

Can Real Social Epistemic Networks Deliver the Wisdom of Crowds?

forthcoming in *Oxford Studies in Experimental Philosophy* (Vol. 3). Eds. Lombrozo, T., Knobe, J., & Nichols, S. OUP

Emily Sullivan, Delft University of Technology

Max Sondag, TU Eindhoven

Ignaz Rutter, University of Passau

Wouter Meulemans, TU Eindhoven

Scott Cunningham, Delft University of Technology

Bettina Speckmann, TU Eindhoven

Mark Alfano, Delft University of Technology & Australian Catholic University

Abstract: In this paper, we explain and showcase the promising methodology of testimonial network analysis and visualization for experimental epistemology, arguing that it can be used to gain insights and answer philosophical questions in social epistemology. Our use case is the epistemic community that discusses vaccine safety primarily in English on Twitter. In two studies, we show, using both statistical analysis and exploratory data visualization, that there is almost no neutral or ambivalent discussion of vaccine safety on Twitter. Roughly half the accounts engaging with this topic are pro-vaccine, while the other half are con-vaccine. We also show that these two camps rarely engage with one another, and that the con-vaccine camp has greater epistemic reach and receptivity than the pro-vaccine camp. In light of these findings, we question whether testimonial networks as they are currently constituted on popular fora such as Twitter are living up to their promise of delivering the wisdom of crowds. We conclude by pointing to directions for further research in digital social epistemology.

Keywords: social epistemology, PageRank, network analysis, network visualization

Word count: 9273

Can Real Social Epistemic Networks Deliver the Wisdom of Crowds?

1. Introduction¹

One of the central ways people gain knowledge is through the testimony of others. Despite this, 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.

Only recently have philosophers moved beyond the hearer-speaker dyad. For example, in epistemology some have considered what it means for *groups* to testify (Tollefsen 2007; Lackey 2015, 2018). In the philosophy of science, others have started to explore which network structures are conducive to sharing knowledge within ideal communities of scientific researchers (Holman and Bruner 2015; Rosenstock et al. 2016; Zollman 2007, 2010). Still lacking is an 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. 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. Minimizing geodesics is an epistemic strategy that can be formally modelled, but do modern advances suggest more epistemically promising strategies? In this paper, we seek to make progress on this front. In particular, we are interested in identifying good sources and receivers of information based on the surrounding network structures of their epistemic community.

Consider a typical search engine. After a user enters a query, a series of links to presumed answers to the query are ranked and displayed to the user. There are several underlying mechanisms that give rise to the resulting list. Some of these mechanisms are inherently networked and constitute a link-based ranking. It is the interconnection among the links that helps determine which ones are displayed first. PageRank, developed by one of Google's founders, is one of these networked mechanisms (Brin & Page 1998). Importantly, one of the motivations behind PageRank is its epistemic benefits. The underlying assumption is that groups of individuals can be more reliable and converge on true answers more so than individuals acting in isolation. Following Masterson et al. (2017), we refer to this assumption as the

¹ This publication was supported by a subaward agreement from the University of Connecticut with funds provided by Grant No. 58942 from John Templeton Foundation. Its contents are solely the responsibility of the authors and do not necessarily represent the official views of UConn or John Templeton Foundation.

wisdom-of-crowds hypothesis.² One of the distinguishing features of PageRank and algorithms like it is that, instead of determining rank solely based on the number of incoming links, it accounts for the *importance* of those links. For example, if a well-connected news source links to a particular webpage, that link is given higher weight than a link to the same webpage from an obscure blog. Such a ranking system has the potential to track epistemic importance. In fact, Masterton et al. (2016, 2017) have shown that PageRank can deliver epistemic benefits and justify a wisdom of crowds thesis, albeit in artificial or ideal cases.

There are many similarities between a large testimonial network and a network of webpages. Like webpages, individuals share novel information. Individuals also decide to share information with others by posting or linking to it. Thus, we can ask whether ranking individuals based on formal properties of their information-sharing network is truth-tracking or justifies a wisdom of crowds thesis of testimony. More precisely, can PageRank identify good sources of information and well-informed receivers of information in a testimonial network?

In order to shed light on this question and test the related wisdom-of-crowds hypothesis, we conducted two case studies involving non-ideal epistemic communities. Both studies involve discussions of vaccine safety on Twitter. The results of these studies indicate that, in non-ideal cases, PageRank is sufficient neither for identifying individuals who are good sources of information nor for identifying individuals who are likely to be well-informed receivers. We conclude by arguing that, in testimonial networks, considerations of source independence and source diversity are needed to give a fuller picture of the network-structures conducive to knowledge transmission.

In what follows, we first expound on the nature of networks, especially how they can help answer epistemological questions by harnessing the wisdom of crowds. In sections 3 and 4 we describe our two case studies. We end with a discussion of our results and point to areas of future research for both philosophers and graph theorists.

2. Networks, epistemology, and PageRank

A network or graph is an abstract mathematical object. The simplest networks represent nothing beyond the bare presence of nodes and whether, for each pair of nodes, there exists a connection between them. Connections between nodes can be directed or undirected. In an undirected network the connections between the nodes (edges) represent a bidirectional relationship. A network of connected cities is an example of a undirected network: the cities are the nodes and the highways or interstates are the edges that connect them. In a *directed* network, by contrast, each edge is unidirectional. A network showing the flow of money to politicians from campaign funders is an example of a directed network. Individual politicians and campaign contributors are the nodes, and the flow of money is represented by directed edges connecting the nodes. In this toy example, the direction of the edges is always *from* the campaign contributors *to* the politicians. Moreover, edges can be *weighted*. The weight of an edge indicates the varying strength of an edge. In a campaign contributions network, edge weights could represent how much money each contributor gives to a particular politician. The research question and the target phenomenon dictate how modelers determine which network parameters, such as edge weights and directions, to use.

² Wisdom-of-crowds hypotheses are not unique to the web. They also underlie arguments for the epistemic value of democracy (List and Goodin 2001), the epistemic value of juries (Hedden 2017), and crowd wisdom more generally (Surowiecki 2004).

As mentioned above, in this paper we are interested in identifying good sources and receivers of information based on the surrounding structures of their testimonial network. In order to successfully model such a testimonial community we need to get clear on what the appropriate network structure is for our target phenomenon. In the remainder of this section, we discuss important testimonial concepts (2.1), their operationalization in a network (2.2), and the wisdom-of-crowds hypothesis regarding how to rank individuals within a network (2.3).³ In what follows, “testimonial community” refers to the target phenomenon and “testimonial network” refers to the network-model of the community.

2.1 Testimonial concepts⁴

In a community, an epistemic agent can play one or more roles. They can be a primary *source* of knowledge for others. Alternatively, they can *receive* knowledge from others. Additionally, they can act as a *conduit* who passes along knowledge from a primary source (or another conduit) to a receiver (or another conduit). An agent is a powerful epistemic source if their messages are likely to be transmitted to others in the network. Likewise, an agent is a well-positioned receiver if messages in the network are likely to be transmitted to them. And an agent is a powerful conduit if they have the ability to control whether and which messages originated by others propagate through the network. In this paper, we focus primarily on sources and receivers, leaving conduits to future research.

In each case of testimonial exchange, there is a single *testifies-to* relation that originates from the source to the receiver. The *testifies-to* relation between one source and one receiver has been explored in great depth in social epistemology. Considerations of trust, reliability, social power, and expertise are all present in a single case of testimony (Fricker 2007; Medina 2013). However, most discussions focus on whether one isolated source is trustworthy, reliable, or has power over the receiver, and how these considerations reflect on epistemic norms or virtues. Importantly, there are macro-level aspects to testimonial trust, power, and source reliability that are sure to be ignored, if we focus only on *one* source and *one* receiver. Consider, for example, epistemic power. A source of information can have great epistemic power not only in virtue of her social identity (i.e., occupying a role society deems valuable or important) but also simply in virtue of her *epistemic reach*: how many individuals in the community will or are likely to receive her testimony first-, second-, or third-hand. This latter sense of epistemic power captures how chains of information -- *epistemic paths* -- propagate through the community. Such paths of testimony can be short, as in the dyadic case of a single source transmitting to a single receiver. They can also be indefinitely long, as when a source transmits through a string of conduits who together manage to convey the information to a receiver.

How much epistemic reach certain sources have is no doubt relevant to the epistemic norms and virtues surrounding testimonial knowledge. For instance, it seems plausible that someone with great reach should be more cautious in what they say, compared to someone with very limited reach. Likewise, it seems plausible that part of what it takes to embody the virtue of

³ Broadly, our research follows the workflow outlined in Alfano & Higgins (forthcoming).

⁴ We begin by conceptualizing the research question and all variables involved because inadequate, imprecise, and ambiguous conceptualization is arguably one of the root causes of the replication crisis currently racking psychology (Lurquin & Miyake 2017).

intellectual humility is to be disposed to notice gaps in one's testimonial network and make efforts to fill them, when those gaps are sufficiently large and important enough.

