

Information Pathways in Online Science Communication: The Role of Platform Actors and News Media

Alexandros Efstratiou¹, Giuseppe Russo², Luca Luceri³

¹University of Washington

²EPFL

³Information Sciences Institute, University of Southern California
aefstra@uw.edu, giuseppe.russo@epfl.ch, lluceri@isi.edu

Abstract

Online discussions of science involve complex interactions among experts, news media, and social media users as they interpret and disseminate scientific findings. While prior work has examined these actors in isolation, their interplay in shaping science communication remains poorly understood. Using the COVID-19 pandemic as a case study, we analyze 1.24M tweets and 211k news articles that reference pandemic-related scientific papers. We find that the most influential Twitter accounts in this discourse are predominantly individuals with medical or research credentials. However, we also identify a coordinated network that disproportionately amplifies a small set of prominent credentialed experts who advance contrarian, anti-consensus positions on vaccines, lockdowns, and related topics. The papers promoted by these influential actors substantially overlap with those covered by news media, but with key differences: pro-consensus experts primarily engage with studies featured by mainstream and medical outlets, whereas contrarian experts align more closely with papers promoted by low-quality, pseudoscientific, or conspiratorial sources. Notably, news outlets tend to report on scientific studies after they have been highlighted by social media superspreaders. Together, these findings reveal multi-level pathways of information flow and coordinated amplification structures that shape science communication across social media and news, offering new insights into the dynamics of the broader information ecosystem.

Introduction

Skepticism toward scientific facts long supported by empirical evidence—such as climate change (Hornsey and Lewandowsky 2022) and the benefits of vaccination (Paoletti et al. 2024)—has increased over the past decade (Kennedy and Tyson 2023), driven in part by growing polarization in public attitudes toward science (Milkoreit and Smith 2025). While this skepticism is often attributed to politically motivated reasoning or organized disinformation campaigns (Lewandowsky et al. 2024), recent research suggests a more nuanced dynamic: even individuals who hold misinformed views on scientific issues attend to scientific credentials (Jalbert, Harris, and Williams 2025), and actors perceived as scientifically authoritative can come to domi-

nate discourse within anti-science online communities (Efstratiou and Caulfield 2021; Harris et al. 2024).

At the same time, science dissemination and interpretation increasingly occur through grassroots activity in online spaces. Prior work shows that scientific papers are shared within online communities for a range of purposes, including coordinated challenges to established findings (Yudhoatmojo, Cristofaro, and Blackburn 2021), selective misrepresentation of scientific consensus (Beers et al. 2023; Efstratiou et al. 2024; Yudhoatmojo, De Cristofaro, and Blackburn 2023), and open scientific debate conducted by researchers in public forums (Efstratiou et al. 2024). Beyond the circulation of peer-reviewed research, scientific misinformation also emerges through substandard community-driven data analyses (Lee et al. 2021) and the production or amplification of pseudoscientific content (Papadamou et al. 2022).

Although the centralized influence of perceived experts is well documented (Harris et al. 2024), including cases where such actors have vested interests in promoting alternative narratives (Lewandowsky et al. 2024; Nogara et al. 2022), much less is known about how grassroots, distributed efforts evolve into centralized expert influence or even into coordinated, malicious disinformation. Moreover, while news media can play a substantial role in disseminating scientific information to broader audiences (West and Bergstrom 2021), their position within the overarching pathways of science-related information flow remains poorly understood.

Research questions. In this paper, we use COVID-19 science as a case study to examine the science dissemination ecosystem from a holistic perspective. We analyze *information pathways* linking different classes of Twitter actors, namely, organic users, bots, superspreaders, and coordinated accounts, as well as news media articles, characterizing their distinct roles while examining how their activities intersect and reinforce one another. We define information pathways as patterns of temporal precedence between different types of entities engaging in similar actions, such as influential users and news media discussing the same scientific papers, or users retweeting science-related posts from other users. Importantly, information pathways under this definition are not inherently causal. For instance, if news outlets report on a scientific paper after it has been highlighted by influential Twitter users, Twitter is said to precede the news in this

pathway, even though the news coverage may have emerged independently or in response to other external factors.

To examine how scientific information circulates across social media and news media, we leverage a dataset capturing mentions of COVID-19–related preprints on Twitter and in news articles. Using this data, we address the following research questions:

RQ1. How do different classes of Twitter actors contribute to the amplification and circulation of COVID-19-related scientific papers?

RQ2. What information pathways characterize the flow of COVID-19 scientific research between Twitter actors and news media?

Methods. To address these questions, we employ a suite of methods to identify and characterize actors involved in online science communication. We construct co-activity networks to detect coordinated users, apply bot-detection tools to flag automated accounts, and use activity-based metrics to quantify influence. We then profile these actors in terms of popularity, topical focus, emotional tone, and alignment with scientific consensus. Throughout the paper, we define *contrarians* as users who consistently express anti-consensus positions (e.g., opposition to COVID-19 vaccines), and *conformists* as users whose views align with established scientific consensus. Finally, we analyze potential information pathways originating from superspreaders—the most influential users—and from news media, tracing how academic research papers diffuse within and beyond the Twitter information landscape.

Main findings. We identify a coordinated network of contrarian (anti-consensus) accounts that repeatedly retweet similar content within narrow time windows. While the primary superspreaders in the dataset are largely conformist users with scientific credentials—alongside a smaller number of contrarian experts—the coordinated network disproportionately amplifies content produced by contrarian superspreaders. We find no evidence that this coordination is automated. Moreover, news media activity exhibits temporal alignment with the activity of superspreaders. Mainstream and higher-trust outlets tend to cover the same scientific papers as conformist superspreaders, whereas low-trust, conspiratorial, or pseudoscientific outlets align more closely with contrarian superspreaders. Importantly, news coverage typically follows superspreader activity on Twitter, suggesting that COVID-19 scientific research often gains visibility on social media before appearing in news reporting—though this temporal ordering should not be interpreted as causal. Taken together, these findings underscore the complex interplay between social media actors and news outlets in science communication, revealing how information pathways across platforms can both amplify and obscure scientific consensus.

Related Work

In this section, we cover existing work on the online science communication ecosystem, as well as work on online coordination more broadly.

Science (Mis)communication

Science communication has taken on a participatory form, where both experts (Williams et al. 2022) and laypeople (Hafid et al. 2022; Yudhoatmojo, De Cristofaro, and Blackburn 2023) often share scientific articles with their audiences and peers in online social media. Additionally, beyond academic or other research venues, science can also be covered in other centralized formats like news media (Pei et al. 2025) or even documentaries (Naiman et al. 2025) and podcasts (Gu, Hickey, and Ryokai 2025). This creates a particularly complex science communication environment. On the one hand, science communicators may need to adopt platform-specific tactics to attract as much engagement as possible (Bagchi, Malmi, and Grabowicz 2025; Pera and Aiello 2024). On the other hand, scientific misinformation is rampant (West and Bergstrom 2021), especially around polarizing topics like climate change (Uyheng, Tyagi, and Carley 2021) or COVID-19 vaccines (Harris et al. 2024; Nogara et al. 2024). This may spread beyond decentralized contexts. For example, Bhat, Romero, and Horvat (2025) report that scientists may withdraw from engaging with news media due to potential misrepresentations of their work.

The problem is compounded by how this science communication is taken up online. It is common for discourse around scientific topics like vaccines and climate change to split into polarized clusters, each representing two sides of a debate (Santoro et al. 2023; Schmidt et al. 2018; Uyheng, Tyagi, and Carley 2021; Zollo et al. 2017), and these biased, collective behaviors may also reinforce counter-science groups like the anti-vaccination movement (Kata 2012). Nonetheless, signals of scientific expertise appear to be universally valuable for forming or supporting science-related attitudes (Jalbert, Harris, and Williams 2025), even among people who do not otherwise hold consensus positions. Online, this may manifest as scientific credentials being used to justify existing scientific opinions, rather than for knowledge attainment around scientific topics. Recent works have shown that the way in which science is disseminated through social networks can lead to false consensus effects (Efstratiou et al. 2024), where science is occasionally selectively cited in public posts by anti-science communities in efforts to claim scientific evidence in their favor (Beers et al. 2023). Moreover, perceived experts who share scientific articles at higher rates—but are, nonetheless, contrarians—are central in anti-vaccination communities (Harris et al. 2024). Taken together, this body of work suggests science contagion mechanisms that place “desired” experts at the center of communities, which may be inclined to use them as sources for their beliefs, although this potential social organization structure has yet to be empirically explored.

Social Collaboration and Coordination

Social organization around science communication can be viewed through the lens of *collaborative sensemaking*, a social process through which people make sense of complex information (Dailey and Starbird 2015). This process is occasionally driven by focused attention on certain information “spotlights” (Zhou et al. 2023), and coordinating the

agents involved can steer the quality and breadth of information that they are exposed to (Hong, Benjamin, and Müller-Birn 2018).

