggest that the best response to this crisis of epistemic vulnerability is to cut oneself off entirely from many sources of information. Rini (2017) likewise argues that individual epistemic dispositions and virtues are inadequate to address the problem, and that we need to focus instead or in addition on the structure of epistemic institutions such as social media platforms.

Regardless of whether one thinks the current state of epistemic affairs represents a new crisis of epistemic vulnerability, the worries underlying these discussions are not new. Epistemologists have traditionally concerned themselves with related issues, such as the epistemic harms of dogmatism (White 2006), the ways beliefs should be responsive to evidence (Chisholm 1957), and epistemic norms associated with testimonial exchange (Coady 1995). Nevertheless, while the epistemic worries themselves are not new, the emergence of these problems in a fast-paced social-network context suggests that new tools may be necessary to diagnose and evaluate the present problems.

When epistemologists think about social exchanges, they tend to idealize away the ways epistemic agents depend on epistemic networks to both acquire and maintain their knowledge. Instead, attention is addressed to the transmission of knowledge from *exactly one* person to *exactly one* other person, neglecting proximal and (even more so) distal social sources of knowledge, error, and ignorance. The few epistemologists who go beyond epistemic dyads either consider only scientific communities and the structure of their epistemic networks (Zollman 2012, 2015) or employ simulated networks with unclear ecological validity (Masterton et al. 2016; Masterton & Olsson 2017, 2018). This makes it impossible to address the question whether real social networks *as they are currently constituted* are more liable to produce the wisdom of crowds or the madness of masses. In other words, contemporary approaches are ill-

equipped to address the extent to which we suffer from epistemic vulnerability. In recent work, Sullivan et al. (2019) show that, at least when it comes to discourse about controversial topics like vaccine safety, social media tends to be highly polarized and to amplify misinformation and disinformation rather than broadcasting expert consensus. Thus, to better evaluate the epistemic vulnerabilities associated with filter bubbles, echo chambers, and polarization, we can no longer abstract away from the fact that epistemic agents are situated in large epistemic communities and receive information from many interconnected sources. What we need is a methodology that makes it possible to assess the epistemic qualities and drawbacks of social networks. This is not to suggest that individual epistemology should be abandoned, or that there are additional ways in which people can be epistemically vulnerable — from low intelligence to poor education to malnutrition to simply being sleepy and tired. However, these individual traits and states have been studied to the neglect of the social dimension of vulnerability, so we focus on the social in this paper.

In this paper, we give pride of place to the epistemic network an agent is situated in.¹ This context, we argue, is an essential starting point for any effort to evaluate whether an agent knows, is well-informed, is a reliable source of information, and so on. In addition, this approach enables us to identify problematic structures in real epistemic communities.

Our aims in this paper are 1) to motivate a network-focused methodology in epistemology, 2) to identify and conceptualize important networked epistemic concepts, 3) to outline how to formally define and identify these concepts in real networks, and 4) to apply our formal theory to a concrete case (namely the discussion of vaccine safety on Twitter). In particular, we construct a network pattern that characterizes epistemic agents who draw on multiple independent sources, which makes them well-placed epistemic *observers*. We establish a method for locating this observer pattern in epistemic networks, and argue that promoting this pattern is well-suited to combat filter bubbles, echo chambers, and group polarization. In so doing, we show that social-network epistemology has much to contribute to contemporary philosophical and societal debates about the extent and severity of epistemic vulnerability.

¹ In developing this approach, we do not discount other epistemic concepts and phenomena. For example, a knowledge-first approach to epistemology takes knowledge as basic and seeks to understand and motivate concepts such as rationality and justification in terms of knowledge (Williamson 2000). By contrast, a function-first approach analyses epistemic concepts in terms of their purposes and functions in epistemic communities (Kelp forthcoming). Still further, virtue epistemology looks to the epistemically ideal agent and identifies skills, abilities, or character traits that contribute to someone doing well epistemically (Zagzebski 1996; Baehr 2011; Alfano 2013). While these approaches primarily differ over which epistemic phenomena take pride of place, they are not mutually exclusive. A virtue epistemologist may refer to the function and purpose of say, open-mindedness, in order to explicate when agents are too open-minded or too close-minded.

2 Motivating Social Network-Epistemology

Imagine an individual in the epistemic state of nature (Craig 1999). This individual, call her Eva, has to fend for herself, epistemically speaking. Eva only learns what she manages to perceive and infer on her own. On the one hand, this makes her epistemically secure against those who might lie to her or distract her. On the other hand, Eva has only very limited knowledge about the world. Her epistemic network is just a single node, herself (Figure 1).

Figure 1
Epistemic Singleton

We can imagine that Eva might feel frustrated in epistemic Eden. She wants to learn about the world outside her garden. To slake her curiosity, she can go exploring, but of course the world is big and Eva can only be at one place at a time. She would benefit from an informant: someone who could report to her about the things he's seen and thought that she hasn't seen or thought herself. Given her epistemic needs and aims, Eva needs a *source*. Let's call him Steve. With the addition of Steve, Eva's network expands and takes the structure of a testimonial dyad (Figure 2).

Figure 2
Testimonial Dyad

However, Eva's network is still quite limited. Her knowledge can only expand as far as Steve's knowledge. In this case, even though Steve has explored parts of Eden that Eva herself has not explored, his experience isn't so different from hers. Eva decides that, while having a single source is nice, having *multiple* sources would be better. So she goes to talk with her neighbors, Theresa, Ursula, and Vera. With the addition of these new sources Eva's epistemic network expands, forming a testimonial star network (Figure 3).