Questions of epistemic power go beyond epistemic sources. Epistemic receivers could also be well-placed in a very direct way, if they receive testimony from a very large number of sources. However, most real epistemic communities are much sparser. In such communities, an agent might be a well-placed receiver, not due to a large number of sources, but simply in virtue of the fact that a large number of messages in the community eventually make their way to her -- sometimes through short epistemic paths and sometimes through much longer ones.

These points, about the relation between epistemic power, on the one hand, and the source and receiver roles, on the other hand, hold *mutatis mutandis* for other epistemic phenomena, such as trust and reliability.⁵ Thus, there are macro-level influences that are operative in large-scale epistemic communities that are not operative in a dyadic community. Attending to these phenomena, using the concepts and distinctions discussed here, makes it possible and worthwhile to explore what Alfano (2016) calls the *topology* of communities of epistemic trust, both from a normative and from a descriptive point of view.

In order to explore these macro-level dimensions of the testifies-to relation, it is paramount to investigate larger, more structured epistemic communities than have typically been studied by social epistemologists. Hence, in this paper we model a non-ideal epistemic community. Such a model will help to inform us how we should think about macro-level aspects of epistemic trust and power, as well as what makes for a good source or receiver of information in an interconnected epistemic community.

2.2 Testimonial communities as directed networks

As discussed above, the designation of nodes, edges, edge weights, and network directedness are central determinations to make when modeling a target phenomenon as a network, such as an epistemic community, as a network.⁶ In our testimonial network the nodes represent epistemic sources and receivers of testimony. The edges represent the *testifies-to* relation discussed above. However, in order to focus on the macro-level structural influences of testimonial networks we abstract away any details of trust or weighted reliability in the testifies-to relation. Thus, the weight of edges simply represents the number of messages from a source to a receiver. By first analyzing these simple structural relations, we will then be able to more clearly see how trust and reliability could function within macro-level structures.

Next, we need to determine whether a testimonial epistemic community is best modeled as a directed or undirected network. The information-sharing network models deployed in the philosophy of science are largely undirected models (Holman and Bruner 2015; Rosenstock et al. 2016; Zollman 2007, 2010). The use case for these models is how information is shared within a scientific community and how groups of researchers converge on a single plan of action, such as prescribing one medication over another. The acting assumption in these models is that testimonial exchanges are *symmetric*: agents jointly trade information with each other.⁷ It is fair

⁵ See Alfano & Huijts (forthcoming) for a discussion on how trust manifests in non-dyadic epistemic networks.

⁶ This is not meant to be an exhaustive list. For example, dynamic network-models that involve agent-based simulation include node parameters that simulate agent decisions that update throughout the simulated run. For an information sharing example see Holman and Bruner (2015), Rosenstock et al. (2016), and Zollman (2007, 2010).

⁷ Holman and Bruner (2015) deploy a model where the testifying relationship is not entirely symmetric. In their agent-based model they simulate that researchers tend to rely on and trust some sources more than others. However,

to say that in these models edges represent a testifies-*with* relation. For any pair of connected agents, each agent provides information to, and receives information from, the other. In contrast, if we consider a testimonial exchange between an expert and a non-expert, symmetry can no longer be presumed. In many cases, the non-expert does not offer testimony in exchange with the expert; the non-expert is merely a receiver of testimony. Thus, when modeling testimonial networks outside of an idealized community of near-peer researchers, where agents have different epistemic goals and projects, it is important to take into account the directional nature of the the testifies-*to* relation. That said, the same individual can take on the source role as well as the receiver role in the larger community. However, each testifies-*to* relation remains an asymmetrical case of sharing information *to* someone who is presumed to lack information, and where the receiver is not presumed to give information back in response.

2.3 (Page)Ranking sources and receivers

Our main goal in this paper is not simply to model a real epistemic community, but to explore how we can identify which individuals in the network are good sources of information by virtue of the surrounding network that links to (or shares) that information, and which individuals in the network are well-informed by virtue of their position and the positions of their sources. PageRank offers a plausible approach.

Recall that PageRank was implemented in web-search technologies to help identify sources that are likely to have epistemic importance. Roughly, a webpage's PageRank is the probability that someone randomly clicking through webpages will end up visiting that page (Brin & Page 1998). The ranking is not based solely on the number of pages that link to it, but also includes the importance of the pages that link to it. Pages with several links "pass on" some of the value of those links to the pages they link to. One perhaps-counterintuitive result of this approach is that some pages with considerably fewer incoming links can be ranked higher than a page with far more links. The rationale is that if important webpages are linking to something, that in itself is an important indicator of the page's merit. For example, if *The New York Times* were to link to another source in its reporting, then that source may have a higher PageRank, even if it is a more obscure page. However, despite the page's obscurity, it does ground the *Times* reporting, and thus the content of the page should seem to weigh heavier than the second-hand reporting from the *Times*. The same intuitions motivate testimonial exchanges. Sources of testimony that attract more trust are more likely sharing epistemically useful information, which is a strong though defeasible reason to give more weight to those sources. Thus, it seems that PageRank houses the theoretical tools not only to track epistemic importance on the web, but also to exploit the wisdom of crowds in testimonial networks (Masterton et al. 2016, 2017).

Furthermore, there are practical benefits to applying link-based rankings to testimonial networks. Algorithms can be designed to increase the epistemic position of those in online testimonial communities, as they have in optimizing search queries. Thus, the time is ripe to ask whether PageRank actually lives up to its wisdom-of-crowds promise in non-ideal testimonial networks. In order to test this, we conducted two case studies consisting of real epistemic communities found on Twitter.

they model this through the dynamics of the agent-based system on how nodes update, not through introducing directed edges.

3. Study 1

In order to test the wisdom-of-crowds hypothesis in a non-ideal network, we needed to identify accessible data sources that we expect to contain a large number of epistemic agents acting as sources and receivers, potentially mediated by conduits. One of the largest social epistemic networks currently in existence is Twitter. Although the quality of discourse on this platform varies greatly, it provides a suitable source for exploring whether PageRank lives up to its promise of delivering the wisdom of crowds for real testimonial communities.

On the Twitter platform, users have the ability to compose new tweets of up to 280 characters; they also can see the tweets published by other users and can interact with these by “liking” them, retweeting them, or replying to them. The retweeting functionality, which is similar to functionalities on other platforms (e.g., “sharing” on Facebook) is especially interesting for our purposes. It enables a user to pass along content (and potentially whatever knowledge could be acquired from that content), making them a conduit. It also indicates (heuristically) that the person read the testimony and thought it was worth sharing. This suggests the testimony was a source of information for the retweeter, whether or not they ultimately accepted that testimony as true.

When user A gets retweeted by user B, that establishes a link *out* from A (and, by the same token, a link *in* to B). The number of links out from a given node is that node’s out-degree; the number of links in to a given node is that node’s in-degree. PageRankIn and PageRankOut are calculated using the PageRank algorithm (Brin & Page 1998) with $d=0.85$ on the network of users and links, and the reverse of this network respectively. We use *high PageRankOut* as a heuristic for identifying a good epistemic *source*; high PageRankOut indicates that the user’s tweets tend to be retweeted to a wide audience. We use *high PageRankIn* as a heuristic for identifying well-placed epistemic *receivers*; high PageRankIn indicates that the user tends to retweet more messages, thus drawing on a wide number of sources. Table 1 shows the epistemic concepts to be modeled along with their operationalizations.

Table 1: definitions and operationalizations of key terms.

Concepts	Network-Definitions	Formal Parameters
well-positioned epistemic source	a node that transmits new information into the network	high PageRankOut
well-positioned epistemic receiver	a node that receives information propagated from elsewhere in the network	high PageRankIn
epistemic path	a path through from a source through zero or more conduits to a receiver	chain of retweets between nodes
epistemic geodesic	the shortest epistemic path between a pair of nodes	shortest chain of retweets between nodes

3.1 Data collection

We conducted a search query on the Twitter stream API that ran from March 5, 2017 to March 11, 2017. We explain how this is done below and provide our code in an appendix. We were interested in seeing a discussion surrounding a controversial topic that would include information that varied along the dimensions of both quality and quantity. For this purpose, we chose to look at discussions of vaccine safety. We searched for 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. These users were identified inductively: researchers searched for sources on both sides of the debate that were prominently involved in ongoing discussions of vaccine safety, then added those sources to a list. The full search query can be found in the appendix. The search resulted in 60,230 tweets from 36,390 users.

3.2 Cleaning and processing the data

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. As discussed above, a retweet network is best-suited to our research questions.

Almost all the 36,390 accounts in our raw dataset are minimally connected to the discourse on vaccine safety. Thus, we proceed by isolating the core of the network in order to avoid imposing artificial boundaries, aid in network visualization, and analyze those who participate in ongoing engagement and discussion. 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). We repeat this process in stages until zero accounts are removed in a stage, indicating that the core of the network is all that remains.

All actors removed from the original network were never engaged in the conversations of the core directly. They 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 important actors. We designate the largest of these networks as our core network. This process filtered out the vast majority of accounts from the raw data, leaving a core network of just 240 nodes.

3.3 Analysis

Our methodology for interpreting a network proceeds along two tracks. The first track is through inferences from the mathematical properties of the network. The second track, visualization, is discussed in the next section.