The idea of coordination, however, is also used in works that aim to detect potentially malicious activity. This is often done by considering signals of shared actions, like sharing the same URLs or news domains (Cinus et al. 2025), sharing the same social media posts (Nizzoli et al. 2021; Pacheco et al. 2021), or posting content whose text is highly similar between users (Pacheco, Flammini, and Menczer 2020). Common link-sharing on Facebook has been used to study coordination around COVID-19 misinformation (Gruzd, Mai, and Soares 2022) and COVID-19 vaccines more specifically (Song et al. 2025). Occasionally, the timing of this shared activity is also considered (Tardelli et al. 2024), and recent work has successfully deployed a combination of such signals in what are called “fused networks” to improve the generalizability of detection methods across information operations (Luceri et al. 2024).

However, these signals can also sometimes be attributed to organic, not inauthentic coordination (Cao et al. 2015), in line with collaborative sensemaking. In such cases, we may expect that participants are drawn to shared entities in a bottom-up, organic manner, as with cherry-picking academic work that aligns with existing beliefs (Beers et al. 2023; Efstratiou et al. 2024), or gravitating towards central, perceived experts (Harris et al. 2024; Jalbert, Harris, and Williams 2025). To this latter end, other scholars have also focused on detecting “superspreaders”; users who post frequently, and whose posts tend to gain a lot of traction, within specific communities or topics (DeVerna et al. 2024).

Present Work

This paper aims to synthesize these insights to assess information pathways around science communication, using COVID-19 as a case study. We focus specifically on scientific paper sharing, and examine the distinct roles of coordinated users, influential accounts (or “superspreaders”), automated activity, and the news media. Moreover, we attempt to link the activities of these distinct entities to each other, offering a first look into “multi-level” information pathways in science communication.

Dataset

We use a dataset from Efstratiou et al. (2024). It includes over 25k COVID-related preprints (and their published versions, where available) released on bioRxiv and medRxiv from the start of the pandemic until November 4, 2022, along with data on their mentions in news and social media. The dataset contains 1.24M tweets about these papers posted by 346k users, and 211k news articles published by 2.34k unique outlets. Moreover, it labels users as contrarians or conformists based on a graph convolution of the retweet network, validated against annotated tweet content of sampled users. In addition, papers are annotated by topic by applying a BERTopic model to their abstracts.




Figure 1: Network of coordinated accounts based on co-retweet activity.

Characterizing On-Platform Actors

In this section, we examine different actors on Twitter. We begin with the detection of coordinated retweet networks, followed by bots and superspreaders. Finally, we consider the interaction between the coordinated retweet network and the superspreaders.

Detection of Coordinated Retweet Networks

Following Luceri et al. (2024), we construct a network based on retweet similarity, or simply co-retweet. First, we drop users who retweeted < 5 times and tweets with < 10 retweets, and we split the data into 30-minute time windows. For every user, we count how many times they retweet each tweet within that window to form a windowed retweet vector. We apply TF-IDF to these vectors and then compute cosine similarity between every pair of users. The result is a user-user network in which edges are weighted by those similarity scores.

We retain the top 1% of eigenvector centrality values as coordinated nodes, as per Luceri et al. (2024). This threshold is chosen as a conservative cutoff, following Luceri et al. (2024)’s intuition that precision is more important than recall in the context of coordination: it is important not to attribute coordinated activity to non-coordinated actors. Nonetheless, we verify that the coordinated network structure remains the same with cutoffs of 0.5% and 2% (see the Appendix). The resulting network with a 1% centrality threshold is shown in Figure 1. There are six distinct subclusters emerging. To provide more context, we show the most-retweeted post from each cluster in Table 1. These posts also align with the highest TF-IDF weights in their respective clusters, indicating the strongest coordination signals. We also use topic labels based on BERTopic modeling of paper abstracts to quantify the types of papers cited in

Subcluster	Most retweeted post
1	New Japanese research pre-print publication ‘SARS-CoV-2 vaccination was associated with higher risk of myocarditis death, not only in young adults but also in ALL age groups including the elderly’ [links to paper]
2-5	You REALLY couldn’t make it up. I repeat ‘We are dealing with the greatest miscarriage of medical science, attack on democracy, damage to population health and erosion of trust in medicine that we will witness in our lifetime’ [quote-tweets post discussing paper that shows “negative vaccine effectiveness”]
6	“Breaking news” on T-cells [links to 6 papers, dated between May 2020 and November 2021, that show T-cell immunity after infection with COVID-19]

Table 1: Most-retweeted posts per coordinated subcluster.

the most shared tweets per cluster. Invariably, clusters 1-5 predominantly share papers on vaccines or boosters, viral mutations and immunity, and excess deaths and mortality, in this order. Cluster 6 predominantly shares papers on vaccines or boosters, T-cells and cell memory, and excess deaths and mortality.

We perform deeper analyses of the users that the coordinated network promotes, but we defer discussion of this to the Superspreader section below. Moreover, we characterize them in terms of their profile features and topical focus.

Profiles. To derive coordinated account characteristics, we perform bootstrapped independent samples t-tests across 10,000 iterations to compare the coordinated accounts with random samples of equal size ($N = 306$ per group). This random sampling is done only from the set of non-coordinated users (i.e., users *not* in the top 1% of centrality values) with ≥ 5 retweets to maintain consistency with the coordinated set. We opt for bootstrapping rather than direct comparisons to the non-coordinated population, as the size of the population can inflate statistical significance; the statistics we report are averages across the bootstrapped iterations. Coordinated accounts, on average, follow and are followed by a similar number of accounts, and have similar percentages of their science-related tweets being retweets (all bootstrapped $p > 0.1$) as the non-coordinated samples. However, coordinated accounts are significantly younger in terms of days since they were created ($t = -3.25, p = 0.001$). We show the distributions of these metrics in Figure 2. Distributions for the non-coordinated accounts are based on the entire non-coordinated population.

Nearly all (96.4%) of the coordinated accounts are contrarian, compared to only about 29% of accounts in the non-coordinated population. This discrepancy is significant in a chi-squared test ($\chi^2_{(1)} = 672.37, p < 0.001$). For this test, we estimate the expected number of contrarians and con-

Group	Expected	Observed
Contrarian	89	295
Conformist	217	11

Table 2: Contingency table for coordination vs stance.

formists in the coordinated network by obtaining their proportions from the non-coordinated population and scaling them down to the coordinated network’s size (rounded to the nearest integer). We then compare these expected values with the observed counts in the coordinated network (Table 2).

Activity. Next, we focus on characterizing coordinated accounts’ retweeting activity, as this is the signal we used to detect coordination; 59.2% of coordinated accounts only ever perform retweet actions on COVID-19 scientific discussions, compared to 30.6% in the overall population with five or more retweets ($\chi^2_{(1)} = 116.22, p < 0.001$, medium effect $\phi = 0.44$). Of the retweeted accounts, 75.9% of them are contrarian, while 90.4% of the tweets retweeted are from these contrarian accounts. Overall, retweets by the coordinated network most often discuss papers on vaccines or boosters (33.8%), virus mutations and natural immunity (18.4%), or excess deaths and mortality (9.5%). These top-retweeted topics are similar to the topics that are most prominent in the overall discussion; however, the distribution is more unequal and concentrated among the top few in the coordinated network (Gini coefficient, $G = 0.80$) as opposed to the distribution of the overall network ($G = 0.69$) or the rest of the contrarian subset ($G = 0.76$). To test for statistical significance in this inequality, we perform 10,000 bootstrapped iterations where we obtain random samples of N equal to the size of the coordinated network from the over-

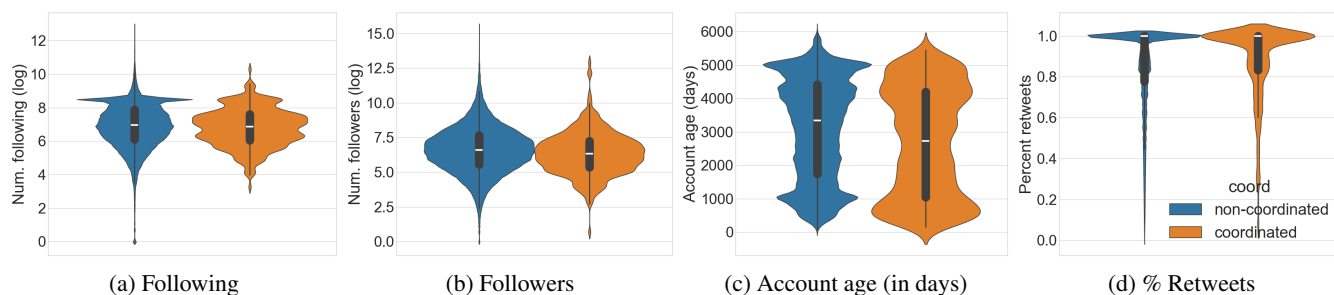


Figure 2: Violin plots for distribution comparisons between coordinated and non-coordinated accounts.

all and contrarian subsets with ≥ 5 retweets. In each iteration, we compute the G difference between the coordinated network and the comparison sample, and obtain the mean difference as well as 95% confidence intervals. In both the overall ($M_{diff} = 0.11$, $SE = 0.0001$, 95% CI: [0.1094, 0.1099]) and contrarian subset ($M_{diff} = 0.054$, $SE = 0.0001$, 95% CI: [0.0534, 0.0538]) comparisons, the coordinated network demonstrates significantly higher inequality (i.e., the 95% CIs do not span zero). That is, coordinated accounts focus on a narrower set of topics than the rest of the network.