Figure 3
Testimonial Star Network

Theresa, Ursula, and Vera have a lot to say. They've been all over the face of the Earth and seen many things. Moreover, they've reached a bunch of conclusions that neither Eva nor Steve would have drawn using only their own evidence and their own powers of inference. When she doesn't have any direct experience related to a given proposition, Eva doesn't just have to take Steve's word about it anymore. Instead, she can poll Steve, Theresa, Ursula, and Vera. Unless she has special reason to distrust one or another of them, she can then let the majority rule. In this way, Eva is able to *aggregate* the testimony that flows to her from her epistemic network.

However, at moments when they don't think she's listening, Eva overhears Theresa, Ursula, and Vera talking to yet another person, Luke. It turns out that, while Theresa, Ursula, and Vera do have their own personal wells of experience, much of what they've been telling Eva is just what Luke told them. Pondering this overheard conversation, Eva realizes that, while having sources helps her learn about the world, it also makes her vulnerable. She becomes painfully aware of her own *epistemic vulnerability*.

Figure 4
Vulnerable Testimonial Network

In addition, Eva realizes that, while she'd thought that she had expanded the number of her sources from one to four, in a sense, she really only had two sources: Steve and Luke. Thus, her epistemic network does not have the structure of a testimonial star (Figure 3), but the structure of a vulnerable testimonial network (Figure 4). Because Theresa, Ursula, and Vera are not *independent*, they serve to *amplify* the stream of messages coming from Luke. Previously, when Steve had disagreed with Theresa, Ursula, and Luke, Eva tended to disregard Steve and believe what the other three had to say. But it now seems to her that she'd been too hasty. Three is better than one, but only when the three have independent evidence, employ different methods of inquiry, or differ in their sensitivity to epistemic reasons. Neglecting the issue of independence essentially makes Luke into Eva's epistemic dictator.

To better cope with her epistemic vulnerability, Eva decides to keep talking both to Steve and to Theresa, Ursula, and Vera. From now on, though, she plans to treat the chorus coming from the latter three as a single source of testimony, not three different sources. Eva is now doing her best to *monitor* the structure of her testimonial network with an eye to the number of independent sources she has. She is also in a position to *adjust* her credence in light of structural defects such as the amplification of Luke by Theresa, Ursula, and Vera. She no longer uses a majority-rules heuristic when aggregating their testimony. Instead, she thinks about whether Steve or Luke is epistemically more reliable about the topic in question, as well as whether the fact that Theresa, Ursula, and Vera all accept Luke's testimony is evidence in its favor.

Eva also takes steps to *restructure* the geometry of her epistemic network — something she can do only because she is successfully monitoring it. She decides that it would be good to hear from someone who is neither a naif like Steve nor the amplifier of a single additional source like Luke. Eva wants to hear from multiple, independent sources. What's more, she wants to hear from multiple, independent, and *diverse* sources. Someone who follows the same path through Eden as Steve but never speaks to him might be independent in some sense, but they won't have much

novel information to offer. Eva prefers to hear from people who've been places and thought thoughts that her other sources have not. So she decides to speak with two additional sources who (as far as she can tell) have different perspectives and who are not in communication with Steve, Theresa, Ursula, Vera, or Luke. The structure of Eva's epistemic network is now more secure (Figure 5)

Figure 5
Secure Testimonial Network

This change in the structure of Eva's epistemic network makes it more challenging for her to wisely aggregate the sometimes-conflicting testimony that she receives, but at least she has a chance now. At the same time, she realizes that listening to more and more sources is starting to take up a lot of her time. Life is short, and Eva would like to have the opportunity to continue exploring the world on her own sometimes. She realizes that learning from testimony takes *effort*, which should be expended efficiently. In addition, she realizes that the best sources to listen to are ones that have proven *reliable*, at least to some extent, in their past testimony.

3 Formal Categorizations

In the previous section, we motivated and introduced the key constructs of social network-epistemology. Many of these constructs refer to the topology or geometry of the network, which can be studied only by moving beyond the simple dyads that receive the lion's share of attention in contemporary social epistemology (Alfano 2017). These constructs help us to make sense of Surowiecki's (2005) analysis of the wisdom of crowds, according to which such wisdom is best harnessed by having multiple, independent, diverse, reliable sources in a decentralized network, then aggregating the information they supply appropriately. To succeed in the epistemic state of nature, Eva needs to *monitor* the structure of her testimonial network with an eye to the *number of sources* in it, along with the *independence*, *diversity*, and *reliability* of those sources. Her success also depends on her ability to *aggregate* the testimony of these sources by *adjusting* her credence in light of imperfections in her testimonial network.

Furthermore, Eva can take an active role in *restructuring* the geometry of her testimonial network, though this takes *effort*. Her best chance of succeeding in the epistemic state of nature thus requires her to embody three families of dispositions or intellectual virtues (monitoring, adjusting, and restructuring) related to the shape of her testimonial network, which is characterized by four parameters (number of sources, independence, diversity, reliability, and effort required to change the structure).²

While reliability has been discussed at great length by social epistemologists, the other parameters have largely gone unaddressed. In this section, we formalize and model these central epistemic concepts regarding epistemic position (current state) and security (threats to and opportunities for improving epistemic position). In the discussion below, we denote by $N = (V, E)$ a network with n nodes $V = \{v_1, \dots, v_n\}$ and directed edges E that are a subset of $V \times V$. An edge (u, v) in E represents the “testifies-to” relation, that is, node u provides testimony to node v .

We use $N \setminus \{v\}$ to denote the network N without the vertex v and its associated edges. A path in the network is a sequence of vertices $\langle v_0, v_1, \dots, v_l \rangle$ with $(v_i, v_{i+1}) \in E$ for all i with $0 \leq i < l$. The length l of such a path is its number of vertices minus one, or equivalently its number of edges. A path in the network represents a sequence of testimony and thus represents potential for flow of information. We use $\delta(u, v)$ to denote the length of the shortest path from u to v in a network.