The measures used to analyze networks are derived from the relational properties of nodes. As we explained above, the basic relational property is a node’s *degree*, or the number of

other nodes it is connected to. In directed networks, we can distinguish between a node's *in-degree* (the number of edges that point to it) and its *out-degree* (the number of edges that point from it). In our retweet network on vaccine safety a node with a high out-degree is one that has been retweeted often. A node with a high in-degree retweets others often. Degree is a local property of nodes: a node with three connections could exist at the periphery of a network or function as a significant bridge or bottleneck in the network, depending on how the rest of the network is structured. In order to better understand the roles played by different nodes, it's therefore helpful to employ more holistic properties. As mentioned above, for this purpose we use PageRankOut and PageRankIn because they informatively summarize holistic information.

As an initial check on the viability of our API query, two of the authors independently hand-coded random samples of tweets both from the unfiltered network (300 tweets) and from the core network (300 tweets). 28.8% of the tweets in the unfiltered network were deemed irrelevant, whereas only 7% of tweets in the core network were deemed irrelevant. Many of the irrelevant tweets were by (or retweets of, replies to, or mentions of) the specific accounts mentioned in section 3.1 above. For this reason, in the second case study, we did not automatically include these accounts in our API query.

In addition, to further establish the viability of our approach, we use latent dirichlet allocation (LDA -- Blei et al 2003) topic modelling on the full collection of tweets, and define $C(u,v)$ as the union of the most dominant topic per tweet that is retweeting along edge (u,v) . 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 topics surrounding alternative medicine and anti-establishment politics, among others. We surmise that this third group was responsible for the higher proportion of irrelevant tweets in the unfiltered network as compared to the core network. There were also more nuanced topic 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).

Next, because the number of nodes in the core network was of a manageable size, we labeled them by hand rather than using automated natural language processing. In line with the semantic broad topics identified by LDA, three independent coders labeled each of the nodes of the core network as being either pro-vaccine, con-vaccine, or neutral. No codebook was used; instead, the coders simply read with an eye to the opinions expressed in the tweets and any urls they linked to. We observed adequate interrater reliability (Fleiss's $\kappa = .734$, $z = 29.7$, $p < .001$). The main source of disagreement was a set of accounts that could not easily be classified as pro-vaccine, con-vaccine, or neutral. Two researchers reviewed cases of disagreement and designated 18 accounts as *irrelevant*; these accounts were removed from the core network, leaving 185 nodes. The average out-degree, in-degree, PageRankOut, and PageRankIn of nodes in each of the three camps is summarized in Table 2. For more detailed plots of the data, we refer to Figure 7 in the appendix.

Table 2: mean degree and PageRank properties of pro-vaccine, neutral, and con-vaccine accounts.

	<i>n</i>	out-degree	in-degree	PageRankOut	PageRankIn
pro-vaccine	91	2.49	2.46	10.4*(10 ⁻⁶)	7.98*(10 ⁻⁶)
neutral	1	1	1	4.35*(10 ⁻⁶)	4.28*(10 ⁻⁶)
con-vaccine	93	8.69	8.73	16.9*(10 ⁻⁶)	13.5*(10 ⁻⁶)

The number of pro- and con-vaccine accounts in the core of the network is nearly identical, and there is just one neutral account. The average out-degree, in-degree, and PageRankOut of con-vaccine accounts is higher, whereas the average pro-vaccine PageRankIn is higher. Because there are so few neutral accounts, we exclude them from significance testing. Kruskal-Wallis tests reveals that differences in out-degree ($\chi^2 = 16.17$, $df = 1$, $p = .001$), in-degree ($\chi^2 = 7.96$, $df = 1$, $p = .046$), and PageRankOut ($\chi^2 = 9.64$, $df = 1$, $p = .02$) are significant, whereas differences in PageRankIn are not ($\chi^2 = 1.80$, $df = 1$, $p = .62$). Together, these results suggest that there are very few neutral accounts tweeting about vaccine safety, and that the reach (out-degree & PageRankOut) of con-vaccine accounts is greater. Regarding receptivity (in-degree & PageRankIn), the results are mixed: the con-vaccine community has significantly greater in-degree but statistically indistinguishable PageRankIn.

3.4 Visualization

In principle, everything that can be learned from a visualization could also be learned by analyzing the bare mathematics of the network itself. However, given the relative strengths and weaknesses of human cognition and perception, it is sometimes (though by no means always) advantageous to employ visualizations.⁸ This is especially acute for analyzing nodes ranked by PageRank or degree. Importantly, visualization is a tool for analysis and understanding.

We first plot the network using a simple force-directed layout as implemented in D3.js.⁹ This gives the impression of two well-separated communities, one more tightly knit than the other, but we should keep in mind that a force-directed layout provides us with no guarantees about perceived and actual distance: nodes that are far apart in the visualization may in fact be fairly close in the network structure. Moreover, it is difficult to correlate PageRank scores with sentiment (being pro- or con-vaccine, etc).

⁸ Visualizations can also be misleading, so it's important to employ best practices as outlined in Munzner (2014).

⁹ We used version 4 of D3.js by Mike Bostock. <http://www.d3js.org>, accessed August 2018.




Figure 2: A force-directed layout of the graph. Nodes represent accounts and arrows represent retweeting.

We therefore alter the layout and use the node positions to communicate these values. As position is the strongest visual variable (Munzner 2014), this allows us to best focus on and assess the values of PageRank and sentiment. Below are visualizations of users' attitudes towards vaccination based on the three independent coders, with all 185 accounts in the core network represented. Position on the x-axis is determined by averaging the sentiment scores of all three raters. We slightly offset these positions horizontally, to get a better sense of the number of accounts and a clearer view of the network: the horizontal position within a column is *not* indicative of strength of opinion. Position on the y-axis represents an account's reach (PageRankOut, Figure 3) or receptivity (PageRankIn, Figure 4). The properties of the most prominent pro and con accounts are summarized in caption boxes.




Figure 3: Active pro-vaccine, neutral, and con-vaccine Twitter accounts ranked by PageRankOut. Nodes represent accounts and arrows represent retweeting. We treat accounts as sources to the extent that they have high PageRankOut. The most prominent con-vaccine and pro-vaccine accounts are, respectively, @[REDACTED1] and @[REDACTED2].¹⁰

¹⁰ Following the ethical considerations outlined in (Townsend and Wallace 2016), we keep any non-public figure anonymous where possible. In this context, public figures are those who either have been verified by Twitter (i.e., have a blue check) or have a wikipedia entry about them.