Superspreaders

To detect superspreaders, we follow the approach outlined by DeVerna et al. (2024), who propose an h-index metric for uncovering influential spreaders. Accordingly, superspreaders are defined as users with a high number of posts that also tend to get a high number of retweets; for example, a user with a superspreader score (i.e., h-index) of 10 is a user with at least 10 tweets that received at least 10 retweets *each*. We consider accounts in the top 1% of scores as superspreaders, in line with DeVerna et al. (2024).

Profiles. As with the coordinated network, we first characterize the superspreader accounts by comparing them against bootstrapped random samples of equal size ($N = 764$) using independent-samples t-tests. We plot the distributions of the account characteristics we compare in Figure 3. Superspreaders tend to follow more accounts overall ($t = 3.96$, $p < 0.001$), this being driven by the tail-ends of the distribution (Figure 3a). They are also slightly older accounts ($t = 3.05$, $p = 0.013$) but do not differ in terms of daily volume of tweets ($t = -1.90$, $p = 0.12$). Moreover, we observe that superspreaders have significantly more followers than the general sample ($t = 42.71$, $t < 0.001$), with a clear mean difference in their distribution (Figure 3b).

Superspreaders also tend to be (legacy) verified at higher rates (28% vs. only 1.8% in the general population), and they are more likely to be conformists (80.6% vs. 70.8% in the general population; $\chi^2 = 35.62$, $p < 0.001$). Their substantially higher number of followers, higher verification rates, and more pro-science stance suggest that superspreaders of COVID-19 science are predominantly science communicators or science professionals. Our manual assessment of these accounts is in line with this observation (see Appendix for our labeling approach and criteria). As

seen in Table 3, the majority of conformist superspreaders are physicians/medical doctors or scientists actively working in COVID-related fields, with a substantial proportion of contrarian superspreaders (39.8% jointly) having similar credentials. The largest class of contrarian superspreaders, however, is people who do not primarily work within the medical science field, though they are often experts in other fields (e.g., law, economics, etc. – see the Appendix for how the scientific credentials in Table 3 are defined).

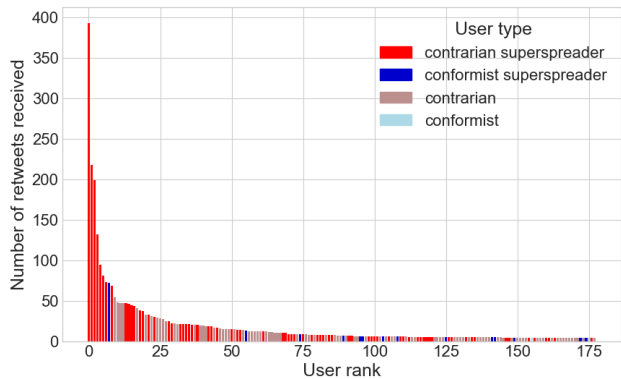
Topics. We analyze the topics of papers discussed by conformist and contrarian superspreaders separately, as these groups may engage distinct audiences: conformist superspreaders may participate in broader science communication, whereas contrarian superspreaders may concentrate on topics that more directly challenge or indict official institutions. Consistent with this interpretation, conformist superspreaders exhibit a more even distribution of topic engagement ($G = 0.58$) than contrarian superspreaders ($G = 0.74$). Moreover, conformists’ topical activity is more evenly distributed—and contrarians’ more concentrated—than that of the overall network ($G = 0.68$). Together, these results indicate that conformist superspreaders engage across a wider range of scientific topics, while contrarian superspreaders focus on a narrower subset.

In terms of the most discussed topics, vaccines/boosters and epidemic models are the top two for both conformists (9.4% and 9.1%, respectively) and contrarians (14.1% and 12.7%, respectively). For conformists, the top five are completed with genomic sequencing/mutations (7.7%), viral mutations/natural immunity (7.3%), and seroprevalence (i.e., how widespread the virus is; 7.2%). For contrarians, the top five are completed by excess mortality (11.6%), vaccination coverage/immunity (8.4%), and seroprevalence (8.3%). The distribution of topical prevalence for the top 20 topics per group is shown in the Appendix.

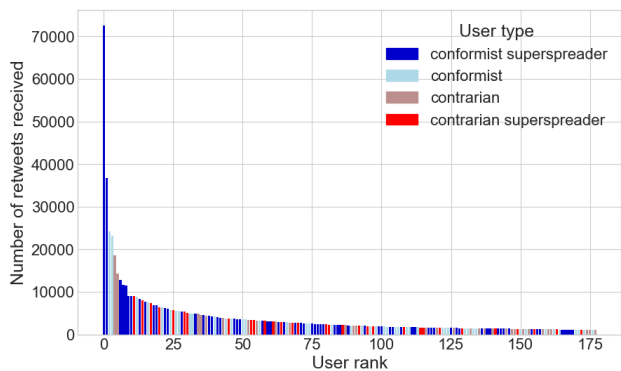
Key terms. To better understand how these topics are discussed across contrarian and conformist superspreaders, we extract the top keywords and top hashtags from the original (i.e., non-retweet) tweets of these groups. We show these as wordclouds in Figure 4; word size is proportional to the relative frequency across the respective group’s tweets. We exclude words like “COVID-19” and other synonyms (e.g., “coronavirus”, “SARS-COV-2” etc.) due to their topical prevalence.

User type	Anger	Disgust	Fear	Joy	Neutral	Sadness	Surprise
Random contrarian	7.66	0.99	37.93	5.60	32.30	6.08	9.44
Random conformist	3.62	0.61	37.50	7.29	40.47	4.20	6.31
Superspreader contrarian	4.22	0.99	35.31	3.34	44.34	4.40	7.39
Superspreader conformist	2.57	0.44	50.41	3.78	35.31	3.16	4.32

Table 4: Percentage of emotion share per user type.



(a) Coordinated network.



(b) Rest of the network.

Figure 5: Distribution of most-retweeted accounts by superspreader and related stance.

this is the number of accounts retweeted within the coordinated network), 576 of them are never retweeted by the coordinated network. This pattern also holds when only considering contrarian accounts. The rank correlation remains low-to-moderate (Kendall $\tau = 0.38$, $p < 0.001$), but 449 of the 884 most-retweeted contrarians overall are not retweeted by the coordinated network. Thus, this shows that the coordinated network focuses on a specific subset of accounts that are not necessarily the most popular in the social network.

Bots

For accounts with ten COVID-19 related tweets and above ($N = 28.9k$), we also collect their bot scores using Botometer’s historical data; this is suitable for our dataset as it was collected before Botometer stopped operating due to Twitter

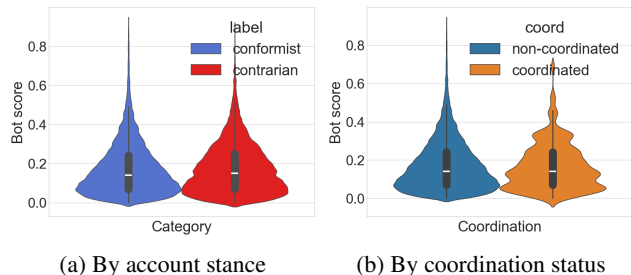


Figure 6: Bot score distributions.

API changes on May 31st, 2023.¹

This analysis does not yield any suspicious patterns. Bot scores do not produce significant differences between conformist and contrarian accounts ($t = -1.55$, $p = 0.12$; Figure 6a) nor between coordinated accounts and bootstrapped non-coordinated samples ($t = -0.52$, $p = 0.53$; Figure 6b). Indeed, based on a manual inspection, many of the bots appear to be overt ones that are automated to share papers, science news, etc. For context, from the accounts in the top 1% of assigned bot scores, just over 1 in 6 (17.8%) include one of the following overt terms in their name: bot, science, paper, hub, medrxiv, biorxiv, preprint, or article. These terms are typical in overt, automated, science-sharing bots. On the other hand, less than 1% of the accounts in the bottom 1% of bot scores contain any of those terms in their usernames. This finding is consistent with prior work showing that coordination does not necessarily imply automation.

Summary (RQ1)

Taken together, these patterns suggest mostly contrarian coordination. That is, while superspreaders are predominantly science communicators who promote COVID-19 science to their wider networks, a cluster of coordinated accounts tends to promote mostly anti-consensus superspreaders who share contrarian science, often focusing on a few specific topics.

Automated activity is mostly restricted to overt bots that typically share scientific papers, preprints, or other resources in a transparent manner. In short, coordination is not driven by automation in this case. Indeed, several prior works have made the case that coordination and automation are two distinct behaviors, since coordination can also be the result of trolling efforts, human-operated information campaigns, or organic cooperation (Mazza et al. 2022; Nizzoli et al. 2021). Similarly, coordination alone does not necessarily constitute

¹<https://botometer.osome.iu.edu/>

Rank				
Coordinated network	Overall network	Contrarian subset	Description	
	1	385	82	Anti-vaccine cardiologist
	2	149	30	Anti-lockdown professor
	3	37	9	Pro-herd immunity pathologist
	4	149	30	Anti-vaccine physician
	5	122	26	Pro-herd immunity physician

Table 5: Ranks refer to the most-retweeted accounts in each set, with 1 being the highest.

evidence of malicious intent (Pacheco et al. 2021), however, it can artificially amplify content and create illusions of public consensus (Luceri et al. 2024). Thus, while we cannot determine the intent behind this coordination, our results suggest a non-automated amplification of specific contrarian accounts.