3.1 Epistemic position

As we observed above, there are four main components that determine how good someone’s epistemic position currently is within a network: the number of sources they have, how independent these sources are from each other, how diverse the viewpoints of these sources are, and how reliable the sources are. Because reliability has been covered quite well by the existing literature, we here focus on the other three parameters. This subsection shows how we quantify these components to arrive at a single score that characterizes how strong a given node’s epistemic position is within a network.³

² Our approach is compatible with non-virtue theoretical accounts of epistemic duties or norms. The important point for our purposes is that in order to be doing well epistemically one needs to monitor the epistemic network one finds oneself in, and take into account defects in that structure.

³ This approach differs in several ways from the agent-based models used in the philosophy of science literature (e.g., Alexander et al. 2015; Muldoon and Weisberg 2011). First, we are not proposing a mechanistic belief-updating scheme with different “types” of agents, such as followers and mavericks. Second, we do not assume that all members of a network are epistemically cooperative and have primarily epistemic goals, as members of the ideal scientific community are often assumed to be. Third, we are interested in the structural epistemic position of each node in the network, whereas agent-based modeling tends to focus on the network as a whole. This of

First consider independence. Let v denote the node whose epistemic position in the network we want to characterize. Each node u in V that has an edge to v (that is, (u, v) is in E) is a potential source of information for v , since each edge represents the testifies-to relation. To capture the independence of two sources x and y with respect to v , we call a node x *m-independent* with respect to v , if there is no node s such that $\delta(s, x) + \delta(s, y) \leq m$ where the distance δ is measured in the network $N \setminus \{v\}$. That is, the absence of a node s that may influence both x and y within a certain distance implies that these sources are testimonially independent. Note that s may also be x (or y), in which case $\delta(s, x) = 0$. For example, any two distinct nodes are always 0-independent, they are 1-independent if there is no edge from one to the other, etc.

Second, we cannot consider the number of sources a given node has isolated from the independence of those sources. Thus, we measure how many sources a node has relative to a given threshold of independence. We start by defining a node as an *(m,k)-observer*. A node v in an epistemic network is an *(m,k)-observer* if it receives testimony from k distinct sources that are pairwise *m-independent* with respect to v . For example, v is a (2,3)-observer if it receives testimony from 3 distinct sources, and each of those sources is a distance of at least 2 from each of the other sources (unless they communicate via v itself). The *(m,k)-observer* status of a node, utilizing independence and number of sources, captures how well it is positioned structurally in the network.

Quantifying better and worse *(m,k)-observer* positions is not straightforward. A node can be an *(m,k)-observer* for multiple values of m and k . We also observe that if a node v is an *(m,k)-observer*, then it is also an *(m',k')-observer* for any $m' \leq m$ and $k' \leq k$. First, generally higher values indicate a better position in the network. However, not all values are immediately comparable: how does a (4,2)-observer compare to a (3,3)-observer? Is it epistemically better to have four sources that are less independent of one another or three sources that are more independent? What is the appropriate tradeoff between the number of sources and the distance between those sources? To answer this question, we use the following formula: we compute the maximum multiplication of m and k , for those pairs of values for which v is an *(m,k)-observer*. This gives a slight preference to, for example, a (3,3)-observer over a (4,2)-observer (since $3 \times 3 = 9$ while $4 \times 2 = 8$). Note, however, that a (4,2)-observer and a (2,4) observer are considered equal by this measure.

Moreover, we restrict our attention to reasonable values of m and k . For instance, any node with at least one source is an $(\infty, 1)$ -observer. This, however, is uninteresting because it does not get beyond the simple dyadic model. In addition, in most real-world social networks the maximum value of m is 6 (Milgram 1967), so it is unreasonable to expect independence values greater than

course does not undermine the agent-based modeling approach, but we want to be clear that what we are doing is distinct.

5. For these reasons, we consider values of m only up to 5 and values of k greater than or equal to 2. A node with zero or one source in total is given a final score of 0 or 1 respectively. Formally, consider the set $MK = \{ (m,k) \mid 1 \leq m \leq 5 \text{ and } 2 \leq k \leq 5 \text{ and } v \text{ is an } (m,k)\text{-observer} \}$. Then the structural position of v is defined as:

$$S = \max_{(m,k) \in MK} (m \times k)$$

Since, the (m,k) -observer is a structural measure, it is agnostic to both the means and the content of testimony. Testimony may be provided orally, in writing, via semaphore, or whatever. In addition, as discussed above, even if sources are structurally independent, they may communicate the same viewpoint. Thus, we must consider diversity of viewpoint to better capture epistemic position. By ‘diversity’ we have in mind not the demographic backgrounds of the sources but the set of positions they are attracted to.⁴ Instead we consider that the content of the testimony given by each source is in some way distinct. We capture this by considering a classification of the nodes in V into viewpoints. Let T denote the set of viewpoints.⁵ Then this classification is a function $C' : V \rightarrow \mathcal{P}(T)$, where $\mathcal{P}(T)$ denotes the powerset of T , that is, the set containing all subsets of T . The diversity of the sources that testify to a given node is then simply the number of viewpoints covered by all nodes that testify to v :

$$D' = | \bigcup_{(u,v) \in E} C'(u) |$$

We explicitly choose to consider all sources in this equation, whether or not they are among the independent sources that contribute to the m and k parameters. In particular, if two sources communicate with one another via a path shorter than m but occupy different viewpoints, then these still add to the diversity of information that v receives. The above assumes that the channels of communication are not filtered: if your source has a certain viewpoint, then they will also communicate this to you. To accommodate for modes of communication where certain edges in the network are not used to convey all viewpoints of the source, we might strengthen our above definition, by instead classifying the edges. That is, $C : E \rightarrow \mathcal{P}(T)$ is a function that now takes an edge as input. We then use the following definition of D for a node v :

$$D = | \bigcup_{(u,v) \in E} C(u,v) |$$

⁴ Of course, we recognize that demographic diversity and viewpoint diversity are likely to be confounded in the real world.