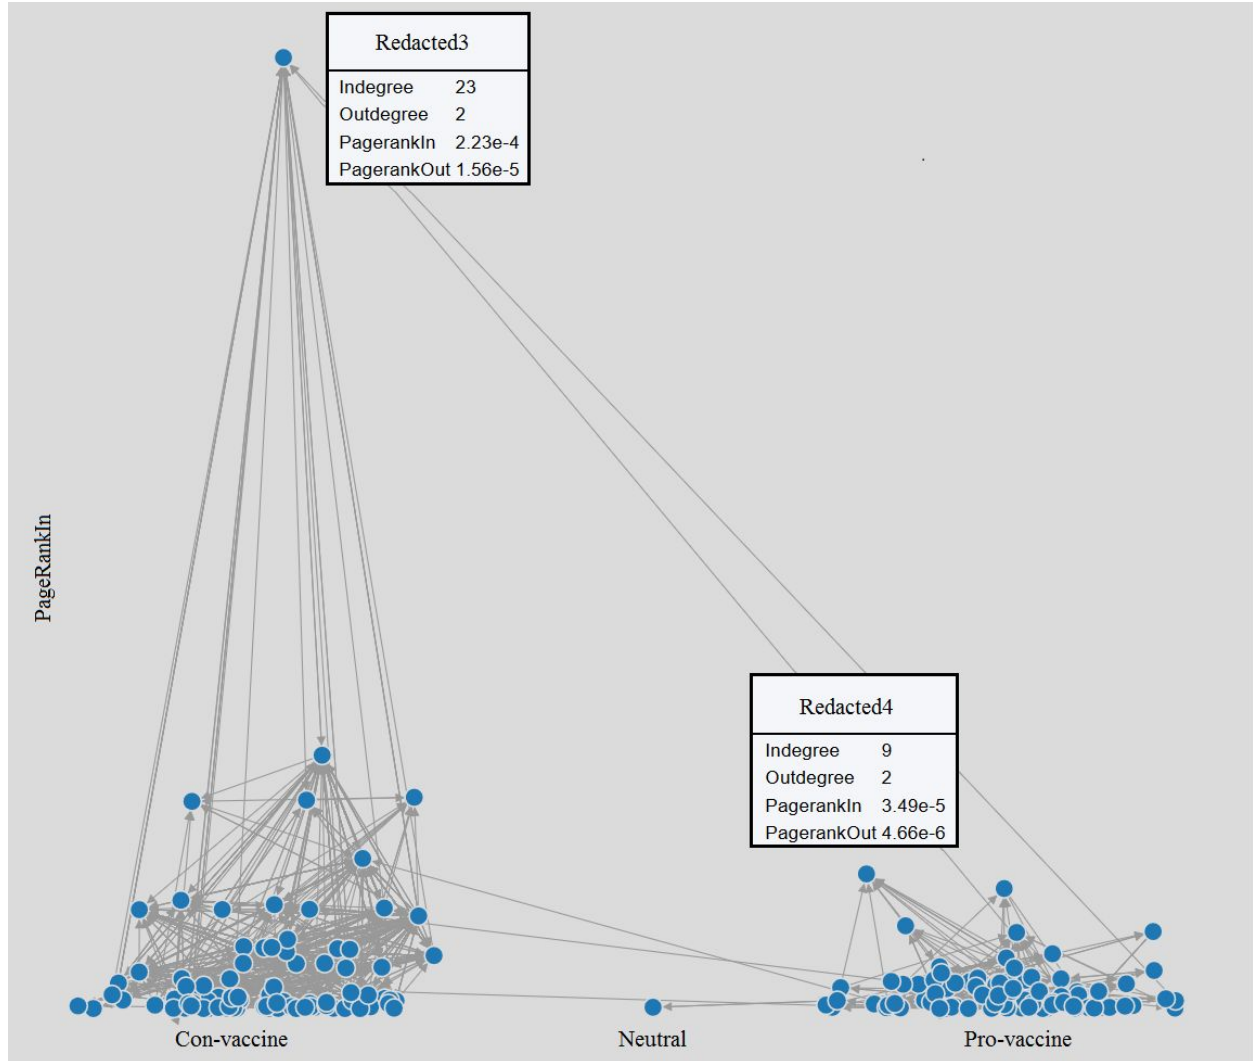


Figure 4: Active pro-vaccine, neutral, and con-vaccine Twitter accounts ranked by PageRankIn. Nodes represent accounts and arrows represent retweeting. We treat accounts as receivers to the extent that they have high PageRankIn. The most prominent con-vaccine and pro-vaccine accounts are, respectively, @[REDACTED3] and @[REDACTED4].

Figures 3 and 4 illustrate the reach and receptivity of both the pro-vaccine and con-vaccine camps. Communication between the camps is rare, and there are just a handful of neutral participants. To test this more formally, we use a *homophily* test, following the exposition of Easley and Kleinberg (2010). The main idea is to compare the fraction of cross-camp links to the expected fraction of such links if links were chosen randomly. To do so, let p denote the fraction of nodes in one camp, and $q = 1 - p$ the fraction of nodes that are not in this camp. A random link has probability p^2 of being within the considered camp, q^2 of being fully outside the camp, and $2pq$ of being between a node in the considered camp and one outside of it. As such, we can assess the homophily of the network by comparing the actual cross-camp links to this threshold. Following Meulemans and Schulz (2015), we can extend this to a *degree of homophily*, defined as $1 - x/(4pq)$ if $x < 2pq$ and $(1 - x)/(2 - 4pq)$ otherwise, where x is the actual fraction of cross-camp links. The degree of homophily is 0 if there are only cross-camp links, 0.5

if there is no homophily at all, and 1 if there is a strong indication of homophily. We assess the homophily of each camp by lumping together the other two camps to assess the strength of preference of communicating within the camp to communicating between camps. For example, when considering the pro-vaccine camp, we treat both con-vaccine and neutral accounts as a single distinct camp. The homophily of the pro- and con-vaccine camps are 0.993 and 0.995 respectively, indicating a strong bias towards retweeting information from accounts that have the same viewpoint.

In addition, while the number of members in each camp is roughly equivalent, the communication pattern in the con-vaccine camp shows greater density, whereas the pro-vaccine community is more diffuse. This corroborates our analysis in the previous section, which showed higher PageRankOut and PageRankIn of accounts in the con-vaccine community. What accounts for this difference? One possibility is that con-vaccine accounts are solely or primarily focused on this one topic (as commentators, activists, etc.), whereas the pro-vaccine accounts are occasional commentators who also engage with a range of other topics. We explore this hypothesis in more detail in section 5 below.

3.5 Discussion

This study suggests that network analysis and visualization can be used to better understand social epistemic communities. We have identified two important formal roles that epistemic agents can have in a community: sources and receivers. We then operationalized these roles and identified those who are seen as good sources of information and those who are active receivers. These are the first steps in assessing the wisdom-of-crowds hypothesis. Our tentative conclusion is that high PageRankOut is epistemically bivalent. Given the scientific consensus on vaccine safety, we operated on the heuristic that the pro-vaccine accounts are more likely to offer reliable testimony compared to anti-vaccine accounts. Our observations, from taking a close look at various accounts, stayed true to this heuristic. Nodes with high PageRankOut that were labeled pro-vaccine seem to offer reliable, expert testimony to others in their community. For example, Redacted2 is a medical doctor and the president of a prominent medical association. However, nodes labeled as anti-vaccine with high PageRankOut (including the node with by far the highest PageRankOut) spread misinformation and disinformation. The formal properties of the network under discussion here do not make it possible to distinguish between these. Likewise, we suggest that high PageRankIn is epistemically bivalent. Again after taking a close look at various accounts we saw that some nodes with high PageRankIn seem to be well-positioned to learn that vaccines are generally safe, but others are liable to end up more misinformed or disinformed than they would have been if they simply disconnected from the testimonial network entirely. And as with the sources, the formal properties of the network under discussion here do not make it possible to distinguish between these cases. If this is right, then the wisdom-of-crowds hypothesis needs to be significantly curtailed. Using PageRankOut as a heuristic to identify trustworthy sources does not succeed in this network. Using PageRankIn as a heuristic to identify well-informed receivers does not succeed in this network. In Study 2, we replicate this study with a larger sample.

4. Study 2

Study 1 suggested that an unvarnished version of the wisdom-of-crowds hypothesis is false. Of course, that was based on one week's discussion of vaccine safety in English on Twitter. It would be rash to generalize from just a single case. To shed further light on this topic, we conducted a second study with a new and larger data set.

4.1 Data collection

As in Study 1, we conducted a search query on the Twitter stream API. This search ran from March 1, 2018 to March 31, 2018. These dates were selected to ensure that any seasonal effects on our data would be minimal, and to expand the amount of data under analysis by a factor of approximately 4. We searched for tweets that used the same hashtags and text strings as in Study 1. These include such hashtags and text strings as #vaxxed, #vaccineswork, #vaccinesafety, 'vaccine', and 'antivax'. However, we did not automatically include all tweets from specific accounts because these seemed to be a source of noise in Study 1. The full search query can be found in the appendix. The search resulted in 167,521 tweets from 156,932 users.

4.2 Cleaning and processing the data

Almost all the 156,932 accounts in our raw dataset are minimally connected to the discourse on vaccine safety. Thus, we again isolate the core of the network by repeatedly removing 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. We repeat this process in stages until zero accounts are removed in a stage, indicating that the core of the network is all that remains. After removing these actors, we again designate the largest remaining network as our core network. This process filters out the vast majority of accounts from the raw data, leaving a core network of 897 nodes.

4.3 Analysis

Two of the authors independently hand-coded random samples of tweets both from the unfiltered network (300 tweets) and from the core network (300 tweets). 12.8% of the tweets in the unfiltered network were deemed irrelevant, whereas only 1.8% of tweets in the core network were deemed irrelevant.¹¹ This indicates that ceasing to include all tweets from specific accounts greatly reduced the amount of noise in the data, and further corroborates our decision to focus the analysis on the core network rather than the unfiltered network.

Next, we assigned nodes to categories, as in Study 1. Some of the nodes from the core network of Study 1 were also part of the core network in Study 2. We did not re-label these nodes; instead, we carried over their labels from Study 1. For all 824 new nodes in the core network, two independent coders labeled each of the nodes of the core network as being either pro-vaccine, con-vaccine, neutral, mixed, or irrelevant. We observed adequate interrater reliability (Fleiss's $\kappa = .811$, $z = 26.4$, $p < .001$). The two labelers then reviewed cases of disagreement and designated a few accounts as *irrelevant*; these accounts were removed from the

¹¹ We further note that several of the tweets in the core network were in French, Italian, or Japanese. They ended up in our network primarily because they also included one or more of the hashtags in our API query. Further investigation may benefit from a multi-language approach.

network before calculating the core, leaving 888 nodes in the core network. The average out-degree, in-degree, PageRankOut, and PageRankIn of nodes in each of the four camps is summarized in Table 3. For more detailed plots of the data, we refer to Figure 7 in the appendix.

Table 3: mean degree and PageRank properties of pro-vaccine, neutral, and con-vaccine accounts.