Media Information Pathways

In our subsequent analyses, we consider the role of news media in propagating COVID-19 science through news articles published in various outlets. Beyond the media’s isolated role, we are also interested in how media outlets’ activity is related to the activity of Twitter, and in particular, superspreaders and “science influencers.” We consider news articles by media outlets that published COVID-related articles in the English language, and drop any news aggregators (e.g., Yahoo News, Google News) from our analyses.


Media Behavior Overview

We start with a broad description of media activity over the observation period. For these analyses, we rely on the domain trustworthiness scores from Lin et al. (2023), who combine six expert ratings from media watchdogs, fact-checkers, and researchers, and use imputation to extend coverage to 11.5k news domains. Moreover, we scrape the daily number of visitors for each outlet’s website from StatsCrap. We plot the 7-day rolling averages of the number, popularity, and trust score of media outlets reporting on COVID-19 scientific articles in Figure 7.


We also observe some syndication patterns. Syndication occurs when different outlets publish the same article at the same time or very close to one another. This is standard in journalism, often happening when outlets report the same story after a media embargo is lifted, or when large national outlets license articles to smaller local outlets for broader community reach. We report an analysis of syndication in the Appendix, which does not indicate any suspicious patterns. For subsequent analyses, we remove syndication by keeping only the earliest version of each news article. When multiple articles are published at the *exact* same time, we retain the one from the most popular outlet (by daily visits), assuming larger outlets license their articles to smaller ones rather than vice versa.

Crossover of Media and Superspreader Activity

To assess how media outlets and superspreaders relate to one another, we ask whether they tend to discuss the same pa-



(a) Number and total daily visits of outlet websites.



(b) Trust score of outlets.

Figure 7: Seven-day rolling averages for number, total daily visits, and average news trustworthiness of news outlets reporting on COVID-19 papers.

pers. We focus on these two actor types because they primarily drive visibility and engagement. We collect all DOIs shared by either superspreaders or outlets—this set defines our “vocabulary.” Each outlet and superspreader is represented as a sparse DOI vector, with values equal to the number of articles (outlets) or tweets (superspreaders) mentioning each paper. After TF-IDF transformation, we compute cosine similarities between every outlet-superspreader pair, treating similarity as alignment in the papers they discuss.

We next identify each superspreader’s K-nearest outlets and compare the neighbors of conformists and contrarians.

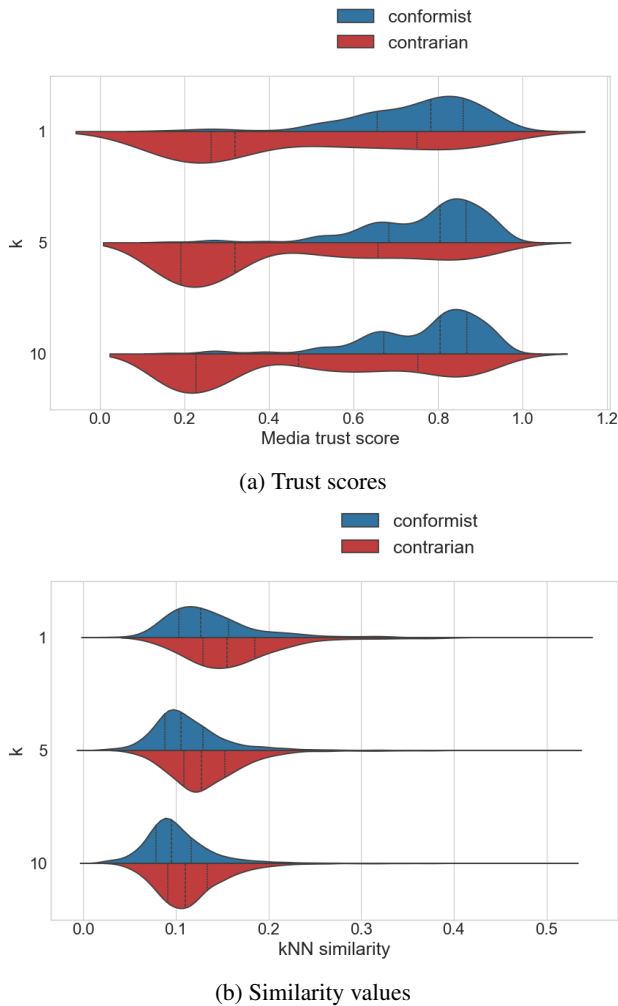


Figure 8: Violin plots representing (a) the trust scores of closest-neighbor media outlets by group, and (b) how similar the activity is between contrarian and conformist superspreaders' closest-neighbor media outlets.

We test $K = 1, 5,$ and 10 for robustness, and show differences between both overall similarities and neighboring outlet trust scores between the two groups in Figure 8.

Across all K , contrarian superspreaders align with significantly lower-trust outlets, and are significantly more similar to their outlet neighbors (all $p < 0.001$ in independent-samples t-tests). In short, contrarian superspreaders show stronger alignment with low-credibility outlets than conformists do with high-credibility outlets.

Breakdown of neighboring outlets. For further context, we show the news outlets that are most often identified as 5-NN neighbors of conformist and contrarian superspreaders in Figure 9, which confirms the trust patterns we observe in Figure 8. Conformist superspreaders tend to share similar DOIs with mainstream or medical news media, whereas contrarian superspreaders mostly neighbor with conspiratorial, low-quality, or pseudoscientific sources (as determined

by Media Bias/Fact Check.²)

Posting times. The preceding analyses show how different groups of media outlets and superspreaders talk about different academic papers, but they do not unpack whether one type of source precedes another. To address this, we examine periods of high activity for each paper DOI.

We first retain only DOIs with at least 100 mentions in both media outlets *and* Twitter. For each DOI, we apply Gaussian kernel density estimation on all mentions by news outlets and superspreaders between the DOI's first and last mention. We then scan the density grid to identify the smallest interval that contains at least 50% of total activity. The intuition is that, because DOI attention typically shows a bursty, log-normal pattern, this interval captures the main triggering event while excluding early noise and long trailing mentions. We visually show this 50% density region distribution for the top 4 most-mentioned papers in the Appendix. To ensure robustness, we also perform backward lookups of 3, 6, and 12 hours before the detected start date and take the earliest instance within these periods as the start of the high-density region. This sensitivity analysis, reported in the Appendix, produces nearly identical results.

Within this high-density region, we identify the density peaks for conformist (pro-consensus) superspreaders (PSPs), contrarian (anti-consensus) superspreaders (ASPs), low-credibility (LC), and high-credibility (HC) outlets. Outlets are binarized into LC and HC using a trust score cutoff computed as Youden's J (see Appendix) for interpretability. For each DOI, we determine the individual density peak of each class, and derive whether the peak of one class precedes the other.

We then construct a directed network showing this precedence across all DOIs among the classes. We draw pairwise edges between classes and, as we iterate through DOIs, increment the edge weight whenever one class's density peak precedes another's. For example, if ASP's peak appears before LC, we increment the $ASP \rightarrow LC$ edge weight by one. We scale the edge weights row-wise so that each class's outgoing edges sum to 100 (i.e., percentages). These relationships are visualized in Figure 10.

These information pathways reveal that superspreaders typically precede news outlets. Taking into account base posting frequencies, contrarian superspreaders are significantly more likely to precede both low-quality ($\chi^2_{(1)} = 127.56, p < 0.001, \text{median } \Delta_{time} = 28h$) and high-quality ($\chi^2_{(1)} = 123.41, p < 0.001, \text{median } \Delta_{time} = 23h$) news. Similarly, conformist superspreaders are more likely to precede both high-quality ($\chi^2_{(1)} = 235.33, p < 0.001, \text{median } \Delta_{time} = 23.5h$) and low-quality ($\chi^2_{(1)} = 166.08, p < 0.001, \text{median } \Delta_{time} = 19.6h$) news. On the contrary, there are no significant differences in the precedence order of low- and high-quality news ($\chi^2_{(1)} = 1.87, p = 0.17$) or contrarian and conformist superspreaders ($\chi^2_{(1)} = 0.51, p = 0.47$), indicating a more generalized Twitter-to-news information pathway. However, we stress that this analysis does not

²<https://mediabiasfactcheck.com/>

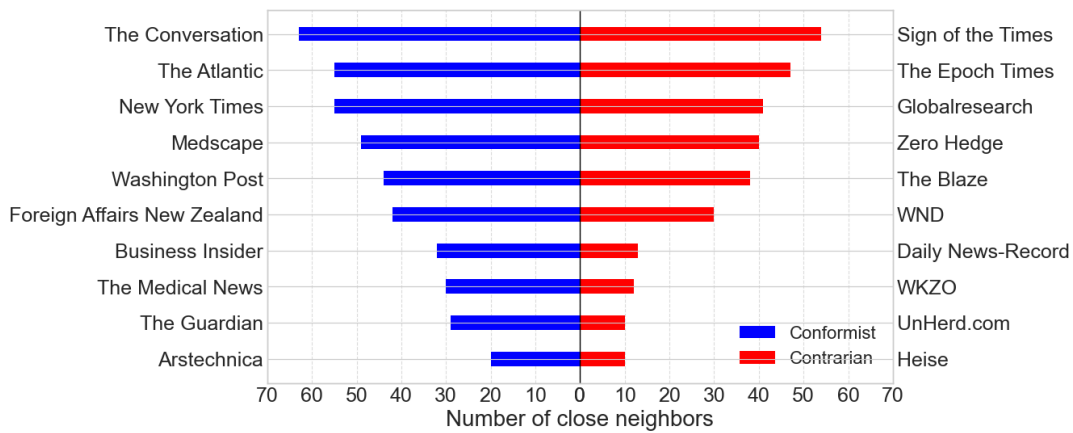


Figure 9: Most frequent closest-neighbor outlets by superspreader group.