⁵ We use T rather than V because the mathematical procedure used to calculate this value is often referred to as *topic modeling*. However, topic modeling can model something closer to viewpoint, not merely topics. Since our data set is focused on the single topic of vaccine safety, our use of topic modeling captures diversity of viewpoint in terms of whether a source testifies that vaccines are safe or not.

Note that this model is strictly stronger than the definition for D' : if we have the function C' , we can define $C(u, v) = C'(u)$, that is, the edges concern all viewpoints of the source, in which case $D = D'$.

We can now quantify the epistemic position of v , by considering its structural position and the diversity of information to which it has access via testimony. We define this as:

$$\pi = D \times S$$

A node's epistemic position is thus the product of the number of different viewpoints among its sources (D) and its structural position (S), which is itself the maximum product of m and k . In this way, we condense three of the essential components of social network-epistemology (multiple sources, independent sources, diverse sources) into a single number.

3.2 Epistemic security and vulnerability

To get a full picture of the epistemic health of an agent in the context of her larger epistemic network, it is important to go beyond quantifying how strong her epistemic position currently is; it is also necessary to quantify how resilient her epistemic position is to change or how much change is needed to improve her current position. With epistemic security, we aim to identify a metric to calculate exactly how robust v 's epistemic position is. We want to be able to say how easily v 's epistemic position could be made better or worse. There are two sides to this: (1) opportunities for v to improve its own position and (2) vulnerabilities posed either by v 's own errors in modifying the network or by others' interventions on the network. Epistemic security is thus a modal notion. It refers not to the actual epistemic position of v but to the set of all v 's possible epistemic positions given a small number ϵ of additions or deletions of edges. In other words, we want to say how easily minor modifications to the network structure could improve or worsen v 's epistemic position. This sort of analysis is similar to the construction of stable equilibria in game theory (Kohlberg & Mertens 1986), where an equilibrium is expected to be robust to small perturbations in players' strategies.

In order to assess epistemic security, it is essential to identify which actions can be undertaken on the structure of the network. Improvements to v 's position can only be achieved by increasing the number of independent sources (k), increasing their level of independence (m), or increasing the diversity of sources (D). Unless v is some sort of epistemic dictator who can determine who talks to whom (e.g., by forbidding other nodes in the network from communicating), increasing the independence among its sources is not a realistic option. For this reason, we assume that v can only influence its direct connections; v only has the power to add edges from other nodes to v , not to add or delete edges between pairs of nodes not including v . Adding a new source may increase v 's structural position S — assuming the new source is independent enough from the others to which v is already connected — or increase diversity D .

As one may expect, cutting off testimony from one's current sources does not have the potential to increase the epistemic position of v .⁶ In the real world, this would mean either refusing even to hear testimony from a source or assigning no credence to a source. We can thus formulate opportunity $O(\varepsilon)$ as the maximum improvement in epistemic standing achievable by adding ε edges from other nodes to v . More formally, let $\pi(0)$ be v 's initial epistemic standing and $\pi(\varepsilon)$ be v 's best possible epistemic standing after adding up to ε edges to v :

$$O(\varepsilon) = \frac{\pi(\varepsilon) - \pi(0)}{\varepsilon}$$

As we pointed out above, finding and listening to new sources takes effort, so we assume that ε is small.

Whereas connecting to new sources may improve v 's epistemic standing, other changes to the network may worsen v 's epistemic standing. Such epistemic threats can manifest in a number of ways. First, the addition of an edge in the network could decrease the independence of sources, thereby reducing the epistemic position of v . Such a change may be the unfortunate side-effect of well-intentioned epistemic actions (e.g., someone trying to improve their own epistemic position), but it may also result from malevolent actions (e.g., people starting to share information to amplify a message, perhaps with the intent to deceive or mislead). Second, removing edges from sources that currently testify to v can decrease v 's epistemic position (note that this is not the case with edges not directly pointing to v). This is a case where someone who previously spoke to v now no longer speaks to v , directly decreasing v 's epistemic position.

We can thus formulate vulnerability or threat $T(\zeta)$ as the maximum decrement in epistemic standing achievable by adding or removing ζ edges anywhere in the network. More formally, let $\pi(0)$ be v 's initial epistemic standing and $\pi(\zeta)$ be v 's worst possible epistemic standing after adding or removing up to ζ edges anywhere in the network:

$$T(\zeta) = \frac{\pi(0) - \pi(\zeta)}{\zeta}$$

As before, adding and removing edges takes effort, so we assume that ζ is small.

⁶ Of course, if the node being cut off is highly unreliable, cutting it off might still be epistemically beneficial, but in this paper we are concerned with the opportunities and vulnerabilities embedded in the *structure* of epistemic networks, not with assessing the reliability of particular sources. As we mention above, reliability is already well-addressed in the literature.

Interestingly, the mode of communication has an immediate impact on what threats are realistic. For example, in personal testimony, a source of information can more directly control who to testify to and who to keep information from. But in an open, online platform such as Twitter, the source can only decide to send or not to send the message in general, not who reads the testimony or receives the information.⁷

Putting opportunities and threats into the same formula, we can capture epistemic security as:

$$\sigma(\varepsilon, \zeta) = O(\varepsilon) / T(\zeta)$$

When epistemic security is greater than 1, opportunities are more salient than threats. When epistemic security is less than 1, threats are more salient than opportunities. We note that in the same network, opportunities may be more salient for some values of ε and ζ while threats are more salient for other values of ε and ζ . Note also that by allowing ε and ζ to vary independently, we make it possible to model both relatively benign situations (in which ζ is much less than ε) and situations of epistemic siege (in which ζ is much greater than ε).