	<i>n</i>	out-degree	in-degree	PageRankOut	PageRankIn
pro-vaccine	549	5.82	5.74	$2.98*(10^{-6})$	$1.71*(10^{-6})$
neutral	3	1.67	3.00	$1.64*(10^{-6})$	$2.02*(10^{-6})$
mixed	9	1.33	1.33	$1.78*(10^{-6})$	$1.51*(10^{-6})$
con-vaccine	327	17.42	17.52	$5.74*(10^{-6})$	$2.31*(10^{-6})$

As before, the number of neutral (and mixed) accounts in the core of the network is very small.¹² Unlike Study 1, however, Study 2 suggests that there are nearly twice as many pro-vaccine accounts as con-vaccine accounts in the core of the network. The two researchers who labeled the data noted a large number of pro-vaccine accounts that seemed to represent pediatricians and other types of physicians (as indicated by their account names). It is unclear whether Study 1 simply missed tweets from such accounts or whether instead an active campaign by pediatricians and other doctors is distinctive of the year 2018. Because there are so few neutral and mixed accounts, we exclude them from significance testing. Kruskal-Wallis tests reveals that differences in out-degree ($\chi^2 = 19.67$, $df = 1$, $p = .0002$), in-degree ($\chi^2 = 41.99$, $df = 1$, $p < .0001$), PageRankOut ($\chi^2 = 12.95$, $df = 1$, $p = .005$) and PageRankIn ($\chi^2 = 11.80$, $df = 1$, $p = .008$) are all significant. Together, these results suggest that there are very few neutral or mixed accounts tweeting about vaccine safety, and that both the reach (out-degree & PageRankOut) and the receptivity (in-degree & PageRankIn) of con-vaccine accounts is greater. It's noteworthy that these differences obtain despite the fact that there are almost twice as many pro-vaccine accounts in the core network. This is in keeping with the findings of Radzikowski et al. (2016), who studied Twitter discussions during a measles outbreak in February 2015. Their data "indicated that a bottom-up campaign [...] far outweighed the presence of official sources such as the top-down efforts of CDC and WHO." Our findings go further: even when individual pediatricians and doctors get involved in debates over vaccine safety (recall that pro-vaccine accounts far outnumber con-vaccine accounts in the core network), their reach and receptivity do not match the reach and receptivity of the con-vaccine accounts.

4.4 Visualization

As before, we also visualize the data to aid in interpretation. These visualizations are based on the core network, with 888 accounts represented. Position on the x-axis is determined by

¹² It is not clear whether distinguishing neutral from mixed views in this way is worth the effort when analyzing testimonial networks about vaccine safety. Perhaps in other testimonial networks on other topics, it would be so.

sentiment, with neutral and mixed accounts grouped together for convenience; again, we offset these horizontally for legibility but remind the reader that this offset within a column carries no meaning. Position on the y-axis represents each account's reach (PageRankOut, Figure 4) or receptivity (PageRankIn, Figure 5). We again see that communication between the camps is rare. We compute the degree of homophily and observe that with 0.988 and 0.983 for the pro- and con-vaccine camps, respectively, these values remain very high (recall that the maximum possible value is 1.0) and are only marginally lower than those for Study 1.

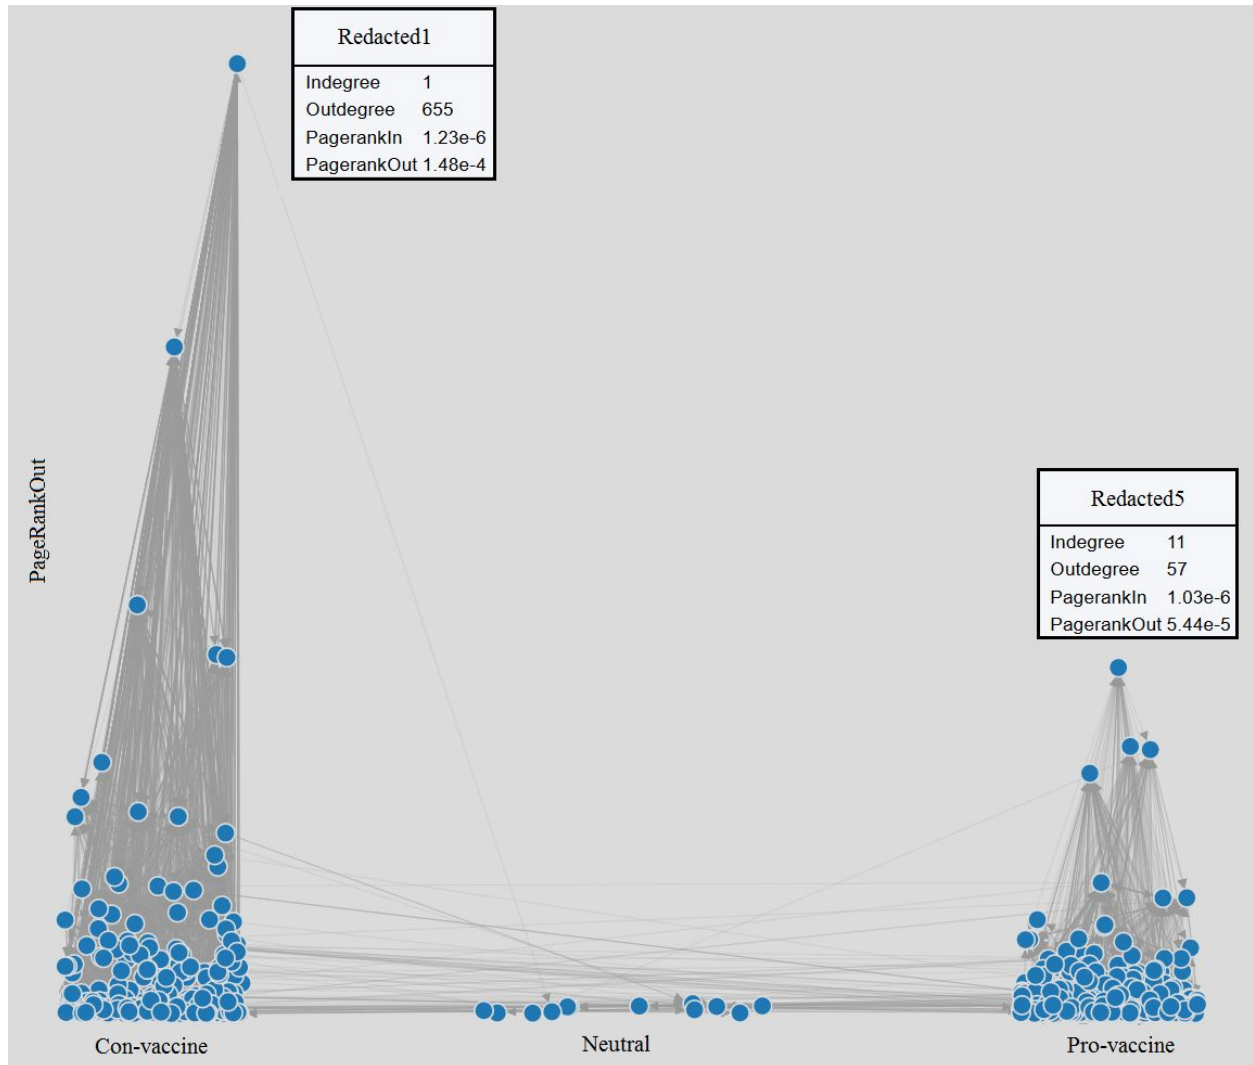


Figure 5: Active pro-vaccine, neutral, and con-vaccine Twitter accounts ranked by PageRankOut. Nodes represent accounts and arrows represent retweeting. We treat accounts as sources to the extent that they have high PageRankOut. The most prominent con-vaccine and pro-vaccine accounts are, respectively, @[REDACTED1] and @[REDACTED5].