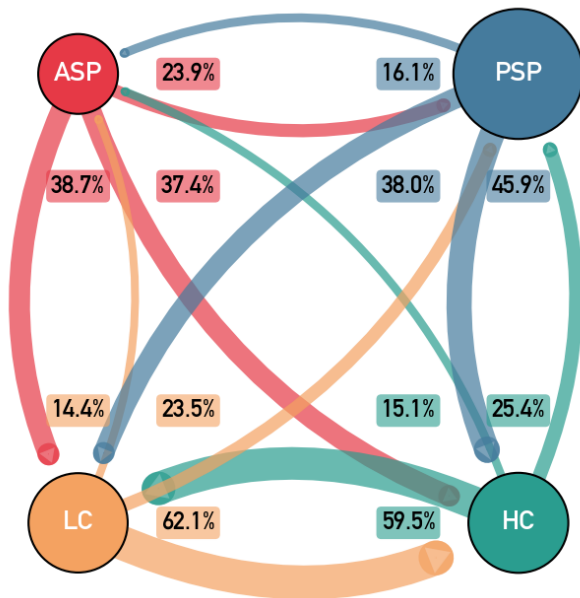


Figure 10: Precedence pathways.

establish any causal information flows. Rather, it demonstrates that scientific papers typically receive traction on Twitter before being picked up by news media, whose articles may subsequently be re-introduced to a wider platform base (Alperin et al. 2024; West and Bergstrom 2021).

Summary (RQ2)

Overall, we find overlap in the science dissemination of superspreaders and news media: while conformist superspreaders discuss papers picked up mostly by mainstream and higher-trust media, contrarian superspreaders align more with low-quality media like conspiratorial or pseudoscientific outlets in the science they share. Moreover, although we cannot claim causality, we find that, on average, science tends to be discussed by superspreaders on Twitter

first before being picked up in news media.

Discussion

In this paper, we examine the roles of different actors involved in the communication of COVID-19 science. We make several important observations, including how the obfuscation of scientific consensus (Beers et al. 2023; Efstratiou et al. 2024) and centralization of perceived experts (Harris et al. 2024) noted in other works may arise. We find that coordinated accounts amplify targeted experts who possess the credentials to introduce (selective) science into their wider networks. Although the majority of influencers driving scientific discussions on Twitter are non-contrarian science communicators, coordination seems to be strongest around those contrarian superspreaders.

Moreover, we find that news media mirror the roles of superspreaders. While mainstream media tend to report on science that is also shared by pro-consensus experts, lower-trust outlets are aligned with the science reporting activity of contrarian superspreaders in terms of the papers they cover. Our findings paint a comprehensive picture of how science communication and its aberrant forms take place across different media and entities of various influence and roles.

Furthermore, we add to the literature on the interplay of news media and platforms. While news articles about science may garner more attention on social media than scientific articles themselves (Alperin et al. 2024), we show that, in addition to social media conversations following these articles (Smeros, Castillo, and Aberer 2019; Wright et al. 2022), news media reporting on specific scientific papers often follows discussions of these papers by influential figures on Twitter. Our findings indicate possible “closed loops”, where platform influence may precede news reporting, with this news reporting then driving further platform attention.

This work highlights the multi-level landscape of online science communication, both within the same platform and between media. We show that seemingly disjointed activity may compound, as in the case of coordinated networks amplifying selected experts, or the boosting of selective pieces of science by high-influence actors like superspreaders and

news outlets with substantial followings. Although multi-platform information diffusion is already a broad area of interest for researchers (Gerard et al. 2025; Wilson and Starbird 2021), we demonstrate the importance of these considerations to science communication as well. At a time when many scientific topics are polarized and draw in actors such as government agencies, political commentators, and other non-scientific influencers, it is essential to understand how these groups collectively shape science communication, whether their participation is genuine or opportunistic.

However, there are some limitations associated with our work. First, we only consider news and tweets about scientific papers specifically. Although this is indeed the intended scope, there may be other avenues through which science communication is conducted. Second, these information pathways concern a specific time period and topic (COVID-19) on a specific platform (Twitter). Therefore, the generalizability of our findings to other platforms or scientific topics is limited. Finally, especially with respect to media and Twitter relationships, the patterns we report cannot be deemed to be causal, but rather as descriptive accounts of the types of events that tend to precede one another.

Acknowledgments

This work has been supported by the University of Washington's Center for an Informed Public, the John S. and James L. Knight Foundation (G-2019-58788), and the William and Flora Hewlett Foundation (2023-02789).

References

- Alperin, J. P.; Fleerackers, A.; Riedlinger, M.; and Hausteine, S. 2024. Second-order citations in altmetrics: A case study analyzing the audiences of COVID-19 research in the news and on social media. *Quantitative Science Studies*, 5(2): 366–382.
- Bagchi, C.; Malmi, E.; and Grabowicz, P. A. 2025. Effects of Research Paper Promotion via ArXiv and X. *Proceedings of the International AAAI Conference on Web and Social Media*, 19: 160–177.
- Baly, R.; Karadzhov, G.; Alexandrov, D.; Glass, J.; and Nakov, P. 2018. Predicting Factuality of Reporting and Bias of News Media Sources. In Riloff, E.; Chiang, D.; Hockenmaier, J.; and Tsujii, J., eds., *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 3528–3539. Brussels, Belgium.
- Baly, R.; Karadzhov, G.; An, J.; Kwak, H.; Dinkov, Y.; Ali, A.; Glass, J.; and Nakov, P. 2020. What Was Written vs. Who Read It: News Media Profiling Using Text Analysis and Social Media Context. In Jurafsky, D.; Chai, J.; Schluter, N.; and Tetreault, J., eds., *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 3364–3374. Online.
- Beers, A.; Nguyen, S.; Starbird, K.; West, J. D.; and Spiro, E. S. 2023. Selective and deceptive citation in the construction of dueling consensus. *Science Advances*, 9(38): eadh1933.
- Bhat, M.; Romero, D.; and Horvat, E.-A. 2025. Scholarly Disengagement as an Epistemic Crisis: Clickbait, Credibility, and the Decline of Public-Facing Science. In *Proceedings of the ACM Collective Intelligence Conference*, CI '25, 286–296. New York, NY, USA.
- Cao, C.; Caverlee, J.; Lee, K.; Ge, H.; and Chung, J. 2015. Organic or Organized? Exploring URL Sharing Behavior. In *Proceedings of the 24th ACM International Conference on Information and Knowledge Management*, CIKM '15, 513–522. New York, NY, USA.
- Cinus, F.; Minici, M.; Luceri, L.; and Ferrara, E. 2025. Exposing Cross-Platform Coordinated Inauthentic Activity in the Run-Up to the 2024 U.S. Election. In *Proceedings of the ACM on Web Conference 2025*, WWW '25, 541–559. New York, NY, USA.
- Dailey, D.; and Starbird, K. 2015. "It's Raining Dispersants": Collective Sensemaking of Complex Information in Crisis Contexts. In *Proceedings of the 18th ACM Conference Companion on Computer Supported Cooperative Work & Social Computing*, CSCW'15 Companion, 155–158. New York, NY, USA.
- DeVerna, M. R.; Aiyappa, R.; Pacheco, D.; Bryden, J.; and Menczer, F. 2024. Identifying and characterizing superspreaders of low-credibility content on Twitter. *PLOS ONE*, 19(5): e0302201.
- Efstratiou, A.; and Caulfield, T. 2021. Misrepresenting Scientific Consensus on COVID-19: The Amplification of Dissenting Scientists on Twitter.
- Efstratiou, A.; Efstratiou, M.; Yudhoatmojo, S.; Blackburn, J.; and De Cristofaro, E. 2024. "Here's Your Evidence": False Consensus in Public Twitter Discussions of COVID-19 Science. *Proc. ACM Hum.-Comput. Interact.*, 8(CSCW2): 471:1–471:33.
- Gerard, P.; Hanley, H. W. A.; Luceri, L.; and Ferrara, E. 2025. Bridging the Narrative Divide: Cross-Platform Discourse Networks in Fragmented Ecosystems.
- Gruzd, A.; Mai, P.; and Soares, F. B. 2022. How coordinated link sharing behavior and partisans' narrative framing fan the spread of COVID-19 misinformation and conspiracy theories. *Social Network Analysis and Mining*, 12(1): 118.
- Gu, Q. C.; Hickey, D.; and Ryokai, K. 2025. When AI Tells Their Story: Researchers' Reactions to AI-Generated Podcasts as a Tool for Communicating Research. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, CHI EA '25, 1–6. New York, NY, USA.
- Hafid, S.; Schellhammer, S.; Bringay, S.; Todorov, K.; and Dietze, S. 2022. SciTweets - A Dataset and Annotation Framework for Detecting Scientific Online Discourse. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, CIKM '22, 3988–3992. New York, NY, USA.
- Harris, M. J.; Murtfeldt, R.; Wang, S.; Mordecai, E. A.; and West, J. D. 2024. Perceived experts are prevalent and influential within an antivaccine community on Twitter. *PNAS Nexus*, 3(2): pgae007.
- Hartmann, J. 2022. Emotion English DistilRoBERTa-base. <https://huggingface.co/j-hartmann/emotion-english-distilroberta-base/>.
- Hong, M.-T.; Benjamin, J. J.; and Müller-Birn, C. 2018. Coordinating Agents: Promoting Shared Situational Awareness in Collaborative Sensemaking. In *Companion of the 2018 ACM Conference on Computer Supported Cooperative Work and Social Computing*, CSCW '18 Companion, 217–220. New York, NY, USA.
- Hornsey, M. J.; and Lewandowsky, S. 2022. A toolkit for understanding and addressing climate scepticism. *Nature Human Behaviour*, 6(11): 1454–1464.
- Jalbert, M.; Harris, M.; and Williams, L. 2025. Who is perceived to be an expert on COVID-19 vaccines on social media? Biomedical credentials confer expertise, even among vaccine-hesitant and conservative observers. *Information, Communication & Society*, 28(4): 669–687.
- Kata, A. 2012. Anti-vaccine activists, Web 2.0, and the postmodern paradigm – An overview of tactics and tropes used online by the anti-vaccination movement. *Vaccine*, 30(25): 3778–3789.
- Kennedy, B.; and Tyson, A. 2023. Americans' Trust in Scientists, Positive Views of Science Continue to Decline.