3.3 Epistemic assessment

After articulating how to quantify the current epistemic position of a node within a network and how secure such a position is, we are now able to assess the general epistemic well-being of a node. We find that the importance of epistemic security varies depending on the agent's current epistemic position. Consider the scenarios envisaged in Table 1, which we rank in terms of their overall epistemic quality.

Table 1: assessment of epistemic position $\times$ epistemic security

	high epistemic security	low epistemic security
high epistemic position	best	second-best
low epistemic position	second-worst	worst

Starting from the bottom, we observe that someone who finds themselves in a low epistemic position with low epistemic security is epistemically vulnerable and the worst off. This node is poorly poised to gain knowledge or good information from the sources it listens to because those sources are few, not diverse, and not independent. Moreover, it faces more epistemic threats than opportunities. So while the node can receive knowledge from time to time, this will mostly be due to luck concerning the content of the information the source shares and not because the

⁷ A user can designate some tweets or all of their tweets as “protected.” These tweets can only be seen by the user's followers, but not all followers are guaranteed to see them. Thus, even in this case there is a lack of control over who of the user's followers reads or receives the testimony.

network structure is conducive to gaining knowledge. Even worse, the structure is such that there is a high risk of becoming even more poorly off. Next, someone with low position but high security is not currently well-positioned to gain knowledge from their sources for the same reasons as the worst off, but with some effort they could quickly improve their situation. Third, someone with high position and low security is currently well-poised to learn from their sources. However, they need to be mindful of changes to the structure of their network, as threats are more salient than opportunities. Finally, someone with high position and high security is in the catbird seat. They have multiple, independent, diverse sources, and further changes to the structure of their network are more liable to improve it than not. In the next section, we use a case study to explore this matrix and assess the networked components that influence echo chambers, filter bubbles, and group polarization, as well as positive epistemic health in a real social network.

4 Measuring Epistemic Position in a Real Online Social Network

As we mentioned above, it would be possible to use simulated data and networks in conjunction with agent-based models to address some of the concerns motivating this paper. However, agent-based models build in several assumptions about belief revision and adoption that may not reflect real world belief formation, especially on social media platforms. This is important because even weak assumptions about network structure may turn out to be false. For example, Masterton & Olsson (2018) showed that networks built up through preferential attachment can still deliver the wisdom of crowds, but Sullivan et al. (2019) found empirically that misinformation and disinformation remain rife in real social epistemic networks.

More than this, we are not so much concerned with how echo-chambers form or what norms people use to quantify testimony or pass on information, but how to measure epistemic position *given an existing network*. Thus, to ensure ecological validity and guarantee that our object of study is relevant to contemporary concerns about the epistemic vulnerabilities associated with filter bubbles, echo chambers, and group polarization, we prefer to use data from a real epistemic network.

Thus far, we have argued that there are several key network structures that are relevant to diagnosing whether someone is doing well epistemically within a social network. Furthermore, we have suggested that these network structures can help us to diagnose when someone is in a problematic filter bubble, along with whether the surrounding network affords them more epistemic threats than opportunities for improving their position. Our approach is different from other approaches found in computer science (Kempe et. al 2003; Li et al. 2018), social science (Dunn et al. 2015, 2017; Qiu et al. 2017) and philosophy Zollman (2012). All these aforementioned approaches conduct analysis on the whole network level and look at how

knowledge can spread throughout the entire network. Our approach focuses on the epistemic standing of each individual node within the network and thus can address the problems of filter bubbles on an individual level, not merely on a group level.

To demonstrate the usefulness of this approach, we apply it to a real use case involving a testimonial network of discussions of vaccine safety on Twitter. This platform gives us access to a large data set of people sharing information. Using this data, we are able to investigate the prevalence of epistemic vulnerabilities. In particular, by looking for epistemic observers in the network we show that discussions of vaccine safety on Twitter involve a high proportion of individuals who are highly vulnerable. Furthermore, the discussion is highly polarized, leaving the diversity metric quite small.

4.1 Data collection and cleaning

We conducted a search query on the Twitter stream API that ran from March 5, 2017, to March 11, 2017. We searched for English language tweets that used hashtags and text strings such as #vaxxed, #vaccineswork, #vaccinesafety, ‘vaccine’, and ‘antivax’. In addition, we collected tweets that were from, to, or mentioned specific users, such as @realnaturalnews and @CDCgov. The full search query can be found in the appendix. Data collected through the Twitter API contains several data points about each tweet, including the text of a tweet, whether it is a retweet or reply, how many likes and followers the user has, and sometimes the geographic location the tweet was made from, among (many) other data points. Our focus was on building a retweet network since we are interested in assessing epistemic position and epistemic security concerning what sources people are listening to. The search resulted in 60,230 tweets from 36,390 users.

Almost all the accounts in our raw dataset are only minimally connected to the discourse on vaccine safety. So we isolated the core of the network in order to avoid imposing artificial boundaries and to select the nodes that participate in ongoing discourse about vaccine safety. The core network is formed by identifying a set of actors and the retweet paths between these actors in the following way. Starting with the raw data, we repeatedly remove nodes (accounts) if they are only minimally connected. A node counts as minimally connected if the sum of its in-degree and out-degree is 0 or 1. In other words, we eliminate users who published one vaccine-related tweet that was retweeted exactly 0 or 1 times (and didn’t retweet any other tweets about vaccines), as well as users who retweeted exactly 0 or 1 tweets about vaccine safety (and didn’t publish any tweets about vaccines that were retweeted by others). This is a sensible approach because it eliminates nodes that do not qualify as observers in the first round (because their associated k parameters are 0 or 1).