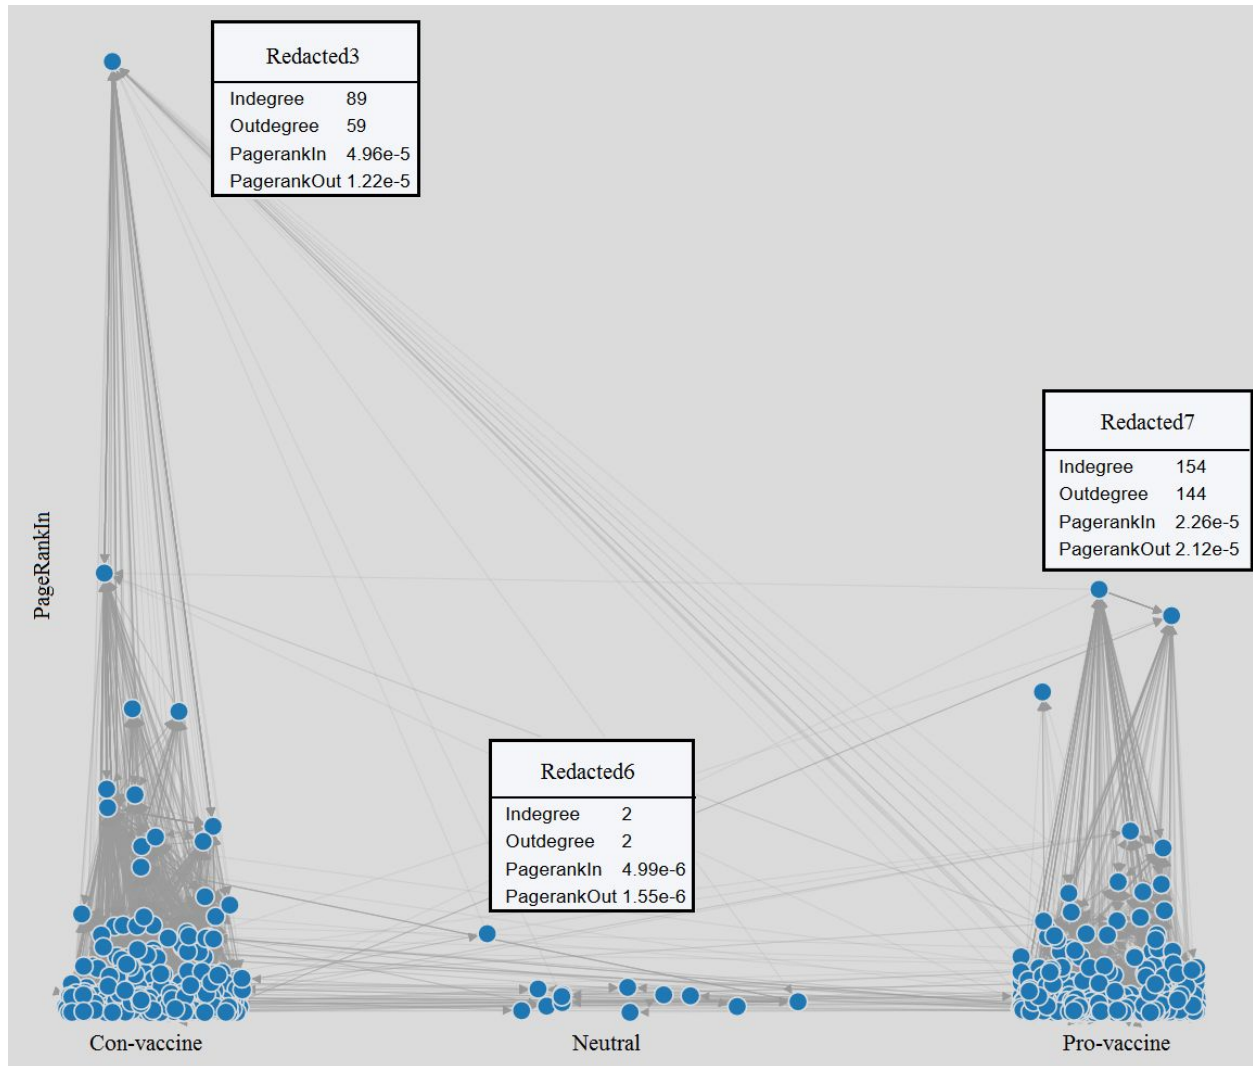


Figure 6: Active pro-vaccine, neutral, and con-vaccine Twitter accounts ranked by PageRankIn. Nodes represent accounts and arrows represent retweeting. We treat accounts as receivers to the extent that they have high PageRankIn. The most prominent con-vaccine, neutral, and pro-vaccine accounts are, respectively, @[REDACTED3], @[REDACTED6], and @[REDACTED7].

In addition, while there are more members of the pro-vaccine camp, the communication pattern in the con-vaccine camp shows greater density, whereas the pro-vaccine community is more diffuse. This corroborates our analysis in the previous section and in Study 1, which showed higher PageRankOut and PageRankIn of accounts in the con-vaccine community. In section 5, we return to this finding.

4.5 Discussion

This study once again suggests that network analysis and visualization can be used to better understand social epistemic communities. As in Study 1, the wisdom-of-crowds hypothesis is not borne out by our data and analyses. Once again, it appears that high PageRankOut is epistemically bivalent. After taking a close look at various accounts in the network, we found that some nodes with high PageRankOut seem to offer reliable, expert testimony to others in their community. However, other nodes with high PageRankOut (including the node with by far the highest PageRankOut) spread misinformation and disinformation. In fact, the same misinformation account that had the highest PageRankOut in our first study, continued to have the highest PageRankOut. The formal properties of the network under discussion here do not make it possible to distinguish between these. Likewise, high PageRankIn is again epistemically bivalent. We found that some nodes with high PageRankIn seem to be well-positioned to learn that vaccines are generally safe, but others are liable to end up more misinformed or disinformed than they would have been if they simply disconnected from the testimonial network entirely. For such actors, avoiding testimony altogether would both save them time and improve their epistemic position; this may hold more generally in the context of fake news and a polluted epistemic ecosystem (Levy 2017). And as with the sources, the formal properties of the network under discussion here do not make it possible to distinguish between these cases. In light of this replication with a new and larger dataset, the wisdom-of-crowds hypothesis needs to be significantly curtailed, especially in regards to polarized topics. Using PageRankOut as a heuristic to identify trustworthy sources does not work in the sorts of non-ideal epistemic networks found in the real world. Using PageRankIn as a heuristic to identify well-informed receivers also does not work in the sorts of non-ideal epistemic networks found in the real world.

5. General discussion and directions for future research

On its own, PageRank does not seem to deliver on the promise of harnessing the wisdom of crowds for agents in real testimonial communities. In both of our case studies, the nodes with highest PageRankOut were not reliable sources, the nodes with highest PageRankIn were not likely to be well-informed, and both PageRankOut and PageRankIn were overall higher in the con-vaccine camp. Perhaps some other metric would do better in such an epistemically treacherous environment. In this section, we point to two such features that deserve further research.

First, most articulations of the wisdom-of-crowds hypothesis require that the crowd be structured in a way that preserves sufficient independence among the members of the community (Surowiecki 2004). One attractive way to spell out the independence criterion is in terms of the distance between the sources that a given receiver relies upon. The metaphors of the “filter bubble” and “echo chamber” (Pariser 2011) that are nowadays common in criticisms of public discourse on social media suggests a lack of independence. Perhaps if a formal definition of relying on multiple independent sources could be developed and operationalized for the study of real testimonial networks, it would fare better than PageRank in harnessing the wisdom of crowds.

Second, in this paper we noted that the discourse on vaccine safety is highly polarized and partisan. The proportion of discussants whose views could plausibly be considered neutral or

mixed was very low. In Study 1, just 1 out of 185 accounts was not partisan. In Study 2, only 12 out of 888 accounts were not partisan. The cross-talk between the camps was also rare. In Study 1, the degree of homophily for the pro- and con-vaccines camps were .993 and .995, respectively. In Study 2, these values were .988 and .983, respectively, extremely close to the maximum of 1.0. This suggests that, in addition to purely structural metrics such as PageRank and source independence, assessments of the health of epistemic communities would do well to attend to diversity of viewpoint or content.¹³ Indeed, Tafuri et al. (2014) advocate fostering diversity among news media to address anti-vaccine misinformation, and Dubois & Blank (2018) argue that citizens with sufficient curiosity operating in a diverse media environment typically manage to avoid getting trapped in echo chambers more generally.

However, measuring and improving viewpoint diversity is challenging for several reasons. While it is possible to delineate broad topical communities using LDA, as we did in Study 1, this is a very coarse way of carving up the space of discussion. It is important to be able to assess not just how various camps disagree with each other but also their *reasons* for disagreeing. In other words, there is not simply diversity in viewpoint regarding diversity in the conclusion reached (either pro or anti-vax), but also diversity in the reasons given for a given conclusion. Furthermore, when taking a close look at the profiles of both the pro-vaccine and con-vaccine accounts with highest PageRankOut and PageRankIn, we found that the most prominent con-vaccine accounts in our data seem to have a single-minded focus on vaccines. We observed that the most prominent sources on the pro-vaccine side of the controversy were medical professionals and the most prominent receivers on this side of the controversy were non-medical accounts that engaged in a wide range of topics. By contrast, the most prominent sources on the con-vaccine side of the controversy include vaccination-related information in their account descriptions, along with references to homeopathy and alternative medicine. The most prominent receivers on this side of the controversy seem to be controlled by activists who believe that their children have been harmed by vaccination. These accounts are also more single-minded in their topic focus, with the vast majority of content referring directly or tangentially to vaccine safety, compared to the pro-vaccine accounts. Plausibly, this may lead them to seek out, attend to, interpret, and remember information about vaccine safety that makes them epistemically worse off. This suggests that beyond viewpoint- and reasons-diversity *between* individuals and groups, it may be important to attend to diversity of interests *within* individuals and groups. Measuring these three types of diversity (viewpoint-, interest-, and reasons-) is a tall order.

Moreover, there are at least two important caveats regarding the epistemic value of diversity that make harnessing a diversity metric in the wisdom of crowds even more challenging. First, regarding interest-diversity: we are well aware that activists are often criticized for lacking diversity in their interests or being obsessed by an *idée fixe*. We do not want to unthinkingly condemn all activism. Moreover, presumably, the vaccine skeptics in our data see themselves as being activists in exactly this way. This makes interest-diversity an ambivalent epistemic heuristic.