- Lee, C.; Yang, T.; Inchoco, G. D.; Jones, G. M.; and Satyanarayan, A. 2021. Viral Visualizations: How Coronavirus Skeptics Use Orthodox Data Practices to Promote Unorthodox Science Online. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21, 1–18. New York, NY, USA.
- Lewandowsky, S.; Ecker, U. K. H.; Cook, J.; van der Linden, S.; Roozenbeek, J.; Oreskes, N.; and McIntyre, L. C. 2024. Liars know they are lying: differentiating disinformation from disagreement. *Humanities and Social Sciences Communications*, 11(1): 986.
- Lin, H.; Lasser, J.; Lewandowsky, S.; Cole, R.; Gully, A.; Rand, D. G.; and Pennycook, G. 2023. High level of correspondence across different news domain quality rating sets. *PNAS Nexus*, 2(9): pgad286.
- Luceri, L.; Pantè, V.; Burghardt, K.; and Ferrara, E. 2024. Unmasking the Web of Deceit: Uncovering Coordinated Activity to Expose Information Operations on Twitter. In *Proceedings of the ACM on Web Conference 2024*, WWW '24, 2530–2541. New York, NY, USA.
- Mazza, M.; Avvenuti, M.; Cresci, S.; and Tesconi, M. 2022. Investigating the difference between trolls, social bots, and humans on Twitter. *Computer Communications*, 196: 23–36.
- Milkoreit, M.; and Smith, E. K. 2025. Rapidly diverging public trust in science in the United States. *Public Understanding of Science*, 34(5): 616–627.
- Naiman, J.; Pessianzadeh, A.; Zhao, H.; Christensen, A.; Nunn, A.; Srikanth, S.; Gami, A.; Maxwell, E.; Zhang, L.; Yeragorla, S. N.; and Rezapour, R. 2025. Beyond the Lens: Quantifying the Impact of Scientific Documentaries through Amazon Reviews. In *Proceedings of the 17th ACM Web Science Conference 2025*, WebSci '25, 413–424. New York, NY, USA.
- Nizzoli, L.; Tardelli, S.; Avvenuti, M.; Cresci, S.; and Tesconi, M. 2021. Coordinated Behavior on Social Media in 2019 UK General Election. *Proceedings of the International AAAI Conference on Web and Social Media*, 15: 443–454.
- Nogara, G.; Pierri, F.; Cresci, S.; Luceri, L.; and Giordano, S. 2024. Misinformation and Polarization around COVID-19 vaccines in France, Germany, and Italy. In *Proceedings of the 16th ACM Web Science Conference*, WEBSCI '24, 119–128. New York, NY, USA.
- Nogara, G.; Vishnuprasad, P. S.; Cardoso, F.; Ayoub, O.; Giordano, S.; and Luceri, L. 2022. The Disinformation Dozen: An Exploratory Analysis of Covid-19 Disinformation Proliferation on Twitter. In *Proceedings of the 14th ACM Web Science Conference 2022*, WebSci '22, 348–358. New York, NY, USA.
- Pacheco, D.; Flammini, A.; and Menczer, F. 2020. Unveiling Coordinated Groups Behind White Helmets Disinformation. In *Companion Proceedings of the Web Conference 2020*, WWW '20, 611–616. New York, NY, USA.
- Pacheco, D.; Hui, P.-M.; Torres-Lugo, C.; Truong, B. T.; Flammini, A.; and Menczer, F. 2021. Uncovering Coordinated Networks on Social Media: Methods and Case Studies. *Proceedings of the International AAAI Conference on Web and Social Media*, 15: 455–466.
- Paoletti, G.; Dall'Amico, L.; Kalimeri, K.; Lenti, J.; Mejova, Y.; Paolotti, D.; Starnini, M.; and Tizzani, M. 2024. Political context of the European vaccine debate on Twitter. *Scientific Reports*, 14(1): 4397.
- Papadamou, K.; Zannettou, S.; Blackburn, J.; Cristofaro, E. D.; Stringhini, G.; and Sirivianos, M. 2022. “It Is Just a Flu”: Assessing the Effect of Watch History on YouTube’s Pseudoscientific Video Recommendations. *Proceedings of the International AAAI Conference on Web and Social Media*, 16: 723–734.
- Pei, J.; Wright, D.; Augenstein, I.; and Jurgens, D. 2025. Modeling Public Perceptions of Science in Media.
- Pera, A.; and Aiello, L. M. 2024. Shifting Climates: Climate Change Communication from YouTube to TikTok. In *ACM Web Science Conference*, 376–381. Stuttgart Germany.
- Santoro, A.; Galeazzi, A.; Scantamburlo, T.; Baronchelli, A.; Quattrociochi, W.; and Zollo, F. 2023. Analyzing the changing landscape of the Covid-19 vaccine debate on Twitter. *Social Network Analysis and Mining*, 13(1): 115.
- Schmidt, A. L.; Zollo, F.; Scala, A.; Betsch, C.; and Quattrociochi, W. 2018. Polarization of the vaccination debate on Facebook. *Vaccine*, 36(25): 3606–3612.
- Smeros, P.; Castillo, C.; and Aberer, K. 2019. Scilens: Evaluating the quality of scientific news articles using social media and scientific literature indicators. In *The world wide web conference*, 1747–1758.
- Song, Y.; Zhang, Y.; Zou, S.; Yang, X.; and Huang, Q. 2025. The spread of pro- and anti-vaccine views by coordinated communities on facebook during COVID-19 pandemic. *Journal of Computational Social Science*, 8(4): 98.
- Tardelli, S.; Nizzoli, L.; Tesconi, M.; Conti, M.; Nakov, P.; Da San Martino, G.; and Cresci, S. 2024. Temporal dynamics of coordinated online behavior: Stability, archetypes, and influence. *Proceedings of the National Academy of Sciences*, 121(20): e2307038121.
- Uyheng, J.; Tyagi, A.; and Carley, K. M. 2021. Mainstream Consensus and the Expansive Fringe: Characterizing the Polarized Information Ecosystems of Online Climate Change Discourse. In *Proceedings of the 13th ACM Web Science Conference 2021*, WebSci '21, 196–204. New York, NY, USA.
- West, J. D.; and Bergstrom, C. T. 2021. Misinformation in and about science. *Proceedings of the National Academy of Sciences*, 118(15): e1912444117.
- Williams, S.; Jones, R.; Reinecke, K.; and Hsieh, G. 2022. An HCI Research Agenda for Online Science Communication. *Proc. ACM Hum.-Comput. Interact.*, 6(CSCW2): 490:1–490:22.
- Wilson, T.; and Starbird, K. 2021. Cross-platform Information Operations: Mobilizing Narratives & Building Resilience through both ‘Big’ & ‘Alt’ Tech. *Proc. ACM Hum.-Comput. Interact.*, 5(CSCW2): 345:1–345:32.
- Wright, D.; Pei, J.; Jurgens, D.; and Augenstein, I. 2022. Modeling Information Change in Science Communication with Semantically Matched Paraphrases.
- Yudhoatmojo, S.; De Cristofaro, E.; and Blackburn, J. 2023. Understanding the Use of e-Prints on Reddit and 4chan’s Politically Incorrect Board. In *Proceedings of the 15th ACM Web Science Conference 2023*, WebSci '23, 117–127. New York, NY, USA.
- Yudhoatmojo, S. B.; Cristofaro, E. D.; and Blackburn, J. 2021. “We won’t even challenge their lefty academic definition of racist.” Understanding the Use of e-Prints on Reddit and 4chan. *arXiv*.
- Zhou, K.; Wilson, T.; Starbird, K.; and Spiro, E. S. 2023. Spotlight Tweets: A Lens for Exploring Attention Dynamics within Online Sensemaking During Crisis Events. *Trans. Soc. Comput.*, 6(1-2): 2:1–2:33.
- Zollo, F.; Bessi, A.; Vicario, M. D.; Scala, A.; Caldarelli, G.; Shekhtman, L.; Havlin, S.; and Quattrociochi, W. 2017. Debunking in a world of tribes. *PLOS ONE*, 12(7): e0181821.