We then repeat this process in stages until zero accounts are removed in a stage, indicating that the core of the network is all that remains. In later rounds, this method may eliminate nodes that

would have been observers in the full original network — but again, this would have been based only on nodes that are on the fringes of the discussion and thus do not structurally contribute to it. In other words, all actors removed from the original network were never engaged in the conversations of the core directly. They only propagated the conversation of the core outwards in the unfiltered network, or supplied input to the core via proxy.

After removing these actors, we have multiple networks consisting of only the most prominent actors. We designate the largest of these networks the core network, as the conversation in this network should be the most interesting to apply our methods of identifying well-placed epistemic agents to. The result of the filtering left us with a network of 185 nodes, or individual users.

4.2 Computing epistemic observers and diversity

One of the characteristics of online epistemic vulnerability embodied in filter bubbles, echo chambers, and group polarization is that people are only listening to sources that share a particular worldview. Looking for individuals in a network that are drawing from independent sources enables us to identify users that are drawing from less densely interconnected sources and drawing on independent sources of evidence. Thus, our primary goal here is to see how many well-placed observers there are in a controversial discussion on Twitter.

To determine whether a node v is a (m,k) -observer in the core network N we create a new network N' on the set of nodes S that have an edge directed to v , that is, the nodes that supply information to v . For each pair of nodes (x,y) in S that are not m -independent in N with respect to v , we add an edge from x to y in N' . Determining whether v is an (m,k) -observer now reduces to determining whether there is an independent set of size k in N' . Since N' is comparatively small, we can use brute-force techniques to answer this question.

This takes care of independence, but we also need to take into account the diversity of each node's sources in terms of the content of the sources' testimony. Thus, we also sought to compute D within our vaccine network. Computing D is straightforward once we have the classification of the edges. To do so, we use latent dirichlet allocation (LDA — Blei et al 2003) topic modelling on the collection of tweets, and define $C(u,v)$ as the union of the most dominant topic/viewpoint per tweet that is retweeting along edge (u,v) . We chose to only analyze the tweets within our retweet network. While it is quite possible that users are reading other tweets from the sources they retweet, we have no way of knowing which tweets a given user reads. However, we can safely assume that users who retweet something have read the tweet they retweeted. In this way we ensure that the content of the tweets were sources for the user. This gives us a conservative estimate of the number of viewpoints that a node v is informed on by its sources.

The LDA topic model reveals three main kinds of tweets — a predominately pro-vaccination set of tweets, a predominantly anti-vaccination set of tweets, and a set that represents alternative

perspectives. The latter set includes viewpoints associated with alternative medicine and anti-establishment politics. There were also more nuanced viewpoint trends present in the data. For example, within the pro-vaccination tweets some users present a positive case for societal vaccination, while others present a negative argument against the viewpoints of those who are anti-vaccination (in other words, they are anti-anti-vaccination). This finding reveals that there are two aspects of the diversity worth considering when evaluating epistemic position: diversity of conclusion and diversity of reasons for a given conclusion. In other words, someone could hear a diversity of viewpoints regarding *whether* vaccines are safe, or they could hear a diversity of reasons *why* vaccines are safe. No doubt both kinds of diversity are important. For the purposes of this paper we focus on *diversity of conclusion*, and leave how best to incorporate diversity of reasons within the same conclusion to further research.

4.3 Data analysis

Our analysis of the vaccine safety debate on Twitter shows that the conversation is polarized and the majority of nodes occupy a low epistemic position. Our analysis also shows that it is possible to identify what look like problematic epistemic vulnerabilities using network analysis based on the epistemic position of individual nodes in the network.

We suggest that a node counts as having *high* epistemic position if its epistemic position is at least as good as a (3,3)-observer that hears from two distinct viewpoints. This means that a node has a high epistemic position so long as $\pi \geq 18$. This threshold is motivated partly due to the size of our core network and the high polarity of users engaging with a single viewpoint, in addition to the considerations discussed in the previous sections. Figure 6 shows one of the nodes in our core network that is a (4,4)-observer and thus has a high value of S . The distribution of the values of π and the contributing factors (S and D) are shown in Figure 7. Of the 185 nodes under evaluation, only 28 (15%) reach or exceed the threshold for having a high epistemic position. The rest we deem to have low epistemic position. Furthermore there are 64 nodes that have an S and thus π value of 0, which means that they are not even a (1,2)-observer in the core network. Therefore, roughly 34% of the nodes have a dangerously low epistemic position.

Figure 6

The red node is a (4,4)-observer. The blue nodes are all pairwise 4-independent with respect to the red node.

Figure 7

Distribution of π and the contribution factors. Each bar represents nodes in the core network with the same values of D , S and π . The width of the bar indicates the number of nodes. The height of the bar encodes the S value, the color of the bar encodes the D value. The π values are shown as a line above each bar.

These results suggest that the debate on Twitter concerning vaccine safety is rife with epistemic vulnerability. We have shown this by using individualized network analysis. Individual nodes are unlikely to hear from multiple independent sources and are likely to only hear from sources that accept the same conclusion concerning vaccinations. That said, there are limits to the type of network that we are considering with this analysis. Since we are analyzing a retweet network, there are undoubtedly other pieces of testimony a user on Twitter is exposed to and reads concerning vaccine safety. Thus, the retweet network is not a perfect analog to a testimonial network. Users may only retweet items they agree with and not share all items they critically engage with. For example, a pro-vaccine user might not retweet any information coming from a con-vaccine account, but may still have read and critically assessed this information, thereby having a better epistemic position than the model indicates. However, this analysis is still quite telling. It shows that the information that individuals find worth sharing is the information within

a single worldview from sources that are not independent. This alone can lead to the creation of filter bubbles and is indicative of epistemic vulnerability.