Second, in the case of vaccine safety, scientific evidence supports one side of the controversy. In such conditions, neutrality, ambivalence, or conciliation is not necessarily a virtue and may be a vice. As mentioned above, Levy (2017) argues that avoiding testimony is the

¹³ It is also worth investigating how PageRank fares in non-ideal testimonial networks involving non-controversial topics that have a lower degree of homophily.

correct response in these sorts of cases. However, inter-camp discussion does not necessarily require partisans to change their views, reduce their confidence, or conciliate epistemically in some other way. As John Stuart Mill (1859 / 1989) argued in chapter two of *On Liberty*, there is much to be said for the free expression of opinions, whether they are true, false, or partially true. According to Mill, false views are epistemically defensible because expressing false opinions leads to debate, which is one of our best methods for arriving at the truth on a given question. Mill famously indicts those who are unwilling to defend in open debate (what they take to be) the truth as being at risk of having their knowledge demoted to “dead dogma” and “phrases retained by rote.” Furthermore, when the topic under discussion is complex -- as is typically the case in the political, moral, and medical domains -- it often happens that any given person’s opinion contains some truth and some falsehood. For instance, health experts do not recommend that all populations receive all vaccines as early as possible; the measles-mumps-rubella (MMR) vaccine is generally not recommended for infants younger than twelve months.¹⁴ In complex cases, universal generalizations (e.g., “All vaccines are safe for all populations”) are almost always false. Debate, when it plays out well, helps to refine the proposition under discussion until it becomes epistemically acceptable.

That said, there are limits to the epistemic value of debate. In particular, the rise of trolling as a form of engagement means that many invitations to debate are not made in good faith. By good faith we mean, roughly, that one is willing to change their views in light of some conceivable evidence or argument that one’s opponent might offer. Furthermore, it’s important to recognize that engaging in debate often brings with it significant opportunity costs in terms of time, effort, and emotional engagement. Incessant, bad-faith invitations to engage in debate are a well-known tactic of trolls called ‘sea lioning’ (Poland 2016, p. 145). Trolls who engage in sea lioning might seem to have John Stuart Mill on their side: after all, can’t they accuse their would-be opponents of adhering to “dead dogma”? On both this question and the question of interest-diversity, simple-minded maximization heuristics (*as many interests as possible! as many viewpoints as possible! as much debate as possible!*) are hopeless.

Thus, much needs to be done on how to establish a diversity of viewpoint metric that can harness the wisdom of crowds that can balance issues of when to cut off testimony and debate and when to engage in order to maintain and secure one’s epistemic position. Articulating and defending operationalizable heuristics is a massive task that will require the expertise and input of philosophers, computer scientists, social scientists, and many others.

¹⁴ See url = <

<http://www.immunise.health.gov.au/internet/immunise/publishing.nsf/Content/Handbook10-home~handbook10part4~handbook10-4-9> >.

References

- Alfano, M. (2016). The topology of communities of trust. *Russian Sociological Review*, 15(4): 30-56.
- Alfano, M. & Higgins, A. (forthcoming). Natural language processing and semantic network visualization for philosophers. In e. Fischer & M. Curtis (eds.), *Methodological Advances in Experimental Philosophy*. Bloomsbury.
- Alfano, M. & Huijts, N. (forthcoming). Trust and distrust in institutions and governance. In J. Simon (ed.), *Handbook of Trust and Philosophy*. Routledge.
- Blei, D., Ng, A., Jordan, M. (2003). Latent dirichlet allocation. *Journal of Machine Learning Research*, 3(4-5): 993-1022.
- Borchers, C. (2017, October 4). Retweets ≠ endorsements? Oh, yes, they do, say the Hatch Act police. *The Washington Post*. URL = <
<https://www.washingtonpost.com/news/the-fix/wp/2017/10/04/retweets-endorsements-hatch-act/>>. Retrieved 18 February 2018.
- Borsboom, D. & Cramer, A. (2013). Network analysis: An integrative approach to the structure of psychopathology. *Annual Review of Clinical Psychology*, 9: 91-121.
- Brandt, M., Sibley, C., & Osborne, D. (2018). What is central to belief system networks? PsyArXiv Preprints. DOI = < 10.17605/OSF.IO/TYR64>. URL = <
<https://psyarxiv.com/tyr64>>.
- Brin, S. & Page, L. (1998). The anatomy of a large-scale hypertextual Web search engine. *Computer Networks and ISDN Systems* 30: 107–117.
- Clark, A. (2002). Towards a science of the bio-technological mind. *International Journal of Cognition and Technology*, 1(1): 21-33.
- Coady, C. A. J. (1992). *Testimony: A Philosophical Study*. Oxford University Press.
- Conover, M., Ratkiewicz, J., Francisco, M., Goncalves, B., Flammini, A., and Menczer, F. (2011). Political polarization on Twitter. In *Proc. 5th International AAAI Conference on Weblogs and Social Media (ICWSM)*.
- Dubois, E. & Blank, G. (2018). The echo chamber is overstated: The moderating effect of political interest and diverse media. *Information, Communication & Society*, 21(5): 729-45.
- Easley, D. & Kleinberg, J. (2010). *Networks, Crowds and Markets: Reasoning About a Highly Connected World*. Cambridge University Press.
- Epskamp, S., Cramer, A., Waldorp, L., Schmittmann, V., Borsboom, D. (2012). Qgraph: Network visualizations of relationships in psychometric data. *Journal of Statistical Software*, 48(4): 1-18.
- Fricke, M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford University Press.
- Hedden, B. (2017). Should juries deliberate? *Social Epistemology*, 31(4): 368-86.
- Holman, B., & Bruner, J. P. (2015). The problem of intransigently biased agents. *Philosophy of Science*, 82(5), 956-968.
- 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.
- Levy, N. (2017). The bad news about fake news. *Social Epistemology Review and Reply Collective*, 6(8): 20-36.

- List, C., & Goodin, R. E. (2001). Epistemic democracy: Generalizing the Condorcet jury theorem. *Journal of political philosophy*, 9(3), 277-306.
- Lurquin, J. & Miyake, A. (2017). Challenges to ego-depletion research go beyond the replication crisis: A need for tackling the conceptual crisis. *Frontiers in Psychology*, 8: 586.
- Masterton, G., Olsson, E. J. & Angere, S. (2016). Linking as Voting : How the Condorcet Jury Theorem in Political Science is Relevant to Webometrics. *Scientometrics*. 106: 3, pp 945-966.
- 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.
- Medina, J. (2013). *The epistemology of resistance: Gender and racial oppression, epistemic injustice, and the social imagination*. Oxford University Press.
- Meulemans, W. & Schulz, A. (2015). A Tale of Two Communities: Assessing Homophily in Node-Link Diagrams. *Graph Drawing and Network Visualization*, LNCS volume 9411, pp 489-501.
- Mill, J. S. (1859 / 1989). *On Liberty*. Edited by S. Collinni. Cambridge University Press.
- Munzner, T. (2014). *Visualization Analysis and Design*. AK Peters.
- Open Science Collaboration. (2015). Estimating the reproducibility of psychology science. *Science*, 349(6251): 943.
- Pariser, E. (2011). *The Filter Bubble, What the Internet is Hiding from You*. Penguin.
- Poland, B. (2016). *Haters: Harassment, Abuse, and Violence Online*. University of Nebraska Press.
- Radzikowski, J., Stefanidis, A., Jacobsen, K. H., Croitoru, A., Crooks, A., & Delamater, P. L. (2016). The measles vaccination narrative in Twitter: a quantitative analysis. *JMIR public health and surveillance*, 2(1).
- Raichle, M. (2015). The brain's default mode network. *Annual Review of Neuroscience*, 38: 433-447.
- Rosenstock, S., Bruner, J., & O'Connor, C. (2017). In Epistemic Networks, Is Less Really More?. *Philosophy of Science*, 84(2), 234-252.
- Senturk, R. (2005). *Narrative Social Structure: Anatomy of the Hadith Transmission Network, 610-1505*. Stanford University Press.
- Surowiecki, J. (2004). *The Wisdom of Crowds*. Anchor.
- Tafari, S., Gallone, M. S., Cappelli, M. G., Martinelli, D., Prato, R., & Germinario, C. (2014). Addressing the anti-vaccination movement and the role of HCWs. *Vaccine*, 32(38), 4860-4865.
- Tollefsen, D. (2007). Group testimony. *Social Epistemology*, 21(3): 299-311.
- Townsend, L., & Wallace, C. (2016). Social media research: A guide to ethics. *University of Aberdeen*, 1-16.
- Wehbe, E. (2011). A catastrophic failure. *The American Journal of Medicine*, 124(3): e7-e9.
- 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.