Paper Checklist

1. For most authors...
 - (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? **Yes**
 - (b) Do your main claims in the abstract and introduction accurately reflect the paper's contributions and scope? **Yes**
 - (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? **Yes**
 - (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? **Yes**
 - (e) Did you describe the limitations of your work? **Yes**
 - (f) Did you discuss any potential negative societal impacts of your work? **Yes**
 - (g) Did you discuss any potential misuse of your work? **Yes**
 - (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? **Yes**
 - (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? **Yes**
2. Additionally, if your study involves hypotheses testing...
 - (a) Did you clearly state the assumptions underlying all theoretical results? **N/A**
 - (b) Have you provided justifications for all theoretical results? **N/A**
 - (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? **N/A**
 - (d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? **N/A**
 - (e) Did you address potential biases or limitations in your theoretical framework? **N/A**
 - (f) Have you related your theoretical results to the existing literature in social science? **N/A**
 - (g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? **N/A**
3. Additionally, if you are including theoretical proofs...
 - (a) Did you state the full set of assumptions of all theoretical results? **N/A**
 - (b) Did you include complete proofs of all theoretical results? **N/A**
4. Additionally, if you ran machine learning experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **N/A**
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **N/A**
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **N/A**
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **N/A**
 - (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? **N/A**
 - (f) Do you discuss what is "the cost" of misclassification and fault (in)tolerance? **N/A**
5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets, **without compromising anonymity**...
 - (a) If your work uses existing assets, did you cite the creators? **Yes**
 - (b) Did you mention the license of the assets? **N/A**
 - (c) Did you include any new assets in the supplemental material or as a URL? **N/A**
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? **N/A**
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **Yes**
 - (f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR? **N/A**
 - (g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset? **N/A**
6. Additionally, if you used crowdsourcing or conducted research with human subjects, **without compromising anonymity**...
 - (a) Did you include the full text of instructions given to participants and screenshots? **N/A**
 - (b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? **N/A**
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? **N/A**
 - (d) Did you discuss how data is stored, shared, and de-identified? **N/A**

Statistic	0.5%	1%	2%
#nodes	153	306	612
%contrarian	95.4	96.4	96.7
Density	0.379	0.15	0.062
Modularity	0.5	0.617	0.614
Avg. path length	1.84	2.16	2.37
Avg. clustering coefficient	0.957	0.814	0.675

Table 6: Network statistics by coordination centrality threshold.

Appendix

Coordination Cutoff Robustness Checks

To ensure that the network structure is not an artifact of our eigenvector centrality threshold, we compare our 1% top-centrality coordinated network to thresholds of 0.5% and 2%. Specifically, we examine visual network structure (Figure 11), as well as statistics like the percentage of contrarians under each configuration, number of nodes, density, etc. (Table 6).

The patterns we observe show that the coordination detection method is robust to the centrality threshold, as the coordinated network simply expands in number of nodes instead of resulting in a different network. In every case, the networks have a diameter of 4 and show a single connected component. The percentage of contrarians remains high (above 95%) for all three thresholds. We also observe a nearly identical modularity for the thresholds of 1% and 2%, which is higher than the modularity for a threshold of 0.5%. Looking at Figure 11, this is because the higher thresholds capture more of the clusters that are adjacent to the coordinated center. Beyond this, however, the differences between the different thresholds are not structural, but rather, completeness-related (i.e., higher thresholds naturally capture more nodes and more edges).

Manual Annotation of Superspreaders

The manual assessment of superspreader accounts was conducted by one of the authors who is closely familiar with scientific discourse of COVID-19 on Twitter, and prominent scientific voices in the discourse. The profile of each account (handle, bio, number of followers etc.) was manually assessed to determine whether it fell into any of the following categories:

1. **Medical doctor** (e.g., MDs, physicians, etc.) For this classification, we look for evidence of a person possessing medical knowledge that enables them to practice medicine on patients.
2. **Scientist** (e.g., researcher, academic, etc.) We restrict this classification to people working in COVID-adjacent fields, e.g., medicine, biology, virology, epidemiology, etc. We do not classify other types of scientists as such if they work in unrelated fields (e.g., physics, computer science, social sciences, etc.), unless they work in those fields in a COVID-related capacity (e.g., physical virology, computational epidemiology, public health, etc.)

For this label, we only consider people who worked in their respective fields in an *official* capacity, e.g., a public health role, academic or industry research role, etc., unless they have otherwise demonstrable expertise (e.g., professional society memberships, recognized accreditations, published peer-reviewed research on COVID-19, etc.)

3. **Scientific organization** (e.g., science-related agencies, scientific journals, scientific or medical companies, science NGOs, academic departments or labs). Note that this does not include non-scientific government agencies or other organizations not related to science.
4. **Science communicator** (e.g., high-profile science personalities, science journalists, scientific news aggregators, etc.) This does not include people who do not explicitly or predominantly focus on science for their content or reporting (e.g., journalists who only covered science during the pandemic).
5. **Other scientific authority**. This includes any other science-related accounts that do not fall in any of the above categories.
6. **Non-scientific**. Any account that cannot be classified into any of the above.

Whenever the annotator encountered accounts that they were unfamiliar with, they verified the account’s identity through the account owner’s online presence before making an annotation decision (e.g., an existing physician practice for MDs, Google Scholar profiles for researchers, other relevant digital footprints like Wikipedia pages or authored books, etc.) We assumed negligible risk of impersonation, as impersonation profiles would have been unlikely to reach the prominence of these superspreader accounts without being detected and taken down. If a determination could not be made, or if an account could not be verified (e.g., due to the absence of real, traceable names), a label of “non-scientific” was applied.

In several cases, accounts could be classified in more than one category (e.g., medical doctors doing active research). In cases where a person held a medical doctorate (MD) or had previously worked as a physician, a medical doctor label took precedence to indicate medical knowledge even if their primary function was that of a researcher (i.e., “medical doctor” was taken in a literal sense). In other cases (e.g., where someone was both a scientist and science communicator), the label was based on the account’s primary function and/or what the account was primarily known for, as determined through bio self-identification, personal or Wikipedia pages, Google profiles, etc.

Extended Topical Distributions

Here, we show the distribution of topical shares for contrarian superspreaders, conformist superspreaders, and the overall network. The topic model detected 53 topics overall (excluding the outlier cluster). Due to some topics having negligible prevalence ($> 1\%$) and for better visualization, we plot the top 20 for each group in Figure 12.

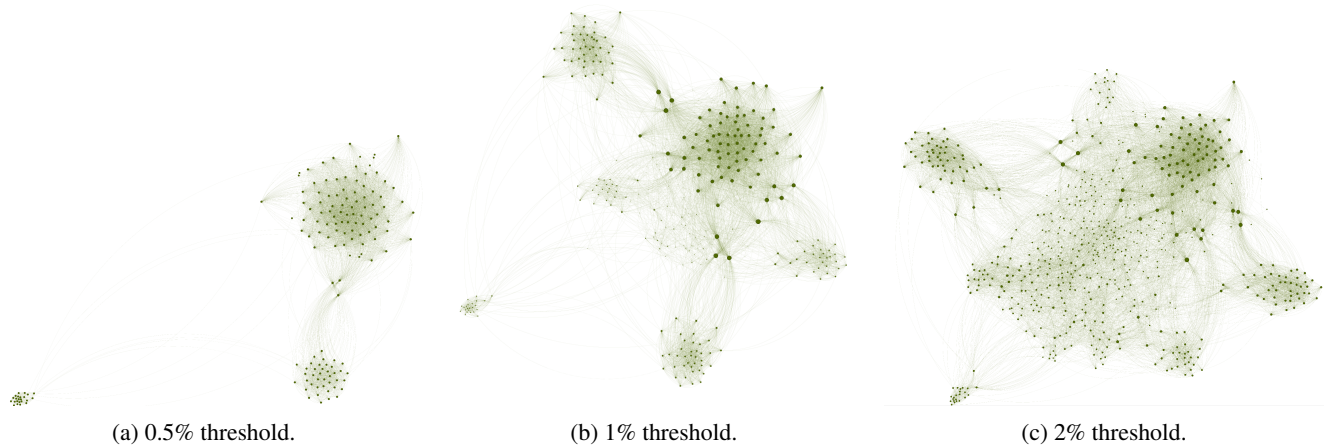


Figure 11: Visual coordinated network structures with different centrality thresholds.