Lastly, as discussed in the previous sections, in order to get a full picture of the epistemic position of agents within a network, we should not only look to the current position a node has, but also how secure the node is to changes in the structure of the network. It is possible that while the Twitter network that we analyzed is currently polarized with few well-placed observers, there could be more epistemic opportunities than threats. Applying our metrics of epistemic security to these real life uses cases could thus help to provide a solution to filter bubbles by exposing what actions individuals should take within a network to improve their position, and how malleable the filter bubbles actually are. We leave this further investigation for future research. In particular, it opens up the door to collaboration with computer scientists to be able to algorithmically identify which connections individuals nodes can efficiently make to improve their position and to make progress on combating filter bubbles.

5 Conclusion and Future Research

In this paper, we introduced a social network approach to the epistemology of filter bubbles and social epistemology more generally. According to this approach, the primary locus of epistemic evaluation is epistemic network structures. In particular, we argued for the importance of agents' drawing on multiple, independent, diverse sources. The (m,k) -observer construct is especially fruitful for evaluating individual nodes within a network in terms of their current epistemic position. Furthermore, we introduced the concept of epistemic security that seeks to locate possible epistemic threats to a given node's position and possible opportunities for its position to be improved. The latter is especially fruitful for addressing epistemic vulnerabilities such as filter bubbles and echo chambers. Importantly, our approach does this based on the evaluation of individual nodes within the network and not merely at the whole-network level. We showed the usefulness of this approach by applying it to an empirical use case: discussions of vaccine safety on Twitter.

The framework and analytical tools developed here point toward two directions for future research. First, from the formal and technical point of view, social epistemologists could collaborate with computer scientists and mathematicians to build on this work. Of particular interest would be an algorithm that calculates the epistemic security of each nodes in the network and outputs recommendations for efficient improvements (as well as warnings about changes to avoid making).

Second, this work points in normative and philosophical directions. We have only scratched the surface of the social network approach to epistemology. For example, in this paper we were primarily concerned with the receivers of testimony, but the sources themselves are also a main component in the epistemology of testimony. There is a direct analogue to the (m,k) -observer

where instead of a node receiving testimony a node gives testimony. Such (m,k) -broadcasters would be nodes with high epistemic power. This phenomenon has attracted the interest of philosophers (Fricker 2007; Alfano & Robinson 2017) but has not been adequately formalized. Our approach is well suited to do just that.

Moreover, we have largely glossed over how the social network approach points to certain epistemic norms or epistemic virtues that hearers (and givers) of testimony in social networks (including online networks) should develop (Heersmink forthcoming). One clear norm is that agents should seek out independent and diverse sources, but this is not especially surprising. We also briefly suggested that a promising complement to a social network approach is to think in terms of social epistemic virtues. Agents would do well to develop monitoring virtues concerning the structure of their networks, aggregating virtues concerning how to weight testimony of sources given strengths and flaws of the network structure, and restructuring virtues concerning how to improve and guard their epistemic position in light of epistemic threats and opportunities. There is great potential for further research regarding the nature of these virtues. For example, these virtues may be either self-regarding (making oneself better off epistemically) or other-regarding (making others better off epistemically). They may include the disposition to gossip (Alfano & Robinson 2017) or leak information as a whistleblower (DesAutels 2009). Furthermore, there are likely to be unavoidable tradeoffs between cutting yourself off from individuals to protect others' epistemic position and keeping your own strong position, indicating a sort of epistemic collective action problem. Thus, social network-epistemology can serve as a promising framework for social epistemology that is both thoroughly philosophical and interdisciplinary in its approach and application.

References

- Alexander, J. M., Himmelreich, J., & Thompson, C. (2015). Epistemic landscapes, optimal search, and the division of cognitive labor. *Philosophy of Science*, 82(3), 424-453.
- Alfano, M. (2013). *Character as Moral Fiction*. Cambridge University Press.
- Alfano, M. (2017). The topology of communities of trust. *Russian Sociological Review*, 15(4): 30-56.
- Alfano, M. & Robinson, B. (2017). Gossip as a burdened virtue. *Ethical Theory and Moral Practice*, 20: 473-82.
- Baehr, J. (2011). *The Inquiring Mind: On Intellectual Virtues and Virtue Epistemology*. Oxford University Press.
- Blei, D., Ng, A., Jordan, M. (2003). Latent dirichlet allocation. *Journal of Machine Learning Research*, 3(4-5): 993-1022.
- Chisholm, R. (1957). *Perceiving: A Philosophical Study*. Cornell University Press.
- Coady, C. (1995). *Testimony: A Philosophical Study*. Clarendon Press.
- Craig, E. (1999). *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford University Press.
- DesAutels, P. (2009). Resisting organizational power. In L. Tessman (ed.), *Feminist Ethics and Social and Political Philosophy: Theorizing the Non-Ideal*. Springer.
- Dunn, A., Leask, J., Zhou, X., Mandl, K., & Coiera, E. (2015). Associations between exposure to and expression of negative opinions about human papillomavirus vaccines on social media: An observational Study. *Journal of Medical Internet Research*, 17(6): e144.
- Dunn, A., Surian, D., Leask, J., Dey, A., Mandl, K., & Coiera, E. (2017). Mapping information exposure on social media to explain differences in HPV vaccine coverage in the United States. *Vaccine*, 35: 3033-40.
- Fricker, M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press.
- Li, Y., Fan, J., Wang, Y. and Tan, K.L (2018) "Influence maximization on social graphs: A survey." *IEEE Transactions on Knowledge and Data Engineering*
- Heersmink, R. (forthcoming). A virtue epistemology of the Internet: Search engines, intellectual virtues, and education. *Social Epistemology*.
- Kelp, C. (2016). Assertion: A function first account. *Nous*.
- Kempe, D & Kleinberg, J & É. Tardos (2003). *Maximizing the spread of influence through a social network*. Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining. 137-146
- Kohlberg, E. & Mertens, J.-F. (1986). On the strategic stability of equilibria. *Econometrica*. 54(5): 1003-1037.
- Levy, N. (2017). The bad news about fake news. *Social Epistemology Review and Reply Collective*, 6(8): 20-36.