Appendix

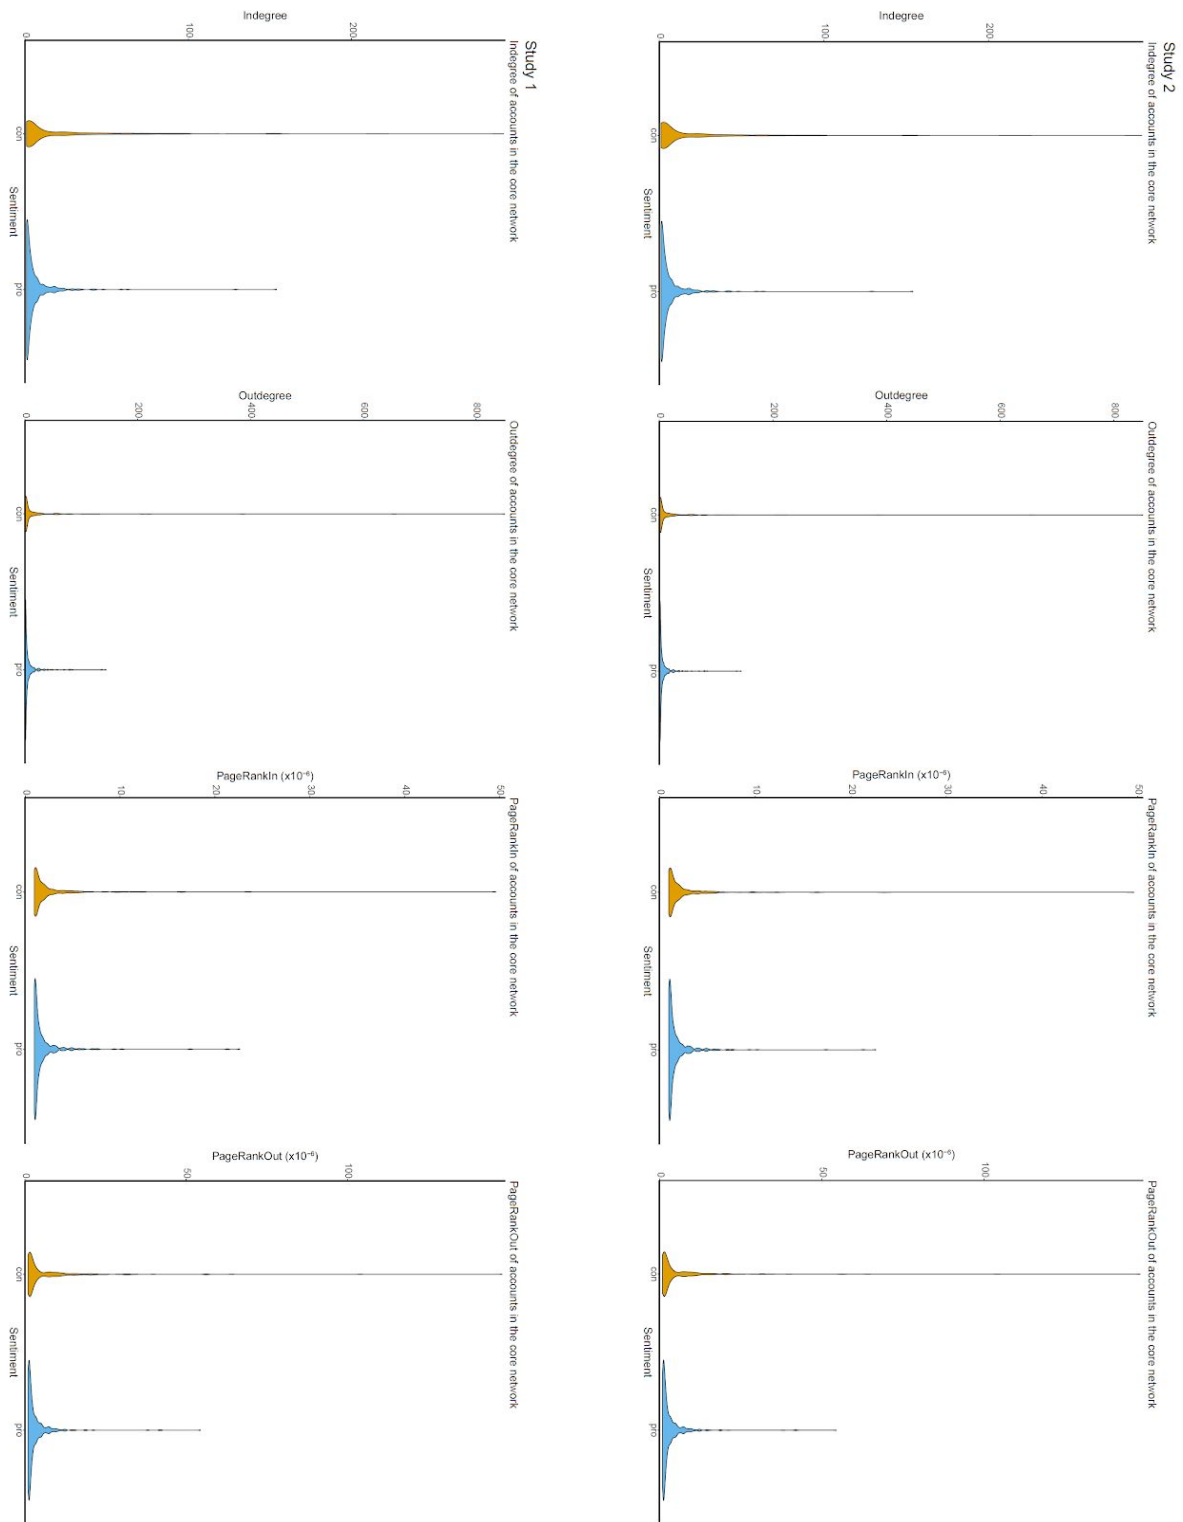


Figure 7: Violin plots to show the distribution of nodes in terms of in- and outdegree, and PageRankIn and -Out for both studies.

Twitter Search Query for Study 1:

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

Twitter Search Query for Study 2:

text: vax, vaxxed, vaccine, vaccination, vaccinations, vaxsafety, vax saftey, vaccineswork, vaccines work, vaccinesaftey, vaccine saftey, vaccines revealed, vaccinesrevealed, novax, no vax, no-vax, antivax, anti-vax, anti vax, immunisation, Vaccin, Vaccinaties, vaccinatiezorg, vaccine injury, vax injury, vaccinatieschade

hashtags: #vax, #vaxxed, #vaccine, #vaccination, #vaccinations, #vaxsafety, #vaccineswork,

```
#vaccinesaftey, #vaccinesrevealed, #novax, #antivax, #immunisation, #Vaccin,
#Vaccinaties, #vaccinatiezorg, #vaccinatieschade, #nvkp, #rvp,
#rijksvaccinatieprogramma, #vaccineinjury, #vaxinjury, #anti-vax
```

Code for Twitter Stream API in python 3:

```
# import libraries
import tweepy
from tweepy import Stream
from tweepy.streaming import StreamListener
import json
import time
import datetime

# Set up the Twitter Stream
class MyListener(StreamListener):
    def __init__(self, api, file_name_base, file_length):
        self.tweet_list=[]
        self.count = 1
        self.file_name_base = file_name_base
        self.file_length = file_length

    def on_data(self, data):
        try:
            j = json.loads(data)
            self.tweet_list.append(j)
            n = len(self.tweet_list)

            if (n % self.file_length == 0):
                ts = time.time()
                stamp = datetime.datetime.fromtimestamp(ts).strftime('%m%d%H%M%S')
                file_name = self.file_name_base + stamp + '.json'

                #writes the file of tweets to your computer
                with open(file_name , 'w') as f:
                    json.dump(self.tweet_list, f)

                print("Output File", self.count)
                self.count = self.count + 1
                self.tweet_list = []
            return True
        except BaseException as e:
            print("Error on_data:", str(e))
            return True
```

```

def on_error(self, status):
    print("Error",status)
    return False

class Credentials:
    def __init__(self, access_token, access_token_secret, consumer_key,
                 consumer_secret):
        self.access_token = access_token
        self.access_token_secret = access_token_secret
        self.consumer_key = consumer_key
        self.consumer_secret = consumer_secret

def mine_tweets(credentials, search_terms, file_name, file_length):
    auth = tweepy.OAuthHandler(credentials.consumer_key, credentials.consumer_secret)
    auth.set_access_token(credentials.access_token, credentials.access_token_secret)

    api = tweepy.API(auth)
    twitter_stream = Stream(auth, MyListener(auth, file_name, file_length))
    print("Starting up!")

    while(True):
        try:
            twitter_stream.filter(track=search_terms)
        except:
            time.sleep(30)
            print("program restart")

# Input your personal Twitter-keys and Path
ACCESS_TOKEN = '** input access token'
ACCESS_TOKEN_SECRET = '** input access token secret'
CONSUMER_KEY = '** input consumer key'
CONSUMER_SECRET = '** input consumer_secret'

file_name_base = '** input personal path'

#sets how many tweets you want per file.
file_length = 100

#Customize your twitter search query
text_list = []
hashtag_list = []
mentions_list = []
from_list = []

```

```
to_list = []
exclude_list = []

full_list = text_list
full_list.extend(hashtag_list)
full_list.extend(mentions_list)
full_list.extend(from_list)
full_list.extend(to_list)
full_list.extend(exclude_list)

# Start the stream
credentials = Credentials(ACCESS_TOKEN, ACCESS_TOKEN_SECRET,
    CONSUMER_KEY, CONSUMER_SECRET)

mine_tweets(credentials, full_list, file_name_base, file_length)
```