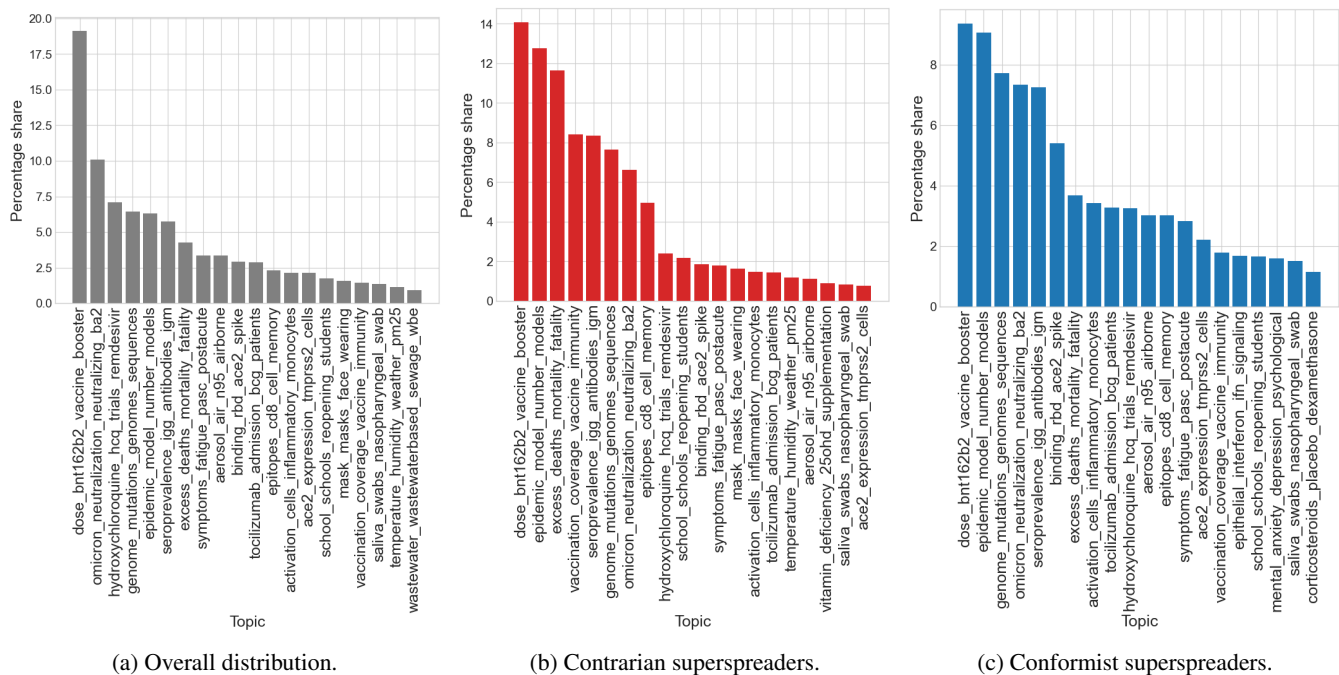


Figure 12: Distribution of the top 20 topics by group. X-axis labels are the most representative topical terms.

Syndication Analysis

We perform a syndication analysis to ensure that any media coordination is not anomalous. We take outlets that shared the exact same title within an hour of each other and form undirected edges between them; the resulting network is shown in Figure 13. Typically, neighboring outlets are those under the same brand (e.g., NBC-affiliated outlets) or local networks coordinating to share national news to local communities. The deviating black cluster represents an Australian high-quality news neighborhood; therefore, the deviation is due to locality and does not suggest suspicious activity.

Distribution of Paper Mentions Over Time

Figure 14 shows the density of mentions over time for the top 4 most-mentioned papers in the dataset, within the 50% density region. As can be seen, papers typically receive a burst of attention, followed by longer tails of more sporadic activity.

Robustness Checks for Precedence Pathways

In Figure 15, we show robustness checks for the media-superspreader precedence links by range of the backward lookup.




Figure 13: Syndicated outlet neighborhoods.

Outlet Binarization

The binarization into LC and HC outlets is done using a cutoff in the 11.5k outlet trust scores reported by Lin et al. (2023). To derive this cutoff, we use labeled datasets from Baly et al. (2018) and Baly et al. (2020), who use “low”, “mixed”, or “high” factuality labels based on Media Bias/Fact Check (MBFC) ratings of 1.42k unique domains. We collapse low- and mixed-factuality sources into the “questionable” class, and use high-factuality sources as ground-truth for the “trustworthy” class. We treat these scores (ranging from 0 to 1) as prediction probability estimates and obtain the ROC curve, from which we compute Youden’s J (i.e., the optimal point at which the true positive minus the false positive rate is maximized). We find that $J = 0.647$, which we use as our cutoff point (F1 score = 0.83).

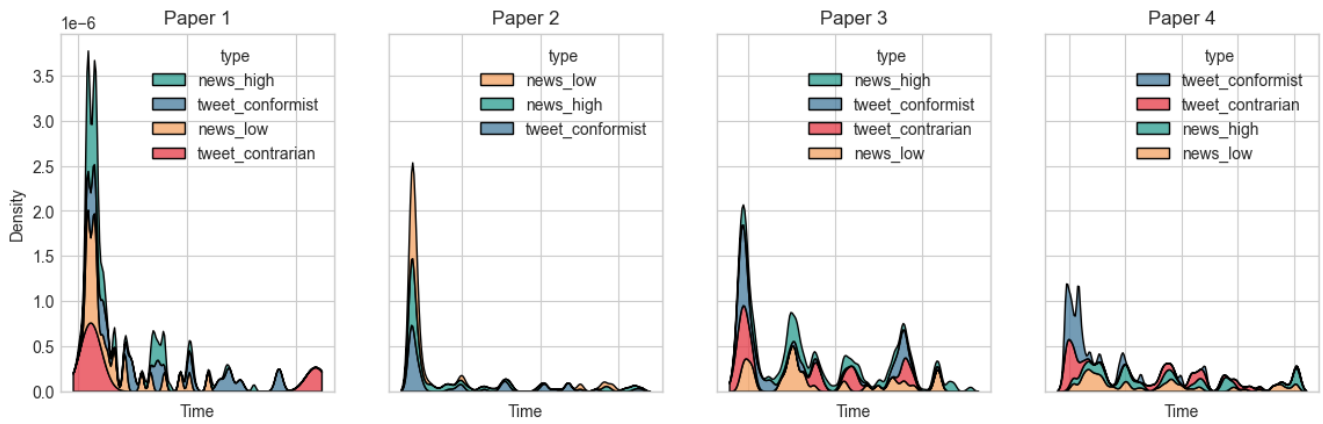


Figure 14: Over-time density distributions for most mentioned papers.

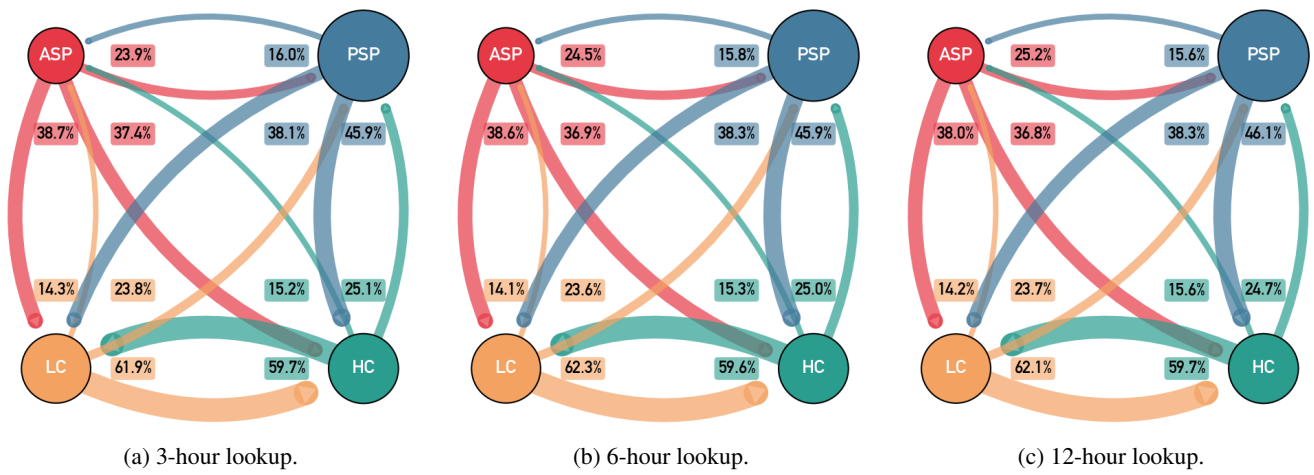


Figure 15: Backward lookup robustness checks.