- Lynch, M. (2016). *The Internet of Us: Knowing More and Understanding Less in the Age of Big Data*. Liveright.
- Masterton, G., Olsson, E. J. & Angere, S. (2016). Linking as voting: How the Condorcet jury theorem in political science is relevant to webometrics. *Scientometrics*. Akademiai Kiado.
- Masterton, G. & Olsson, E. J. (2017). From impact to importance: the current state of the wisdom-of-crowds justification of link-based ranking algorithms. *Philosophy and Technology*. Springer.
- Masterton, G. & Olsson, E. J. (2018). PageRank's ability to track webpage quality: Reconciling Google's wisdom-of-crowds justification with the scale-free structure of the web. *Helvion*, 4(11): 1-34.
- Milgram, S. (1967). The small world problem. *Psychology Today*, 1(1): 60-67.
- Muldoon, R., & Weisberg, M. (2011). Robustness and idealization in models of cognitive labor. *Synthese*, 183(2), 161-174.
- Nguyen, C. T. (forthcoming). Cognitive islands and runaway echo chambers: Problems for epistemic dependence on experts. *Synthese*.
- Pallavicini, J., Hallsson, B., & Kappel, K. (forthcoming). Polarization in groups of Bayesian agents. *Synthese*.
- Pariser, E. (2011). *The Filter Bubble: What the Internet is Hiding from You*. New York: Penguin Press.
- Qiu, X., Oliveira, D., Shirazi, A. S., Flammini, A., & Menczer, F. (2017). Limited individual attention and online virality of low-quality information. *Nature Human Behavior*. doi = 10.1038/s41562-017-0132.
- Rini, R. (2017). Fake news and partisan epistemology. *Kennedy Institute of Ethics Journal*, 27(S2): 43-64.
- Roberts, D. (2017a, November 2). America is facing an epistemic crisis. *Vox Media*. Url = <www.vox.com/policy-and-politics/2017/11/2/16588964/america-epistemic-crisis>. Accessed 17 February 2018.
- Roberts, D. (2017b, May 19). Donald Trump and the rise of tribal epistemology. *Vox Media*. Url = <www.vox.com/policy-and-politics/2017/3/22/14762030/donald-trump-tribal-epistemology>. Accessed 17 February 2018.
- 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*, vol. 3. Oxford University Press.
- Sunstein, C. (2017). *#Republic: Divided Democracy in the Age of Social Media*. Princeton University Press.
- Surowiecki, J. (2005). *The Wisdom of Crowds*. Anchor.
- White, R. (2006). Problems for dogmatism. *Philosophical Studies*, 131(3): 525-57.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press.

- Zagzebski, L. (1996). *Virtues of the Mind*. Cambridge University Press.
- Zollman, K. (2012). Network epistemology: Communication in epistemic communities. *Philosophy Compass*, 8: 15-27.
- Zollman, K. (2015). Modeling the social consequences of testimonial norms. *Philosophical Studies*, 172 (9). doi:10.1007/s11098-014-0416-7.

Appendix

Full Twitter Search Query

text: vax, vaxxed, vaccine, vaxsafety, vaccineswork, vaccinesafety, vaccinesrevealed, novax, antivax, vaccination, vaccinations, immunization

hashtags: #vax, #vaxxed, #vaccine, #vaxsafety, #vaccineswork, #vaccinesafety, #vaccinesrevealed, #novax, #antivax, #vaccination, #vaccinations

mentions: @CDCgov, @drwakefield, @realnaturalnews, @drpanmd, @conservabotia, @WHO, @HHSgov, @DrRandPaul, @MicheleBachmann, @kwakzalverij, @rivm, @gezondheidsraad, @RolandPierik, @Gert_van_Dijk, @nvkp_nl, @VaccinatieRaad, @AnthonySc6, @LotusOak, @jelani9

from: @CDCgov, @drwakefield, @realnaturalnews, @drpanmd, @conservabotia, @WHO, @HHSgov, @DrRandPaul, @MicheleBachmann, @kwakzalverij, @rivm, @gezondheidsraad, @RolandPierik, @Gert_van_Dijk, @nvkp_nl, @VaccinatieRaad, @AnthonySc6, @LotusOak, @jelani9

to: @CDCgov, @drwakefield, @realnaturalnews, @drpanmd, @conservabotia, @WHO, @HHSgov, @DrRandPaul, @MicheleBachmann, @kwakzalverij, @rivm, @gezondheidsraad, @RolandPierik, @Gert_van_Dijk, @nvkp_nl, @VaccinatieRaad, @AnthonySc6, @LotusOak, @jelani9

exclude: @realdonaldtrump, @RandPaul

Acknowledgments

This paper is the winner of the PERiTIA Prize of the 2019 IJPS Robert Papazian essay competition on the subject of Vulnerability and Trust. The PERiTIA Prize is funded by UCD Centre for Ethics in Public Life. The project Policy, Expertise and Trust in Action — PERiTIA has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 870883.